Lost in Transmission:
UHV Infrastructure and the Markup Response of Chinese Industrial Firms

Songhao Guo[1], Hongcong Chu[1], Xiaoyin Yang[2], Xinmin Zhang[34]

Abstract

China's energy resources and industrial demand are separated by 1,000–4,000 kilometers, yet causal evidence on how the cost savings from bridging this gap are shared between firms and buyers is scarce. We exploit the staggered commissioning of nine ultra-high voltage (UHV) transmission lines during 2009–2015 as a quasi-natural experiment. Using approximately 2.45 million firm-year observations from China's Annual Survey of Industrial Firms, we estimate a triple-difference model that identifies the effect of lower effective electricity costs on industrial markups. High electricity-intensity firms in receiving-end regions raise markups by 1.70 percentage points (heterogeneity-robust estimate). The mechanism is incomplete cost pass-through: unit costs fall by 3.3%, yet only 34% of this decline reaches output prices. A sufficient-statistics welfare calculation shows that, although firms retain roughly two-thirds of the direct cost savings, buyers capture about 80% of the net welfare gain because the modest price decline cumulates over a large purchase base. These findings offer new micro-level evidence on infrastructure-driven cost pass-through and welfare distribution, with implications for grid modernization in developing economies.

Keywords: UHV transmission; firm markups; triple-difference; cost pass-through; welfare distribution

[1] School of Business Administration, China University of Petroleum (Beijing) at Karamay, Karamay 834000, China
[2] School of Northeast Asia Studies, Shandong University，Weihai, 264209, China
[3] Corresponding E-mail: 1457007724@qq.com
[4] School of Economics and Business Administration, Chongqing University, Chongqing 400030, China

1. Introduction

Many developing economies confront a common structural problem: energy resources lie far from industrial centers, and unreliable electricity supply forces firms to absorb heavy costs from production stoppages and backup generation. During China's 2002–2004 power shortages, large numbers of manufacturing firms were compelled to run self-generation equipment at costs far exceeding grid rates, or simply curtail output in response to rationing (Fisher-Vanden et al., 2015). Drawing on Indian data, Allcott et al. (2016) estimates that electricity shortages reduce industrial output by roughly 5%. For these firms, the total electricity-related cost of producing one more unit—the posted tariff plus rationing-induced losses plus backup generation expenses—far exceeds the tariff alone. We call this all-in concept the "effective marginal electricity cost." Cross-regional transmission investment is widely viewed as a policy lever for easing these constraints and lowering effective costs.

Yet infrastructure-driven cost savings do not automatically reach buyers. In imperfectly competitive markets, firms face downward-sloping demand curves, and the optimal response to an exogenous cost reduction is to pass only part of the savings through to output prices, retaining the rest as profit (Weyl & Fabinger, 2013). The pass-through rate—the share of a cost change reflected in the output price—governs how savings are split between firms and buyers. Assessing the social return to transmission investment therefore requires a distributional answer: who benefits, and by how much? Existing work focuses mainly on the aggregate output or productivity effects of electricity infrastructure (Allcott et al., 2016; Fried & Lagakos, 2023; Kassem, 2024; Lipscomb et al., 2013; Rud, 2012); causal evidence on this micro-level distributional question is lacking.

This paper exploits the staggered roll-out of China's ultra-high voltage (UHV) transmission network during 2009–2015 to estimate, simultaneously, the cost pass-through rate from transmission infrastructure and the welfare split between producers and buyers. China's primary energy reserves are concentrated in its western and northern interior, while industrial demand clusters along the eastern seaboard, 1,000–4,000 km away. Between 2009 and 2015, nine UHV lines were progressively commissioned (Table 1), serving 60 county-level areas across 12 provinces and enabling large-scale, low-loss power transfer. Under the regulated tariff regime in force during the sample period, industrial posted prices were virtually unchanged (Section 2.2); commissioning therefore did not alter posted tariffs but reduced firms' effective marginal electricity costs by alleviating rationing and cutting backup generation needs. The staggered timing provides a favorable quasi-natural experiment: the induced cost shocks are large enough to detect at the firm level, while the long planning-to-commissioning lag makes precise timing difficult for firms to anticipate. Figure 1 maps their geographic distribution, and Figure 2 illustrates the commissioning timeline relative to our sample period.

We merge the Annual Survey of Industrial Firms with UHV commissioning data to build a panel of roughly 2.45 million firm-year observations and estimate a triple-difference (DDD) model. The first difference contrasts receiving-end areas with non-receiving areas; the second compares pre- and post-commissioning periods; the third exploits predetermined industry-level electricity intensity. Identification does not require random substation siting—substations are indeed built near load centers—but rather that, within the same receiving-end area, high- and low-intensity firms trend alike before commissioning. Pre-treatment event-study coefficients are individually and jointly insignificant (Section 3.4), the effect appears abruptly at commissioning, and coefficients decline monotonically with distance from substations.

Our core findings are as follows. UHV commissioning raises markups of high-intensity firms in receiving-end regions by 1.70 percentage points (heterogeneity-robust estimate), roughly 2.1% of the sample mean (0.802). The mechanism is incomplete cost pass-through: unit costs fall by about 3.3%, but only 34% of this decline shows up in output prices; the remainder is retained as profit. Profit margins calculated directly from accounting data—requiring no production-function assumptions—are also significantly positive, cross-validating the result. On the market-structure dimension, the HHI rises significantly and the number of firms falls, consistent with (Melitz, 2003)'s prediction that favorable cost shocks screen out marginal firms and raise concentration. On the efficiency dimension, TFP of high-intensity firms rises significantly. The three channels are complementary, not competing. Heterogeneity analysis confirms that the effect concentrates among firms with greater market power. Export-oriented firms, facing more elastic international demand, exhibit significantly smaller markup increases, providing demand-side corroboration of incomplete pass-through.

Applying a sufficient-statistics welfare framework (Chetty, 2009; Weyl & Fabinger, 2013), we translate these reduced-form estimates into welfare magnitudes. Firms retain about two-thirds of the cost savings ($1 - \rho \approx 66\%$), while buyers receive one-third through lower prices ($\rho \approx 34\%$). The welfare distribution, however, differs from this direct split. Although the price decline is modest in magnitude, it applies to the entire purchase volume; cumulative buyer surplus therefore far exceeds the producer surplus gain. Buyers capture roughly 80% of net welfare and producers about 20%. This ratio is highly robust to the demand-elasticity calibration, being driven primarily by the pass-through rate. The calculation covers only the receiving end; sending-end effects are left to future work (Section 8.3). Appendix Table D provides the full $\rho \times \varepsilon$ joint sensitivity grid for buyer benefit share and net social welfare.

Our analysis also speaks to the cost pass-through literature. Empirical evidence in this area concentrates on cost *increases* in advanced economies (Ganapati et al., 2020; Nakamura & Zerom, 2010) ; we provide evidence on cost *decreases* under regulated tariffs in a developing economy. Within the China infrastructure literature, existing work mainly examines the output and trade effects of transportation networks (Banerjee et al., 2020; Faber, 2014) ; the distributional effects of electricity infrastructure at the firm level have not been studied.

The paper proceeds as follows. Section 2 presents the institutional background and derives testable predictions. Section 3 describes the data and identification strategy. Sections 4 and 5 report baseline results and robustness checks. Section 6 examines mechanisms. Section 7 discusses heterogeneity. Section 8 conducts the welfare analysis. Section 9 concludes.

2. Institutional Background and Theoretical Predictions

2.1 Spatial Mismatch of Electricity Supply and Demand

China's energy endowments and electricity demand are severely misaligned geographically. Roughly 76% of coal reserves lie in Shanxi, Inner Mongolia, Shaanxi, and Xinjiang; 80% of exploitable hydropower is in Sichuan, Yunnan, and Tibet. Yet Guangdong, Jiangsu, Zhejiang, and Shandong account for over 40% of national industrial electricity consumption while holding less than 5% of primary energy reserves (National Bureau of Statistics, 2009). For decades this gap was bridged mainly by rail-hauled coal—in 2008 alone, 1.5 billion tons of coal moved by rail, nearly half of all freight—a mode that was costly and frequently bottlenecked.

For coastal industrial hubs, local generation fell short and cross-regional transmission was limited. Firms faced frequent power rationing—up to one or two days per week during peak months (Fisher-Vanden et al., 2015). Rationing halted production and forced firms to self-generate at costs far above grid rates. Total electricity-related costs thus far exceeded the posted tariff; expansion of transmission infrastructure directly affects firms' production costs and pricing behavior.

2.2 UHV Transmission Network

UHV technology achieves transmission losses and per-unit-capacity costs far below those of conventional extra-high voltage lines, making bulk transfer over distances exceeding 1,000 km economically viable (National Energy Administration, 2015). The first UHV AC line—Jindongnan–Nanyang–Jingmen—entered service in January 2009. By end-2015 nine lines were operational (six DC, three AC), serving 60 county-level areas in 12 provinces (Table 1, Figures 1 and 2). Commissioning dates span 2009 to 2015; receiving-end regions differ in provincial location and industrial mix, providing the spatial and temporal variation needed for a staggered difference-in-differences design.

Table 1 UHV Transmission Lines Commissioned 2009–2015

Line	Type	Year	Sending End	Receiving End
Jindongnan–Nanyang–Jingmen	AC	2009	Shanxi Changzhi	Hubei Jingmen
Yunnan–Guangdong	DC	2010	Yunnan Chuxiong	Guangdong Guangzhou
Xiangjiaba–Shanghai	DC	2010	Sichuan Yibin	Shanghai Fengxian
Jinping–South Jiangsu	DC	2012	Sichuan Xichang	Jiangsu Suzhou
Huainan–North Zhejiang–Shanghai	AC	2013	Anhui Huainan	Shanghai
Nuozhadu–Guangdong	DC	2013	Yunnan Pu'er	Guangdong Jiangmen
Hami South–Zhengzhou	DC	2014	Xinjiang Hami	Henan Zhengzhou
Xiluodu–Zhejiang	DC	2014	Sichuan Yibin	Zhejiang Jinhua
North Zhejiang–Fuzhou	AC	2015	Zhejiang Huzhou	Fujian Fuzhou

Note: The North Zhejiang–Fuzhou line was commissioned December 26, 2014; treatment year coded as 2015. Yunnan–Guangdong DC achieved full bipole operation June 2010. After deduplication, 60 county-level areas are treated.

