Electricity Supply Shocks and Manufacturing: Evidence from the Fukushima Nuclear Accident*

Daiji Kawaguchi† Takanao Tanaka‡

February 2026

Why can shocks to a small-share input—such as electricity—generate large output losses? We study the unanticipated and persistent contraction in electricity supply to Japanese manufacturing following the 2011 Fukushima accident, using confidential plant-level administrative data, a shift-share instrument for regional exposure, and a structural model. Electricity use falls by 6%, yet manufacturing revenue declines by 1.7%, despite electricity's 2% expenditure share, implying an annual loss of 4.9 trillion JPY. We find little evidence of adjustment in technology, product mix, within-firm reallocation, or fuel substitution. The model rationalizes these disproportionate losses through low input substitutability and endogenous capital idleness, which amplify the marginal value of electricity under scarcity. Finally, we quantify the welfare gains of counterfactual market reforms that lower marginal generation costs.

Keywords: Electricity, Supply Shocks, Infrastructure, Manufacturing, Nuclear
JEL Codes: H54, L60, O13, O14, Q40

*We thank Joseph Shapiro, James Sallee, Meredith Fowlie, and Guojun He, as well as seminar and conference participants at UC Berkeley, AERE, WEAI, GARESC, JADE–JTREE, SWET, IDE, Osaka University, Sophia University, and the University of Tsukuba, for insightful and constructive comments. We acknowledge financial support from the Japan Center for Economic Research and JSPS KAKENHI (19H00589; 22H00057). All errors are our own.

†Graduate School of Economics, University of Tokyo. Email: kawaguchi@e.u-tokyo.ac.jp

‡Department of Agricultural and Resource Economics, University of California, Berkeley. Email: takanao.tanaka.0323@gmail.com

1 Introduction

Can shocks to small-share inputs generate large output losses? In canonical settings, the aggregate impact of an input disruption should roughly scale with that input's expenditure share. Yet historical energy shocks often appear to violate this benchmark. From repeated global oil crises to recent electricity disruptions, major energy shocks are routinely followed by deep economic downturns (Hamilton, 1983; Kilian, 2008). This paper resolves this long-standing puzzle by combining confidential administrative microdata, quasi-experimental shock variation, and a structural model to uncover mechanisms through which energy supply shocks, and electricity supply shocks in particular, propagate into disproportionately large manufacturing output losses.

We show that electricity supply shocks generate output contractions far exceeding what electricity's expenditure share, typically below 2%, would predict. In principle, plants could mitigate disruptions through technology upgrading, substitution toward alternative fuels, product switching, outsourcing, or within-firm reallocation. In practice, we find limited adjustment along these margins. To rationalize these large losses, we develop a structural model with two mechanisms: strong complementarity between electricity and other primary inputs, and endogenous capital idleness that renders installed capital temporarily unproductive when electricity is curtailed. Given these limits to plant-level adjustment, we use the model to evaluate policy counterfactuals that operate through electricity market structure.

We study an unanticipated, nationwide, and long-lasting contraction in electricity supply to Japanese manufacturing following the 2011 Fukushima nuclear accident. Triggered by the country's largest recorded earthquake and tsunami, the disaster caused nuclear core meltdowns and substantial radioactive releases. The resulting political backlash led to the suspension of all nuclear reactors, and within two years, nuclear power was fully phased out. Because nuclear generation had accounted for approximately 30% of Japan's total electricity supply, this policy response produced a contraction in supply without modern precedent. To avoid widespread blackouts, the government curtailed demand through mandates and information campaigns. Concurrently, replacing nuclear capacity with higher marginal cost fossil fuel generation led to sharp and persistent increases in electricity prices.

We exploit quasi-experimental regional variation in the severity of electricity supply shocks using restricted plant-level data from the Census of Manufactures (CMF) for 2005–2018, combined with manually collected policy records. During our study period, Japan's electricity market was characterized by regional monopolies and limited interregional transmission, so local shortfalls could not be offset elsewhere. Exposure to the shock was

therefore highly heterogeneous: regions that relied more heavily on nuclear generation and had lower reserve margins experienced the most severe imbalances. We leverage this spatial heterogeneity and use shift-share instruments based on the pre-shock energy mix and reserve margins to identify the causal effects of electricity supply disruptions on manufacturing outcomes.

Our empirical analysis proceeds in four steps. First, we show that electricity supply shocks generate revenue losses far exceeding what electricity's expenditure share—about 2.0% of revenue—would predict. During the post–phase-out period (2013–2018), these shocks reduce manufacturing electricity use by roughly 6%. Yet manufacturing revenue falls by 1.7%, corresponding to about 4.9 trillion JPY annually—an order of magnitude larger than the small-share benchmark. Employment and materials also decline substantially; because these inputs account for most production costs, their contraction is a key channel through which electricity shortfalls translate into large revenue losses. These effects are disproportionately concentrated among small plants and exporting industries, while we find limited evidence of differential responses by baseline electricity intensity.

Second, we test for behavioral responses across several potential margins of plant-level adaptation. In principle, plants could dampen the impact of a supply shock through technology upgrading, shifts in product mix, or the reallocation of production across plants and locations. Empirically, however, we find limited adjustment along these channels. We find no systematic evidence of productivity improvements, no meaningful reductions in product-level electricity intensity, and little adjustment through outsourcing, relocation, or within-firm production shifting.

This rigidity extends to the energy input mix. We find little evidence that plants substitute toward alternative energy inputs; instead, production fuel use declines, suggesting short-run complementarity between electricity and production fuels in modern manufacturing (e.g., motive power versus thermal processes). A notable exception is in-house electricity generation: plants with pre-existing capacity expand self-generation and largely avoid output losses. However, adjustment is constrained by a thin extensive margin, with only about 10% of plants using in-house generation. Consequently, the vast majority of manufacturers remain exposed to the electricity supply contraction, helping to rationalize the large aggregate output losses.

Third, we develop and discipline a plant-level production model that treats electricity as an input. We estimate CES gross-output production functions using the approach of Gandhi et al. (2020). To assess the model's empirical plausibility, we follow Todd and Wolpin (2006) and compare model-implied responses with our reduced-form causal estimates.

Two mechanisms are central. First, the model requires limited substitutability between

electricity and other primary inputs (labor and materials) to generate the joint contractions documented in our empirical analysis. Our preferred parameterization implies a low elasticity of substitution ($\sigma = 0.5$); higher elasticities instead predict substantial substitution away from electricity and markedly smaller output losses. For example, setting $\sigma = 0.99$, which is effectively Cobb–Douglas technology, allows plants to substitute more aggressively away from electricity, implying an electricity contraction that is 57% larger, along with smaller contractions in labor and materials, yielding revenue losses 38% smaller than in our baseline. Second, we incorporate capital idleness induced by electricity curtailment. Because electricity and capital are highly complementary in the short run, supply disruptions render part of the installed capital stock temporarily unproductive. We formalize this mechanism through rigidities in predetermined capital–electricity ratios (Atkeson and Kehoe, 1999); absent this friction, predicted output losses fall by 76% relative to the baseline.

Fourth, given limited scope for plant-level adjustment, we use the model to evaluate supply-side counterfactuals that operate through electricity market structure. After the Fukushima accident, Japan met the shortfall through a combination of quantity-based conservation measures and higher marginal generation costs, as nuclear output was replaced by costly fossil fuels. Our decomposition shows that higher electricity prices, rather than rationing mandates, account for nearly all of the observed contractions in inputs and revenue (about 95%). In contrast, rationing binds for only a small fraction of plants (8%), limiting its aggregate contribution.

Motivated by the dominant role of prices, we evaluate two counterfactuals central to Japan's energy debate that lower generation costs: partial nuclear restart and regional market integration. Restarting nuclear reactors to 33% of pre-disaster capacity reduces marginal generation costs and electricity prices, mitigating producer surplus losses by 23%. Regional market integration reduces damages by 27.5% by expanding interregional trade and enabling reserve sharing, thereby reducing the need for costly backup capacity within each monopoly area. Together, these results suggest that the economic fallout reflected not only the nuclear phase-out itself, but also the costs of limited interregional coordination in a fragmented grid.

This paper makes several key contributions. First, we provide micro-level evidence that helps resolve the classic "oil puzzle," the empirical regularity that major energy disruptions generate output losses far larger than what energy's small expenditure share would predict (Hamilton, 1983; Kilian, 2008). Existing explanations emphasize monetary policy responses (Bernanke et al., 1997), sectoral reallocation frictions (Davis and Haltiwanger, 2001), and supply-side amplification through endogenous markups, production networks, and factor underutilization (Rotemberg and Woodford, 1996; Baqaee and Farhi, 2019; Atke-

son and Kehoe, 1999). Using restricted administrative plant-level data, we provide direct evidence on a complementary mechanism that is difficult to observe in aggregate data: limited plant-level adaptation along technological, organizational, and input-adjustment margins, which transmits energy shocks into broad contractions in inputs and output.

Second, this paper sheds light on the non-linear effects of energy prices on firm performance by providing some of the first micro-level evidence on the consequences of large, disruptive negative energy supply shocks. Existing studies typically find modest (Marin and Vona, 2021; Fontagné et al., 2024) or null effects (Gerster and Lamp, 2024; von Graevenitz and Rottner, 2024) in settings where energy prices move by small, transitory amounts due to taxes, network costs, or incremental shifts in the energy mix. By contrast, we study an electricity supply disruption described as a "major turning point in Japan's energy policy" (Agency for Natural Resources and Energy, 2022). In this setting, plants exhibit substantially higher responsiveness to electricity prices: across our reduced-form estimates and structural model, we find own-price elasticities of electricity demand between -0.6 and -0.7 and output elasticities between -0.15 and -0.2.[1]

Third, we contribute to the energy reliability literature by showing that the economic consequences of infrastructure failures differ systematically between developed and developing economies. Prior work focuses on blackouts and greyouts in developing-country settings (e.g., Fisher-Vanden et al., 2015; Allcott et al., 2016; Cole et al., 2018; Fried and Lagakos, 2023; Cisse, 2025), where supply interruptions often shut down production altogether. In a high-income economy, we study partial constraints in which plants remain in operation but at sharply higher marginal cost. Output losses are substantial yet smaller than in shutdown settings, highlighting that economic damage depends on whether reliability shocks bind at the intensive margin or force production to cease.

Finally, we contribute a discipline-and-validate framework that links quasi-experimental evidence on electricity supply shocks to a structural production model with electricity as a primary input, following a validation logic in the spirit of Todd and Wolpin (2006). We estimate industry-specific gross-output production functions using the approach of Gandhi et al. (2020), addressing endogeneity in input choices and allowing for heterogeneity across industries. This approach pins down key substitution parameters and the extent of capital underutilization required to match the joint responses of electricity, other inputs, and output, providing a transparent basis for counterfactual policy analysis.

[1]Consistent with this interpretation, evidence from shale gas expansions indicates that large energy price changes can generate substantial local economic effects (Feyrer et al., 2017; Allcott and Keniston, 2018); see also Arezki et al. (2017).

2 Institutional Background

During our study period, Japan's electricity sector was sharply segmented by regulation, so supply shocks were absorbed largely within regional service areas rather than smoothed through national trade. The country was divided into nine regional markets, each dominated by a vertically integrated utility supplying roughly 95% of local electricity. Wholesale trading was negligible and cross-border adjustment mechanisms were limited. As a result, common nationwide disturbances generated substantial cross-regional heterogeneity in disruption severity.

Two institutional features drove this heterogeneity. First, interregional transmission and trade were constrained by regulatory rules. Under the Electricity Business Act, each utility held a local monopoly and was legally responsible for reliability within its service area, which discouraged exports whenever local adequacy could be affected. Consequently, even in the severe post-Fukushima period, interregional transfers remained extremely small, often below 1% of peak-day demand, and some shortage regions continued to export electricity.[2]

Second, utilities operated under cost-of-service regulation, causing the supply shock to translate into asymmetric regional price adjustments. Retail electricity prices were set to recover realized costs of service plus a regulated return. Because generation portfolios differed markedly across regions, most notably in their reliance on nuclear power, the forced substitution toward fossil fuel generation after the Fukushima accident raised costs at different rates. These increases were passed through to consumers only after lengthy regulatory approval processes, often taking several months to more than a year.

Prior to the Fukushima Daiichi nuclear accident, this segmented system operated without major disruptions. In 2010, total electricity generation in Japan was approximately 930 TWh, and industrial users accounted for 71% of electricity consumption. The generation mix was diversified: 60% thermal, 31% nuclear, and 9% hydroelectric. No region experienced major supply disruptions during this period.

This stability was abruptly disrupted by the Great East Japan Earthquake and the subsequent Fukushima Daiichi nuclear accident in March 2011. The accident, which led to core meltdowns at multiple reactors, fundamentally altered public perceptions of nuclear safety in Japan and triggered an unprecedented policy response. In the months that followed, the government initiated a nationwide suspension of nuclear power operations, effectively removing a major source of baseload capacity. By May 2012, all nuclear reactors

[2]While physical constraints, notably the incompatibility between eastern Japan's 50 Hz grid and western Japan's 60 Hz grid, existed, the primary barrier to equalizing supply shocks was institutional rather than technological.

were offline, and most remained suspended for several years.[3] Panel A of Figure 1 illustrates the sharp and persistent collapse in nuclear capacity utilization that followed.

Because nuclear power accounted for nearly one-third of electricity generation prior to the accident, its removal precipitated a nationwide supply crisis. The strain was particularly severe during summer peak periods, when elevated air-conditioning demand sharply increased the risk of blackouts. The severity of the shock varied substantially across regions, driven by differences in pre-accident nuclear reliance and reserve margins. Regions with high nuclear dependence and limited reserve margins experienced especially acute supply disruptions, while others were less affected.

Policymakers responded to the electricity supply shock with a combination of demand-side rationing and supply-side adjustments. On the demand side, the government prioritized avoiding blackouts through "physical rationing." Most notably, in the summer of 2011, the government issued a formal mandate requiring large establishments in the most affected regions to reduce peak electricity usage by 15%. This marked the first time such emergency powers had been invoked since the 1973 oil shocks.[4] The government also relied on conservation campaigns and moral suasion, encouraging manufacturers to improve energy efficiency, dim lighting, reduce air-conditioning use, and shift operations away from peak hours.

On the supply side, utilities reactivated retired and inefficient fossil fuel plants and constructed new thermal generators. These measures increased available capacity but at a markedly higher marginal cost and fell short of fully replacing lost nuclear baseload generation. As shown in Panel B of Figure 1, the net effect was a contraction in aggregate electricity supply over the medium run. Although renewable generation expanded during this period, its contribution remained limited throughout our study horizon.

The shift from low-marginal-cost nuclear power to high-marginal-cost fossil fuel generation led to a substantial increase in electricity prices. Under Japan's cost-of-service regulatory regime, these higher generation costs were largely passed through to consumers. Panel C of Figure 1 illustrates the resulting evolution of regional electricity prices. While prices broadly tracked global fuel markets, rising through 2014 and declining thereafter, the defining feature of this period was the widening dispersion across regions, with regional differences in price growth reaching up to 20%. This divergence exposed manufacturers in different regions to markedly different input cost shocks, creating the cross-sectional variation exploited in our empirical analysis.

[3]The timing of the shutdown was largely predetermined by the expiration of operating licenses and mandatory safety inspections, which prevented reactors from restarting without regulatory approval.

[4]Based on Article 27 of the Electricity Business Act...

3 Data

3.1 Peak-Day Electricity Demand and Supply

To measure electricity supply shocks, we assemble regional-level data on peak-day electricity demand and supply from 2005 to 2018. Our primary source is the Electricity Supply–Demand Verification Subcommittee,[5] a government body that convenes annually to assess summer and winter electricity supply conditions and to formulate policy responses to potential shortages. The subcommittee collects detailed information on regional supply capacity and demand on peak days, which the national government uses to design and implement interventions. Importantly, these assessments are conducted at the peak-day level rather than the peak-hour level, reflecting institutional practice in Japan.

3.2 Census of Manufactures

We use establishment-level administrative data from the Census of Manufactures (CMF) for 2005–2018, compiled by the Ministry of Economy, Trade, and Industry (METI). The census annually surveys all manufacturing establishments with 30 or more employees, covering approximately 80% of total manufacturing output in Japan. In scope and structure, the CMF is comparable to manufacturing census data used in studies of the United States and other major economies.

The Census of Manufactures provides detailed plant-level information on outputs, inputs, and production characteristics. We measure output as revenue adjusted for inventory changes. We construct the capital stock as the sum of the reported book values of land, buildings, and machinery, net of depreciation. We measure electricity, materials, and labor using expenditure data. All monetary values are deflated using industry-specific output and input price deflators, the consumer price index, and capital stock deflators.[6] Because the data do not report plant-level electricity prices, we use average regional electricity prices, which are appropriate under regulated pricing during our study period. Each establishment is assigned a 2- to 4-digit industry code.[7] Throughout the paper, we use the three-digit industry classification as our baseline definition of an "industry" unless otherwise noted.

[5] We supplement these data with information from the Agency of Natural Resources and Energy (ANRE) and utilities' power supply planning documents.

[6] Industry-specific output and input deflators are from the Bank of Japan and are used to deflate output and intermediate inputs, including materials, fuel, and electricity. We deflate total wages and salaries using the CPI, and deflate capital stock using capital deflators based on national capital stock series published by the Cabinet Office.

[7] The data include 24 two-digit industries, 153 three-digit industries, and 546 four-digit industries.

We impose three sample restrictions. First, we exclude establishments located in municipalities severely affected by the Great East Japan Earthquake and the Fukushima nuclear accident. Second, for each outcome regression, we restrict the sample to establishments observed in the baseline year (2010) and the outcome year so that our main dependent variable, the change in the log outcome, is defined. Third, we restrict the CMF sample to establishments with non-missing values for revenue and primary input expenditures (labor, materials, electricity, and capital).

The CMF includes several additional features central to our analysis. Firm identifiers allow us to examine production reallocation within multi-plant firms. A subset of establishments reports product-level information, enabling an analysis of product switching.

3.3 Energy Consumption Statistics

We complement the CMF with plant-level administrative data from the Energy Consumption Statistics (ECS), provided by ANRE.[8] These surveys report physical consumption of fuel inputs, including natural gas (city gas, LPG, LNG), gasoline, kerosene, heavy oil, diesel oil, and steam.[9] We convert all fuels into a common energy unit based on heat content to construct a standardized measure of energy input.

A key advantage of the ECS is that it distinguishes fuel used for production from fuel used for on-site self-generation, including electricity produced by boilers, cogeneration systems, or diesel generators. Relative to the CMF, the ECS differs in industry coverage and timing.[10] In our econometric analysis, we therefore align the shock and outcome definitions to the ECS fiscal-year structure.

3.4 Summary Statistics

Table 1 presents summary statistics. In the CMF (N = 442,099), the average plant generates approximately 5,040 million yen in annual revenue (about $40–50 million in 2010 USD) and employs 127 workers. Approximately 2.8% of establishments export, and 42% belong to multi-plant firms. On average, labor, capital, and materials account for 22%, 3.2%, and

[8] In practice, we use both the Energy Consumption Statistics and the Petroleum and Other Fuels Monthly Statistics. These surveys collect similar variables but differ in industry coverage: the latter focuses on a subset of energy-intensive industries, while the former covers a broader range.

[9] A small number of establishments report expenditures instead of physical quantities; in these cases, we infer quantities using average fuel prices.

[10] The ECS covers fewer industries overall, but it includes plants that operate self-generators even when employment is below 30 workers. The ECS is reported on a fiscal-year basis; as a result, the Fukushima Daiichi nuclear accident in March 2011 is recorded in fiscal year 2010 rather than calendar year 2011, so we use 2009 as the reference year. Finally, the ECS reports employment but not total payroll.

51% of revenue, respectively, while electricity represents a relatively small share at 2.0%. Among establishments covered by the ECS (N = 381,891), approximately 11% own self-generation capacity. These plants purchase 30 TJ of electricity from the grid and generate 13 TJ on site, implying that self-generation accounts for roughly 30% of total electricity use.

4 Empirical Strategy

4.1 Constructing a Measure of Electricity Supply Shocks

We construct a measure of regional electricity supply shocks based on projected shortfalls in planned supply capacity on the peak day relative to peak demand. During the post-accident period, the government's primary objective was to avoid widespread blackouts, which required meeting demand on the system peak, when consumption is highest and weather conditions are most extreme. We classify a region as facing a binding supply constraint when its projected reserve margin falls below 5% of peak demand, the minimum buffer used to accommodate stochastic demand fluctuations.

Formally, we define the supply shock intensity for region r in year t as

$$\mathcal{S}_{rt} = \left(\frac{1.05\, D^{\text{peak}}_{r,2010} - S^{\text{plan,peak}}_{rt}}{D^{\text{peak}}_{r,2010}} \right) \mathbb{1}\left\{ S^{\text{plan,peak}}_{rt} < 1.05\, D^{\text{peak}}_{r,2010} \right\}, \tag{1}$$

where $D^{\text{peak}}_{r,2010}$ is realized peak-day demand in region r in 2010, and $S^{\text{plan,peak}}_{rt}$ is planned available supply capacity for the peak day in year t. The indicator function equals one when planned capacity implies a reserve margin below 5% relative to the 2010 peak-demand baseline. When it is active, $\mathcal{S}_{rt}$ measures the supply gap, expressed as a share of 2010 peak-day demand, required to restore a 5% reserve margin.