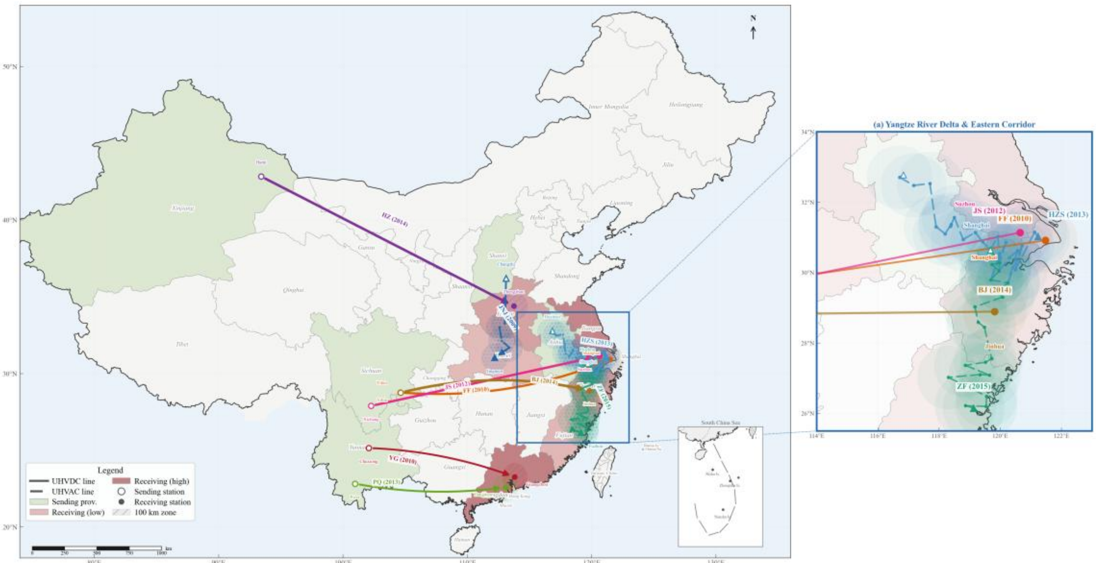

Figure 1 Geographic Distribution of UHV Transmission Lines (2009–2015)

Note: The map shows the routing of nine UHV lines and the locations of receiving-end substations. Red markers denote receiving-end substations; blue markers denote sending-end converter/substations. The base map shows provincial administrative boundaries.

Regulated tariff regime and the cost channel. Provincial price authorities set industrial posted tariffs by cost-plus rules; adjustments require NDRC approval and occur infrequently (Cao et al., 2021; Davidson & Pérez-Arriaga, 2020). Cross-provincial power transfers do not trigger tariff recalculation in the receiving province. Market-oriented reform did not formally begin until 2015 (State Council Document No. 9), and posted tariffs were not fully abolished until 2021. Throughout the 2005–2015 sample period, administrative regulation was never loosened. The benefits of UHV therefore show up mainly on the quantity margin—alleviating rationing, cutting production losses, and reducing backup generation costs (Bao et al., 2024; Fisher-Vanden et al., 2015) . We identify the overall change in the effective marginal electricity cost; improved supply reliability is itself a cost decrease, even if observed posted tariffs do not move. Section 6.1 provides empirical confirmation from provincial price data.

Siting and exogeneity. Receiving-end substations are sited according to grid planning criteria—load-center location and interconnection conditions. Our DDD exploits the differential response of high- versus low-intensity firms within the same receiving-end area. Confounders correlated with siting but not varying systematically with industry electricity intensity are absorbed by design. The key assumption is that, within receiving-end areas, high- and low-intensity firms do not display differential pre-trends. From planning approval to commissioning typically takes five to eight years, and planning preceded the sample start (2005), making precise anticipation of commissioning timing difficult. Section 3.4 evaluates the parallel-trends and no-anticipation conditions.

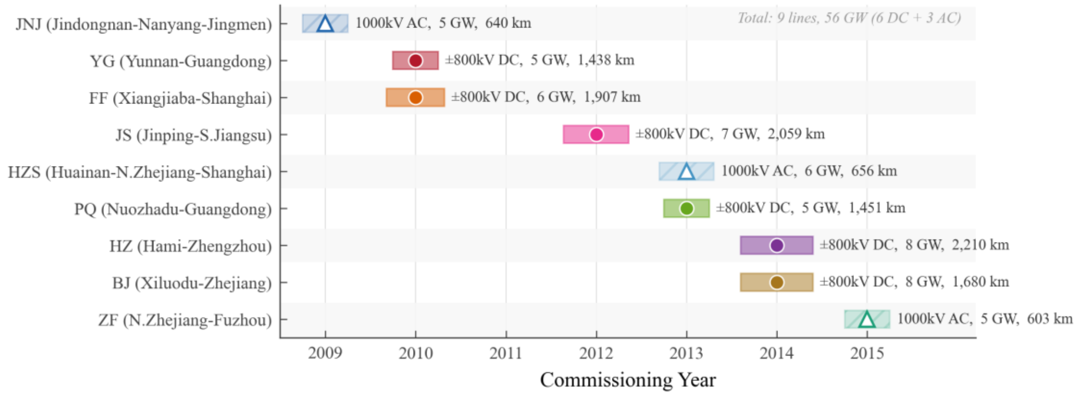

Figure 2 UHV Line Commissioning Timeline and Sample Period

Note: The horizontal axis is year (2005–2015). Vertical lines mark each line's commissioning date. The shaded area indicates the sample coverage period.

2.3 Theoretical Framework: Cost Shocks, Pass-Through, and Markups

This section sets up a parsimonious framework linking cost shocks, pricing behavior, and welfare, and derives testable predictions. The framework draws on (Weyl & Fabinger, 2013)'s pass-through analysis and (De Loecker & Warzynski, 2012)'s markup measurement approach.

Basic setup. Consider an imperfectly competitive firm facing a downward-sloping demand curve q(p) with marginal cost c. The profit-maximizing first-order condition yields the pricing rule:

$$p = \frac{\varepsilon}{\varepsilon - 1} \cdot c = \mu \cdot c \tag{1}$$

where ε is the demand elasticity (absolute value) and $\mu = \varepsilon/(\varepsilon - 1) > 1$ is the theoretical markup. In practice, markups are estimated indirectly via a production function approach (Section 3.2); the estimated level may deviate from the theoretical value, but the direction of change is consistent with the theoretical prediction. Under CES demand, ε is constant and μ does not vary with c; under more general demand systems, ε may vary with the price level, causing μ to depend on c.

Cost shock and pass-through rate. Suppose UHV commissioning reduces marginal cost from c to $c' = c - \Delta c$ ($\Delta c > 0$). The firm adjusts price to the new optimum $p' = \mu' \cdot c'$. Define the cost pass-through rate as:

$$\rho = \frac{\Delta p}{\Delta c} = \frac{p - p'}{c - c'} \tag{2}$$

Markup decomposition. Using the definition μ = p/c:

$$\frac{\Delta \mu}{\mu} = \frac{\Delta p}{p} - \frac{\Delta c}{c} = \frac{\Delta c}{c}\left(\frac{\rho}{\mu} - 1\right) \tag{3}$$

Since μ > 1 and typically ρ < 1, the term in parentheses is negative; with Δc > 0 (cost decline), Δμ > 0: incomplete pass-through of a cost decline raises markups.

Testable predictions. The framework generates four main predictions:

Prediction 1 (Markup increase). If pass-through is incomplete (ρ < 1), a cost decline raises markups. This is the primary testable implication, corresponding to the DDD core coefficient.

Prediction 2 (Cost decline exceeds price decline). The unit cost decline should exceed the output price decline, with the difference retained by firms as profit.

Prediction 3 (Market consolidation). Cost shocks screen out marginal firms and raise market concentration(Melitz, 2003).

Prediction 4 (Heterogeneous response). Firms with greater market power have stronger retention capacity (ρ lower), yielding larger markup increases. Export firms face more elastic international demand, have smaller pricing margins, and show smaller markup increases.

3. Data and Identification Strategy

The staggered timing of commissioning, technically determined siting, spatial variation between receiving and non-receiving areas, and predetermined differences in industry electricity intensity provide favorable conditions for causal identification.

3.1 Data Sources

Firm-level data come from China's Annual Survey of Industrial Firms (2005–2015), maintained by the National Bureau of Statistics, covering all above-scale industrial enterprises (annual revenue above 5 million yuan, adjusted to 20 million in 2011). The database provides financial indicators, production information, ownership type, industry codes, and geographic location. Cleaning and cross-year matching procedures follow (Brandt et al., 2017) , yielding a final sample of approximately 2.45 million firm-year observations across 457 prefecture-level city clusters.

The sample period is chosen on three grounds. First, the first UHV line was commissioned in 2009, so 2005–2008 provides ample pre-treatment data for parallel-trends tests. Second, by end-2015 nine lines were operational, providing sufficient policy variation. Third, a change in statistical reporting standards after 2015 reduces data quality, making 2015 a natural cutoff for consistency. In the event study, more distant pre-treatment periods (e.g., pre5) are contributed mainly by later cohorts; the earliest cohort (2009) has at most four pre-treatment years. The 2011 above-scale threshold increase from 5 to 20 million yuan may affect sample composition; Appendix Table C1 reports a robustness check restricting the sample to firms above 20 million yuan throughout (DDD coefficient 0.0216 vs. baseline 0.0225), and conclusions are unaffected.

Line information comes from publicly available materials of the State Grid Corporation and China Southern Power Grid, including line routing, converter/substation coordinates, and commissioning dates. We calculate the geographic distance from each firm's city to the nearest UHV substation; areas within 100 km are defined as the treatment group. Industry electricity intensity data come from the China Energy Statistical Yearbook (National Bureau of Statistics, 2009)defined as electricity consumption per unit of output value (kWh/10,000 yuan) at the two-digit industry level; industries above the median are classified as high electricity-intensity.

3.2 Variable Construction

Markup. Firm-level markups are estimated following (De Loecker & Warzynski, 2012). Under cost minimization and Cobb-Douglas production function assumptions:

$$\mu_{it} = \theta_M^j \times \frac{P_{it}Y_{it}}{P_M M_{it}} \tag{4}$$

where μ_{it} is the markup estimate of firm i in year t, $\theta_M{}^j$ is the intermediate input output elasticity of industry *j*, P_{it} Y_{it} is sales revenue, and P_M M_{it} is intermediate input expenditure. Note that μ_{it} in Equation (4) is a production-function-based estimator that is conceptually consistent with, but differs in measurement from, the ideal markup $\mu = p/mc$ in the theoretical framework of Section 2.3: measurement error in production function parameter estimation, deviations of the intermediate input variable from its theoretical concept, and winsorization may all cause estimated levels to depart from theoretical expectations—in particular, $\theta_M{}^j$ being generally below 1 implies that the revenue-to-intermediate-input ratio must be substantially above 1 for μ_{it} > 1.

Consequently, the sample mean of μ_{it} may fall below 1, a phenomenon commonly observed in the literature using the same method and database (Brandt et al., 2017; Lu & Yu, 2015). Crucially, the DDD exploits within-group cross-time variation rather than cross-sectional levels. θ_M^j varies substantially across industries: estimates for 41 two-digit industries range from 0.48 (textiles) to 0.61 (chemical raw materials), with a mean of 0.55 (SD = 0.12). Industry × year fixed effects absorb cross-industry variation in θ_M^j, so the specific level of output elasticity does not affect the core coefficients.

θ_M^j is estimated industry by industry using the (Levinsohn & Petrin, 2003) (hereafter LP) method. (Ackerberg et al., 2015) (hereafter ACF) points out that the LP first stage may fail to nonparametrically identify the variable-input elasticity due to collinearity. Two observations mitigate this concern. First, as noted above, the DDD exploits within-group time variation, and industry × year fixed effects absorb cross-industry differences in θ_M^j, so the specific level of output elasticity has limited bearing on the core coefficient; Appendix Table B2 shows that uniformly adjusting θ_M by ±5%, the DDD coefficient changes by less than 5%. Second, (Brandt et al., 2017) and (Lu & Yu, 2015) estimate markups using the LP method with the same database; following this practice preserves comparability with the existing literature. Appendix Table A1 reports the industry-level production function estimates.

Temporal stability of production function parameters. The intermediate input variable is unavailable after 2008, so we estimate production function parameters using 2004–2007 data (the period in which this variable is fully available). Mean output elasticities for labor, capital, and intermediate inputs are 0.17, 0.06, and 0.55, respectively, with mean returns to scale of 0.78, consistent with (Brandt et al., 2017)'s estimates. Applying the 2004–2007 θ_M^j to the full 2005–2015 sample requires time-invariance. We conduct three checks. First, estimating θ_M^j separately on 2004–2006 and 2005–2007 subsamples, the Pearson correlation is 0.959 and the mean absolute difference is only 0.041 (7.5% of the mean 0.55); returns to scale show no significant difference (paired t-test p = 0.259). Second, the ±5% sensitivity test in Appendix Table B2 shows that even with systematic bias in θ_M^j, the core coefficient barely moves. Third, industry × year fixed effects absorb the level impact of θ_M^j on markups; only if θ_M^j changed differentially after treatment in a way that interacts with both treatment status and electricity intensity would identification be threatened—a very specific confounding pattern that is economically implausible. This assumption is a methodological limitation of the paper.

As a complement to the production function approach, the mechanism analysis also reports DDD results using profit margins (operating profit/revenue), which are directly derived from accounting data without production function assumptions, providing independent cross-validation.

Other key variables. The treatment indicator Treated equals one when a firm's city is within 100 km of the nearest UHV facility, and zero otherwise. The 100 km cutoff is motivated by two considerations: first, in grid dispatching practice, a substation's direct load-management range is typically within 100 km; beyond this, additional step-down transformations attenuate supply improvements. Second, the spatial gradient test shows coefficients declining smoothly from 0.0251 at 50 km to 0.0225 at 100 km, but decaying sharply to 0.0132 at 125–150 km; 100 km lies within the main effect range. Appendix C reports full results at alternative cutoffs; the main findings are robust. The time indicator Post switches on from the commissioning year of the nearest UHV line; control-group cities are assigned according to the commissioning year of their nearest line to ensure a common temporal structure. HighIntensity equals one for industries above the median electricity intensity. Controls include ln(assets), leverage, export intensity, ln(firm age), and ownership type.

3.3 Descriptive Statistics and Balance Tests

Table 2 reports descriptive statistics. The mean of the markup estimator is 0.802 (SD = 0.267), below the theoretical lower bound of $\mu = p/mc > 1$; this is a known feature of the production function approach (see the discussion in Section 3.2). The core identification exploits cross-time variation in markups rather than their levels, and the substantial cross-firm dispersion (P25 = 0.624, P75 = 0.940) provides ample variation for DDD estimation. Mean firm size is ln(10.23) ≈ 27,700 thousand yuan in assets; mean leverage is 53.1%; export intensity 12.9%. Private firms account for 75.8% of the sample, foreign-invested firms 18.2%, and state-owned or collective firms 6.0%.

Table 3 reports pre-treatment (2006) balance tests. All eight standardized differences are below 0.1 in absolute value, well below the 0.25 threshold recommended by (Imbens & Wooldridge, 2009). The joint F-test is statistically significant ($p = 0.007$), but the uniformly small standardized differences suggest this reflects the large sample size rather than economically meaningful imbalance. Appendix Table C3 reports DDD results using entropy-balancing weights (matching pre-treatment means of leverage, ln(assets), export intensity, and ln(firm age)); the coefficient is virtually unchanged (0.0230 vs. baseline 0.0225).

Table 2 Descriptive Statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max
Markup	2,450,464	0.802	0.267	0.361	2.156
ln(Assets)	2,450,464	10.230	1.369	7.003	15.176
Leverage	2,450,464	0.531	0.281	0.014	1.278
Export intensity	2,450,464	0.129	0.298	0.000	1.056
ln(Firm age)	2,450,464	1.916	0.818	0.000	6.026

Note: Sample period 2005–2015. Markups estimated following De Loecker and Warzynski (2012); their levels are affected by the precision of production function parameters, and the identification strategy exploits within-group cross-time variation (see text for discussion).

Table 3 Balance Test (Pre-treatment 2006)

Variable	Treated mean	Control mean	Difference	Std. diff.
Markup	0.751	0.764	−0.012	−0.064
ln(Assets)	9.845	9.748	0.097	0.076
Labor	168.034	185.193	−17.159**	−0.071
Leverage	0.560	0.545	0.015	0.057
TFP	3.397	3.393	0.004	0.006
Labor productivity	438.158	429.019	9.139	0.016
Export intensity	0.133	0.122	0.012	0.041
Firm age	1.775	1.736	0.038	0.045

Note: Treated: firms within 100 km of nearest UHV facility (N=73,126); Control (N=141,050). *** p<0.01, ** p<0.05, * p<0.1.

3.4 Triple-Difference Identification Strategy

The baseline regression specification is:

$$Markup_{it} = \beta(Treated_c \times Post_t \times HighIntensity_j) + X_{it}'\gamma + \alpha_i + \delta_{ct} + \theta_{jt} + \varepsilon_{it} \tag{5}$$

where β captures the differential change in markups of high electricity-intensity firms in covered areas relative to uncovered areas and low electricity-intensity industries after UHV commissioning. α_i is firm fixed effects, δ_{ct} is city × year fixed effects, θ_{jt} is industry × year fixed effects, and standard errors are clustered at the city level. The economic interpretation of β follows from the theoretical framework in Section 2.3: it captures the net effect of a cost decline on markups under incomplete pass-through.

The key advantage of DDD over standard DD lies in the third difference. City × year fixed effects δct already control for all time-varying factors in receiving-end regions (industrial policy changes, regional business cycles, spillovers from public investment, etc.), provided these factors affect high- and low-intensity industries symmetrically. If some omitted region–time shocks happen to interact with electricity intensity (e.g., receiving-end regions simultaneously adopt policies favoring high-intensity industries), β could still be biased. Industry × year fixed effects θ_{jt} further absorb nationwide common trends for high-intensity industries. A confounder threatening β must therefore satisfy all of the following: (i) it coincides temporally with UHV commissioning; (ii) it occurs only in receiving-end regions, not nationwide; and (iii) it affects high-intensity industries systematically differently from low-intensity industries. We consider it very difficult in practice to construct an alternative explanation that meets all three conditions.

Consistency of the DDD relies on three key assumptions.

Assumption 1: Parallel trends. The core assumption requires that, absent treatment, the markup gap between high- and low-intensity firms in receiving-end areas would not change systematically. The event study provides a direct test: pre-treatment coefficients (pre5 to pre2) are individually and jointly insignificant ($F(4, 456) = 1.02$, $p = 0.397$), and the effect appears abruptly at commissioning (Section 4.2, Figure 3). The $\tau = -2$ coefficient is marginally significant at the 10% level (−0.0057, SE = 0.0032), but its absolute value is less than one-quarter of the post-treatment effect. (Rambachan & Roth, 2023)sensitivity analysis shows that, even allowing the pre-trend to continue at its estimated slope δ-hat = 0.00117, all post-treatment coefficients remain significantly positive. Appendix B reports a Probit model predicting substation siting from pre-treatment firm characteristics; the limited predictive power (pseudo $R^2 = 0.085$) further supports the absence of a direct link between siting and firms' markup trajectories.

Assumption 2: No anticipation. If firms accurately anticipate the commissioning date and adjust behavior in advance, post-treatment coefficients will understate the true effect. Two institutional features limit anticipation. As discussed in Section 2.2, the lag from planning approval to commissioning is long, and the actual commissioning date is subject to substantial uncertainty from construction progress, regulatory approvals, and technical commissioning. The flat pre-treatment trajectory followed by the abrupt post-treatment jump in the event study constitutes the most direct evidence.

Assumption 3: SUTVA. The stable unit treatment value assumption requires that control-group firms are unaffected by spatial spillovers from treatment. DDD coefficients decline monotonically from 0.0251 at 50 km to 0.0132 at 150 km. This spatial gradient is consistent with infrastructure effects attenuating with distance and is difficult to reconcile with a regional confounding story (under which coefficients would be roughly constant across distances). If converter-station output were pooled into the regional grid and dispatched uniformly within administrative boundaries, coefficients should not vary systematically with distance; the observed monotonic decline rules out such "complete mixing" and indicates that supply improvements concentrate near the substation due to network topology and capacity constraints. City × year fixed effects further absorb any uniform supply changes within a city, so β reflects only the differential response between high- and low-intensity firms. Section 5.4 provides a more stringent within-control-group distance gradient test.

4. Baseline Results

4.1 Main Results

Table 4 reports DDD estimates with progressively richer fixed effects. The preferred specification (column 4) yields a coefficient of 0.0225 ($SE = 0.0027$, $p < 0.01$), implying that UHV commissioning raises markups of high-intensity firms in receiving-end regions by about 2.25 percentage points relative to controls.

Comparing across specifications clarifies the sources and directions of bias. The pooled OLS with only year, industry, and city FE (column 1) gives 0.0515; adding firm FE (column 2) cuts it to 0.0264, a 49% reduction. This drop indicates that unobserved firm heterogeneity generates substantial upward bias in cross-sectional settings: economically vibrant load centers host both more receiving-end substations and higher-markup firms, and firm FE absorb this selection bias. Adding controls (column 3) barely changes the coefficient (0.0258), confirming that observable firm characteristics have no incremental explanatory power once firm FE are included. The final move from column (3) to column (4)—adding city × year and industry × year FE—compresses the coefficient from 0.0258 to 0.0225, eliminating residual bias from regional macro shocks and industry-wide trends. The coefficient stabilizes in column (4) and is insensitive to further additions.