Electricity supply shocks exhibit substantial spatial and temporal variation. Panel A of Figure 2 plots the geographic distribution of these shocks, revealing widespread exposure across Japan. Panel B displays the time series of the average shock measure. Following the 2011 disaster, shock intensity spiked sharply; while shocks persisted throughout the study period, their severity fluctuated significantly across regions and over time.

This spatial variation is driven largely by the fragmentation of Japan's electricity grid. Appendix Table A1 reports peak-day supply, demand, and interregional transmission volumes for four major regions. When Tokyo faced its most severe supply deficits, it

received virtually no electricity from other regions or non-utility suppliers.[11] Instead, Tokyo remained a net exporter, sending roughly 0.1% to 0.5% of its pre-accident capacity to other regions (Panel B). Similarly, even during severe shocks in 2016, Kansai exported power rather than importing it (Panel D).

4.2 Instrumenting for Electricity Supply Shocks

OLS estimates of electricity supply shocks may be biased by endogenous policy responses. Following the nuclear phase-out, utilities in some regions reactivated dormant thermal plants or brought nuclear reactors back online. These supply-side adjustments may be correlated with unobserved local economic conditions, such as industrial demand recovery or political influence.

To address this concern, we construct a shift-share (Bartik) instrument that isolates variation in regional supply shocks driven by the national nuclear phase-out (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). The instrument leverages pre-determined regional generation shares, interacted with national growth rates by fuel type. We also control for the initial reserve margin, as regions with larger pre-existing buffers were mechanically less exposed to abrupt supply reductions.

The vector of instruments, $\mathbf{Z}_{rt}$, is defined as:

$$\begin{aligned} \mathbf{Z}_{rt} &= \left\{ \omega^{k}_{r,2010} \times g^{k}_{t,-r} \right\}_{k \in K} \cup \left\{ \omega^{RM}_{r,2010} \right\}, \\ & k \in \{\textit{Nuclear, Thermal, Hydro, Others}\}, \end{aligned} \tag{2}$$

where $\omega^{k}_{r,2010}$ denotes the share of electricity generated by source k in region r in 2010, and $g^{k}_{t,-r}$ is the leave-one-out national growth rate of generation for source k. We also include the region's initial reserve margin ratio, $\omega^{RM}_{r,2010}$. All variables are normalized by 2010 peak-day demand. This instrument captures variation in local supply constraints predicted by historical fuel mix and the national nuclear phase-out, conditional on pre-existing reserve margins.

4.3 Econometric Specification

To quantify the effects of electricity supply shocks on manufacturing outcomes, we regress log changes in plant-level outcomes on the intensity of regional supply shocks. We es-

[11] Here, "transmission" refers to electricity imports from other regional utility monopolies or non-utility suppliers (e.g., independent power producers). Physical constraints, including the 50 Hz/60 Hz frequency mismatch and geographic bottlenecks, further limited interregional flows.

timate the model separately for two distinct post-accident periods: the phase-out period (2011–2012), during which nuclear reactors were gradually suspended, and the post-phase-out period (2013–2018), following the complete cessation of nuclear power. Throughout the paper, we focus on the post-phase-out period (2013–2018).

For each period, we estimate:

$$\Delta \ln Y_{ijrt} = \rho \, \mathcal{S}_{rt} + \mu_{jt} + \epsilon_{ijrt}, \tag{3}$$

where $\Delta \ln Y_{ijrt} \equiv \ln(Y_{ijrt}) - \ln(Y_{ijr,2010})$, and $\mathcal{S}_{rt}$ denotes the electricity supply shock intensity defined in Section 4.1. The coefficient ρ captures the semi-elasticity of Y_{ijrt} with respect to $\mathcal{S}_{rt}$, allowing it to differ between the phase-out and post-phase-out regimes. All specifications include industry-by-year fixed effects, μ_{jt}, which control for time-varying industry-specific shocks. The error term ϵ_{ijrt} captures idiosyncratic plant-level variation.

We compute standard errors using the wild cluster bootstrap, clustered at the region level. Because treatment assignment varies across only nine regions, standard asymptotic approximations may be unreliable and prone to over-rejection (Cameron and Miller, 2015). Following Canay et al. (2021) and Roodman et al. (2019), we implement the wild bootstrap using Rademacher weights and report the resulting p-values and confidence intervals.

5 Electricity Supply Shocks and Plant Outcomes

5.1 Trends in Inputs and Outputs of Manufacturing Plants

Figure 3 plots trends in manufacturing-plant inputs and outputs from 2005 to 2018, grouping plants by whether they are located in high-shock or low-shock regions. Panel A shows that electricity use evolved similarly across the two groups prior to 2011. Immediately after 2011, however, the series diverge: electricity consumption falls sharply in high-shock regions, while it changes only modestly in low-shock regions. The resulting gap persists through 2018.

Panels B–D show analogous patterns for employment, the total salary, and materials, though the divergences are smaller in magnitude than for electricity. Revenue (Panel E) mirrors these input adjustments: after 2011, plants in high-shock regions exhibit a sustained shortfall relative to those in low-shock regions.

Figure 4 complements these time-series patterns with cross-sectional evidence. Using outcomes residualized with 3-digit industry fixed effects and aggregated to the nine regions, Panels A and C relate regional electricity supply shocks to changes in electricity consumption and revenue (both measured relative to 2010). Regions experiencing larger

shocks exhibit larger post-accident declines in electricity use, with fitted slopes of −0.83 in the medium term (2013–2015) and −0.77 in the long term (2016–2018). Revenue displays a similar relationship, with slopes of −0.26 and −0.24, respectively. In contrast, Panels B and D show little association in the pre-accident period, when we plot pre-2011 changes against regions' average post-phase-out shocks, suggesting that the post-accident declines are unlikely to be explained by differential pre-trends correlated with shock exposure.

5.2 First-stage Results: Testing the Instruments

Table 2 reports first-stage estimates for our instrument for regional electricity supply shocks, using region-by-year data from 2011 to 2018. Identification stems from the sharp nationwide contraction in nuclear generation following the 2011 accident, interacted with cross-sectional variation in pre-accident regional reliance on nuclear power. In 2010, nuclear power accounted for an average of 18% of peak-day demand ($Sh_{r,2010}^{\text{Nuclear}}$). During the post-phase-out period (2013–2018), nationwide nuclear generation fell by 97% ($g_{t\geq 2013}^{\text{Nuclear}}$).

Our preferred specification (Column (2)) yields a coefficient of −0.25 on $Sh_{r,2010}^{\text{Nuclear}} \times g_t^{\text{Nuclear}}$. This estimate implies that the nuclear phase-out increased the intensity of regional supply shocks by about 4.4 percentage points ($0.18 \times 0.97 \times 0.25 \approx 0.044$). Interpreted in demand-equivalent terms, meeting peak load to avoid widespread outages would have required peak-demand reductions of roughly 4.4% in the average region and about 7% in the most exposed regions, holding other margins fixed.

Other supply components contribute less to the observed variation. Although the coefficients on thermal and hydro power are statistically significant, their quantitative importance is limited: these sources declined by only 2.4% and 1.8% of 2010 peak-day demand on average, respectively. Renewables increased by just 0.7% of peak-day demand. Finally, the initial reserve margin ($Sh_{r,2010}^{\text{RM}}$) is negatively associated with supply shocks, consistent with higher pre-existing capacity buffers mitigating the severity of the disruption.

The first stage is strong. The first-stage F-statistic is 147 for the phase-out period (2011–2012) and 47 for the post-phase-out period (2013–2018).[12]

[12]These values exceed the conventional $F > 10$ threshold suggested by Stock and Yogo (2003). Lee et al. (2022) argues that substantially larger F-statistics may be desirable in some settings to ensure robust instrument relevance.

5.3 The Effects of Supply Shocks on Manufacturing Outcomes

Impacts on Electricity Inputs Table 3 reports OLS and 2SLS estimates of equation (3) for plant-level electricity use. Supply shocks reduce electricity consumption (Column (1)), and the effect strengthens over time. During the phase-out period (2011–2012), the IV estimates (Panel B) imply that a 1 ppt increase in the supply shock lowers electricity use by about 0.3%. The effect rises to 0.81% in the post-phase-out period (2013–2018). Given a mean supply shock of 7.4% after the nationwide shutdown, these estimates imply an average decline in manufacturing electricity use of roughly 6% during the post-phase-out period.

This temporal pattern is consistent with the evolution of policy and tariff adjustments. In the immediate aftermath of the accident, the policy response relied primarily on seasonal, quantity-based conservation measures that targeted peak-load months. If plants adjusted only through these seasonal constraints, the implied annual reduction would be approximately one quarter of our supply-shock measure,[13] since the shock season spans roughly one quarter of the year, which is close to the short-run estimate.

The larger post-2013 response is consistent with a growing role for price-based adjustment. Under cost-of-service regulation, utilities were permitted to pass through higher generation costs as nuclear output was replaced by fossil generation, but tariff revisions occurred with lags. Some utilities raised retail tariffs in 2012 (e.g., Tokyo), whereas others implemented comparable increases in 2013 (five utilities) or later.

OLS estimates (Panel A) are slightly smaller in magnitude than 2SLS estimates (Panel B), -0.63 versus -0.81 in the post-phase-out period. This divergence is consistent with endogeneity that attenuates OLS estimates—for example, mitigation efforts or other unobserved adjustments that are more prevalent in severely affected regions. At the same time, the broadly similar patterns across specifications underscore the limited scope for fully offsetting the supply disruption. Moreover, the event-study estimates in Figure 5 show little evidence of differential pre-trends,[14] supporting the view that post-2011 changes are not driven by pre-existing regional trends.

Impacts on Other Primary Inputs Columns (2)–(4) of Table 3 and Panels B–D of Figure 5 show that electricity supply shocks also reduce other primary inputs, including labor and materials, though the magnitudes are smaller than for electricity. During the post-phase-out period (2013–2018), a 1 ppt increase in the supply shock reduces employment by

[13] This calculation assumes that a 1% supply shock translates one-for-one into a 1% reduction in manufacturers' electricity use during the shock season.

[14] Event-time coefficients are estimated via separate regressions of equation (3) for each year. In pre-accident years, each region is assigned its post-accident average shock as a placebo.

approximately 0.2%. Material inputs decline by about 0.18%, but the estimate is imprecise, likely reflecting the limited number of regional clusters.

The labor results indicate disproportionate impacts across worker types. Appendix Table B1 shows little effect on wages per worker (Column (1)), suggesting that adjustment operates primarily through employment rather than wages. Estimated effects are similar for male and female workers (Columns (2) and (3)). By contrast, effects on non-regular workers (Column (5)) are more than twice as large as those on regular workers (Column (4)).

Appendix Table B2 shows that investment responds sharply, while capital-stock responses are slower-moving and less precisely estimated. A 1 ppt increase in the supply shock reduces investment by 0.52% in 2011–2012 and 1.16% in 2013–2018 (Appendix Table B2, Panel B, Column (1)). In contrast, capital-stock estimates are near zero in the short run and remain noisy in the post-phase-out period (Column (2)); if anything, their magnitudes are comparable to the employment and materials responses, consistent with partial capital fixity and measurement error. Overall, the evidence points to adjustment primarily through reduced capital formation rather than scrapping existing assets.

Impacts on Outputs and Economic Magnitudes Electricity supply shocks reduce manufacturing revenue (Column (5) of Table 3). In the post-phase-out period, the IV estimate is -0.23, which implies an average revenue decline of about 1.7% at the mean supply shock.

Revenue effects are close to zero and statistically insignificant during the initial phase-out period (2011–2012). One interpretation is that policy uncertainty muted short-run adjustment: because Japan did not formally commit to a long-run phase-out path until mid-2012, plants may initially have viewed the disruptions as transitory. Revenue losses emerge more clearly after 2013, once the decline in electricity availability appeared persistent and retail tariff increases had largely taken effect.

A back-of-the-envelope calculation illustrates the economic magnitude. In 2010, manufacturing employed 7.7 million workers and generated 289 trillion JPY in annual revenue (METI, 2012). Applying our estimated post-2013 effects implies that supply shocks reduced manufacturing employment by roughly 110,000 workers and annual revenue by about 4.9 trillion JPY.[15]

[15]These figures are for manufacturing plants with at least four employees (METI, 2012). We assume the post-phase-out estimates are representative and compute these totals by multiplying the estimated percentage effects by the 2010 aggregates. For reference, 4.9 trillion JPY is on the order of $40–60 billion, depending on the exchange rate.

5.4 Alternative Explanations and Robustness Checks

Supply Chain Disruptions We first examine whether our main results are driven by supply-chain disruptions. A primary concern is that our estimates reflect indirect damages propagated through input–output linkages rather than the direct effects of electricity supply contractions, consistent with evidence that the 2011 earthquake and nuclear accident generated sizable spillovers through supply chains (Carvalho et al., 2021).

To address this concern, we augment our baseline specification with two measures of supply-chain exposure constructed from the prefecture-by-industry (2-digit) input–output tables in RIETI (2005). The first measure captures direct exposure: the share of intermediate inputs sourced from the three most affected prefectures (Iwate, Miyagi, and Fukushima). The second measure captures indirect (Leontief-type) exposure by incorporating upstream linkages, for example, reliance on suppliers that themselves source from the disaster zone. We compute these measures using pre-disaster relationships and interact them with year dummies to allow the effects to vary over time. As shown in Table B3, the estimates are essentially unchanged, indicating that supply-chain spillovers are unlikely to drive our results.

Demand-Side Shocks We next investigate whether the observed contractions reflect localized demand shocks rather than supply-side constraints. Local demand may decline due to direct and indirect damages from the earthquake, tsunami, and nuclear accident, and it may also respond endogenously to electricity supply shocks. We therefore control for time-varying prefecture-level local economic conditions, including log gross prefectural product in non-manufacturing sectors (primary and tertiary), total and working-age populations, construction investment, aggregate consumption, and the job-to-application ratio (MIC, 2025). Table B3, columns (3)–(4), shows that the estimated effect on electricity use is modestly attenuated (from -0.81 to -0.68), while the estimated effect on revenue is little changed. These results indicate that local demand conditions do not account for our main findings.

Alternative Specifications and Weighting Finally, we assess robustness to plant-level controls, alternative estimators, and weighting schemes. We add controls for pre-shock (2010) log revenue, input cost shares (electricity, capital, and labor), export intensity, and multi-plant firm status. The estimated effects remain stable when these controls are included (column (4)). We also re-estimate the model using a two-way fixed effects specification with plant and year fixed effects rather than first differences (column (5)). In addition, the estimates are similar in unweighted regressions (column (6)) relative to our baseline specification that weights by log total salary.

5.5 Heterogeneous Impacts

Figure 6 explores heterogeneity by splitting the sample at the median of pre-determined plant and industry (3-digit) characteristics. Panel A shows that the electricity response is similar across most subgroups. The full-sample 2SLS estimate of −0.81 is broadly representative of the group-specific estimates, which are typically close in magnitude. This pattern is consistent with the post-Fukushima institutional environment, in which conservation efforts and regulated tariff increases applied broadly. One exception is export exposure. Plants in high-export industries reduce electricity consumption more than those in low-export industries (−1.09 versus −0.52).

Panel B shows greater heterogeneity in revenue outcomes. Revenue losses are larger for smaller plants: a split by pre-shock revenue yields an estimate of −0.33 for below-median plants and −0.16 for above-median plants. Revenue losses are also substantially larger in high-export industries (−0.39) than in low-export industries (−0.08), consistent with greater exposure to international competition and more elastic demand, in line with evidence on output responses to labor-cost shocks in Harasztosi and Lindner (2019). For other dimensions, including electricity intensity, productivity, and market concentration, differences are modest.

6 Plants' Responses to Electricity Supply Shocks

6.1 Impacts on Revenue Productivity

This section examines how and to what extent manufacturing plants adapt to electricity supply shocks. A central question is whether plants can mitigate revenue losses by increasing revenue productivity (TFPR), i.e., generating more revenue from the same observed input bundle. We describe the estimation of TFPR in Section 7.

There are competing theoretical predictions. On one hand, the Porter Hypothesis (Porter, 1991) suggests that shocks can induce operational efficiency improvements and adoption of energy-saving technologies. Japanese manufacturers, in particular, engaged in extensive energy-conservation activities after Fukushima, ranging from optimizing compressor pressure and intake temperatures to operational adjustments such as raising chiller output temperatures. On the other hand, TFPR could decline if electricity tightness distorts the input mix or induces underutilization of quasi-fixed capital, consistent with evidence from past oil shocks (Kilian, 2008).

We find no evidence of productivity improvements. Panel A of Figure C1 shows that TFPR trends in high-shock and low-shock regions are similar throughout our study period. Panel B, which reports event-study estimates, shows no statistically significant

effects. Point estimates are close to zero in the post-phase-out period (Table 4, Column (1)). Overall, these results suggest limited evidence that plants achieved productivity gains through technological innovation or operational improvements that would offset the inefficiencies created by electricity scarcity.

A related pattern is that electricity use falls more than revenue. In our baseline estimates, a mean shock implies roughly a 6% decline in electricity consumption, while revenue declines by about 1.7% in the post-phase-out period. This increases the revenue-to-electricity ratio and could be misinterpreted as an improvement in "electricity-specific productivity." However, changes in this ratio can also arise mechanically from relative input price changes and re-optimization of input intensity, even without technological improvement. In Appendix C-2, we show that accounting for relative price movements during our study period explains most of the increase in revenue per unit of electricity, leaving little evidence of residual productivity growth.

6.2 Impacts on Observed Plant Behaviors

We next examine observable adjustment margins in response to electricity supply shocks. Beyond scaling down inputs, plants may change product mix, adjust procurement (including outsourcing), reallocate activity across establishments within multi-plant firms, or adjust along the extensive margin through relocation and exit. To assess these channels, we compile a set of plant- and firm-level outcomes from multiple data sources (Appendix C-3).

Product mix and outsourcing. We test whether plants adjust along product-market and procurement margins in response to electricity supply shocks. A natural hypothesis is that plants shift toward less electricity-intensive products or increase outsourcing to economize on electricity when local supply tightens. Columns (2)–(3) of Table 4 show no evidence of increased product switching or declines in product-level electricity intensity. Columns (4)–(5) likewise indicate little change in outsourcing, measured both in levels and as a share of materials. Overall, these margins do not appear to be first-order channels of adjustment in our setting.

Intra-firm reallocation. We next examine whether multi-plant firms reallocate activity across establishments to buffer localized electricity supply shocks (Chen et al., 2025). In our data, 42% of plants belong to multi-plant firms, accounting for 57% of employment and 66% of revenue. Appendix Table C1 provides limited evidence of such reshuffling: shocks are associated with a rise in electricity use at the firm's other establishments

(Column (1)), but we detect no corresponding response in revenue (Column (2)). This pattern suggests that intra-firm reallocation, while present on some margins, was not quantitatively important for offsetting the overall production impact.

Extensive margin: entry and exit. Appendix Table C2 shows no statistically significant effects of supply shocks on plant turnover. Using baseline entry and exit rates of 1.7% and 4.8%, a mean shock of 7.4 ppt, and the IV coefficients in Columns (2)–(3), the implied annual changes are 0.14 ppt for entry ($1.7\% \times 1.13 \times 7.4$) and 0.37 ppt for exit ($4.8\% \times 1.04 \times 7.4$). Because the data do not distinguish relocation from shutdown, we interpret these estimates as effects on overall plant turnover rather than on any single margin.

6.3 Substitution Between Electricity and Other Fuels

We examine whether plants mitigated electricity supply shocks by substituting away from grid electricity along two margins. First, they may switch to other fuels in production, feasible for some thermal processes (e.g., boilers) but costly for electricity-specific applications (e.g., precision motor-driven machinery or electric arc furnaces). Second, plants may substitute indirectly by using alternative fuels for on-site self-generation, as documented by Allcott et al. (2016) and von Graevenitz and Rottner (2024). We distinguish these channels using the Energy Consumption Statistics, which separately reports fuel inputs used in production and those used for self-generation.

Direct fuel substitution. We find little evidence of substitution away from electricity toward fuels used directly in production. Column (1) of Table 5 shows that supply shocks reduce production-fuel use in the post-phase-out period; the 2SLS point estimate is about -0.27, though imprecise (Panel B). The decline is at least as large as those for labor and materials (Section 5.3), consistent with electricity (motive power) and production fuels (thermal inputs) moving together in aggregate rather than serving as close substitutes.

Self-generation. In contrast, we find a clear increase in self-generation. Plants exposed to larger shocks raise fuel use for on-site generation (Column (2)) and increase electricity generated on site (Column (3)). The response is particularly pronounced among plants that already own generators. For this group, the intensive-margin 2SLS estimate is 3.43 (Appendix Table C4, Column (2)), implying about a 25% increase in self-generated electricity at the mean shock of 7.4 ppt ($\approx 3.43 \times 0.074$). Because only about 10% of plants own generators, this group can offset a substantial share—and, in some cases, plausibly all—of the reduction in grid-provided electricity by expanding self-generation.

Consistent with this mechanism, when we split the sample by baseline generator ownership, plants with generators exhibit no statistically detectable revenue decline (Table C3, Column (2)), whereas revenue losses are concentrated among plants without generators (Column (4)), despite similar reductions in grid-purchased electricity (Columns (1) and (3)).