Staggered TWFE may be biased by negative weights under effect heterogeneity (de Chaisemartin & D'Haultfœuille, 2020; Goodman-Bacon, 2021). Section 5 reports three heterogeneity-robust estimators; Stacking DID and Callaway–Sant'Anna both yield 0.0170, roughly 24% below TWFE. We adopt 0.0170 as the primary benchmark throughout. Under this benchmark, the 1.70 pp markup increase represents 2.1% of the sample mean (0.802). For a firm at the sample median revenue (~28 million yuan), the markup rises from 0.802 to 0.819, widening the price–cost wedge by about 2.1% and adding roughly 120,000 yuan to annual profit. The ~24% gap between TWFE and the robust estimate is consistent with moderate upward bias from negative weights in a staggered design, but the qualitative conclusion is unaffected: the effect is economically and statistically significant under both estimates.

Column (5) replaces the dependent variable with the profit margin (operating profit/revenue); the DDD is 0.0029 (SE = 0.0008, $p < 0.01$). This result is valuable for its methodological independence. Markups come from production-function estimation and are subject to the θ_M assumption and functional-form choice; profit margins come directly from accounting statements and require no production-function assumptions. The fact that two independent approaches yield directionally consistent and significant results indicates that the markup effect is unlikely to be an artifact of specific production-function assumptions.

Table 4 Effect of UHV on Firm Markups and Profit Margins

	(1) Pooled OLS	(2) TWFE	(3) TWFE+Controls	(4) High-dim FE	(5) Profit margin
DDD	0.0515***	0.0264***	0.0258***	0.0225***	0.0029***
	(0.0049)	(0.0078)	(0.0078)	(0.0027)	(0.0008)
Observations	2,450,464	2,450,464	2,450,464	2,450,464	2,449,088
Adj. R^2	0.434	0.603	0.605	0.646	0.530
Controls	No	No	Yes	Yes	Yes
Firm FE	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	No
City×Year FE	No	No	No	Yes	Yes
Industry×Year FE	No	No	No	Yes	Yes

Note: Dependent variable: markup in (1)–(4); profit margin in (5). DDD = Treated × Post × High Intensity. City-clustered SE in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

4.2 Event Study and Parallel Trends

The event study uses the year before commissioning ($\tau = -1$) as the reference period. Pre-treatment coefficients ($\tau = -5$ to -2) are close to zero and individually insignificant; the joint F-test cannot reject the null of no differential pre-trends ($F(4, 456) = 1.02$, $p = 0.397$). The $\tau = -2$ coefficient is marginally significant at the 10% level (−0.0057, SE = 0.0032), but two features limit its threat to causal inference. First, it is negative while the post-treatment effect is positive—opposite in sign. If a differential pre-trend were driving the post-treatment markup increase, pre-treatment coefficients should drift upward; a downward drift, if anything, biases our positive estimate toward zero. Second, its absolute value is less than half the smallest post-treatment effect (0.0103), so its economic magnitude is limited.

Rambachan & Roth (2023)'s sensitivity analysis provides a more systematic assessment. Taking the estimated linear slope of pre-treatment coefficients (δ-hat = 0.00117), and subtracting the worst-case trend extrapolation, all post-treatment coefficients remain significantly positive: the adjusted $\tau = 0$ effect is 0.0091 ($t = 3.13$), rising to 0.0232 at $\tau = 1$ ($t = 7.13$) and 0.0273 at $\tau = 2$ ($t = 6.69$). Figure 4 further shows that even when the allowed trend-slope upper bound M-bar far exceeds δ-hat, the effect remains robustly different from zero.

The dynamic pattern of post-treatment coefficients itself offers clues about the transmission mechanism. The commissioning-year coefficient is already 0.0103, consistent with the immediacy of electricity supply improvements: rationing constraints are relaxed almost as soon as the line is energized. The effect then climbs over one to two years to a peak (post1 = 0.0256, post2 = 0.0308), reflecting the time firms need to adjust—retiring backup generators, raising capacity utilization, rearranging production schedules. From post3 onward the coefficient stabilizes at 0.02–0.03, neither rising further nor dissipating, consistent with the market reaching a new equilibrium and ruling out a transient effect as an alternative explanation.

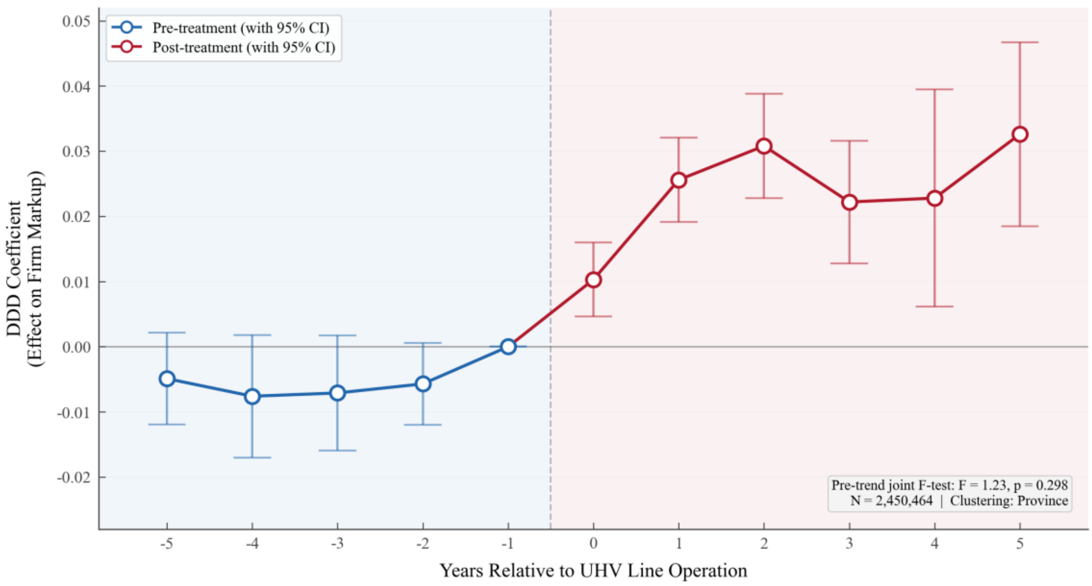

Figure 3 Event Study: Dynamic DDD Coefficients

Note: The dashed line marks zero. The horizontal axis is time relative to commissioning (τ = −5 to +5). Solid dots are point estimates; vertical bars are 95% confidence intervals. All regressions include firm, city×year, and industry×year FE plus controls. SE clustered at city level.

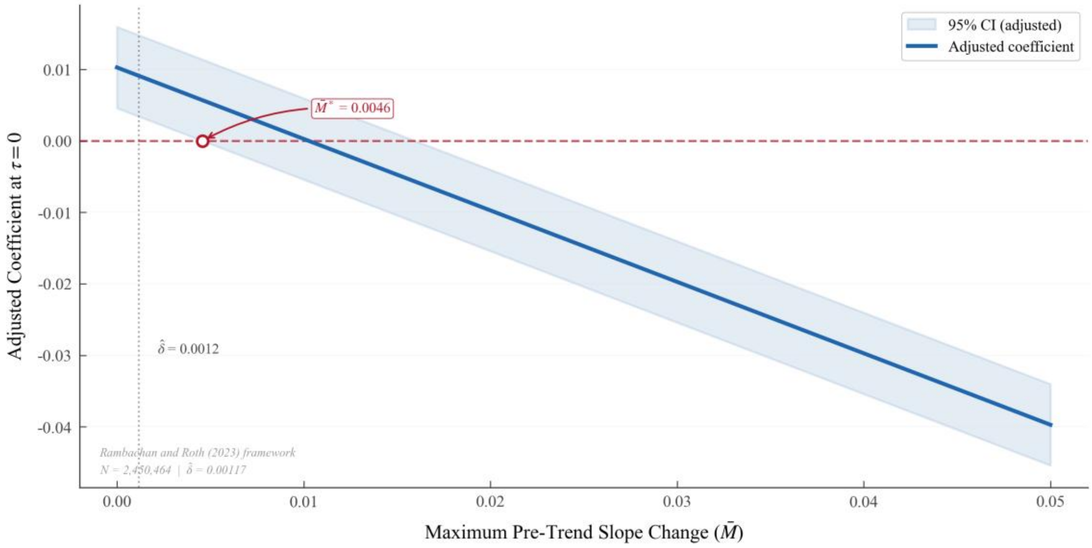

Figure 4 Rambachan–Roth Sensitivity Analysis: Robustness of the τ = 0 Effect
Note: M-bar is the upper bound on the change in the linear slope of pre-trends (Rambachan and Roth, 2023). Delta-hat = 0.00117 is the estimated pre-trend slope. The dashed line marks M-bar = delta-hat.

5. Robustness Checks

The baseline effect is significant in both statistical and economic terms, but several identification threats warrant scrutiny. We examine placebo tests, heterogeneity-robust DID estimators, alternative clustering structures, SUTVA, and subsample stability. Appendix C reports additional checks on sample composition, balanced panels, and entropy-balanced DDD; all are consistent. Appendix Table C2 confirms robustness to balanced-panel restrictions.

5.1 Placebo Tests

Estimating the DD coefficient (Treated × Post) for low-intensity industries alone yields results close to zero and insignificant. Industries that do not depend on electricity gain nothing from transmission improvements, consistent with the cost-channel hypothesis and ruling out the possibility that some broad-based economic boost in receiving-end regions drives the baseline result (Figure 5).

A randomization-inference exercise provides a more systematic check. In each of 500 iterations, we randomly permute treatment status, timing, or both, re-estimate the DDD coefficient, and build an empirical null distribution. The actual TWFE coefficient (0.0225) falls in the extreme right tail; fewer than 1% of randomized coefficients exceed this value (Figure 6).