Two caveats limit the aggregate relevance of this margin. First, generator ownership is rare in our setting: only 11% of plants owned generators at baseline,[16] far below the 44% reported for India by Allcott and Keniston (2018). Second, extensive-margin adjustment is limited. The 2SLS coefficient of 0.36 (Appendix Table C3, Column (1)) implies an increase in generator ownership of only about 0.3 ppt at the mean shock ($0.11 \times 0.36 \times 7.4 \approx 0.003$).

Taken together, self-generation provided an effective shield for a small minority of plants, but the vast majority of manufacturers remained reliant on the grid and therefore exposed to the supply disruptions documented in our main analysis.

7 Production Framework and Model Validation

This section develops and validates a production framework to explain why electricity supply disruptions generate large revenue losses despite electricity's small expenditure share. The model features limited input substitutability and endogenous capital utilization. We estimate the production function and validate the model by comparing its simulated responses with the reduced-form 2SLS evidence.

7.1 Setup

Plants generate output Y_{it}, measured as deflated revenue (real shipments), using capital K_{it}, labor L_{it}, materials M_{it}, and electricity E_{it} under a constant-returns-to-scale CES technology with revenue productivity Ω_{it}. As physical quantities are unobserved, we follow the standard approach in the production-function literature and estimate a revenue-based production function. The production function for plant i at time t is

$$Y_{it} = \Omega_{it} \left(\alpha_K K_{it}^{\rho} + \alpha_L L_{it}^{\rho} + \alpha_M M_{it}^{\rho} + \alpha_E E_{it}^{\rho} \right)^{\frac{1}{\rho}}, \qquad \sigma \equiv \frac{1}{1-\rho}, \tag{4}$$

where σ denotes the elasticity of substitution and α_j are share parameters satisfying $\sum_j \alpha_j = 1$.

[16] This figure is likely an upper bound because the ECS records generator ownership for some plants with fewer than 30 employees. In the CMF, about 8% of plants own generators.

We distinguish between two subperiods within a year: a shock season and an off-season. Plants face electricity supply constraints for a fraction ϕ of the year, a duration known ex-ante to both regulators and plants. Let Y^S denote output in the shock season and Y^O denote output in the off-season. In the absence of shocks, output is identical across seasons (i.e., $Y = Y^S = Y^O$). After the nuclear phase-out, annual output is defined as the weighted average $Y = \phi Y^S + (1 - \phi)Y^O$.

Inputs are classified as either quasi-fixed or seasonally flexible. The installed capital stock K_{it} is quasi-fixed and effectively fixed over the post-accident period due to adjustment costs, though its utilization can vary across years and seasons (Section 7.4). Labor, materials, and electricity are chosen at the seasonal frequency; in particular, labor can adjust through working hours and non-regular employment, consistent with our data, with larger impacts concentrated in non-regular employment.

7.2 Small-Share Benchmark and Reduced-form Evidence

7.2.1 Electricity Shocks as Shadow-Price Changes

We model the shock as two constraints on plants' electricity demand. First, during the shock season for $t \geq 2011$, electricity use is subject to a quantitative restriction:

$$E_t^S \leq (1 - \mathcal{S})E_{2010}^S, \tag{5}$$

where $\mathcal{S}$ indexes the regional severity of the shock. Second, electricity prices rise due to higher generation costs and regulatory pass-through. We capture this by allowing the regulated price to increase from its pre-shock level p_{2010}^E to p_t^E.

With these changes, the effective (shadow) price of electricity in the shock season is $\tilde{p}_t^E = p_t^E + \lambda$, where $\lambda \geq 0$ is the scarcity rent under binding rationing, capturing the wedge between the retail prices and the marginal value of electricity.

7.2.2 Growth Accounting and the "Small-Share" Benchmark

To characterize the impact of these shocks, we rely on a standard growth-accounting decomposition. With constant returns to scale and price-taking in factor markets, the cost-minimization first-order conditions imply that output elasticities equal expenditure shares. Differentiating $\ln Y$ therefore gives

$$d \ln Y = s_E \, d \ln E + s_M \, d \ln M + s_L \, d \ln L + d \ln \Omega, \tag{6}$$

where $s_E \equiv \frac{P^E E}{PY}$, $s_M \equiv \frac{P^M M}{PY}$, and $s_L \equiv \frac{P^L L}{PY}$ denote expenditure shares in nominal revenue, where P is the output price index (so PY equals nominal revenue). We treat the capital stock as quasi-fixed over our horizon, so $d \ln K \approx 0$ (see Section 7.4).

Normalizing by shock severity yields

$$\frac{d \ln Y}{d\mathcal{S}} = s_E \frac{d \ln E}{d\mathcal{S}} + s_M \frac{d \ln M}{d\mathcal{S}} + s_L \frac{d \ln L}{d\mathcal{S}} + \frac{d \ln \Omega}{d\mathcal{S}}. \tag{7}$$

A common benchmark in the literature is the "small-share" approximation, which abstracts from endogenous movements in non-electricity inputs and measured productivity (Rotemberg and Woodford, 1996). Holding L, M, and Ω fixed, the first-order effect of the shock operates only through electricity:

$$\left.\frac{d \ln Y}{d\mathcal{S}}\right|_{L=L^*,\, M=M^*,\, \Omega=\Omega^*} \approx s_E \frac{d \ln E}{d\mathcal{S}}. \tag{8}$$

7.2.3 Comparison with Reduced-Form Results

Our reduced-form estimates imply revenue losses that substantially exceed those predicted by the standard small-share benchmark. In our sample, electricity accounts for only 2% of revenue. During the post-phase-out period, the shocks reduce electricity use at the average plant by about 6%. Under the small-share benchmark, a first-order growth-accounting approximation implies a revenue loss of only $0.02 \times 0.06 \approx 0.12\%$. In contrast, our estimates imply an average revenue contraction of 1.7%—roughly an order of magnitude larger—echoing evidence from past oil shocks (Rotemberg and Woodford, 1996; Bernanke et al., 1997; Kilian, 2008).

Equation (7) clarifies where this wedge can arise. We find little evidence of offsetting productivity responses ($d \ln \Omega / d\mathcal{S} \approx 0$) and limited adjustment through observable technological or organizational channels. By contrast, our reduced-form estimates show sizable declines in both labor and materials ($d \ln L / d\mathcal{S} < 0$, $d \ln M / d\mathcal{S} < 0$). Because labor and materials account for large expenditure shares in revenue (22% and 51%, respectively), their contraction is a key channel through which electricity shortfalls translate into large revenue losses.

7.3 Input Complementarity and Amplification

We now explain how input complementarity provides a simple propagation mechanism from electricity tightness to output losses. When electricity complements other primary inputs, an increase in its effective (shadow) price reduces not only electricity use but also

high-share inputs such as labor and materials, as documented in our data. This broad-based input response amplifies revenue effects relative to the small-share benchmark.

To formalize the propagation from electricity tightness to other inputs, consider a generic non-electricity input $X \in \{L, M\}$. Using the cost-share identity $s_X^C \equiv \frac{p^X X}{cY}$, where c denotes unit cost, log-differentiation yields

$$\frac{d \ln X}{d\mathcal{S}} = \underbrace{\frac{d \ln Y}{d\mathcal{S}}}_{\text{scale}} + \underbrace{\frac{d \ln s_X^C}{d\mathcal{S}}}_{\text{reallocation}} + \underbrace{\frac{d \ln c}{d\mathcal{S}}}_{\text{unit cost}}, \tag{9}$$

holding p^X locally fixed. The scale and unit-cost terms capture movements in input demand through output and unit cost, while the share term captures changes in relative factor demands.

Under CES cost minimization (equivalently, from the CES unit-cost function), cost shares respond to movements in the effective (shadow) price of electricity. Holding p^X and other factor prices fixed,

$$\frac{d \ln s_X^C}{d\mathcal{S}} = (\sigma - 1)\, s_E^C \, \frac{d \ln \tilde{p}^E}{d\mathcal{S}}. \tag{10}$$

Electricity supply shocks raise the effective electricity price $\tilde{p}^E$. When $0 < \sigma < 1$, equation (10) implies $\frac{d \ln s_X^C}{d\mathcal{S}} < 0$: electricity is complementary with other primary inputs, so higher $\tilde{p}^E$ lowers the cost shares of labor and materials and induces broad input contraction rather than operating solely through electricity reductions. By contrast, when $\sigma = 1$ (Cobb–Douglas), cost shares are invariant and the reallocation channel vanishes. Because labor and materials account for most expenditures, propagation to these high-share inputs provides a natural amplification mechanism and can generate revenue losses far larger than electricity's expenditure share would suggest.

7.4 Electricity-Constrained Capital Utilization

We model electricity-constrained capital utilization as an amplification channel. Electricity is a flow input but also a prerequisite for operating machinery. When capital adjustment is costly, a contraction in electricity availability can leave part of the installed capital stock idle, reducing effective capital services and output.

Let K denote installed capital services and let $\tilde{K}_t$ denote effective capital services. Following the putty–clay logic of Atkeson and Kehoe (1999), we assume that operating legacy (pre-shock) capital requires electricity in fixed proportion, so the amount of legacy

capital that can be operated is limited by the electricity allocated to it.

We treat total installed capital as fixed at its pre-shock level. Let $K \equiv K_{2010}$, and assume $K_t \equiv K$ for $t \geq 2011$. We allow the legacy component to decline at rate δ as it is phased out over time; in particular, for $t \geq 2011$,

$$\frac{K_t^{\text{old}}}{K} = (1-\delta)^{t-2010}, \qquad K_t^{\text{new}} \equiv K - K_t^{\text{old}}.$$

Absent vintage-specific electricity-use data, we allocate electricity across vintages in proportion to their capital shares, so $E_t^{\text{old}} = (K_t^{\text{old}}/K)E_t$.

We model effective capital services as the sum of fully operable new capital and the operable portion of legacy capital, which is limited by electricity availability:

$$\tilde{K}_t = K_t^{\text{new}} + \min\left\{K_t^{\text{old}},\, vE_t^{\text{old}}\right\}, \qquad v \equiv \frac{K}{E_{2010}}. \tag{11}$$

Here $v > 0$ is the amount of operable legacy capital services activated per unit of electricity allocated to legacy capital. We normalize v so that the constraint holds with equality in 2010, when capital is fully utilized. The $\min\{\cdot\}$ operator makes the extensive-margin logic explicit: legacy capital is either fully operable or capped by available electricity.

7.4.1 Structural Amplification

We embed effective capital services in a gross-output CES production function:

$$Y_t = \Omega_t \left(\alpha_K \tilde{K}_t^{\rho} + \alpha_L L_t^{\rho} + \alpha_M M_t^{\rho} + \alpha_E E_t^{\rho}\right)^{1/\rho}, \qquad \rho \equiv (\sigma-1)/\sigma. \tag{12}$$

When the legacy-capital utilization constraint binds, operable legacy services satisfy $\tilde{K}_t^{\text{old}} = vE_t^{\text{old}}$, so electricity effectively activates an additional quasi-fixed input. In this regime, a marginal reduction in E_t lowers both electricity services and usable legacy capital services, raising the effective importance of electricity precisely when electricity is scarce.

In the binding regime, the electricity-related component can be re-expressed to absorb the constrained legacy-capital component:

$$\alpha_E E_t^{\rho} + \alpha_K \left(\tilde{K}_t^{\text{old}}\right)^{\rho} = \left(\alpha_E + \alpha_K \left(v\frac{E_t^{\text{old}}}{E_t}\right)^{\rho}\right) E_t^{\rho} \equiv \tilde{\alpha}_{E,t}\, E_t^{\rho}, \tag{13}$$

where we use $\tilde{K}_t^{\text{old}} = vE_t^{\text{old}}$.

Under proportional allocation across vintages, $\frac{E_t^{\text{old}}}{E_t} = \frac{K_t^{\text{old}}}{K} = (1-\delta)^{t-2010}$ for $t \geq 2011$,

so

$$\tilde{\alpha}_{E,t} = \alpha_E + \alpha_K \big(v(1-\delta)^{t-2010}\big)^{\rho}. \tag{14}$$

Analogously, define $\tilde{\alpha}_{K,t}$ as the coefficient on K^{ρ} that reproduces the contribution of new (fully operable) capital services: $\alpha_K (K_t^{\text{new}})^{\rho} \equiv \tilde{\alpha}_{K,t} K^{\rho}$ with $\tilde{\alpha}_{K,t} \equiv \alpha_K \left(1 - (1-\delta)^{t-2010}\right)^{\rho}$.

These rearrangements imply that electricity enters with an amplified effective weight, $\tilde{\alpha}_{E,t} > \alpha_E$, whenever legacy utilization binds. Interpreted through Equation (7), this mechanism raises the effective sensitivity of $\ln Y_t$ to $\ln E_t$ relative to what the observed expenditure share s_E alone would suggest. A contraction in E_t reduces not only electricity services but also operable legacy capital services. As $\tilde{K}_t$ falls, the marginal revenue products of labor and materials decline, inducing contractions in these high-share inputs and amplifying revenue losses relative to the small-share benchmark.

7.5 Production Function Estimation

We estimate the structural parameters of the CES production function separately for each of the 153 three-digit manufacturing industries, allowing technology to differ across industries while imposing a common production technology within industry. We recover the CES parameters in two steps. First, we estimate industry-specific revenue elasticities using the approach of Gandhi et al. (2020). Second, conditional on a candidate elasticity of substitution σ, we map the estimated elasticities into the CES distribution parameters α evaluated at the industry-mean input bundle.

This approach offers three advantages. First, it identifies revenue elasticities of flexible inputs from static first-order conditions, avoiding value-added specifications and the functional-form restrictions embedded in standard control-function estimators (e.g., Olley and Pakes, 1996; Ackerberg et al., 2015). Second, because the first step delivers reduced-form elasticities, the implied CES share parameters can be recovered for alternative values of σ without re-estimating the elasticities. Third, estimating the model separately by three-digit industry accommodates substantial technological heterogeneity across manufacturing, which is essential for disciplined welfare comparisons of pricing and rationing policies.

7.5.1 Setup

We assume standard timing of input choices and perfectly competitive factor markets. The log revenue production function is

$$y_{it} = f(k_{it}, l_{it}, m_{it}, e_{it}) + \omega_{it} + \epsilon_{it}, \tag{15}$$

where lowercase variables denote logs, $\omega_{it} = \ln \Omega_{it}$ is plant-level productivity, and ϵ_{it} is an ex-post revenue shock satisfying $\mathbb{E}[\epsilon_{it} \mid k_{it}, l_{it}, m_{it}, e_{it}] = 0$. Productivity follows a first-order Markov process,

$$\omega_{it} = h(\omega_{it-1}) + \eta_{it}, \tag{16}$$

where the innovation η_{it} is orthogonal to the quasi-fixed input, $\mathbb{E}[\eta_{it} \mid k_{it}] = 0$. Plants choose flexible inputs (l_{it}, m_{it}, e_{it}) after observing ω_{it}, while the quasi-fixed input k_{it} is set prior to the realization of the current productivity innovation.

7.5.2 Flexible inputs

In the first stage, we estimate revenue elasticities of flexible inputs: $\theta_E \equiv \partial f / \partial e_{it}$ (electricity), $\theta_M \equiv \partial f / \partial m_{it}$ (materials), and $\theta_L \equiv \partial f / \partial l_{it}$ (labor). Under price-taking in factor markets, the static first-order condition implies that the expenditure share of a flexible input identifies its revenue elasticity up to the multiplicative revenue shock. For electricity,

$$\ln s_{E,it} = \ln \theta_E - \epsilon_{it}, \tag{17}$$

where $s_{E,it} \equiv \frac{P_t^E E_{it}}{P_{it}^Y Y_{it}}$ and $P_{it}^Y Y_{it}$ is nominal revenue. Equation (17) implies that θ_E is identified by the conditional mean of the log electricity expenditure share.[17] We obtain θ_M and θ_L analogously using materials and labor shares.

7.5.3 Quasi-fixed input

In the second stage, we estimate the revenue elasticity of capital, θ_K. Using the first-stage estimates, we construct the partial residual

$$\widehat{\boldsymbol{y}}_{it} \equiv y_{it} - \hat{\theta}_E e_{it} - \hat{\theta}_M m_{it} - \hat{\theta}_L l_{it}, \tag{18}$$

and define implied productivity as

$$\hat{\omega}_{it}(\theta_K) \equiv \widehat{\boldsymbol{y}}_{it} - \theta_K k_{it} - \hat{\epsilon}_{it}. \tag{19}$$

Following Gandhi et al. (2020), the revenue shock ϵ_{it} can be recovered from the first-stage share equations; we denote this estimate by $\hat{\epsilon}_{it}$. We then estimate θ_K by GMM using the Markov structure in (16) and the moment condition $\mathbb{E}[\eta_{it} \cdot k_{it}] = 0$.

[17] We impose the normalization $\mathbb{E}[\exp(\epsilon_{it}) \mid k_{it}, l_{it}, m_{it}, e_{it}] = 1$, so that $\mathbb{E}[\ln s_{E,it} \mid k_{it}, l_{it}, m_{it}, e_{it}] = \ln \theta_E$.

7.5.4 Recovering CES share parameters

Conditional on a candidate elasticity of substitution σ (with $\rho \equiv (\sigma - 1)/\sigma$), we recover the CES distribution parameters α by mapping the estimated industry-average elasticities and observed input levels into the implied local shares:

$$\hat{\alpha}_j = \frac{\hat{\theta}_j X_j^{-\rho}}{\sum_{k \in \{K,L,M,E\}} \hat{\theta}_k X_k^{-\rho}}, \qquad j \in \{K, L, M, E\}, \tag{20}$$

where X_j denotes the geometric mean of input j within the industry. This normalization ensures that the implied CES technology locally matches the estimated elasticities at the industry-mean input bundle and satisfies constant returns to scale, $\sum_j \hat{\alpha}_j = 1$.

7.5.5 Production Function Results

Table 6 reports estimated revenue elasticities using pre-accident data (2005–2010). Relative to industry-level OLS regressions of log revenue on log inputs, the structural estimates assign a larger elasticity to the quasi-fixed input (capital) and smaller elasticities to flexible inputs, consistent with simultaneity bias. Because flexible input choices covary with unobserved productivity, OLS tends to overstate their elasticities and understate the contribution of quasi-fixed capital. The structural elasticities sum to 0.784 on average. Under constant returns to scale in physical production, this implies an average markup of $\mu \approx (\sum_j \theta_j)^{-1} \approx 1.28$.

7.6 Model Validation and Parameter Discipline

This subsection assesses the empirical plausibility of the production framework and disciplines its structural parameters. Following Todd and Wolpin (2006), we simulate plant-level outcomes under the observed post-phaseout environment (2013–2018) and compare the model-implied responses to our reduced-form 2SLS estimates.

7.6.1 Simulation Setup

The simulation uses the same plant sample as the reduced-form analysis to ensure comparability. Each plant is characterized by its observed 2010 productivity and a predetermined electricity–capital ratio that governs capital idleness under electricity scarcity, together with industry-specific share parameters. We close demand with $Y = BP^{-\eta}$. We set the plant-level shifter B_i using each plant's observed 2010 deflated revenue. We set

$\eta = 6$, implying a 20 percent markup, and unless otherwise noted impose a common elasticity of substitution $\sigma = 0.5$ across industries.

The simulation alters two exogenous policy variables: quantity rationing and electricity price increases. Under rationing, we impose the observed supply shock $\mathcal{S}$ (varying by region and year) as the shock-season constraint in Equation (5), where the peak season lasts approximately 90 days (about one quarter of the year). In addition, plants face year-round increases in regulated electricity prices. Appendix Table D1 reports the full set of simulation inputs and data sources (see Appendix D-2).

The simulation also incorporates heterogeneity in self-generation capacity. As shown in Section 6.3, generator-owning plants reduce purchases from grid utilities while exhibiting no meaningful declines in other inputs or in total revenue. Consistent with this evidence, we treat generator-owning plants as largely insulated from the supply shock and assume they maintain their pre-shock levels of non-electricity inputs and revenue.

7.6.2 Validation: Comparison with Reduced-Form Responses

We generate plant-level predictions for inputs and revenue under the realized electricity price p and rationing regime $\mathcal{S}$, relative to a no-shock baseline $(p_{2010}, \mathcal{S}_{2010})$, holding productivity and other state variables fixed (Appendix D-1). Using these simulated outcomes, we construct the model-implied change in an outcome Y, $\Delta Y \equiv \ln Y(p, \mathcal{S}) - \ln Y(p_{2010}, \mathcal{S}_{2010})$, and re-estimate Equation (3) on the simulated data using the same 2SLS specification and sample restrictions as in the reduced-form analysis.

The model closely matches the empirical benchmarks over the post-phase-out period (2013–2018). As reported in Table 7, the model-implied coefficients for electricity consumption and revenue are -0.80 and -0.20, respectively (Panel B-1), compared with reduced-form estimates of -0.81 and -0.23 (Panel A). The corresponding coefficients for labor and materials are also similar in magnitude. This quantitative correspondence supports the credibility of the estimated production technology and motivates its use for counterfactual analysis.

7.6.3 Disciplining Input Substitutability and Capital Utilization

We use the validation exercise to discipline two amplification mechanisms: limited input substitutability (captured by the elasticity of substitution σ) and electricity-constrained capital utilization. To obtain a plausible substitution elasticity, we vary σ while holding other parameters fixed; for each candidate value, we recalibrate the CES share parameters α so that the technology matches the estimated elasticities at the industry-mean input bundle.To incorporate electricity-constrained capital operation, we replace (α_K, α_E) with the

effective weights $(\tilde{\alpha}_K, \tilde{\alpha}_E)$ implied by the constraint linking operable capital to electricity (via the electricity–capital ratio v).