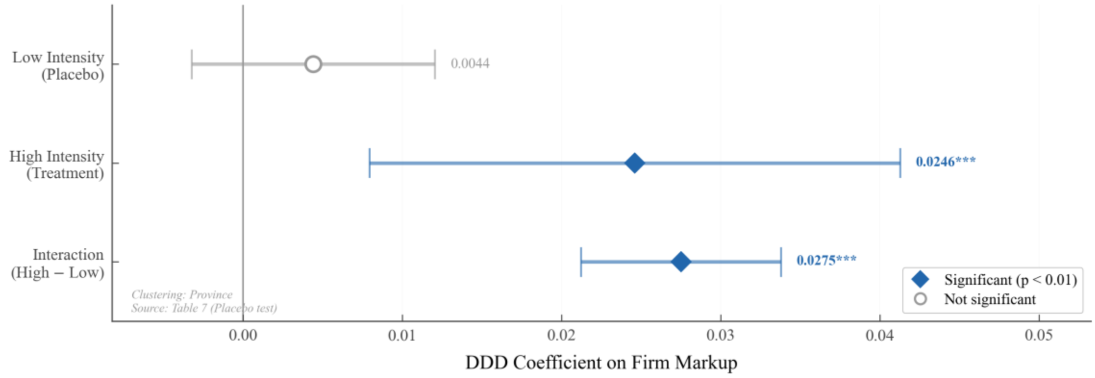

Figure 5 Placebo Test: Low vs. High Electricity-Intensity Industries

Note: Blue denotes high-intensity industries (treated); gray denotes low-intensity industries (placebo). Horizontal segments are 95% confidence intervals. The vertical dashed line marks zero.

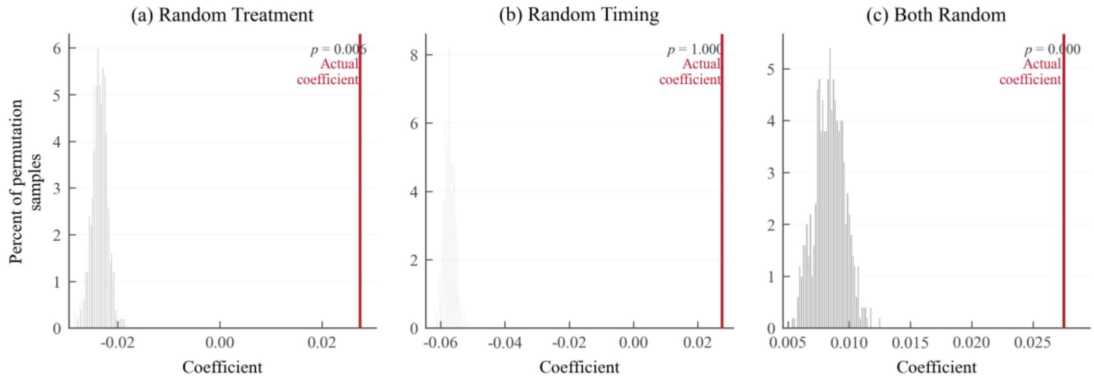

Figure 6 Placebo Test: Permutation Distribution of DDD Coefficients

Note: The distribution is generated from 500 random permutations of treatment status, timing, or both. The solid vertical line marks the actual TWFE DDD coefficient (0.0225). Empirical p < 0.01.

5.2 Heterogeneity-Robust DID Estimates

In a staggered design, TWFE may generate negative-weight bias when already-treated units serve as controls for later cohorts (de Chaisemartin & D'Haultfœuille, 2020; Goodman-Bacon, 2021). With nine lines commissioned over seven years, this is precisely the scenario in which the problem may arise. The third difference in DDD constrains the risk to some extent but cannot fully eliminate it.

Three alternative estimators are reported in Appendix Table C4. Stacking DID (Cengiz et al., 2019) and (Callaway & Sant'Anna, 2021) independently yield an identical point estimate of 0.0170, about 24% below TWFE.(Sun & Abraham, 2021)'s interaction-weighted estimator gives a larger value (0.0404), driven primarily by the three earliest lines (2009–2010, covering Hubei, Guangdong, and Shanghai—major industrial regions) receiving high cohort-size weights; firms in these areas inherently show stronger markup responses.

All four estimates are significant at the 1% level, ranging from 0.017 to 0.040. We adopt the most conservative (0.0170) as the primary benchmark; the TWFE value

(0.0225) serves as an upper-bound reference in selected mechanism analyses and welfare calculations.

5.3 Clustering Inference Robustness

The baseline standard errors are clustered at the city level (457 clusters), corresponding to the level at which Treated is assigned. However, treatment variation also relates to provincial electricity dispatch, and cities within the seven receiving-end provinces may exhibit within-group correlation. Appendix Table C5 reports results at three clustering levels.

Province-level clustering (31 clusters) yields $SE = 0.0038$, $t = 5.95$. Two-way clustering (city + industry) gives $SE = 0.0075$, $t = 3.02$ ($p = 0.003$). To address the small number of province-level clusters, (Cameron et al., 2008) wild cluster bootstrap returns $p = 0.001$. Conclusions are invariant to the inference method.

5.4 SUTVA: Distance Gradient in Controls

If supply improvements in receiving-end areas spill over through the regional grid to nearby non-covered areas, control-group firms close to substations may also benefit, causing the DDD to understate the true effect. China's provincial dispatch system institutionally limits cross-province spillovers, but empirical verification is still needed.

Within the control group, we stratify by distance from the nearest UHV facility and estimate the DD coefficient (HighIntensity × Post). Coefficients across three distance bands are: 100–200 km, 0.0446; 200–400 km, 0.0387; > 400 km, 0.0427—pairwise differences are insignificant and show no monotonic gradient. A continuous distance interaction is −0.0013 ($p = 0.078$), economically negligible. The control group shows no systematic variation with distance from the substation, consistent with SUTVA (Figure 7).

The positive ~0.04 DD within the control group merits brief comment. It reflects a nationwide common change in high-intensity industries after 2009; the macro recovery and upstream price rebound asymmetrically boosting energy-intensive sectors is a plausible explanation. The DDD differences this component out via the Treated dimension; the positive control-group DD does not threaten identification but underscores the necessity of the triple difference.

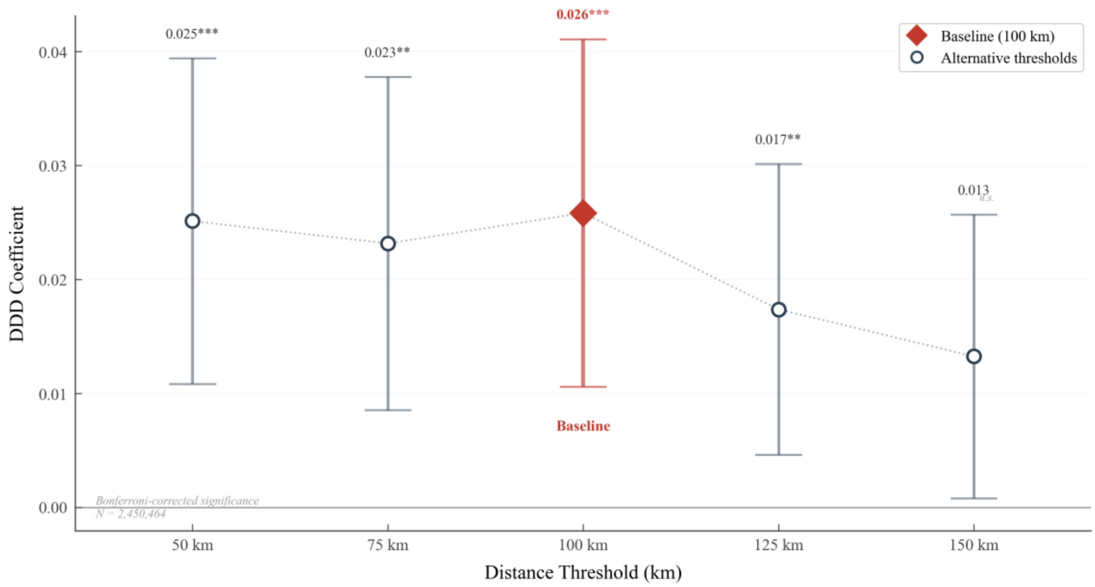

Figure 7 Distance Sensitivity: DDD Coefficients by Distance Cutoff
Note: The horizontal axis is the distance cutoff from the nearest UHV substation (km); the vertical axis is the DDD coefficient estimate. Solid dots are point estimates; vertical bars are 95% confidence intervals.

5.5 Subsample Test: 2004–2010

The full sample spans the 2011 threshold change from 5 to 20 million yuan. Restricting the sample to 2004–2010 yields a positive but insignificant DDD coefficient.

The insignificance most likely reflects insufficient power rather than absence of effect. This subsample covers at most one to two post-treatment years for the two earliest lines, and the event study shows that effects take one to two years to peak. Moreover, 2009–2010 coincides with the early post-crisis recovery, when macro noise further erodes the signal-to-noise ratio. The full sample exploits nine lines commissioned over seven years, offering far richer identifying variation. Appendix Table C1 separately reports results restricting the full sample to firms above 20 million yuan (DDD 0.0216 vs. baseline 0.0225), directly ruling out the threshold change as a driver.

Taken together, the five robustness checks confirm that the baseline finding is robust to industry time trends, negative-weight bias, clustering level, spatial spillovers, and sample-period definition.

6. Transmission Mechanisms

6.1 Provincial Electricity Prices

Using a province–year panel from the China Electricity Statistical Yearbook (2003–2015), the DID coefficient for receiving-end provinces is −0.0023 (SE = 0.0133, p = 0.862)—correct in sign but far from significant (Appendix Table C6).

This null result is unsurprising in the institutional context. Throughout the sample period, industrial posted tariffs were set by the NDRC via cost-plus rules, with adjustments requiring central approval and typically implemented uniformly nationwide (Davidson & Pérez-Arriaga, 2020; Lin & Jiang, 2011); cross-provincial power transfers did not trigger tariff recalculation in the receiving province. documents the same pattern during China's 2002–2004 power shortage: under price-cap regulation, supply shocks transmit through the quantity channel (rationing frequency and duration) rather than through prices. The logic of UHV benefits is analogous: they should manifest mainly as reduced rationing and higher capacity utilization. A simple back-of-the-envelope calculation further rules out a pure price channel: assuming electricity accounts for roughly 10% of total costs, attributing the entire markup effect (0.0170) to a tariff decline would require a ~20% price drop—implausible under the regulated regime.

The DDD identification does not require observable movements in provincial posted tariffs. The third difference captures the differential benefit that high- versus low-intensity firms derive from improved electricity supply within the same area, whether the improvement manifests as greater supply reliability, expanded quantity, or negotiated discounts—all subsumed by the change in the "effective marginal electricity cost."

6.2 Cost Channel and Pass-Through Rate

Provincial prices did not move, but firm-level production costs may have fallen through the quantity channel. Under imperfect competition, an exogenous decline in marginal cost prompts re-optimization: part of the savings is passed through to output prices, part retained as profit. This logic generates three testable predictions: unit costs decline, output prices decline by less than costs, and profit margins rise. Table 5 columns (1)–(3) test each in turn.