Among the values considered, $\sigma = 0.5$ provides the closest match to the reduced-form evidence (Figure 7). When substitution is more flexible (e.g., $\sigma \in \{0.99, 1.5\}$), the model implies substantially larger reductions in electricity use as plants substitute away from electricity toward other inputs. For example, setting $\sigma = 0.99$ (approximately Cobb–Douglas) implies an electricity contraction that is 57% larger, along with smaller contractions in labor and materials, yielding revenue losses 38% smaller than in our baseline and thereby counterfactually overstating substitution away from electricity. Conversely, lower substitutability (e.g., $\sigma = 0.2$) predicts relatively small reductions in electricity use but excessively large contractions in labor, materials, and revenue, as plants shed high-share inputs to maintain near-fixed input proportions.

The validation also underscores the quantitative importance of electricity-constrained utilization as an amplification mechanism. Shutting down the utilization constraint reduces predicted revenue losses by 76% relative to the baseline.

7.6.4 Elasticity Benchmarks

We report reduced-form and model-implied elasticities. The reduced-form estimates imply an own-price elasticity of electricity demand of −0.6 (Table D2, Column (1)). Over 2013–2018, a 1 ppt increase in the regional supply shock raises the effective electricity price by 1.35% and reduces electricity use by 0.81%. We define the effective electricity price as the observed price plus a scarcity component equal to the Lagrange multiplier on the rationing constraint when it binds (and zero otherwise). Under this definition, the estimated cross-price elasticities are −0.17 for labor and −0.13 for materials, and the revenue elasticity with respect to the effective electricity price is −0.17.

The structural model yields elasticities that closely track these reduced-form benchmarks. Table D2, Column (2), reports a model-implied own-price elasticity of −0.67, given by $-\sigma + (\sigma - \eta)s_E^C$ (Appendix D-1.2). The model further implies a cross-price elasticity for non-electricity inputs of $(\sigma - \eta)s_E^C = -0.16$ and a revenue elasticity of $-\eta s_E^C = -0.17$.[18]

[18]Our own price elasticity estimates are in line with quasi-experimental evidence and are closest to von Graevenitz and Rottner (2024) (about −0.4 to −0.6). Prior studies report industrial own-price elasticities ranging from −0.09 to −0.20 (Gerster and Lamp, 2024) to as large as −1.5 under substantial price shocks (Martin et al., 2014).

8 Policy Counterfactuals

In this section, we use our estimated production function to evaluate policy counterfactuals that operate through market structure, given the limited plant-level adaptability we documented in the reduced-form analysis.

8.1 Decomposing the Effects of Rationing and Pricing

We decompose the total effect of the electricity supply shock during the post-phase-out period (2013–2018) into two channels, pricing and rationing. This distinction matters for policy design because mitigation depends on which channel drives plants' responses.

We implement the decomposition using the following log-difference identity:

$$\underbrace{\ln Y(p,\mathcal{S}) - \ln Y(p_{2010},\mathcal{S}_{2010})}_{\text{Total Impact}} = \underbrace{\ln Y(p,\mathcal{S}) - \ln Y(p,\mathcal{S}_{2010})}_{\text{Rationing Contribution}} + \underbrace{\ln Y(p,\mathcal{S}_{2010}) - \ln Y(p_{2010},\mathcal{S}_{2010})}_{\text{Price Contribution}}. \tag{21}$$

We use each component of (21) as the dependent variable in Equation (3). The rationing term varies rationing severity ($\mathcal{S}_{2010} \to \mathcal{S}$) holding prices fixed at p, while the price term varies prices ($p_{2010} \to p$) holding rationing fixed at $\mathcal{S}_{2010}$.

The decomposition indicates that pricing accounts for most of the aggregate decline. Panel B of Table 7 reports a total effect on electricity consumption of -0.80. The price component is -0.77, about 96% of the total, while the rationing component is small at -0.03. A similar pattern holds for labor, materials, and output, where the price component explains roughly 95% of the simulated decline. The small rationing component reflects that physical constraints were rarely binding for manufacturing plants. In our data, only 8.3% of plants exhibit strictly positive shadow values on the rationing constraint.

These results imply that the costs of the nuclear shutdown operated primarily through higher electricity generation costs from substituting nuclear with more expensive fossil generation, rather than through binding quantity constraints; this contrasts with the residential sector, where non-price interventions can meaningfully reduce electricity demand during scarcity (e.g., Reiss and White, 2008; He and Tanaka, 2023).

8.2 Counterfactuals: Nuclear Restart and Market Integration

We use the estimated production framework to evaluate supply-side interventions that mitigate the economic costs of electricity supply disruptions. Because most revenue losses

operate through higher electricity prices, we focus on market reforms that lower marginal generation costs.

We study two policy-relevant scenarios central to Japan's energy debate. The first accelerates nuclear restarts, consistent with the 7$^{\text{th}}$ Strategic Energy Plan's shift from reducing nuclear dependence toward greater utilization of existing reactors. The second considers nationwide market integration. Historically limited interregional transmission has segmented Japan's electricity market, creating dispatch inefficiencies—high-marginal-cost generation runs in importing regions while lower-cost capacity elsewhere remains idle (Hausman, 2025). Fragmentation also limits reserve sharing, requiring redundant reserves within each region and raising system-wide costs. We quantify the gains from relaxing these constraints through interregional trade that lowers marginal generation costs and through more efficient reserve sharing.

Simulation framework. We measure welfare using producer surplus. The aggregate change is

$$\Delta\Pi \equiv \sum_i \left[\Pi_i\left(p^E_{\text{cf}}\right) - \Pi_i\left(p^E_{\text{base}}\right)\right],$$

where $\Pi_i(p^E)$ denotes plant i's producer surplus as a function of the electricity price p^E. Log-linearizing around the baseline and using the model-implied approximation (Appendix D-1), we obtain

$$\Delta\Pi \approx \sum_i \Pi_i\left(p^E_{\text{base}}\right)(1-\eta)\, s^C_{E,i}\, \Delta \ln p^E, \tag{22}$$

where $s_{E,i}$ is the plant-specific electricity cost share and $\Delta \ln p^E$ is the policy-induced change in electricity prices.

Constructing counterfactual prices. We first isolate the component of post-Fukushima retail price changes attributable to the nuclear suspension (Appendix E-1.1). Because retail prices also tracked global fossil-fuel prices—rising through 2014 and falling in 2015–2016 (Figure 1, Panel C)—we use regional and temporal variation in our shock exposure to attribute differential price movements to the nuclear suspension. We estimate that the shock raised retail electricity prices by 8.9% on average (Panel A in Table 8), with larger effects in regions more dependent on nuclear generation prior to the accident (Table E1, Column (1)).

For the nuclear-restart scenario, we scale the nuclear-related shock component by region-specific restart rates and convert the implied reduction in the shock into counterfactual prices using the estimated shock–price mapping (Appendix E-1.2). We classify

regions into feasibility tiers to reflect heterogeneity in seismic risk, regulatory requirements, and local acceptance. Restoring 33.6% of pre-accident national capacity implies a 2.0% decline in the national average retail price relative to the status quo, concentrated in regions with high restart potential (Table E1, Column (2)).

For the grid-integration scenario, we construct counterfactual prices under relaxed transmission constraints and reserve sharing (Appendix E-1.3). We estimate regional supply curves and solve for integrated-grid equilibrium prices that allow lower-cost generators to serve a wider market and enable reserve sharing, thereby relaxing effective capacity constraints. Full integration reduces the national average electricity price by 3.2%.[19] The net price increase relative to the no-shock baseline falls from 8.9% to 5.7%, with large declines in high-cost demand centers such as Tokyo and offsetting increases in low-cost exporting regions (Table E1, Column (3)).

Results. Table 8 summarizes the simulation results. Under the baseline structural simulation (Column (2)), electricity consumption declines by 6.0%, implying reductions in real output and producer surplus of 1.6% and 1.4%, respectively (Panel B). Evaluated at 2010 levels, the output decline corresponds to an annual loss of 4,726 billion JPY, closely aligning with reduced-form revenue losses of 4,920 billion JPY.[20] Producer surplus falls by 656 billion JPY. Although these magnitudes are close to the reduced-form estimates (Column (1)), the slightly smaller structural impacts reflect partial price-mediated adjustment and—consistent with rationing rarely binding in the data—the absence of scarcity rents.

Accelerating nuclear restarts substantially mitigates these costs. Under the nuclear restart scenario (Column (3)), retail electricity price increases are smaller and the decline in electricity use is correspondingly reduced. Annual output losses fall to 3,649 billion JPY and producer surplus losses to 507 billion JPY, a 23% reduction relative to the baseline. Grid integration with reserve sharing yields comparable gains (Column (4)): full integration reduces annual output losses to 3,428 billion JPY and producer surplus losses to 476 billion JPY, a 27.5% reduction.

Taken together, these counterfactuals quantify the costs of slow nuclear reactivation and limited interregional coordination. Either partial nuclear restarts or grid integration recovers roughly one quarter of the baseline losses, indicating that the post-Fukushima

[19]The implied price effects are broadly consistent with existing evidence on market integration. Gonzales et al. (2023) estimate that Chile's 2017 integration reform reduced generation costs by 3–7% in the short run (and by 8–18% when accounting for investment responses). Wolak (2015) finds that relieving transmission congestion in New Zealand lowers wholesale prices by roughly 3% on average. Our estimates are smaller than those in Ryan (2021), who find larger gains in settings with severe transmission shortages.

[20]These correspond to approximately 50–60 billion USD, using an exchange-rate range of 80–95 JPY per USD (around 2010).

contraction reflected not only reduced nuclear generation but also constraints that limited efficient system-wide dispatch and reserve sharing.

8.3 Allocative Inefficiency of Rationing

We compare the welfare effects of pricing and rationing to quantify the costs of misallocation. Although rationing played a secondary role in our setting, the comparison is informative because quantity-based regulation is common in environmental and resource management (Ryan and Sudarshan, 2022; Rafey, 2023; Mansur and Olmstead, 2012). Rationing violates the equi-marginal principle: uniform mandates force plants with high shadow values to curtail electricity use excessively, while low-shadow-value plants continue to consume electricity that could be deployed more productively elsewhere. In our setting, heterogeneity in shadow values reflects industry-specific technologies and plant-level capital underutilization arising from predetermined capital–electricity ratios.

We simulate aggregate final output under pricing and rationing while holding fixed the total quantity of electricity used within each region-year. Because electricity is an intermediate input, welfare comparisons reduce to differences in final output (Appendix E-2).

We estimate a national mean ratio of 1.16, implying that rationing generates output losses 16% larger than a uniform price increase that achieves the same aggregate conservation. This magnitude is comparable to the surplus gains reported in agricultural settings (Ryan and Sudarshan, 2022; Rafey, 2023), though smaller than estimates for urban water (Mansur and Olmstead, 2012). Figure E1, Panel A, shows a right-skewed distribution, with the ratio exceeding 1.5 in several region-years. Panel B shows that the ratio is higher in regions with larger electricity cost shares, consistent with larger allocative distortions in electricity-intensive regions.

9 Conclusion

This paper studies how manufacturing responds to a large and persistent contraction in electricity supply following the 2011 Fukushima nuclear accident. Energy-supply disruptions remain a recurring risk for energy-intensive modern economies and are likely to become more salient amid geopolitical tensions and the clean energy transition (IEA, 2025). Yet evidence on how modern industries adjust to large and persistent energy shocks remains limited.

Using restricted plant-level administrative data and shift-share instruments for regional exposure, we show that a 6% decline in manufacturing electricity use lowered

manufacturing revenue by 1.7%, implying annual losses of roughly 4.9 trillion JPY despite electricity's small (2.0%) expenditure share. A central contribution is a systematic assessment of adjustment margins: across technology upgrading, product-mix changes, fuel substitution, and reallocation, we find limited adaptation for the typical plant, with meaningful expansion in self-generation concentrated among a subset of producers.

To interpret why a small-share input can generate large output losses, we develop and estimate a structural production model in which electricity is strongly complementary with other inputs and scarcity induces endogenous capital idleness. The model matches the magnitude of the reduced-form revenue losses and provides a disciplined framework for evaluating mitigation policies. Consistent with limited plant-level adjustment, our counterfactuals indicate sizable gains from electricity-market reforms that lower marginal generation costs and improve interregional allocation under supply disruptions.

References

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Agency for Natural Resources and Energy (2022). Annual report on energy (energy white paper) 2022: Part i, chapter 1. Technical report, Ministry of Economy, Trade and Industry. Japanese original: .

Allcott, H., Collard-Wexler, A., and O'Connell, S. D. (2016). How do electricity shortages affect industry? evidence from india. *American Economic Review*, 106.

Allcott, H. and Keniston, D. (2018). Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. *The Review of Economic Studies*, 85(2):695–731.

Arezki, R., Fetzer, T., and Pisch, F. (2017). On the comparative advantage of us manufacturing: evidence from the shale gas revolution. *Journal of International Economics*, 107:34–59.

Atkeson, A. and Kehoe, P. J. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, 89.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A., and Friedman, B. M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings papers on economic activity*, 1997(1):91–157.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster- robust inference. *Journal of Human Resources*, 50.

Canay, I. A., Santos, A., and Shaikh, A. M. (2021). The wild bootstrap with a "small" number of "large" clusters. *Review of Economics and Statistics*, 103.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

Chen, Q., Chen, Z., Liu, Z., Suárez Serrato, J. C., and Xu, D. Y. (2025). Regulating conglomerates: Evidence from an energy conservation program in china. *American Economic Review*, 115(2):408–447.

Cisse, A. (2025). The value of electricity reliability: Evidence from a natural experiment in senegal. Working paper.

Cole, M. A., Elliott, R. J., Occhiali, G., and Strobl, E. (2018). Power outages and firm performance in sub-saharan africa. *Journal of Development Economics*, 134.

Davis, S. J. and Haltiwanger, J. (2001). Sectoral job creation and destruction responses to oil price changes. *Journal of Monetary Economics*, 48.

Feyrer, J., Mansur, E. T., and Sacerdote, B. (2017). Geographic dispersion of economic shocks: Evidence from the fracking revolution. *American Economic Review*, 107(4):1313–1334.

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

Fontagné, L., Martin, P., and Orefice, G. (2024). The many channels of firm's adjustment to energy shocks: evidence from france. *Economic Policy*, 39(117):5–43.

Fried, S. and Lagakos, D. (2023). Electricity and firm productivity: A general-equilibrium approach. *American Economic Journal: Macroeconomics*, 15(4):67–103.

Gandhi, A., Navarro, S., and Rivers, D. A. (2020). On the identification of gross output production functions. *Journal of Political Economy*, 128(8):2973–3016.

Gerster, A. and Lamp, S. (2024). Energy tax exemptions and industrial production. *The Economic Journal*, 134(663):2803–2834.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110.

Gonzales, L. E., Ito, K., and Reguant, M. (2023). The investment effects of market integration: Evidence from renewable energy expansion in chile. *Econometrica*, 91(5):1659–1693.

Hamilton, J. D. (1983). Oil and the macroeconomy since world war ii. *Journal of Political Economy*, 91.

Harasztosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, 109(8):2693–2727.

Hausman, C. (2025). Power flows: Transmission lines, allocative efficiency, and corporate profits. *American Economic Review*, 115(8):2574–2615.

He, G. and Tanaka, T. (2023). Energy saving may kill: evidence from the fukushima nuclear accident. *American Economic Journal: Applied Economics*, 15(2):377–414.

IEA (2025). As risks multiply in a world thirsty for energy, diversification and cooperation are more urgent than ever.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of economic literature*, 46(4):871–909.

Lee, D. S., McCrary, J., Moreira, M. J., and Porter, J. (2022). Valid t-ratio inference for iv. *American Economic Review*, 112(10):3260–3290.

Mansur, E. T. and Olmstead, S. M. (2012). The value of scarce water: Measuring the inefficiency of municipal regulations. *Journal of Urban Economics*, 71(3):332–346.

Marin, G. and Vona, F. (2021). The impact of energy prices on socioeconomic and environ-

mental performance: Evidence from french manufacturing establishments, 1997–2015. *European Economic Review*, 135.

Martin, R., de Preux, L. B., and Wagner, U. J. (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, 117.

METI (2012). 2010 census of manufacture: Summary report. https://www.meti.go.jp/statistics/tyo/kougyo/result-2/h22/gaiyo/index.html. Accessed: 2025-05-31.

MIC (2025). Statistical observations of prefectures 2025. Statistics Bureau, Ministry of Internal Affairs and Communications (Japan). Accessed: 2025-06-25.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Porter, M. E. (1991). America's green strategy. In *Scientific American*, page 168.

Rafey, W. (2023). Droughts, deluges, and (river) diversions: Valuing market-based water reallocation. *American Economic Review*, 113(2):430–471.

Reiss, P. C. and White, M. W. (2008). What changes energy consumption? prices and public pressures. *RAND Journal of Economics*, 39:636–663.

RIETI (2005). R-jip2005 prefectural input–output tables. Research Institute of Economy, Trade and Industry (RIETI), Japan. Accessed: 2025-06-25.

Roodman, D., MacKinnon, J. G., Ørregaard Nielsen, M., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *Stata Journal*, 19.

Rotemberg, J. J. and Woodford, M. (1996). Imperfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit & Banking*, 28(4):549–578.

Ryan, N. (2021). The competitive effects of transmission infrastructure in the indian electricity market. *American Economic Journal: Microeconomics*, 13(2):202–242.

Ryan, N. and Sudarshan, A. (2022). Rationing the commons. *Journal of Political Economy*, 130(1):210–257.

Stock, J. H. and Yogo, M. (2003). Testing for weak instruments in linear iv regression (book chapter: 6. asymptotic distributions of instrumental variables statistics with many instruments). *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 2001:80–108.

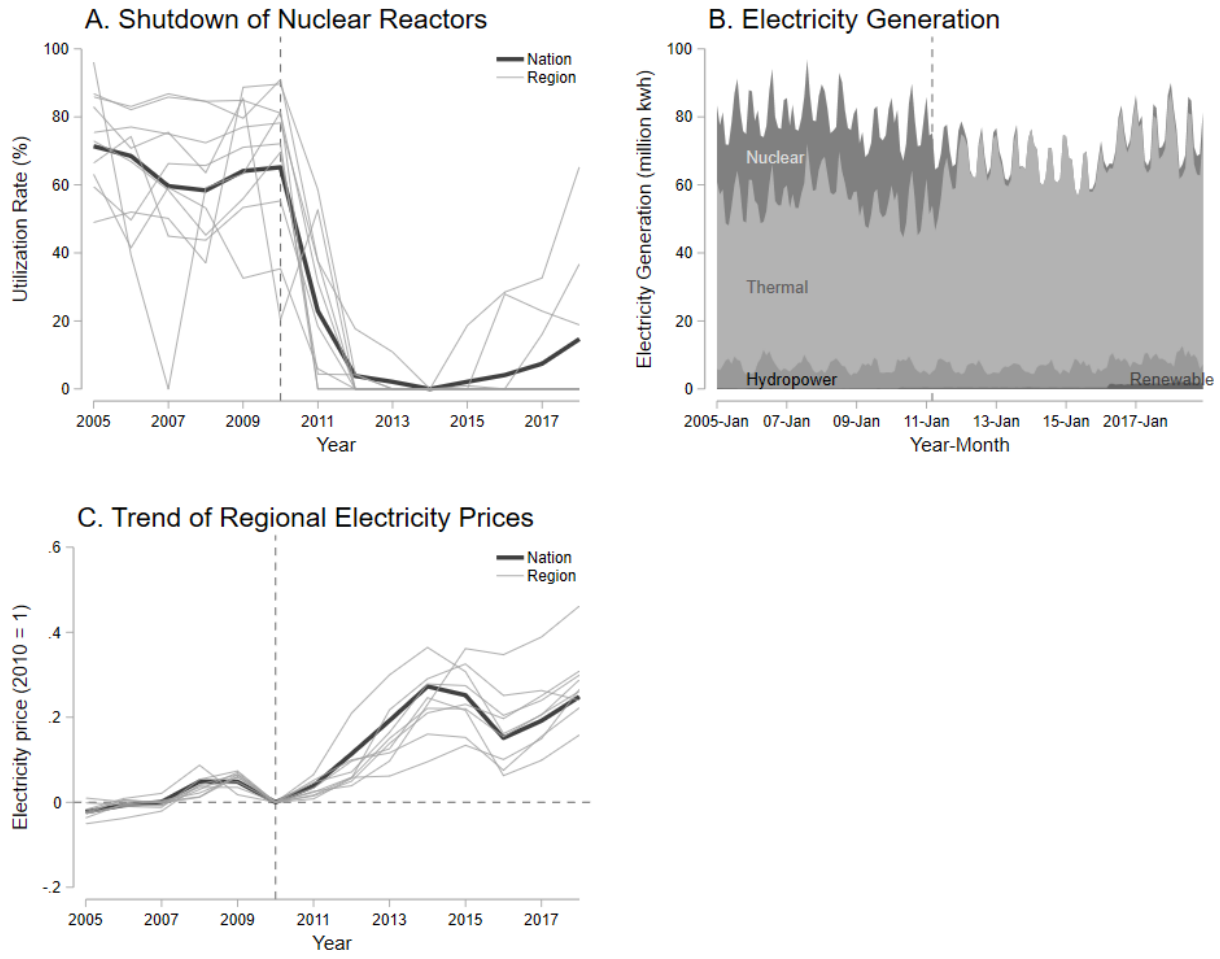
Todd, P. E. and Wolpin, K. I. (2006). Assessing the impact of a school subsidy program in mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American economic review*, 96(5):1384–1417.

von Graevenitz, K. and Rottner, E. (2024). Climate policies and electricity prices: To abate or to generate? Technical report, University of Bonn and University of Mannheim, Germany.

Wolak, F. A. (2015). Measuring the competitiveness benefits of a transmission investment policy: The case of the alberta electricity market. *Energy policy*, 85:426–444.

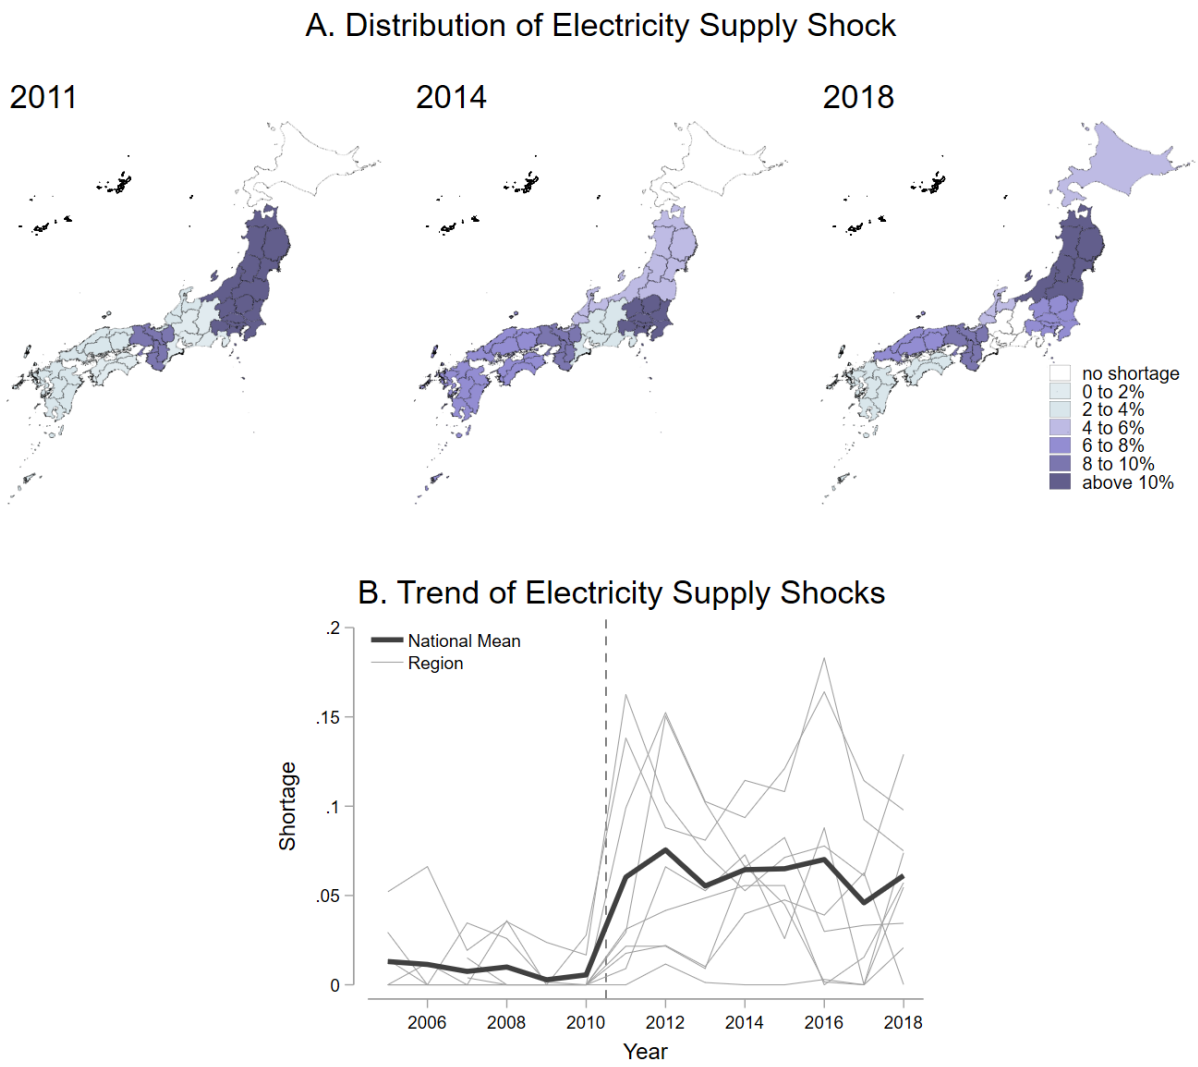
Figures and Tables

Figure 1: Nuclear shutdown and electricity supply shocks



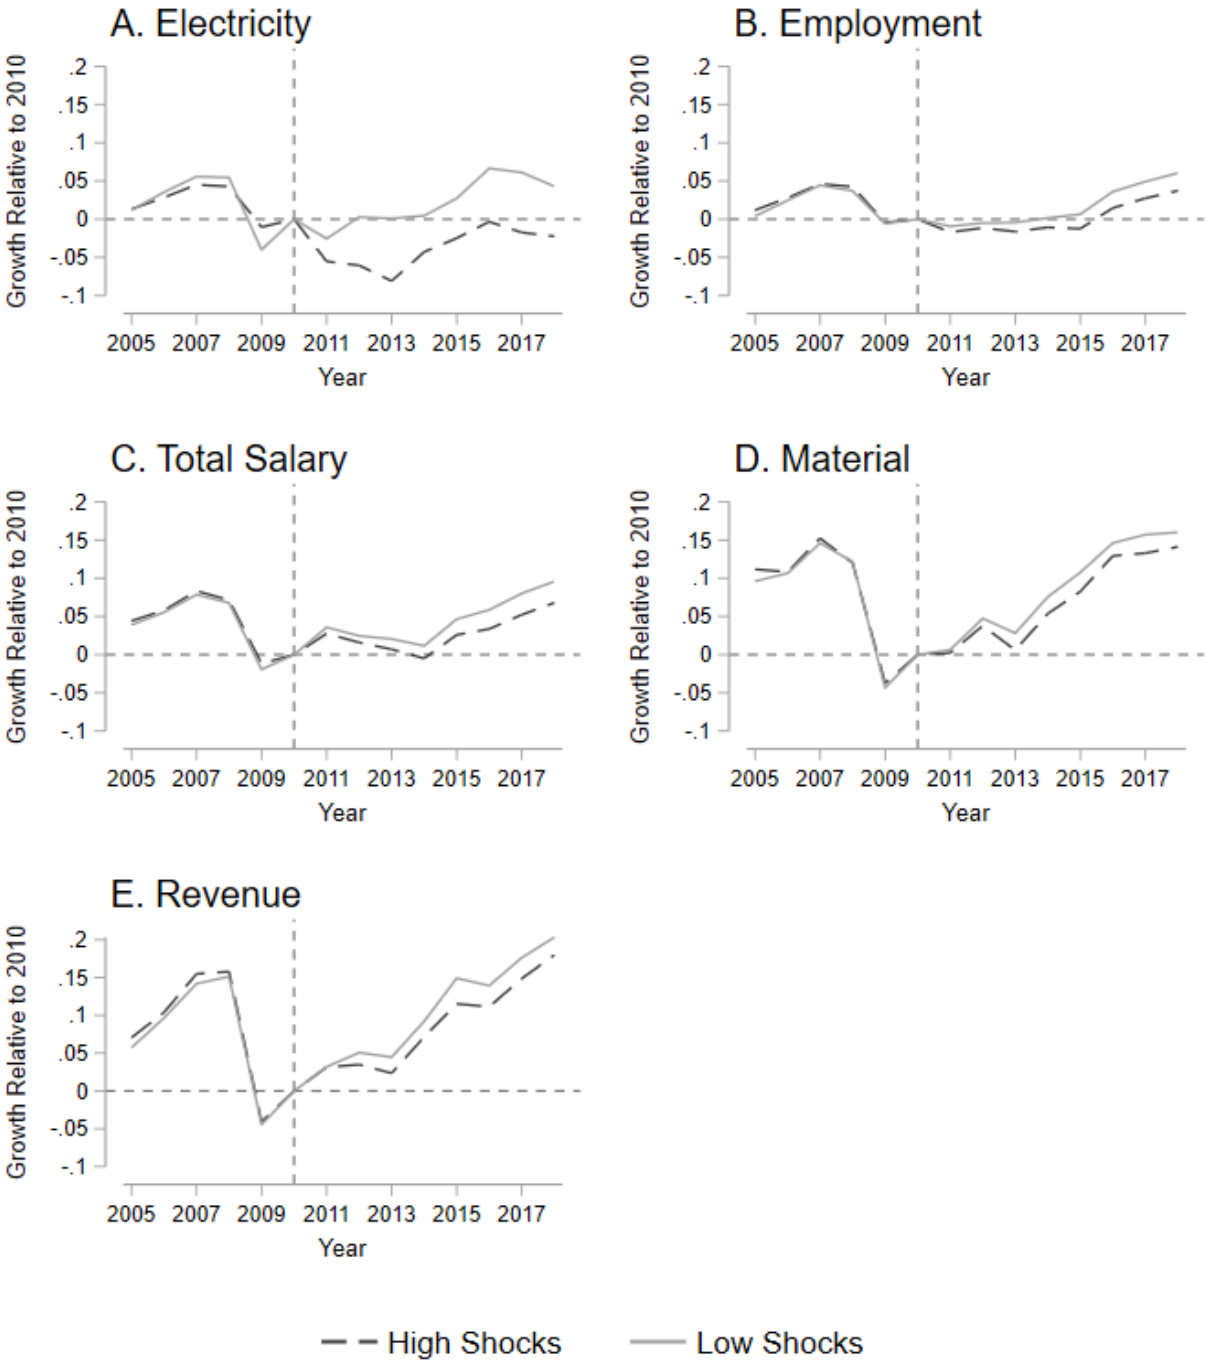
Notes: Panel A shows the utilization of nuclear reactors, defined as power generation divided by supply capacity. Panel B depicts the trend in monthly national electricity demand by energy source. Panel C presents the trend in average electricity prices. In all panels, the vertical dotted line indicates the timing of the Fukushima nuclear accident. In Panels A and C, the bold line represents the national trend, while the thin gray lines represent individual regions.

Figure 2: Distribution and trend of electricity supply shocks



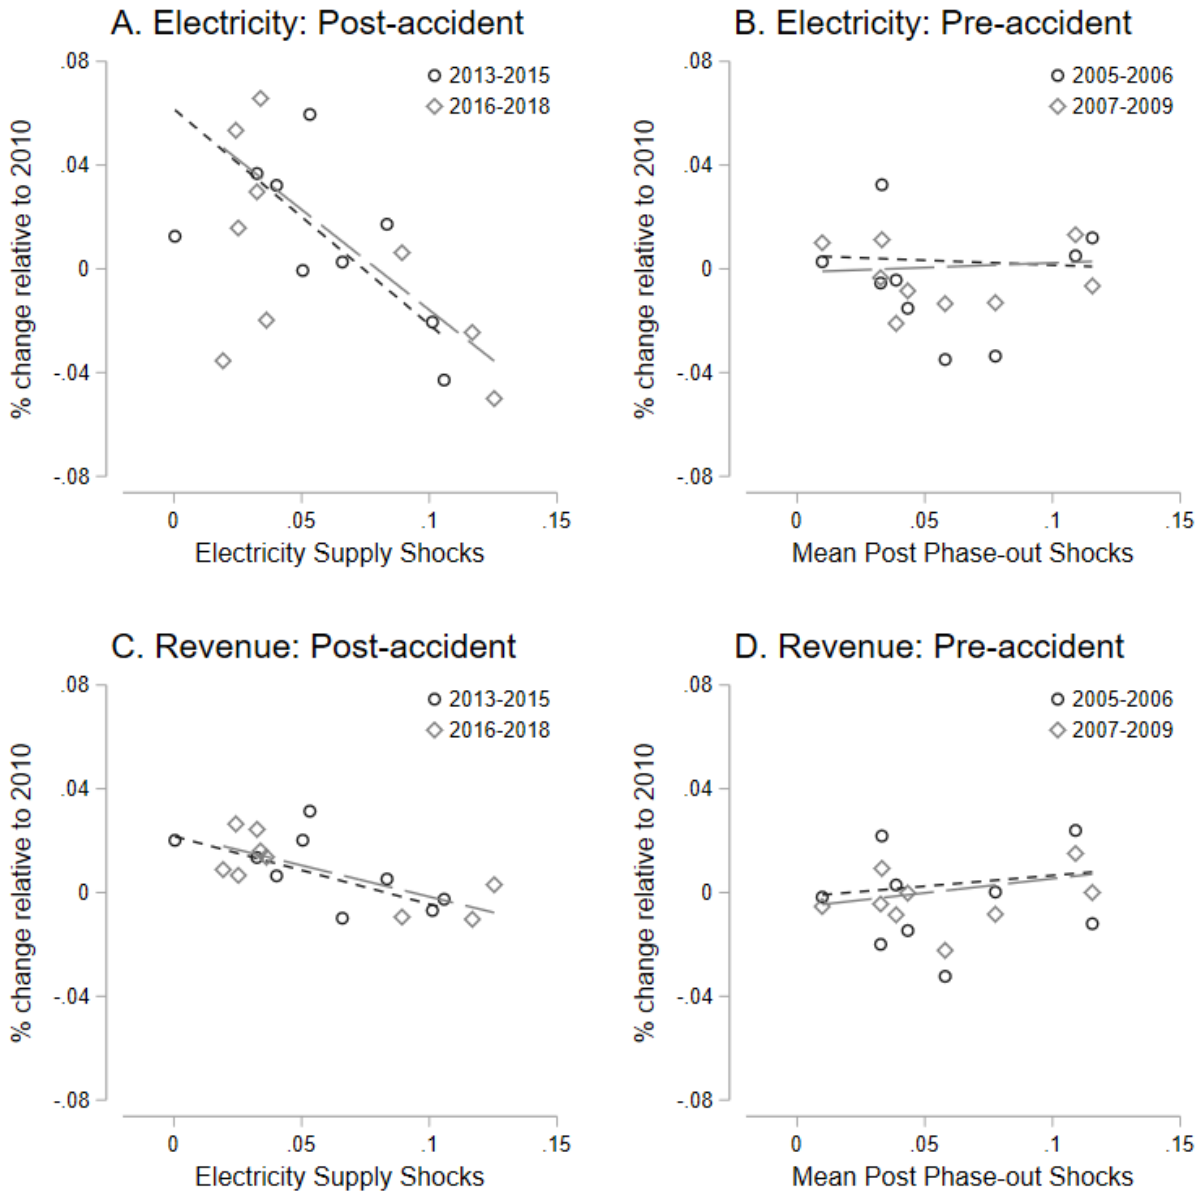
Notes: Panel A presents the geographic distribution of our measure of electricity supply shocks from 2011 to 2018. Panel B shows the corresponding trends, where the bold line represents the national average and the thin gray lines represent individual regions. See the main text for the definition of electricity supply shocks.

Figure 3: Trends of various measures of manufacturing performances between high-shock and low-shock regions



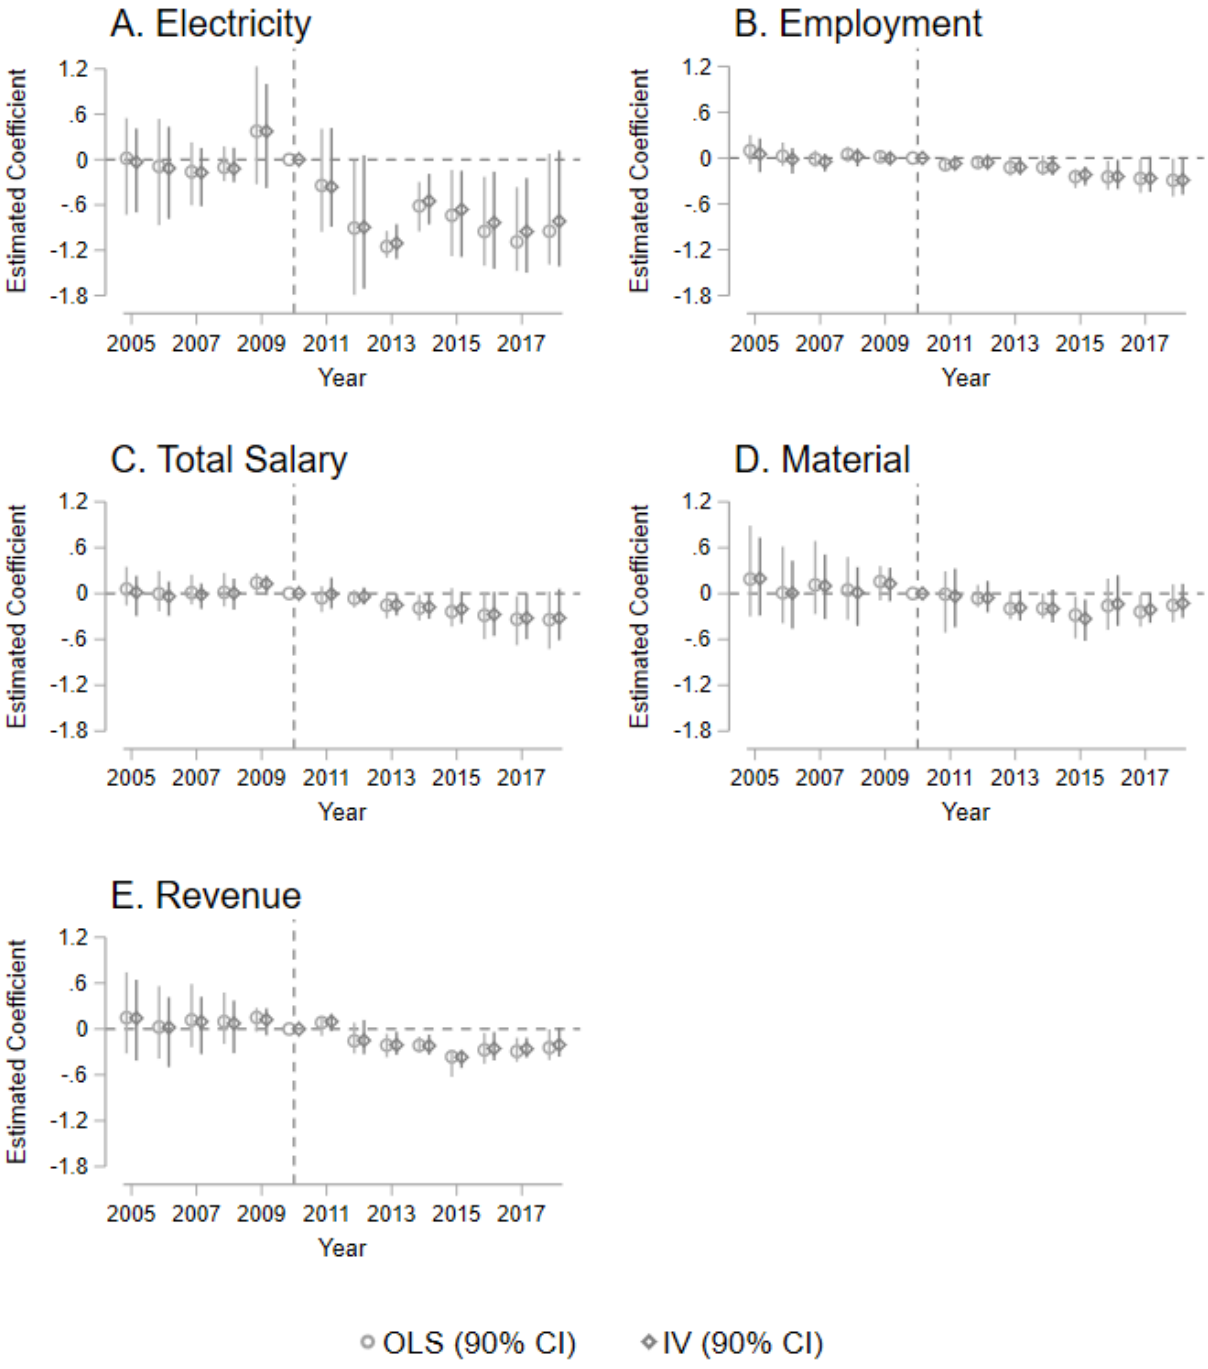
Notes: We classify nine regions into high- and low-shock regions using the median average supply shocks during the post-accident periods. The value is normalized using the 2010 value, so the y-axis represents the percentage change relative to the 2010 value.