Column (1) uses ln(real unit cost) as the dependent variable; the DDD is −0.0328 (SE = 0.0028)—unit costs fall by about 3.3%. Column (2) uses ln(output price); the coefficient is −0.0113 (SE = 0.0025)—prices fall by about 1.1%. Column (3) uses profit margin; the coefficient is 0.0029 (SE = 0.0008)—profit margins rise significantly. All three predictions are confirmed.

Together these three columns paint a picture of incomplete pass-through. Costs fall 3.3%, prices fall only 1.1%, and the gap is retained by firms as profit. The cost pass-through rate follows directly:

$$\rho = \frac{|\beta_p|}{|\beta_c|} = \frac{0.0113}{0.0328} \approx 34.4\% \tag{6}$$

Firms pass roughly one-third of the cost savings through to buyers via lower prices and retain about two-thirds. This rate provides the key empirical input for the welfare calculation in Section 8. Under CES demand, full pass-through ($\rho = 1$) would leave markups unchanged; pure cost retention ($\rho = 0$) would leave buyers with no benefit. The intermediate value of 34.4% is consistent with general expectations for imperfectly competitive markets and falls within a comparable range to Ganapati et al. (2020)'s estimate of ~65% energy-cost pass-through in U.S. manufacturing (though that study uses energy cost rather than total cost as the denominator, making the two figures not directly comparable).

6.3 Market Structure Channel

The impact of cost shocks need not be confined to incumbents' profits; market structure may also shift. Following (Melitz, 2003), a favorable cost shock benefits all firms but disproportionately helps efficient ones; marginal firms, whose relative disadvantage widens, exit, and market share concentrates among survivors. If this channel operates, we should observe rising concentration alongside a declining number of firms. Table 5 columns (4)–(5) test this prediction at the city × industry × year level.

The HHI rises by 0.100 (column 4) and ln(firm count) falls by 0.537 (column 5), both highly significant. Exit of marginal firms releases market share to survivors, strengthening their bargaining power and further supporting markup increases.

6.4 Efficiency Channel

Improved electricity supply may also release resources previously constrained by rationing, enabling firms to optimize factor allocation. In the (De Loecker & Warzynski, 2012) framework, when output prices are held constant, a TFP increase directly widens the wedge between revenue and variable cost—i.e., raises the markup.

TFP of high-intensity firms rises by roughly 4.35% (column 6). Including TFP as a control reduces the markup DDD from 0.0225 to 0.0081 (a 64% decline), suggesting that efficiency gains are an important component of the markup increase. However, TFP is itself affected by treatment and therefore constitutes a "bad control" (Angrist & Pischke, 2009); the reduction is directional evidence only.

Table 5 Mechanisms: Cost, Market Structure, and Efficiency

	(1)	(2)	(3)	(4)	(5)	(6)
		Firm-level		Market-level		Firm-level
	ln(Unit cost)	ln(Price)	P-margin	HHI	ln(Firms)	TFP
DDD	−0.0328***	−0.0113***	0.0029***	0.0997***	−0.5373***	0.0435***
	(0.0028)	(0.0025)	(0.0008)	(0.0110)	(0.0457)	(0.0046)
Obs.	2,450,464	2,450,464	2,449,088	101,251	101,251	2,450,464
Adj. R^2	0.442	0.646	0.530	0.759	0.896	0.933

Note: Columns (1)–(3): firm-level. Columns (4)–(5): city×industry×year level. All regressions include corresponding FE and controls. City-clustered SE. *** p<0.01, ** p<0.05, * p<0.1.

7. Heterogeneity Analysis

If the effect operates through the cost–market-power channel, firms' ability to retain savings should vary systematically across observable characteristics. We interact the DDD term with five baseline (demeaned) firm attributes. Among the five dimensions, export status and firm size provide the most discriminating evidence; the remaining three offer complementary support.

Export-oriented firms face more elastic international demand, which constrains their pricing space (column 2). The interaction is significantly negative (−0.0219, $p < 0.01$): exporters' markup increase is about 2.2 pp smaller than that of domestic-sales firms. Facing a more elastic demand curve, exporters pass a larger share of cost savings through to buyers, leaving less room to widen markups. This is consistent with Prediction 4 in Section 2.3 and coheres with the 34.4% average pass-through rate from Section 6.2: the average is a weighted mean, with exporters above and domestic-sales firms below.

The firm-size interaction is significantly positive (0.0218, $p < 0.01$, column 3): larger firms retain more from cost declines. The main DDD in this column is 0.0074, lower than in other columns, because the size variable (demeaned ln(assets)) absorbs much of the variation through the interaction; each one-log-unit increase in size adds roughly 2.2 pp to the effect.

The TFP interaction is positive (0.0135, $p < 0.01$, column 4). In the (Melitz, 2003) framework, favorable cost shocks cause high-productivity firms to expand while low-productivity firms are marginalized; high-productivity firms should therefore benefit more. From the perspective of (Hsieh & Klenow, 2009), UHV systematically favoring more efficient firms may improve within-industry resource allocation, though a rigorous assessment is beyond the scope of this paper. This result is also consistent with the market consolidation evidence in Section 6.3: the rise in HHI and fall in firm count reflect the expansion of efficient firms and exit of inefficient ones.

The market-concentration interaction is positive (0.0174) but insignificant ($p > 0.10$, column 1), possibly because the third difference in DDD already absorbs most cross-sectional variation in concentration along the industry dimension, leaving limited identifying variation for the interaction. The firm-age interaction is positive (0.0015, $p < 0.01$, column 5) but economically modest, consistent with longer-lived firms maintaining price premiums through customer relationships and brand recognition.

Taken together, the five dimensions point to two complementary threads of logic (Figure 8). On the supply side, firms with greater market power (larger, more productive, longer-lived) retain more and see larger markup increases. On the demand side, firms facing more elastic demand (exporters) are compelled to pass more through and see smaller markup increases. The supply-side heterogeneity explains why the same cost shock generates unequal profit effects across firms; the demand-side heterogeneity provides the micro-foundation for the pass-through rate. These distributional patterns set the stage for the welfare calculation in Section 8.

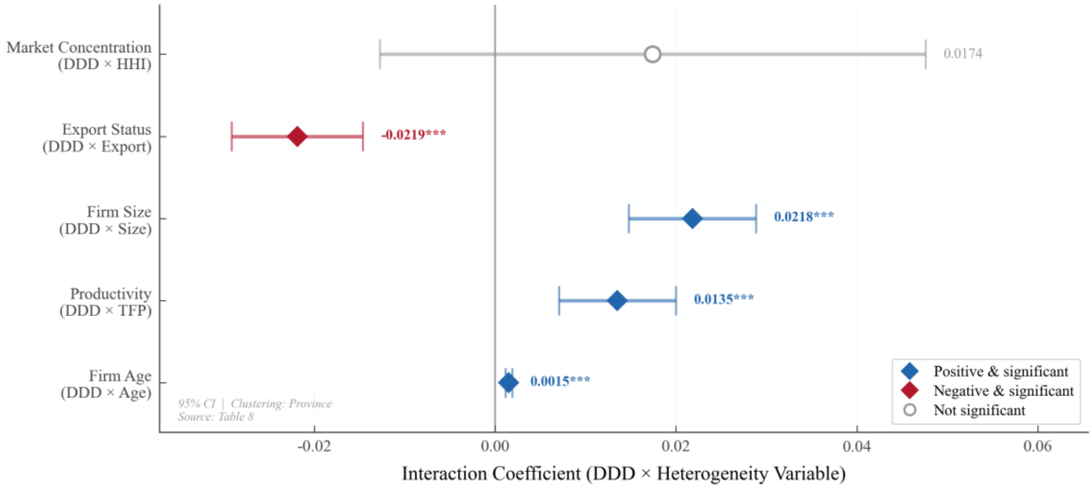

Figure 8 Heterogeneity Analysis: DDD Coefficients and Interaction Terms (Forest Plot)
Note: Dots are point estimates; horizontal lines are 95% confidence intervals. Blue indicates significantly positive; red indicates significantly negative; gray indicates insignificant. The vertical dashed line marks zero.

8. Welfare Analysis

8.1 From Pass-Through to Welfare

Section 6.2 estimated a cost pass-through rate of $\rho = 34.4\%$; Section 7 showed how the pass-through rate varies across firm types from the demand and supply sides. This section applies a sufficient-statistics welfare framework (Chetty, 2009; Weyl & Fabinger, 2013) to translate these reduced-form estimates into welfare magnitudes and assess the distribution of cost savings between producers and buyers.

Under CES demand, the change in buyer surplus is (Weyl & Fabinger, 2013):

$$\Delta CS = \frac{(1+g)^{1-\varepsilon}-1}{\varepsilon-1} \times Revenue \tag{7}$$

where $g = \beta_p = -0.0113$ is the output-price change and ε is the demand elasticity. Setting ε = 3 as the baseline (Broda & Weinstein, 2006; Feenstra, 1994), Table 6 reports the welfare estimates.

Table 6 Welfare Analysis

	Estimate	Std. Error
Panel A: Cost pass-through		
Cost decline (βc)	−3.28%	(0.28%)
Price decline (β_p)	−1.13%	(0.25%)
Pass-through rate (ρ)	34.4%	(8.2%)
Panel B: Welfare effects (% of revenue, ε = 3)		
Buyer surplus (CS)	1.14%	(0.26%)
Producer surplus (PS)	0.29%	(0.08%)
Net social welfare	1.43%	(0.27%)
Panel C: Welfare distribution		
Buyer share	79.8%	
Producer share	20.2%	

Note: ε = 3.0 (Feenstra, 1994; Broda and Weinstein, 2006). Pass-through SE via delta method. Bootstrap SE (B=999): 0.27%, 95% CI [0.89%, 1.92%].

8.2 Sensitivity Analysis

The buyer share is driven almost entirely by the pass-through rate. As ρ rises from 20% to 50%, the buyer share rises roughly in proportion. The demand elasticity ε has negligible influence: varying ε from 1.5 to 6.0, the buyer share shifts by less than one percentage point within any row (Figure 9). This robustness strengthens the distributional conclusion.