Figure 4: Correlation between regional supply shock and power usage and revenue



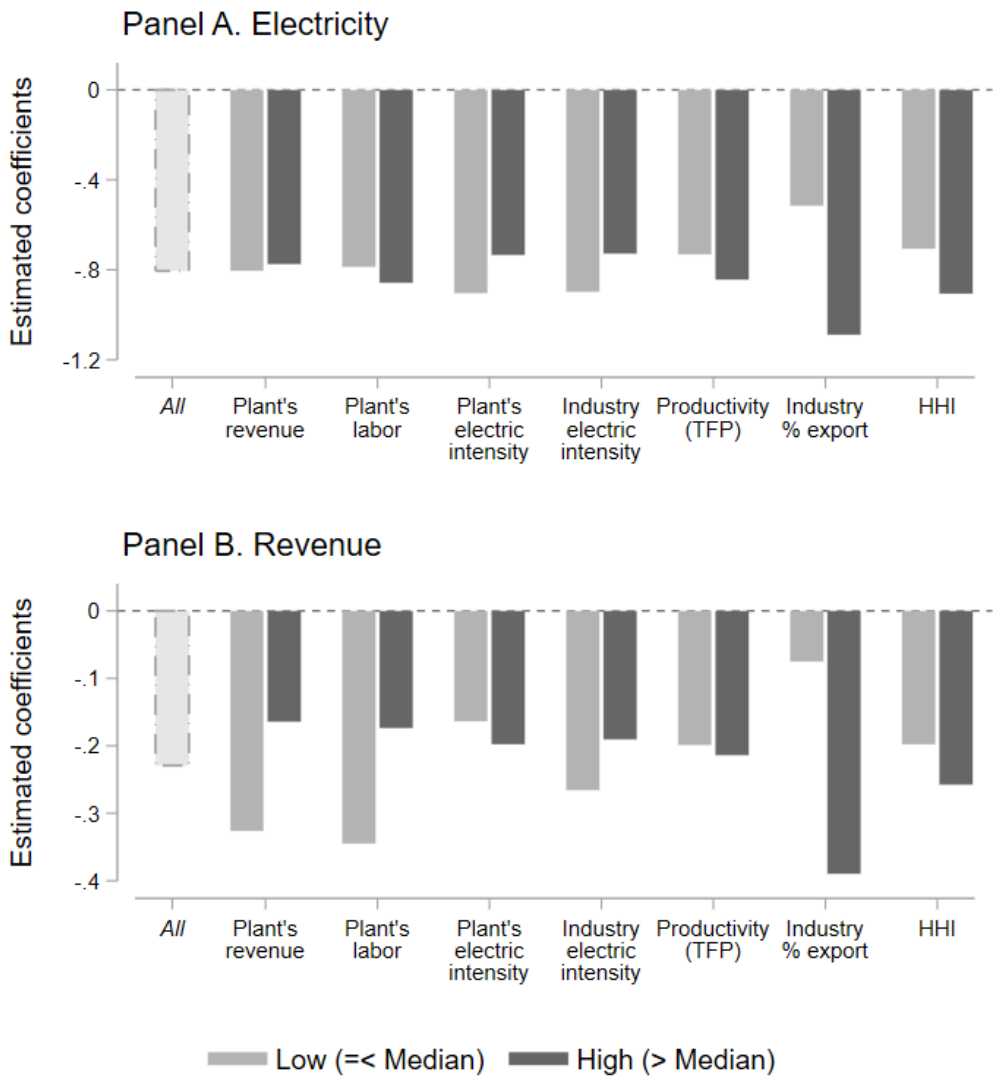
Notes: We first take residuals in outcomes controlling for year-by-industry (3 digits) fixed effects. We then plot the differences in log residual relative to 2010 on the Y-axis and our measures of supply shock on the X-axis. Panel A and B use electricity, and Panels C and D use revenue as outcomes. Panels B and D use mean supply shocks during the post-phase-out period (2013-2018) on the x-axis.

Figure 5: Event study results of effects of supply shock on various manufacture's activities



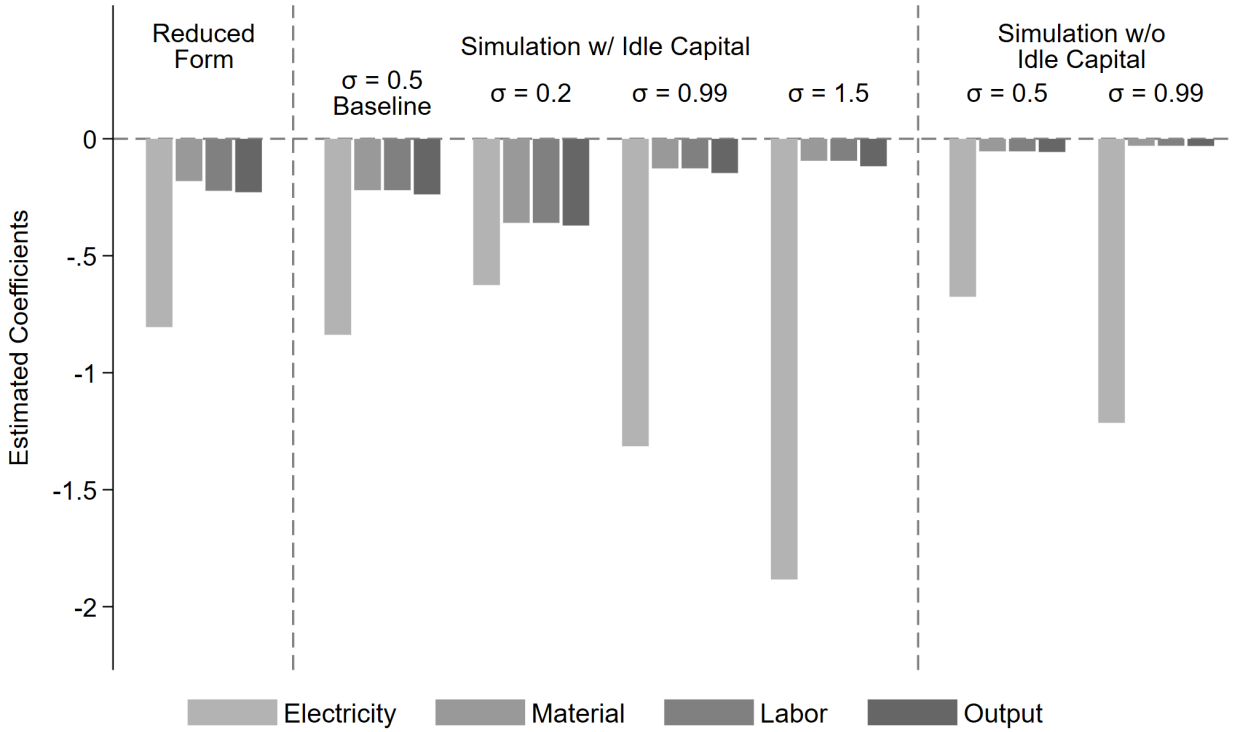
Notes: We plot the regression results by fitting equation (3). Each dot corresponds to a separate regression, in which grey circles denote OLS estimators and black diamonds denote 2SLS estimators. During the pre-accident periods, we use the average shocks in each region during the post-phase-out periods. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap to obtain P-values using regions as clusters.

Figure 6: Heterogeneity impacts of electricity supply shocks



Notes: The figure plots coefficients from separate regressions of equation (3). The leftmost bars report the full-sample estimates, and the light and dark bars report estimates for plants below and above the median of the indicated 2010 characteristic. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap to obtain P-values using regions as clusters.

Figure 7: Reduced-Form Estimates and Structural Simulations



Notes: The figure compares the baseline reduced-form elasticity estimates (leftmost panel) with coefficients derived from structural model simulations. The bars represent the estimated response of electricity, material, labor, and output to an electricity supply shock. The "Simulation w/ Idle Capital" panels display results from the model allowing for variable capacity utilization (Clay) across varying values of the elasticity of substitution ($\sigma \in \{0.2, 0.5, 1.0, 1.5\}$), with $\sigma = 0.5$ serving as the baseline. The "Simulation w/o Idle Capital" panel (far right) displays results from the model with flexible capital (Putty) using the baseline parameter $\sigma = 0.5$. Simulation estimates are obtained by applying the same instrumental variable (IV) specification used in the empirical analysis to the simulated datasets.

Table 1: Summary Statistics

	Mean (1)	S.D. (2)
Panel A. Plant Characteristics		
Revenue (million yen)	5,040	11,427
Employment (headcount)	127	165
Capital stock (million yen)	1,013	2,447
Export share (%)	2.8	11
Multi-plants (yes = 1)	0.42	0.49
Panel B. Revenue shares		
Labor	0.22	0.14
Capital	0.032	0.041
Materials	0.51	0.20
Electricity	0.02	0.023
Panel C: Plant Energy Usage		
Electricity Bought (TJ)	30	639
Self Generation (TJ)	13	285
Own Electricity Generator (yes = 1)	0.11	0.31
Fuel for Production (TJ)	76	1,684
Steam for Production (TJ)	24	388

Notes: Panels A and B include all plants in our sample constructed from the Census of Manufactures (N = 442,099). Expenditure shares in Panel B are calculated as the ratio of input expenditure to plant revenue. Panel C uses data from the Structural Survey of Energy Consumption (N = 381,891). Monetary values are deflated to 2010 prices.

Table 2: First-Stage Regression: Impact of Supply Composition on Supply Shocks

	Dep Var: Supply Shocks	
	Phase-out Terms (1)	Post-Phase-out (2)
$Sh^{Nuc}_{r,2010} \times g^{Nuc}_t$ (Nuclear)	-0.52	-0.25
	(0.00)***	(0.00)***
$Sh^{Fossil}_{r,2010} \times g^{Fossil}_t$ (Fossil)	-0.42	-0.26
	(0.03)**	(0.01)***
$Sh^{Hydro}_{r,2010} \times g^{Hydro}_t$ (Hydro)	0.74	-0.42
	(0.56)	(0.05)*
$Sh^{Renew}_{r,2010} \times g^{Renew}_t$ (Renewable)	1.51	0.14
	(0.09)*	(0.17)
$Sh^{RM}_{r,2010}$ (Reserve margin)	-1.15	-0.50
	(0.00)***	(0.00)***
Year FE	Yes	Yes
Observations	18	54

Notes: The table reports the first-stage regression estimates. The dependent variable is the electricity supply shocks. The independent variables are shift-share instruments defined as the 2010 regional generation shares for each source interacted with the leave-one-out national growth rates. Column (1) covers the nuclear phase-out period (2011–2012), and Column (2) covers the post-adjustment period (2013–2018). P-values are in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Effects of Supply Shocks on Inputs and Outputs

	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$				
	Electricity (1)	Employ-ment (2)	Total Salary (3)	Material (4)	Revenue (5)
Panel A: The effects of supply shocks using OLS					
Supply Shocks	-0.25	-0.02	-0.01	-0.04	-0.03
(Phase-out Terms)	[0.10]	[0.46]	[0.71]	[0.41]	[0.29]
Supply Shocks	-0.63**	-0.15**	-0.18**	-0.14	-0.18***
(Post-Phase-out)	[0.03]	[0.02]	[0.04]	[0.15]	[0.00]
Panel B: The effects of supply shocks using 2SLS					
Supply Shocks	-0.30	-0.02	0.01	-0.03	-0.03
(Phase-out Terms)	[0.22]	[0.50]	[0.85]	[0.63]	[0.40]
Supply Shocks	-0.81**	-0.19**	-0.22*	-0.18	-0.23***
(Post-Phase-out)	[0.02]	[0.04]	[0.05]	[0.16]	[0.00]
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. The "Phase-out Terms" row uses data from 2005–2010 and 2011–2012. The "Post-Phase-out" row uses data from 2005–2010 and 2013–2018. All regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 4: Effects of Supply Shocks on Plants' Behaviors and Productivity

	Productivity	Product-level		Outsourcing	
	TFPR	Switching	Intensity	Outsourcing	Share in Mat.
	(1)	(2)	(3)	(4)	(5)
Panel A: The effects of supply shocks using OLS					
Supply Shocks	0.01	0.09	0.02	0.00	0.00
(Phase-out Terms)	[0.91]	[0.24]	[0.50]	[0.99]	[0.79]
Supply Shocks	-0.01	0.04	0.03	-0.05	-0.00
(Post-Phase-out)	[0.83]	[0.62]	[0.29]	[0.89]	[0.87]
Panel B: The effects of supply shocks using 2SLS					
Supply Shocks	-0.00	0.09	0.02	0.05	0.00
(Phase-out Terms)	[0.97]	[0.23]	[0.61]	[0.77]	[0.95]
Supply Shocks	0.00	0.02	0.05	-0.12	-0.01
(Post-Phase-out)	[0.99]	[0.92]	[0.26]	[0.84]	[0.85]
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. Column (1) uses log differences in $\ln \Omega_{it}$ relative to 2010. Column (2) uses a dummy variable for product switching. Columns (3) and (4) use log differences. Column (5) uses the materials share relative to 2010. The "Phase-out Terms" row uses data from 2005–2010 and 2011–2012. The "Post-Phase-out" row uses data from 2005–2010 and 2013–2018. All regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 5: Effects of Supply Shocks on Energy Usage

	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2009})$		
	Fuel Consumption for:		
	Production (1)	Self-generation (2)	Electricity Generated (3)
Panel A: The effects of Supply Shocks using OLS			
Supply Shocks (Phase-out Terms)	0.04 [0.83]	2.05*** [0.00]	1.45*** [0.00]
Supply Shocks (Post-Phase-out)	-0.17 [0.58]	1.78* [0.06]	0.64* [0.10]
Panel B: The effects of Supply Shocks using 2SLS			
Supply Shocks (Phase-out Terms)	0.03 [0.90]	2.25** [0.01]	1.65** [0.01]
Supply Shocks (Post-Phase-out)	-0.27 [0.54]	2.63** [0.01]	1.05 [0.10]
Year-by-Industry FE	Yes	Yes	Yes

Notes: This table uses the ECS data and reports regression estimates. Columns (1) and (2) report fuel consumption for production and self-generation, respectively. Column (3) reports total electricity generation. The Phase-out Terms row uses data from fiscal years 2005–2009 and 2011–2012. The Post-Phase-out row uses data from fiscal years 2005–2009 and 2013–2018. All regressions include year-by-industry (2-digit) fixed effects and use 2009 employment weights. We compute p-values using a wild cluster bootstrap with regions as clusters; p-values are reported in brackets. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 6: Distribution of Estimated Output Elasticities

	Structural Estimates			Benchmark
	Mean	p25	p75	OLS
	(1)	(2)	(3)	(4)
Labor (θ_L)	0.181	0.138	0.214	0.429
Materials (θ_M)	0.452	0.397	0.512	0.497
Energy (θ_E)	0.014	0.008	0.017	0.061
Capital (θ_K)	0.137	0.083	0.172	0.038

Notes: This table summarizes the distribution of output elasticities estimated using the Gandhi, Navarro, and Rivers (2020) method. Columns (1)–(3) report the mean and the 25th and 75th percentiles of the structural estimates across 153 industries. Column (4) provides a benchmark based on industry-specific OLS regressions that regress log revenue on log inputs.

Table 7: Simulation Results: Reduced-Form vs. Structural Estimates

	Dep Var: $\triangle \ln(Y)$			
	Electricity (1)	Total Salary (2)	Materials (3)	Output (4)
Panel A: Reduced-Form Analysis				
Supply Shocks	-0.81	-0.22	-0.18	-0.23
	[0.02]**	[0.05]*	[0.16]	[0.00]***
Panel B: Structural Model Decompositions				
B1: Total Impact ($p, \mathcal{S}$ vs. $p_{2010}, \mathcal{S}_{2010}$)				
Supply Shocks	-0.80	-0.18	-0.18	-0.20
B2: Rationing Contribution ($p, \mathcal{S}$ vs. $p, \mathcal{S}_{2010}$)				
Supply Shocks	-0.03	-0.01	-0.01	-0.01
B3: Price Contribution ($p, \mathcal{S}_{2010}$ vs. $p_{2010}, \mathcal{S}_{2010}$)				
Supply Shocks	-0.77	-0.18	-0.18	-0.19
Industry-year FE	Yes	Yes	Yes	Yes

Notes: Panel A presents reduced-form estimates using observed data. Panel B presents structural counterfactuals where the dependent variable is the simulated log difference under specific pricing (p) and rationing ($\mathcal{S}$) regimes relative to the 2010 baseline. P-values are reported in brackets. *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Counterfactual Simulation: Manufacturing Performance Losses

		Structural Model		
	Reduced-form (1)	Baseline (2)	Nuclear Reutilization (3)	Market Integration (4)
Panel A: Price Changes relative to 2010				
Electricity ($\Delta \ln P^E$)	–	8.9%	6.9%	5.7%
Panel B: Input and Output Changes relative to 2010				
Electricity ($\Delta \ln E$)	-6.0%	-6.0%	-4.6%	-4.2%
Employment ($\Delta \ln L$)	-1.7%	-1.5%	-1.2%	-1.1%
Material Input ($\Delta \ln M$)	-1.3%	-1.5%	-1.2%	-1.1%
Real Output ($\Delta \ln Y$)	-1.7%	-1.6%	-1.3%	-1.2%
Producer Surplus ($\Delta \ln \Pi$)	–	-1.4%	-1.1%	-1.0%
Panel C: Aggregate Economic Losses (Annual)				
Output Loss (Billion JPY)	-4,920	-4,726	-3,649	-3,428
Profit Loss (Billion JPY)	–	-656	-507	-476

Notes: This table reports estimated aggregate losses for the manufacturing sector under alternative electricity supply scenarios. Column (1) reports reduced-form estimates. Column (2) reports the structural baseline simulation of the observed post-Fukushima shock. Column (3) reports the counterfactual nuclear reutilization scenario. Column (4) reports the counterfactual with full market integration and reserve sharing. All monetary values are expressed in constant 2010 Japanese yen.

ONLINE APPENDIX

Firm Responses to Electricity Shortages: The Case of the Fukushima Nuclear Accidents

Daiji Kawaguchi and Takanao Tanaka*

*Kawaguchi: Graduate School of Economics, The University of Tokyo (email: kawaguchi@e.u-tokyo.ac.jp); Tanaka: Department of Agricultural and Resource Economics, UC Berkeley (email: takanao_tanaka@berkeley.edu).

A Background

Table A1: Peak Day Electricity Supply and Demand (GW)

	Demand 2010	Supply 2010	Planned Supply		
			2012	2014	2016
	(1)	(2)	(3)	(4)	(5)
A. Tohoku Region					
Total	15.6	16.6	14.8	15.5	15.1
Thermal	-	11.9	12.5	14.0	13.5
Nuclear	-	2.5	0.0	0.0	0.0
Hydro	-	1.9	1.4	1.7	1.6
Renewable	-	0.8	0.9	0.9	1.3
Transmission	-	-0.5	-0.1	-1.0	-1.2
Supply Shock	-	-	10.3%	5.3%	7.8%
B. Tokyo Region					
Total	59.9	64.1	57.7	56.1	52.0
Thermal	-	41.5	46.4	43.8	39.2
Nuclear	-	10.7	0.0	0.0	0.0
Hydro	-	3.4	3.0	3.0	2.7
Renewable	-	8.3	8.6	9.5	10.2
Transmission	-	0.3	-0.3	-0.2	-0.1
Supply Shock	-	-	8.8%	11.4%	18.3%
C. Chubu Region					
Total	27.1	29.9	27.9	27.4	27.4
Thermal	-	21.2	23.4	23.4	21.3
Nuclear	-	2.7	0.0	0.0	0.0
Hydro	-	1.5	1.4	1.4	1.4
Renewable	-	4.1	4.1	4.4	5.1
Transmission	-	0.3	-1.1	-1.9	-0.4
Supply Shock	-	-	2.2%	4.0%	3.9%
D. Kansai Region					
Total	31.0	32.7	27.8	29.6	27.4
Thermal	-	16.8	19.2	20.5	20.3
Nuclear	-	8.4	0.0	0.0	0.0
Hydro	-	2.3	2.5	2.6	2.6
Renewable	-	4.5	2.4	4.8	4.7
Transmission	-	0.7	1.2	1.7	-0.2
Supply Shock	-	-	15.2%	9.4%	16.4%

Notes: This table reports peak-day supply and demand (in GW) for four major regions. The peak day differs across regions. Transmission refers to electricity exchanged with utilities in other regions and with non-utility suppliers; thus, transmission does not necessarily sum to zero. Table constructed by the authors.

B Main Results

B-1 Additional Results

Table B1: Effects of Supply Shocks on Labor

	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$				
		Employment			
	Salary per worker (1)	Male (2)	Female (3)	Regular (4)	Non-regular (5)
Panel A: The effects of Supply Shocks using OLS					
Supply Shocks (Phase-out Terms)	0.01 [0.43]	-0.02 [0.69]	-0.06 [0.25]	-0.04 [0.55]	-0.13 [0.48]
Supply Shocks (Post-Phase-out)	-0.03 [0.60]	-0.18** [0.01]	-0.18*** [0.01]	-0.18* [0.08]	-0.36 [0.21]
Panel B: The effects of Supply Shocks using 2SLS					
Supply Shocks (Phase-out Terms)	-0.01 [0.88]	-0.15* [0.06]	-0.20** [0.01]	-0.12** [0.03]	-0.20 [0.13]
Supply Shocks (Post-Phase-out)	-0.05 [0.59]	-0.25* [0.07]	-0.25*** [0.01]	-0.21 [0.16]	-0.49** [0.03]
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. Columns (2) and (3) refer to male and female employment, respectively. Columns (4) and (5) refer to regular and non-regular employment. The "Phase-out Terms" row uses data from 2005–2010 and 2011–2012. The "Post-Phase-out" row uses data from 2005–2010 and 2013–2018. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01$,** $p < 0.05$,* $p < 0.1$.

Table B2: Effects of Supply Shocks on Investment and Capital Stock

Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$	Capital Stock (1)	Investment (2)
Panel A: The effects of Supply Shocks using OLS		
Supply Shocks (Phase-out Terms)	0.14 [0.32]	-0.12 [0.55]
Supply Shocks (Post-Phase-out)	-0.14 [0.48]	-0.80** [0.02]
Panel B: The effects of Supply Shocks using 2SLS		
Supply Shocks (Phase-out Terms)	0.03 [0.88]	-0.52*** [0.00]
Supply Shocks (Post-Phase-out)	-0.24 [0.46]	-1.16*** [0.00]
Year-by-Industry FE	Yes	Yes

Notes: The table reports regression estimates. Column (1) refers to Capital Stock and Column (2) refers to Investment. The "Phase-out Terms" row uses data from 2005–2010 and 2011–2012. The "Post-Phase-out" row uses data from 2005–2010 and 2013–2018. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

B-2 Robustness Checks

Table B3: Robustness Checks

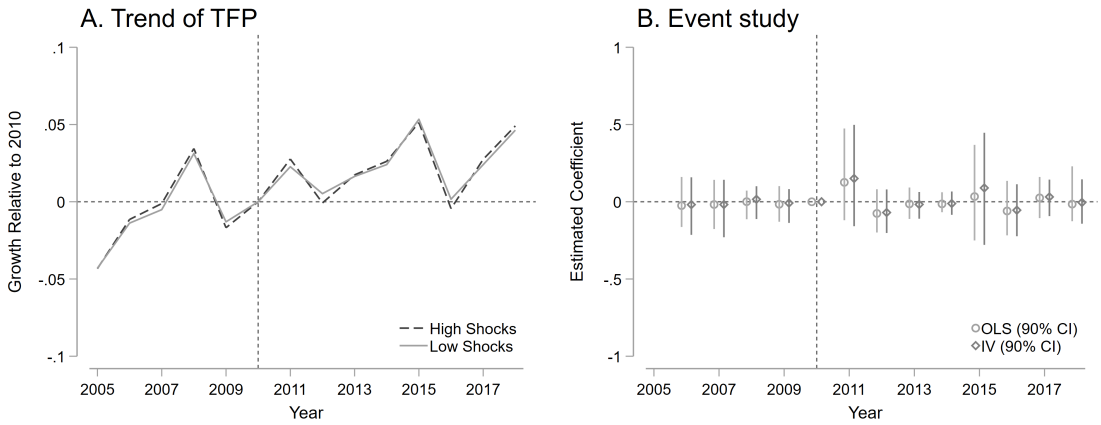
	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$					
	Direct Supply (1)	Total Supply (2)	Local Demand (3)	Plant Control (4)	Twoway FE (5)	No Weight (6)
	Dep Var: $\Delta \ln(E) = \ln(E) - \ln(E_{2010})$					
Panel A: The effects of Supply Shocks using OLS						
Supply Shocks (Phase-out Terms)	-0.26 [0.12]	-0.26 [0.12]	-0.20 [0.19]	-0.27* [0.09]	-0.17 [0.23]	-0.23 [0.12]
Supply Shocks (Post-Phase-out)	-0.67** [0.02]	-0.67** [0.03]	-0.50** [0.02]	-0.66** [0.03]	-0.48** [0.02]	-0.62** [0.03]
Panel B: The effects of Supply Shocks using 2SLS						
Supply Shocks (Phase-out Terms)	-0.33 [0.19]	-0.33 [0.19]	-0.24 [0.32]	-0.32 [0.22]	-0.21 [0.47]	-0.29 [0.21]
Supply Shocks (Post-Phase-out)	-0.87*** [0.01]	-0.87*** [0.01]	-0.68** [0.02]	-0.84** [0.02]	-0.80*** [0.01]	-0.80** [0.01]
	Dep Var: $\Delta \ln(R) = \ln(R) - \ln(R_{2010})$					
Panel C: The effects of Supply Shocks using OLS						
Supply Shocks (Phase-out Terms)	-0.02 [0.51]	-0.02 [0.45]	-0.01 [0.53]	-0.05 [0.39]	-0.05 [0.51]	-0.04 [0.22]
Supply Shocks (Post-Phase-out)	-0.20** [0.04]	-0.21** [0.03]	-0.15* [0.06]	-0.16 [0.21]	-0.18 [0.16]	-0.19*** [0.00]
Panel D: The effects of Supply Shocks using 2SLS						
Supply Shocks (Phase-out Terms)	-0.01 [0.63]	-0.02 [0.58]	-0.01 [0.92]	-0.06 [0.43]	-0.01 [0.93]	-0.04 [0.34]
Supply Shocks (Post-Phase-out)	-0.24** [0.03]	-0.24** [0.02]	-0.21* [0.07]	-0.19 [0.21]	-0.25 [0.25]	-0.24*** [0.00]
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Columns (1)–(6) correspond to the specifications: Direct Supply Chain, Total Supply Chain, adding Local Demand controls, adding Plant controls, Two-way Fixed Effects, and No Weights, respectively. All the regressions include year-by-industry (3-digits) fixed effects. For each robustness check, see the main text. We use wild cluster bootstrap to obtain P-values using region-year as clusters. P-values are in parentheses. $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

C Plants' Responses to Supply Shocks

C-1 Total Factor Productivity

Figure C1: TFPR Responses to Electricity Supply Shocks



Notes: This figure reports evidence on TFPR, constructed as $\omega_{it} = \ln \Omega_{it}$ and expressed as log changes relative to 2010. Panel A plots mean TFPR by year separately for high-shock and low-shock regions, where regions are classified using the median of average supply shocks over 2011–2018. Panel B reports year-by-year estimates from equation (3); each point comes from a separate cross-sectional regression for that year, using region-level supply shocks (and, for pre-2011 years, regions' average post-phase-out shocks), and controlling for 3-digit industry fixed effects. Circles denote OLS and diamonds denote 2SLS; vertical bars show 90% confidence intervals based on wild cluster bootstrap with regions as clusters.

C-2 Electricity-Biased Technical Change

We ask whether the post-accident decline in electricity intensity reflects (i) cost-minimizing factor substitution driven by relative input prices, or (ii) a shift in electricity-augmenting technology (directed technical change). Under a CES aggregator over electricity and labor, relative input demand can be written as

$$\ln\left(\frac{E_{it}}{L_{it}}\right) = -\sigma \ln\left(\frac{p^E_{it}}{p^L_{it}}\right) + (1-\sigma)\,\omega^E_{it}, \tag{C1}$$

where σ is the elasticity of substitution and ω^E_{it} captures non-price shifts in electricity-augmenting technology.[1] Increases in p^E/p^L reduce E/L, while changes in ω^E_{it} shift E/L beyond what prices predict.

We assess whether relative price movements can account for the observed evolution of E/L after Fukushima. Using pre-accident data (2005–2010), we estimate

$$\ln\left(\frac{E_{it}}{L_{it}}\right) = \alpha \ln\left(\frac{p^E_{it}}{p^L_{it}}\right) + \theta_i + \gamma_j \times \text{year}_t + \varepsilon_{it}, \tag{C2}$$

where θ_i are plant fixed effects and $\gamma_j \times \text{year}_t$ are industry-specific linear trends. Under (C1), $\alpha = -\sigma$; because relative prices may be endogenous, we treat $\hat{\sigma}$ as descriptive.

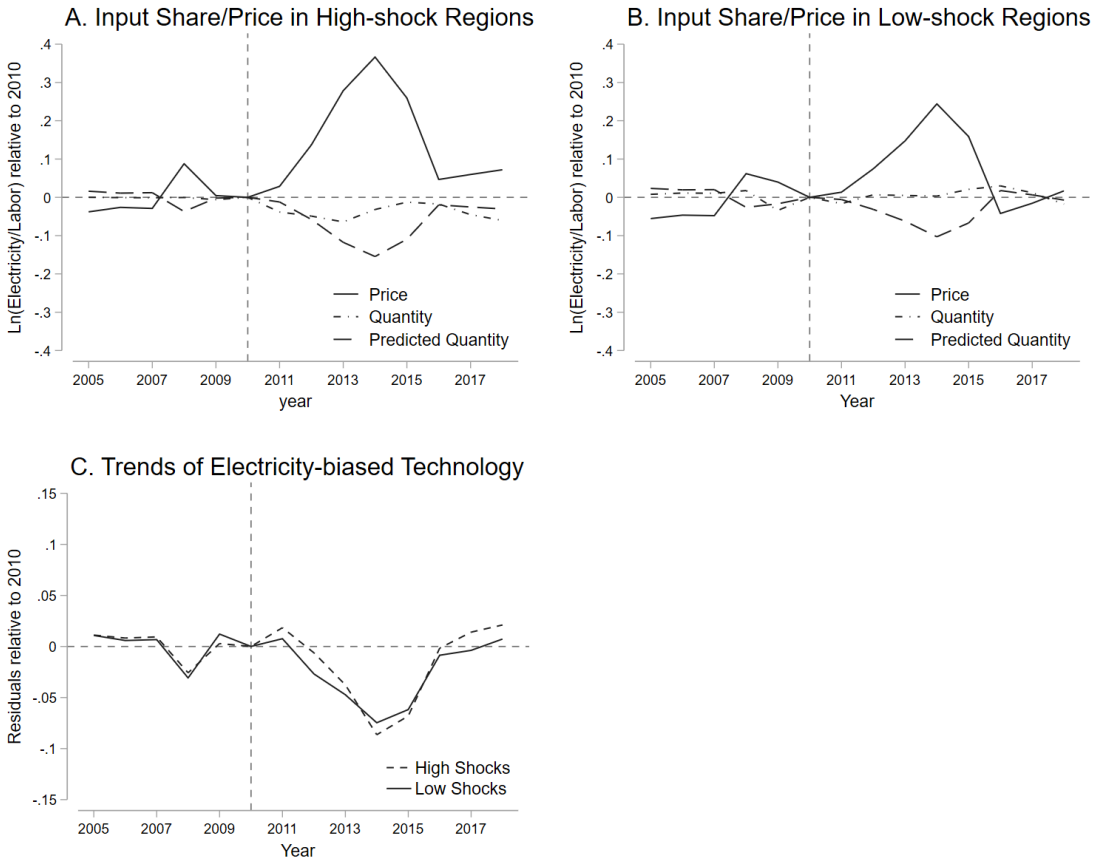
We then use $\hat{\alpha}$ and observed post-2011 relative prices to construct a "price-only" predicted path for $\ln(E/L)$ in 2011–2018. The residual component, $\hat{\varepsilon}_{it}$, corresponds to non-price shifts in relative demand and can be mapped into an implied technology shifter via

$$\hat{\omega}^E_{it} = \frac{\hat{\varepsilon}_{it}}{1-\hat{\sigma}}. \tag{C3}$$

Appendix Figure C2 reports the decomposition for high- and low-shock regions. The observed electricity-to-labor ratio closely tracks the price-predicted series in both groups; consistent with this, the recovered technology shifter $\hat{\omega}^E$ exhibits similar trends across groups (Panel C), providing little evidence of differential electricity-biased technical change induced by the supply shocks. Overall, the post-accident decline in electricity intensity appears to reflect movements along isoquants driven by relative price changes and input reallocation, rather than persistent electricity-saving innovations that shift the production frontier. This conclusion is consistent with the absence of systematic TFPR gains in the main analysis.

[1] Equation (C1) follows from CES demand with factor-augmenting terms: changes in p^E/p^L move the plant along an isoquant, while changes in ω^E shift relative demand holding prices fixed.

Figure C2: Relative Input Share, Prices, and Electricity-biased Technology



Notes: Panels A and B plot the trends of the log relative price of electricity to labor (solid line), the actual log electricity-to-labor ratio (dash-dotted line), and the predicted ratio based on relative price changes (dashed line). Panel A displays regions with supply shocks above the median ("High-shock"), while Panel B displays those below the median ("Low-shock"). Panel C plots the estimated electricity-biased technology (ω^E). All variables are log-differences relative to 2010.

C-3 Observed Plants' Responses

Product Switching We use product-level data from Japan's CMF, in which manufacturing plants report sales by product. There are 2,397 products in total, and the average plant produces 3.0 products. In Column (1), the dependent variable is an indicator for product switching, which equals one if the plant's product set differs from its 2010 set due to adding or dropping at least one product. Between 2010 and 2014, 39% of plants switch at least one product; by 2018, this figure rises to 54%. Column (2) uses the log change in product-level electricity intensity as the outcome. We measure electricity expenditure per unit of output using single-product plants in 2010, and then construct a plant-level measure as the sales-weighted average electricity intensity of the products the plant produces in each year.

Outsourcing Japan's CMF reports expenditures on outsourcing of intermediate inputs, which account for 10% of total costs and 18% of material costs. Column (3) uses the log change in outsourcing expenditure, and Column (4) uses the change in outsourcing's share of material costs as the outcome.

Production Reshuffling Japan's CMF includes firm identifiers, which allows us to test whether multi-plant firms reallocate production across regions in response to localized electricity supply shocks. We aggregate the data to the firm-by-region-by-year level and regress the firm's activity in all other regions on the supply shock in the focal region. This specification captures cross-regional substitution. For example, it tests whether a negative supply shock in the Tokyo region is associated with increased activity at the same firm's establishments in less affected regions, such as Hokkaido.

Entry and Exit We obtain counts of plants, entrants, and exits from the Economic Census. These data are aggregated to the prefecture-by-industry (2-digit)-by-year level. During our study period, establishment registration was updated in 2009, 2014, and 2016. While our main analysis focuses on plants with at least 30 employees, these measures cover all plants. Because the data do not distinguish true entry and exit from relocations of existing plants across regions, our estimates should be interpreted as effects on total turnover. We regress log changes relative to 2009—the pre-Fukushima wave—on electricity supply shocks.

Table C1: Spillover Effects: Production Reallocation Across Regions

Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$	Electricity (1)	Revenue (2)
Panel A: The effects of Supply Shocks using OLS		
Supply Shocks (Phase-out Terms)	0.25*	-0.06
	[0.09]	[0.53]
Supply Shocks (Post-Phase-out)	0.46**	-0.09
	[0.04]	[0.46]
Panel B: The effects of Supply Shocks using 2SLS		
Supply Shocks (Phase-out Terms)	0.26	-0.06
	[0.20]	[0.57]
Supply Shocks (Post-Phase-out)	0.62**	-0.20
	[0.03]	[0.44]
Year-by-Industry FE	Yes	Yes

Notes: The dependent variables are log changes in outcomes for the firm's plants in all other regions (excluding the focal region), measured relative to 2010. The regressor is the supply shock in the focal region. A positive coefficient indicates that a negative shock in the focal region is associated with increased activity elsewhere within the firm. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table C2: Effects of Supply Shocks on Entry and Exit

	Number Plants (1)	Entry (2)	Exit (3)
Panel A: The effects of Supply Shocks using OLS			
Supply Shocks	-0.30	0.74	0.44
	[0.18]	[0.56]	[0.21]
Panel B: The effects of Supply Shocks using 2SLS			
Supply Shocks	-0.34	1.08	0.39
	[0.27]	[0.56]	[0.44]
Year-by-Industry FE	Yes	Yes	Yes

Notes: The table reports estimates of equation (3) for plant counts and turnover outcomes. The dependent variables are log differences relative to the 2009 (pre-Fukushima) wave. The sample includes the 2009, 2014, and 2016 waves. P-values, reported in brackets, are obtained from a wild cluster bootstrap clustered at the region level. $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$.

C-4 Electricity Supply Shocks and Self Generators

Table C3: Effects of Supply Shocks by Generator Ownership

	w/ Generators		w/o Generators	
	Electricity (1)	Revenue (2)	Electricity (3)	Revenue (4)
Panel A: The effects of Supply Shocks using OLS				
Supply Shocks (Phase-out Terms)	-0.29 [0.15]	0.09 [0.43]	-0.23 [0.14]	-0.04 [0.40]
Supply Shocks (Post-Phase-out)	-0.76** [0.02]	0.15 [0.43]	-0.61** [0.03]	-0.22*** [0.00]
Panel B: The effects of Supply Shocks using 2SLS				
Supply Shocks (Phase-out Terms)	-0.40 [0.12]	0.11 [0.43]	-0.29 [0.26]	-0.04 [0.50]
Supply Shocks (Post-Phase-out)	-1.01** [0.02]	0.26 [0.37]	-0.77** [0.03]	-0.28*** [0.00]
Year-by-Industry FE	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. "w/ Generators" refers to firms that owned power generators in 2009 (financial year); "w/o Generators" refers to those that did not. The "Phase-out Terms" row uses data from 2005–2010 and 2011–2012. The "Post-Phase-out" row uses data from 2005–2010 and 2013–2018. All the regressions include year-by-industry (3-digits) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap using regions as clusters. P-values are reported in brackets. $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table C4: Effects of Supply Shocks on Self-generation

Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$	Own Generator (Yes=1) (1)	Intensive Margin (2)
Panel A: The effects of Supply Shocks using OLS		
Supply Shocks (Phase-out Terms)	0.23^{***} [0.00]	7.35^{***} [0.00]
Supply Shocks (Post-Phase-out)	0.23^{**} [0.01]	2.18 [0.13]
Panel B: The effects of Supply Shocks using 2SLS		
Supply Shocks (Phase-out Terms)	0.27^{***} [0.00]	8.29^{***} [0.00]
Supply Shocks (Post-Phase-out)	0.36^{***} [0.01]	3.43 [0.14]
Year-by-Industry FE	Yes	Yes

Notes: This table uses the ECS data and reports regression estimates. Column (1) uses a binary variable indicating generator ownership. Column (2) uses the log difference of electricity generated relative to 2009 (fiscal year). The Phase-out Terms row uses data from fiscal years 2005–2009 and 2011–2012. The Post-Phase-out row uses data from fiscal years 2005–2009 and 2013–2018. All regressions include year-by-industry (2-digit) fixed effects and use 2009 employment weights. We compute p-values using a wild cluster bootstrap with regions as clusters; p-values are reported in brackets. $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

D Production Framework and Model Validation

D-1 Predicting Outputs and Inputs in the Simulation

This section describes how we generate plant-level counterfactual outcomes under (i) electricity-price increases and (ii) seasonal rationing. The simulation is conducted plant-by-plant using industry-specific CES parameters and plant-specific baseline scale.

We use the production and demand environment from Section 7. In this appendix, we suppress tildes and write K for effective capital services whenever utilization is relevant.

D-1.1 Unconstrained equilibrium (baseline and pricing counterfactuals)

Each season without rationing, the plant chooses inputs to maximize profits:

$$\begin{aligned} &\max_{K,L,M,E} \ \Pi = P\,Y - p^K K - p^L L - p^M M - p^E E, \\ &\text{subject to} \quad Y = \Omega(\alpha_K K^\rho + \alpha_L L^\rho + \alpha_M M^\rho + \alpha_E E^\rho)^{\frac{1}{\rho}}. \end{aligned} \tag{D1}$$

Under constant returns to scale, the unit-cost function is

$$c(p) = \frac{1}{\Omega}\left(\alpha_K^\sigma (p^K)^{1-\sigma} + \alpha_L^\sigma (p^L)^{1-\sigma} + \alpha_M^\sigma (p^M)^{1-\sigma} + \alpha_E^\sigma (p^E)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}. \tag{D2}$$

Profit-maximizing pricing implies $P = \mu c(p)$, so equilibrium output is

$$Y^*(p) = B\,[\mu\, c(p)]^{-\eta}. \tag{D3}$$

Conditional factor demands imply the standard CES input demands:

$$X_j^*(p) = \frac{Y^*(p)}{\Omega}\left(\frac{c(p)\,\Omega\,\alpha_j}{p^j}\right)^\sigma, \qquad j \in \{K, L, M, E\}. \tag{D4}$$

Implementation. Given a price vector p, we compute $c(p)$ from (D2), then $Y^*(p)$ from (D3), and finally (K^*, L^*, M^*, E^*) from (D4). Pricing counterfactuals replace p^E and other factor prices with post-shock values and re-solve the unconstrained system.

D-1.2 Analytical elasticities with respect to electricity prices

Define the electricity cost share in unit cost as

$$s_E^C(p) = \frac{\alpha_E^\sigma (p^E)^{1-\sigma}}{\alpha_K^\sigma (p^K)^{1-\sigma} + \alpha_L^\sigma (p^L)^{1-\sigma} + \alpha_M^\sigma (p^M)^{1-\sigma} + \alpha_E^\sigma (p^E)^{1-\sigma}}. \tag{D5}$$

By Shephard's Lemma, $\partial \ln c(p)/\partial \ln p^E = s_E^C$. Since $P = \mu c(p)$ with constant μ, we have $d \ln P / d \ln p^E = s_E^C$, and demand implies $d \ln Y / d \ln p^E = -\eta s_E^C$.

Combining with (D4) yields

$$\begin{aligned} \frac{d \ln Y}{d \ln p^E} &= -\eta \, s_E^C, \\ \frac{d \ln X_j}{d \ln p^E} &= (\sigma - \eta) \, s_E^C, \qquad j \in \{K, L, M\}, \\ \frac{d \ln E}{d \ln p^E} &= -\sigma + (\sigma - \eta) \, s_E^C. \end{aligned} \tag{D6}$$

Producer surplus is $\Pi = (P - c)Y = (\mu - 1)cY$; because Y depends on $c(p)$ through demand,

$$\frac{d \ln \Pi}{d \ln p^E} = \frac{d \ln c}{d \ln p^E} + \frac{d \ln Y}{d \ln p^E} = (1 - \eta) \, s_E^C. \tag{D7}$$

Unless otherwise noted, these elasticities are evaluated at the pre-shock baseline equilibrium.