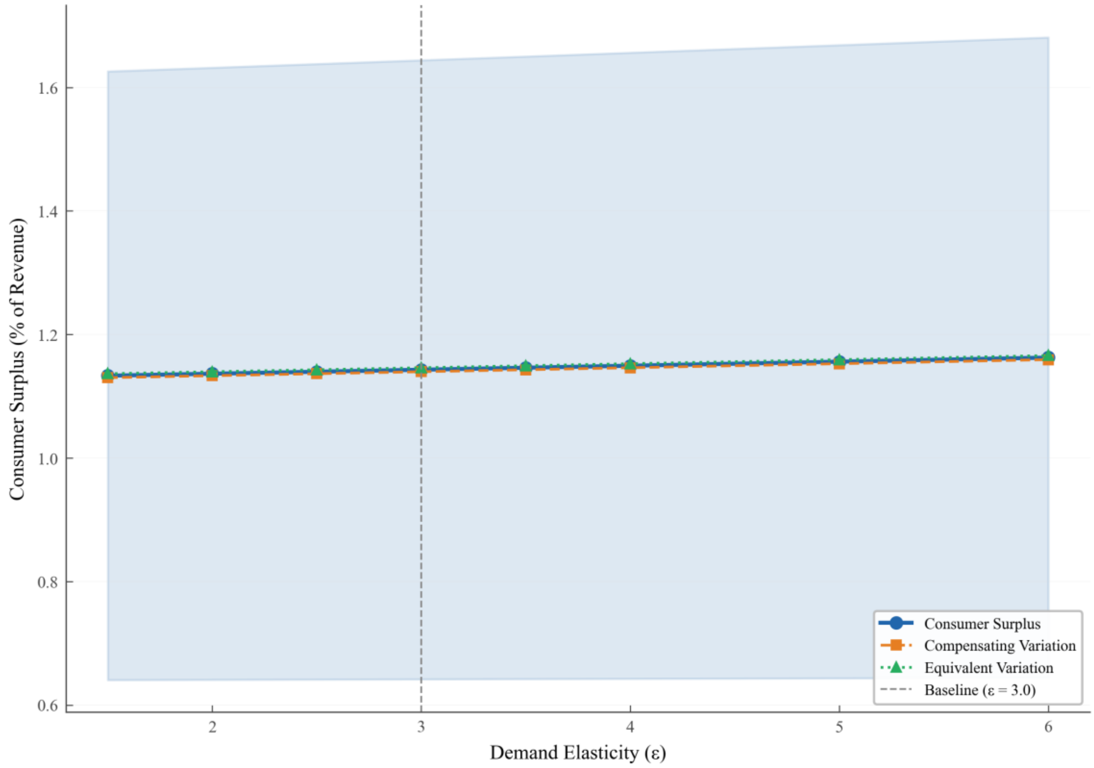

Figure 9 Sensitivity of Welfare Distribution to Pass-Through Rate and Demand Elasticity
Note: The horizontal axis is the pass-through rate ρ; the vertical axis is the buyer welfare share. Different curves correspond to different demand elasticities ε. The baseline point ρ = 0.344, ε = 3.0 is highlighted.

8.3 Limitations

The calculation covers the receiving end only. If sending-end firms face higher electricity costs because power is diverted to coastal regions, the receiving-end net welfare gain would overstate the national-level effect. The CES assumption affects the absolute level of buyer surplus but, as the sensitivity analysis shows, barely affects the distribution ratio. The framework is partial-equilibrium; "buyers capture 80% of net welfare" refers to direct buyers and does not account for general-equilibrium transmission through supply chains.

9. Conclusion

This paper exploits the staggered commissioning of China's UHV transmission lines during 2009–2015 to estimate the causal effect of large-scale transmission infrastructure on industrial markups within a triple-difference framework, and to quantify cost pass-through rates and the welfare split using a sufficient-statistics approach.

Three numbers summarize the core findings. Markups rise by 1.70 percentage points (heterogeneity-robust estimate); the pass-through rate is about 34%; and buyers capture roughly 80% of net welfare gains. Together they tell a coherent story: transmission infrastructure lowers firms' effective electricity costs; firms pass about one-third of the savings to buyers as lower prices while keeping two-thirds as profit; yet the price decline, cumulating over a large buyer base, generates welfare gains for buyers that far exceed producer gains.

Two layers of policy implications follow. First, on evaluation methodology: the "markup → pass-through rate → welfare distribution" framework offers a replicable toolkit for assessing transmission investments in other economies facing similar spatial mismatches (e.g., India, Brazil, Indonesia). Second, on distributional equity: larger and more efficient firms benefit disproportionately from cost declines (Hsieh & Klenow, 2009). Policymakers designing transmission programs should therefore recognize that benefits are not uniformly distributed; the redistributive effects of infrastructure investment deserve as much attention as its aggregate effects.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Allcott, H., Collard-Wexler, A., & O'Connell, S. D. (2016). How Do Electricity Shortages Affect Industry? Evidence from India. *American Economic Review*, *106*(3), 587–624.

Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Banerjee, A., Duflo, E., & Qian, N. (2020). On the road: Access to transportation infrastructure and economic growth in China. *Journal of Development Economics*, *145*, 102442.

Bao, B., Fu, D., Yu, J., & Zhang, Y. (2024). Lights dim, exports down: Examining the trade effects of power shortages on Chinese manufacturing firms. *China Economic Review*, *88*, 102270.

Brandt, L., Van Biesebroeck, J., Wang, L., & Zhang, Y. (2017). WTO Accession and Performance of Chinese Manufacturing Firms. *American Economic Review*, *107*(9), 2784–2820.

Broda, C., & Weinstein, D. E. (2006). Globalization and the Gains From Variety*. *The Quarterly Journal of Economics*, *121*(2), 541–585.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Themed Issue: Treatment Effect 1*, *225*(2), 200–230.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *The Review of Economics and Statistics*, *90*(3), 414–427.

Cao, J., Ho, M. S., Ma, R., & Teng, F. (2021). When carbon emission trading meets a regulated industry: Evidence from the electricity sector of China. *Journal of Public Economics*, *200*, 104470.

Cengiz, D., Dube, A., Lindner, A., & Zipperer, B. (2019). The Effect of Minimum Wages on Low-Wage Jobs*. *The Quarterly Journal of Economics*, *134*(3), 1405–1454.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. In *Annual Review of Economics* (Vol. 1, Issue Volume 1, 2009, pp. 451–488). Annual Reviews.

Davidson, M. R., & Pérez-Arriaga, I. (2020). Avoiding Pitfalls in China's Electricity Sector Reforms. *The Energy Journal*, *41*(3), 119–142.

de Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, *110*(9), 2964–2996.

De Loecker, J., & Warzynski, F. (2012). Markups and Firm-Level Export Status. *American Economic Review*, *102*(6), 2437–2471.

Faber, B. (2014). Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System. *The Review of Economic Studies*, *81*(3), 1046–1070.

Feenstra, R. C. (1994). New Product Varieties and the Measurement of International Prices. *The American Economic Review*, *84*(1), 157–177. JSTOR.

Fisher-Vanden, K., Mansur, E. T., & Wang, Q. (Juliana). (2015). Electricity shortages and firm productivity: Evidence from China's industrial firms. *Journal of Development Economics*, *114*, 172–188.

Fried, S., & Lagakos, D. (2023). Electricity and Firm Productivity: A General-Equilibrium Approach. *American Economic Journal: Macroeconomics*, *15*(4), 67–103.

Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes. *American Economic Journal: Applied Economics*, *12*(2), 303–342.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Themed Issue: Treatment Effect 1*, *225*(2), 254–277.

Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics*, *124*(4), 1403–1448.

Imbens, G. W., & Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature*, *47*(1), 5–86. https://doi.org/10.1257/jel.47.1.5

Kassem, D. (2024). Does electrification cause industrial development? Grid expansion and firm turnover in Indonesia. *Journal of Development Economics*, *167*, 103234. https://doi.org/10.1016/j.jdeveco.2023.103234

Levinsohn, J., & Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, *70*(2), 317–341.

Lin, B., & Jiang, Z. (2011). Estimates of energy subsidies in China and impact of energy subsidy reform. *Energy Economics*, *33*(2), 273–283.

Lipscomb, M., Mobarak, A. M., & Barham, T. (2013). Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil. *American Economic Journal: Applied Economics*, *5*(2), 200–231.

Lu, Y., & Yu, L. (2015). Trade Liberalization and Markup Dispersion: Evidence from China's WTO Accession. *American Economic Journal: Applied Economics*, *7*(4), 221–253.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, *71*(6), 1695–1725.

Nakamura, E., & Zerom, D. (2010). Accounting for Incomplete Pass-Through. *The Review of Economic Studies*, *77*(3), 1192–1230.

National Bureau of Statistics. (2009). *China Energy Statistical Yearbook 2009*. China Statistics Press.

National Energy Administration. (2015). *China Electric Power Development Report*. China Electric Power Press.

Rambachan, A., & Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, *90*(5), 2555–2591.

Rud, J. P. (2012). Electricity provision and industrial development: Evidence from India. *Journal of Development Economics*, *97*(2), 352–367.

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Themed Issue: Treatment Effect 1*, *225*(2), 175–199.

Weyl, E. G., & Fabinger, M. (2013). Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy*, *121*(3), 528–583.

Appendix A: Data Processing and Markup Estimation

Data Sources and Cleaning

Firm-level data come from China's Annual Survey of Industrial Firms (2004–2015). Cleaning steps include (1) dropping observations with negative or zero total assets, sales revenue, or employment, (2) dropping observations with leverage ratios below 0 or above 1, (3) dropping firms with fewer than 8 employees, and (4) winsorizing continuous variables at the 1st and 99th percentiles.

Cross-Year Firm Matching

Because firm identification codes may change across years, we adopt the following matching procedure. First, the legal person code serves as the primary key; second, for firms with missing legal person codes, we use fuzzy matching on "firm name + address." Over 95% of firm–year observations are successfully matched.

Industry Code Harmonization

The 2011 revision of the National Economic Industry Classification (GB/T 4754) changed some industry codes. We map all post-2011 codes to the pre-2011 system. Key mappings include C30 (non-metallic mineral products) → old C31, and C38 (electrical machinery) → old C39. Cross-year firm tracking data verified the accuracy of the mapping.

Production Function Estimation

We use the Levinsohn and Petrin (2003) method and 2004–2007 data (the period with complete intermediate input variables) to estimate production function parameters. Parameter constraints are set to $0.5 < \theta M < 1.2$, $0.01 < \theta L < 0.25$, $0.005 < \theta K < 0.20$.