D-1.3 Seasonal rationing via a shadow-price representation

We allow for rationing only in the shock season S. Electricity use in season S is subject to

$$E^S \leq \bar{E}^S = (1 - \mathcal{S}) \, E_{\text{base}}^{S*}, \tag{D8}$$

where $\mathcal{S}$ is the region-year shock and E_{base}^{S*} denotes baseline (pre-shock) electricity use in season S.

Non-binding case. If the unconstrained optimum under post-shock prices satisfies $E^{S*}(p_{\text{post}}) \leq \bar{E}^S$, rationing is slack and outcomes in season S are given by the unconstrained solution (D2)–(D4) evaluated at post-shock prices.

Binding case. If (D8) binds, we approximate the rationing equilibrium using a shadow-price representation. Let $\tilde{p}^E$ denote the shadow price of electricity in season S (price plus scarcity rent). We choose $\tilde{p}^E$ such that the unconstrained CES solution at $(p^K, p^L, p^M, \tilde{p}^E)$ reproduces the binding electricity quantity. Operationally, we use a first-order log approximation around the unconstrained post-price optimum. Let $E_{\text{ref}}^S \equiv E^{S*}(p_{\text{post}})$ denote

electricity demand under the post-shock price absent rationing. Then

$$d \ln E^S \approx \ln\left(\frac{\bar{E}^S}{E^S_{\text{ref}}}\right). \tag{D9}$$

Using (D6) with respect to $\tilde{p}^E$, the implied shadow-price change is

$$d \ln \tilde{p}^E = \frac{d \ln E^S}{-\sigma + (\sigma - \eta) s^C_E}. \tag{D10}$$

Throughout the rationing approximation, we evaluate s^C_E at the reference point $E^S = E^S_{\text{ref}}$. Given $d \ln \tilde{p}^E$, the associated changes in season-S output and non-electricity inputs are

$$\begin{aligned} d \ln Y^S &= -\eta \, s^C_E \, d \ln \tilde{p}^E, \\ d \ln X^S_j &= (\sigma - \eta) \, s^C_E \, d \ln \tilde{p}^E, \quad j \in \{K, L, M\}, \\ d \ln \Pi^S &= (1 - \eta) \, s^C_E \, d \ln \tilde{p}^E. \end{aligned} \tag{D11}$$

Combining price increases and rationing. We proceed in two steps. First, solve the unconstrained system under the post-shock price to obtain $E^S_{\text{ref}} \equiv E^{S*}(p_{\text{post}})$. Second, if $E^S_{\text{ref}} > \bar{E}^S$, compute $d \ln E^S \approx \ln(\bar{E}^S / E^S_{\text{ref}})$ and apply (D10)–(D11). For the off-season O, outcomes are always given by the unconstrained solution at post-shock prices.

Annual aggregation. Annual outcomes are computed as weighted averages of seasonal log changes,

$$\Delta \ln Z \equiv \phi \, \Delta \ln Z^S + (1 - \phi) \, \Delta \ln Z^O, \tag{D12}$$

a first-order approximation to log changes in annual outcomes defined as seasonal averages in levels.

D-1.4 Plants with generators

Generator-owning plants can substitute from grid purchases to self-generation. Consistent with reduced-form evidence showing large declines in grid purchases but no statistically significant changes in real output or other non-electricity inputs, we impose no scale response for generator owners: (Y, K, L, M) remain at baseline levels, and the shock affects only the composition of electricity between grid purchases and self-generation. In practice, we allow grid electricity purchases to respond to the (effective) electricity price through an own-price elasticity, but set cross-price elasticities with respect to Y, K, L, and M to zero.

D-2 Simulation Parameters

We recover the CES distribution parameters α from the estimated revenue elasticities θ, the industry-mean input bundle, and a candidate σ. Given (α, σ), we recover plant-level TFPR ω and conduct the counterfactual exercises described above. To incorporate electricity-constrained utilization in the post-accident period, we use the electricity–capital fixed ratio v and depreciation rate δ to construct effective capital services and map (α_K, α_E) into time-varying effective weights $(\tilde{\alpha}_{K,t}, \tilde{\alpha}_{E,t})$.

Table D1: Simulation Inputs

	Description	Level	Pre-policy	Post-policy
Panel A: Implied objects				
α	CES distribution parameters	industry	α_{2010}	α_{2010}
ω	TFPR	plant	ω_{2010}	ω_{2010}
σ	Elasticity of substitution	constant	0.5	0.5
Panel B: Capital utilization				
v	Electricity–capital fixed ratio	plant	–	v_{2010}
δ	Depreciation rate	constant	–	5%
Panel C: Market demand				
η	Demand elasticity	constant	6	6
B	Demand shifter (scale)	plant	R_{2010}	R_{2010}
Panel D: Rationing and pricing policies				
ϕ	Shock-season length (days)	constant	–	90/365
$\mathcal{S}$	Supply shock	region-year	–	$\mathcal{S}$
p^E	Electricity price	region-year	p_t^E	p_t^E
p^M	Materials price	year	p_t^M	p_t^M
p^L	Labor price	year	p_t^L	p_t^L

D-3 Results for Elasticities

Table D2: Comparison of Elasticity Estimates

	Estimation method	
	Reduced-form (1)	Structural model (2)
Own-price elasticity (electricity)	-0.60	-0.67
Cross-price elasticity (labor)	-0.17	-0.16
Cross-price elasticity (materials)	-0.13	-0.16
Output elasticity	-0.17	-0.17

Notes: Reduced-form elasticities are computed as the ratio of the outcome response to the effective-price impulse as described in Section 7.6.4 (Column (1)). Structural elasticities are implied by the production-function parameters (Column (2)).

E Counterfactual Analysis

E-1 Counterfactual Electricity Prices

E-1.1 Baseline: shock-induced price increase

This subsection constructs a baseline measure of the electricity price increase attributable to the post-Fukushima nuclear phase-out. We estimate pass-through from regional supply shocks to electricity prices using

$$\Delta \ln p^E_{rt} \equiv \ln p^E_{rt} - \ln p^E_{r,2010} = \beta \, \mathcal{S}_{rt} + \gamma_t + \varepsilon_{rt}, \tag{E1}$$

where p^E_{rt} is the electricity price in region r and year t, $\mathcal{S}_{rt}$ is the region–year supply shock, and γ_t are year fixed effects that absorb nationwide price movements, including common fossil-fuel cost fluctuations.

To summarize the shock-induced price increase by region, we evaluate (E1) at the region's mean exposure in the post-phase-out period (2013–2018),

$$\widehat{\Delta \ln p^E}_r = \widehat{\beta} \, \bar{\mathcal{S}}_r, \tag{E2}$$

where $\bar{\mathcal{S}}_r$ denotes the 2013–2018 mean of $\mathcal{S}_{rt}$. In the IV specification, $\widehat{\beta} \approx 1.20$, so a 1 ppt increase in the supply shock is associated with about a 1.2% increase in electricity prices (in logs).[2] We combine $\widehat{\beta}$ and $\bar{\mathcal{S}}_r$ to obtain the baseline price changes reported in Table E1, Column (1). The demand-weighted national average implied increase is 8.9%, with substantial heterogeneity: low-shock regions such as Hokkaido and Chugoku exhibit increases of about 1–4%, while high-shock regions such as Tokyo and Kansai exceed 10%.

E-1.2 Counterfactual 1: nuclear restarts

Construction. We construct counterfactual electricity prices under partial nuclear restarts by combining (i) the first-stage mapping from the nuclear shift-share component to supply shocks (Table 2) and (ii) the shock–price pass-through in (E1). Let the nuclear component of the instrument be

$$Z^{\text{Nuclear}}_{rt} = Sh^{\text{Nuclear}}_{r,2010} \times g_t, \tag{E3}$$

where $Sh^{\text{Nuclear}}_{r,2010}$ is the nuclear share of regional electricity supply on the 2010 peak day and g_t is the national percent change in peak-day nuclear generation relative to 2010.

[2] For the reduced-form elasticity calculations, we use an *effective* electricity price that incorporates rationing through a shadow-price adjustment. In contrast, the counterfactual prices in this appendix map shocks into observed electricity prices and therefore abstract from scarcity rents.

We model partial restarts via a region-specific reutilization rate $c_r \in [0, 1]$ and define the counterfactual nuclear shortfall as

$$Z_{rt}^{\text{Nuclear,cf}} = (1 - c_r)\, Z_{rt}^{\text{Nuclear}}, \tag{E4}$$

so that $c_r = 0$ reproduces the observed shutdown and $c_r = 1$ eliminates the nuclear shortfall relative to 2010.

We convert the implied change in the nuclear shifter into a change in the supply shock using the first-stage coefficient $\widehat{\theta}$:

$$\Delta \mathcal{S}_{rt}^{\text{cf}} = \widehat{\theta}\, \Delta Z_{rt}^{\text{Nuclear}}, \tag{E5}$$

and map the resulting change in supply shocks into counterfactual electricity prices using $\widehat{\beta}$:

$$\Delta \ln \widehat{p}_{rt}^{E} = \widehat{\beta}\, \Delta \mathcal{S}_{rt}^{\text{cf}}. \tag{E6}$$

Feasibility-based scenario design. We discipline c_r using a three-tier feasibility classification based on publicly observable constraints: progress in Nuclear Regulation Authority safety reviews and compliance actions, seismic/fault-line concerns, and local acceptance, proxied by the status of major regional plants. Kansai, Shikoku, and Kyushu are classified as High feasibility because multiple units have already restarted (e.g., Takahama/Ohi/Mihama; Ikata; Sendai/Genkai). Tohoku, Hokuriku, and Chugoku are classified as Medium feasibility, reflecting partial progress (e.g., Onagawa progressing while Higashidori remains delayed; Shika and Shimane facing remaining uncertainty despite approval/progress). Hokkaido, Tokyo, and Chubu are classified as Low feasibility due to persistent regulatory or seismic constraints and prolonged shutdowns (e.g., Tomari; Kashiwazaki–Kariwa; Hamaoka). In the policy-feasible restart scenario, we set $c_r = 0.70$ (High), $c_r = 0.40$ (Medium), and $c_r = 0.10$ (Low).

Results. Table E1, Column (2), reports implied price changes under this restart scenario. Price relief is largest in regions with both high baseline nuclear exposure and high restart feasibility (Kansai, Shikoku, Kyushu), while implied reductions are small in low-feasibility regions (e.g., Tokyo). The demand-weighted national average price increase falls from 8.9% in the baseline to 6.9% under the restart scenario.

E-1.3 Counterfactual 2: market integration and reserve sharing

Integrated-market pricing. We construct counterfactual electricity prices under market integration by specifying short-run regional supply curves, aggregating supply across regions, and solving for the integrated-market price. We parameterize the inverse supply curve in constant-elasticity form:

$$p_{rt} = \left(\frac{Q_{rt}}{c_{rt}}\right)^{\beta}, \tag{E7}$$

where p_{rt} is the electricity price, Q_{rt} is quantity supplied, $c_{rt} > 0$ is a region–year cost shifter, and $\beta \equiv d\ln p / d\ln Q$ (so $d\ln Q/d\ln p = 1/\beta$).[3]

Under full integration, all regions face a common price p_t^* in year t, implying regional supply

$$Q_{rt}^* = c_{rt}\left(p_t^*\right)^{1/\beta}. \tag{E8}$$

Holding fixed national demand $\bar{Q}_t$, the integrated equilibrium satisfies

$$\bar{Q}_t = \sum_r Q_{rt}^* = \left(p_t^*\right)^{1/\beta} \sum_r c_{rt}, \tag{E9}$$

which yields

$$p_t^* = \left(\frac{\bar{Q}_t}{\sum_r c_{rt}}\right)^{\beta}. \tag{E10}$$

Reserve sharing (effective capacity). Integration can reduce the need for region-specific reserve margins. We capture this mechanism by scaling aggregate effective capacity by $(1+\kappa)$:

$$p_t^* = \left(\frac{\bar{Q}_t}{(1+\kappa)\sum_r c_{rt}}\right)^{\beta}. \tag{E11}$$

We set $\kappa = 0.007$, corresponding to a 0.7% increase in effective capacity due to improved dispatch and reserve sharing.

Estimation. We set $\beta = 5$, corresponding to a supply elasticity $d\ln Q/d\ln p = 0.2$.[4] Rewrite (E7) in logs:

$$\ln p_{rt} = A_{rt} + \beta \ln Q_{rt}, \qquad A_{rt} \equiv -\frac{1}{\beta}\ln c_{rt}.$$

[3] See Wolak (2000), Bushnell and Mansur (2005), Holland and Mansur (2006), and Cicala (2022).

[4] Existing estimates of very short-run electricity supply elasticities are typically smaller (0.02–0.10) (Wolak, 2000; Bushnell and Mansur, 2005; Holland and Mansur, 2006; Cicala, 2022). Because our variation is at the region–year level, we adopt an intermediate value that allows for limited within-year adjustment.

We recover $A_{rt} = \ln p_{rt} - \beta \ln Q_{rt}$ and model it as a function of generation shares and fixed effects:

$$A_{rt} = \gamma_1 Sh_{rt}^{\text{Nuclear}} + \gamma_2 Sh_{rt}^{\text{Hydro}} + \gamma_3 Sh_{rt}^{\text{Thermal}} + \gamma_4 Sh_{rt}^{\text{Others}} + \tau_t + \delta_r + \varepsilon_{rt}. \tag{E12}$$

We map the fitted value $\widehat{A}_{rt}$ back into the cost shifter via

$$\widehat{c}_{rt} = \exp\left(-\widehat{A}_{rt}/\beta\right), \tag{E13}$$

and substitute $\widehat{c}_{rt}$ into (E11) to obtain the integrated-market counterfactual price p_t^*.

Results. Table E1, Column (3), reports implied price changes under market integration. High-cost regions experience sizable counterfactual price reductions, while low-cost regions experience price increases, reflecting reallocation of generation toward lower-cost supply. The demand-weighted national average price increase falls from 8.9% in the baseline to 5.7% under integration.

Table E1: Regional price Changes: Baseline, Nuclear Utilization, and Trade

		$\Delta p^E = p^E - p^E_{2010}$		
#	Region	Baseline (1)	Nuclear restart (2)	Trade (3)
1	Hokkaido	1.2%	0.7%	-12.5%
2	Tohoku	9.3%	7.6%	2.3%
3	Tokyo	13.1%	12.6%	4.0%
4	Hokuriku	4.6%	1.8%	19.0%
5	Chubu	3.9%	3.6%	5.3%
6	Kansai	13.8%	8.6%	8.9%
7	Chugoku	3.8%	3.8%	13.0%
8	Shikoku	5.4%	-0.4%	2.7%
9	Kyushu	6.8%	1.6%	10.5%
	National	8.9%	6.9%	5.7%

Notes: Entries report $100 \times \Delta \ln p^E$. Column (1) reports the shock-induced price increase (baseline). Column (2) and Column (3) report the counterfactual price changes under the nuclear-restart (reutilization) scenario and the market-integration (trade) scenario, respectively. National values are demand-weighted averages using 2010 regional electricity demand.

E-2 Welfare Comparison: Pricing vs. Rationing

This subsection compares a uniform rationing rule to a uniform price mechanism while holding fixed the aggregate reduction in electricity consumption, under the welfare accounting in Section 8.3.

Under rationing, each plant i faces a uniform proportional reduction $\mathcal{S}$ in electricity use relative to its baseline, implying an aggregate reduction

$$\Delta E^{\text{rat}} = \sum_i (-\mathcal{S}\, E_{i,\text{base}}) = -\mathcal{S} \sum_i E_{i,\text{base}}. \tag{E14}$$

Let $\varepsilon_i \equiv d \ln E_i / d \ln p^E$ denote the plant-level electricity-demand elasticity (Appendix D-1.2). The shadow-price change that implements the quota for plant i satisfies

$$d \ln \tilde{p}_i^E = \frac{d \ln E_i}{\varepsilon_i} = -\frac{\mathcal{S}}{\varepsilon_i}. \tag{E15}$$

Under pricing, all plants face a common price change $d \ln p^{E*}$ that delivers the same aggregate reduction:

$$\Delta E^{\text{price}} \approx \sum_i E_{i,\text{base}}\, d \ln E_i = \sum_i E_{i,\text{base}}\, \varepsilon_i\, d \ln p^{E*}. \tag{E16}$$

Let $\bar{\varepsilon} \equiv \sum_i \omega_i \varepsilon_i$ denote the electricity-consumption-weighted average elasticity, where $\omega_i \equiv E_{i,\text{base}} / \sum_j E_{j,\text{base}}$. Equating (E14) and (E16) yields

$$d \ln p^{E*} = \frac{-\mathcal{S}}{\bar{\varepsilon}}. \tag{E17}$$

Output responds to electricity-price changes according to $d \ln Y_i = -\eta s_{E,i}^C\, d \ln p^E$ (Appendix D-1.2), where $s_{E,i}^C$ is the electricity cost share in unit cost. Under pricing, all plants face the common price change $d \ln p^{E*}$. Under rationing, plant i faces the heterogeneous shadow-price change $d \ln \tilde{p}_i^E$ in (E15).

The ratio of plant-level output losses under rationing relative to pricing is therefore

$$\frac{\left| d \ln Y_i^{\text{rat}} \right|}{\left| d \ln Y_i^{\text{price}} \right|} = \frac{\left| -\eta s_{E,i}^C\, d \ln \tilde{p}_i^E \right|}{\left| -\eta s_{E,i}^C\, d \ln p^{E*} \right|} = \frac{\bar{\varepsilon}}{\varepsilon_i}. \tag{E18}$$

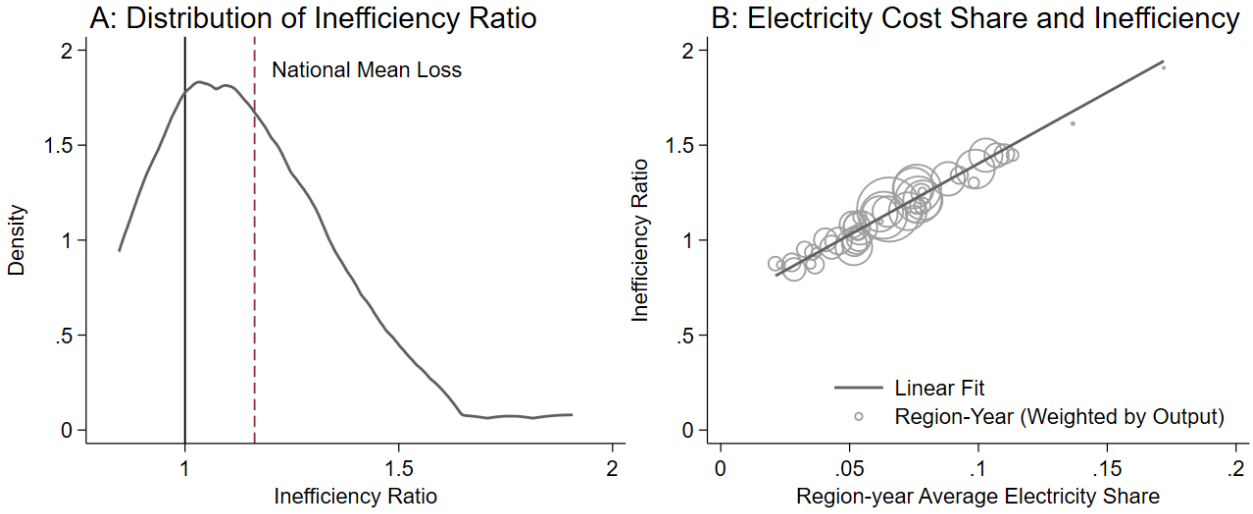
Plants with relatively inelastic electricity demand (small $|\varepsilon_i|$) experience larger output losses under rationing because implementing a fixed proportional quantity cut requires a larger shadow-price increase.

We summarize allocative inefficiency by the ratio of aggregate output losses,

$$\Lambda \equiv \frac{\sum_i \left| d \ln Y_i^{\text{rat}} \right|}{\sum_i \left| d \ln Y_i^{\text{price}} \right|} \geq 1, \tag{E19}$$

so $\Lambda > 1$ indicates that quantity mandates generate larger aggregate output losses than a uniform price increase that achieves the same aggregate electricity conservation.

Figure E1: **The Allocative Inefficiency of Quantity Mandates**



Notes: Panel A displays the kernel density of Λ across region-years. The solid vertical line at 1 corresponds to equivalence between the two mechanisms; the dashed line denotes the weighted national mean. Panel B plots Λ against the output-weighted average electricity cost share for each region-year and overlays a linear fit. The sample includes years $t \geq 2013$.

References for Appendix

Bushnell, James B, and Erin T Mansur. 2005. "Consumption under noisy price signals: a study of electricity retail rate deregulation in San Diego." *The Journal of Industrial Economics* 53 (4): 493–513.

Cicala, Steve. 2022. "Imperfect markets versus imperfect regulation in US electricity generation." *American Economic Review* 112 (2): 409–441.

Holland, Stephen P, and Erin T Mansur. 2006. "The short-run effects of time-varying prices in competitive electricity markets." *The Energy Journal* 27 (4): 127–156.

Wolak, Frank A. 2000. "An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market." *International Economic Journal* 14 (2): 1–39.