Table A1 Production Function Estimates by Industry

Code	Industry	θM	θL	θK	RTS
B06	Coal Mining	0.697	0.079	0.037	0.813
B07	Oil and Gas Extraction	0.211	0.092	0.188	0.491
B08	Ferrous Metal Mining	0.718	0.119	0.054	0.891
B09	Non-ferrous Metal Mining	0.902	0.103	0.029	1.035
B10	Non-metallic Mining	0.585	0.132	0.053	0.770
B11	Other Mining	0.697	0.103	0.053	0.813

(Continued)

C13	Food Processing	0.720	0.102	0.036	0.858
C14	Food Manufacturing	0.603	0.145	0.061	0.809
C15	Beverage Manufacturing	0.770	0.089	0.020	0.878
C16	Tobacco	0.070	0.282	0.143	0.495
C17	Textile	0.483	0.206	0.066	0.754
C18	Textile Garments	0.325	0.269	0.083	0.676
C19	Leather Products	0.436	0.235	0.064	0.735
C20	Wood Processing	0.549	0.168	0.065	0.781
C21	Furniture	0.426	0.245	0.064	0.736
C22	Paper Products	0.403	0.248	0.078	0.730
C23	Printing	0.313	0.233	0.071	0.618
C24	Stationery and Sports Goods	0.407	0.227	0.056	0.690
C25	Petroleum Refining and Coking	0.668	0.153	0.069	0.890
C26	Chemical Materials	0.610	0.149	0.056	0.815
C27	Pharmaceuticals	0.612	0.163	0.038	0.813
C28	Chemical Fibers	0.672	0.154	0.070	0.896
C29	Rubber Products	0.523	0.154	0.067	0.744
C31	Non-metallic Mineral Products	0.587	0.155	0.050	0.792
C32	Ferrous Metal Smelting	0.735	0.156	0.041	0.933
C33	Non-ferrous Metal Smelting	0.769	0.114	0.038	0.921
C34	General Equipment	0.539	0.165	0.060	0.764
C35	Special Equipment	0.555	0.170	0.053	0.778
C36	Automobiles	0.513	0.190	0.056	0.759
C37	Other Transport Equipment	0.560	0.215	0.053	0.828

(Continued)

C39	Electrical Machinery	0.563	0.178	0.055	0.796
C40	Computers and Electronics	0.465	0.252	0.082	0.800
C41	Instruments	0.451	0.233	0.055	0.740
C42	Crafts and Other	0.528	0.177	0.050	0.755
C43	Waste Recycling	0.705	0.117	0.062	0.885
D44	Electricity and Heat	0.401	0.147	0.086	0.633
D45	Gas Supply	0.515	0.156	0.034	0.706
D46	Water Supply	0.430	0.111	0.047	0.588

Note: Note: θM, θL, and θK denote output elasticities of intermediate inputs, labor, and capital, respectively. Returns to scale = sum of the three. Estimated using the Levinsohn–Petrin method on 2004–2007 data. B11 uses the mining industry median due to insufficient sample size.

Appendix B: Identification Supplements

Probit Selection Model

Table B1 reports a Probit model predicting treatment status from pre-treatment (2006) firm characteristics. The dependent variable is Treated (=1 if within 100 km of the nearest UHV facility). Observable characteristics have limited predictive power for site selection (pseudo R^2 = 0.085), supporting the conclusion that site selection is unrelated to firm markup trends.

Table B1 Probit Selection Model (Pre-Treatment 2006 Cross-Section)

	Coefficient	Marginal Effect
ln(Assets)	0.032**	0.011
	(0.014)	
Leverage	0.085	0.029
	(0.063)	
Export intensity	0.041	0.014
	(0.048)	
ln(Firm age)	0.028*	0.010
	(0.016)	
Markup	-0.045	-0.015
	(0.052)	
TFP	0.003	0.001
	(0.019)	

(Continued)

Foreign-invested	-0.022	-0.008
	(0.031)	
SOE/Collective	0.064*	0.022
	(0.038)	
Pseudo R^2	0.085	—
Obs.	214,176	—

Note: Note: Dependent variable is Treated (=1 if within 100 km of the nearest UHV facility). Sample is the 2006 cross-section (pre-treatment); reference group is private enterprises. Standard errors clustered at the city level. *** p < 0.01, ** p < 0.05, * p < 0.1.

Output Elasticity Sensitivity

Table B2 Output Elasticity Ratio Sensitivity

	(1)	(2)	(3)
DDD	0.0250***	0.0262***	0.0274***
	(0.0074)	(0.0077)	(0.0080)
Obs.	2,450,464	2,450,464	2,450,464
Adj. R^2	0.607	0.610	0.617

Note: Note: Dependent variable is markup computed using the corresponding θM ratio. Column (1) reduces all industry θM by 5%; column (3) increases by 5%. Coefficient range is 0.0250–0.0274, deviating from the baseline by no more than 5%. All regressions include firm, city×year, and industry×year fixed effects and controls. Standard errors clustered at the city level. *** p < 0.01, ** p < 0.05, * p < 0.1.

Appendix C: Robustness Checks

Sample Robustness

Table C1 Sample Robustness

	(1)	(2)	(3)	(4)	(5)
Panel A: Markup					
DDD	0.0225***	0.0216***	0.0221***	0.0230***	0.0229***
	(0.0027)	(0.0026)	(0.0029)	(0.0028)	(0.0028)
Panel B: Profit margin					
DDD	0.0029***	0.0040***	—	—	—
	(0.0008)	(0.0009)			
Obs.	2,450,464	1,836,906	1,849,488	2,170,732	2,261,926
Adj. R^2	0.646	0.674	0.675	0.638	0.654

Note: Note: Column (2) restricts to firms with revenue ≥ 20 million yuan (the post-2011 above-scale threshold), dropping approximately 25% of observations; the DDD coefficient is virtually unchanged (0.0216 vs. 0.0225). All regressions include firm, city×year, and industry×year fixed effects and controls. Standard errors clustered at the city level. *** p < 0.01, ** p < 0.05, * p < 0.1.

Balanced Panel Tests

Table C2 Balanced Panel Robustness

	(1)	(2)	(3)	(4)
DDD	0.0264***	0.0300***	0.0324***	0.0295
	(0.0031)	(0.0033)	(0.0043)	(0.0346)
Obs.	2,450,464	1,576,513	730,865	21,570
Cluster (city)	457	444	437	230

Note: Note: "Fully balanced" requires a firm to be observed in every sample year. All regressions include firm, city×year, and industry×year fixed effects and controls. Standard errors clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Entropy-Balanced DDD

Table C3 Entropy-Balanced DDD

	(1)	(2)
DDD	0.0225***	0.0230***
	(0.0027)	(0.0025)
Obs.	2,450,464	2,450,464

Note: Note: Column (2) uses entropy-balancing weights matching treated and control groups on pre-treatment leverage, ln(assets), export intensity, and ln(firm age) means. All regressions include firm, city×year, and industry×year fixed effects and controls. Standard errors clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Distance Cutoff Sensitivity

Main results are presented in the text Table 4 and Figure 7. The coefficient declines monotonically from 0.0251 at 50 km to 0.0132 at 150 km.

Heterogeneity-Robust DID Comparison

Table C4 Heterogeneity-Robust DID Estimates

	(1)	(2)	(3)	(4)
DDD	0.0225***	0.0170***	0.0170***	0.0404***
	(0.0027)	(0.0024)	(0.0022)	(0.0053)
t / z	8.29	7.16	7.83	7.65
Obs.	2,450,464	2,450,464	2,450,464	2,450,464

Note: Note: Column (1) is the baseline TWFE (Table 5 column (4)). Column (2) follows Cengiz et al. (2019), constructing a clean control group for each cohort and stacking. Column (3) is the Callaway and Sant'Anna (2021) estimator. Column (4) is the Sun and Abraham (2021) interaction-weighted estimator. Standard errors clustered at the city level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Cluster Inference Robustness

Table C5 Cluster Robustness

Cluster Level	SE	t-statistic	p-value
City (baseline, 457 clusters)	0.0027	8.29	<0.001
Province (31 clusters)	0.0038	5.95	<0.001
City + Industry (two-way)	0.0075	3.02	0.003
Wild bootstrap (province, 999 reps)	—	—	0.001

Note: Note: The TWFE DDD coefficient is 0.0225 in all rows (N = 2,450,464). The first three rows use analytical cluster-robust standard errors. The last row reports the Cameron et al. (2008) wild cluster bootstrap p-value using Rademacher weights via Frisch–Waugh residualization testing H_0: βDDD = 0. All regressions include firm, city×year, and industry×year fixed effects and controls.

Provincial Industrial Electricity Price DID

Table C6 Provincial Electricity Price DID

	(1)	(2)	(3)	(4)
Receiving×Post	-0.0023	-0.0085	0.0162*	—
	(0.0133)	(0.0136)	(0.0093)	
Obs.	403	325	403	403
Adj. R^2	0.844	0.856	0.816	0.822
Provinces	31	25	31	31
Province FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Note: Note: Province–year panel regressions. The dependent variable in columns (1), (2), and (4) is ln(industrial catalogue electricity price); column (3) uses the price level (yuan/kWh). Treatment group: 7 receiving-end provinces (Henan, Hubei, Guangdong, Jiangsu, Shanghai, Zhejiang, Fujian). Column (2) excludes the 6 sending-end provinces. Column (4) is an event-study specification; the pre-treatment joint F-test: $F(4, 349) = 1.23$, $p = 0.298$. Robust standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix D: Welfare Analysis Supplements

Joint Sensitivity Grid: Buyer Share of Cost Savings and Net Welfare

Table D Buyer Benefit Share and Net Social Welfare: ρ × ε Joint Sensitivity Grid

ρ \ ε	ε=1.5	ε=2.0	ε=3.0	ε=4.0	ε=5.0	ε=6.0
Panel A: Buyer Direct Benefit Share (%)						
ρ=0.20	20.1	20.1	20.2	20.2	20.3	20.3
ρ=0.25	25.1	25.2	25.2	25.3	25.4	25.5
ρ=0.30	30.2	30.2	30.3	30.4	30.5	30.6
ρ=0.344	34.6	34.7	34.8	34.9	35.0	35.2
ρ=0.40	40.2	40.3	40.5	40.6	40.8	41.0
ρ=0.50	50.3	50.4	50.6	50.8	51.0	51.2
Panel B: Net Social Welfare (% of Revenue)						
ρ=0.20	3.283	3.285	3.287	3.290	3.292	3.294
ρ=0.25	3.286	3.288	3.292	3.296	3.299	3.303
ρ=0.30	3.288	3.292	3.297	3.303	3.308	3.313
ρ=0.344	3.291	3.295	3.303	3.310	3.317	3.324
ρ=0.40	3.295	3.301	3.311	3.321	3.330	3.340
ρ=0.50	3.304	3.313	3.330	3.345	3.360	3.374

Note: Note: The bolded row is the baseline (ρ = 0.344). Within any row, the buyer share varies by less than one percentage point over the full range of ε, indicating high robustness to the demand elasticity. Net welfare = ΔCS + ΔPS − DWL, where DWL ≈ 0.5 × (ε − 1) × g^2/ε (Harberger approximation). Net welfare varies by less than 3% (3.28%–3.37%) across the entire parameter space.