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

Daiji Kawaguchi† Takanao Tanaka‡

July 2026

As governments around the world pursue clean energy transitions, electricity reliability and affordability have become central policy concerns. How do adverse shocks to electricity markets affect industrial production? We study the persistent, large-scale contraction in electricity supply experienced by Japanese manufacturers following the 2011 Fukushima nuclear accident. Using confidential plant-level administrative data and an instrumental-variables strategy for localized electricity supply shocks, we show that these shocks reduced electricity use by 6.7% and output by 1.7%. The implied annual output loss is 4.9 trillion JPY, or approximately 50–60 billion USD. We find little evidence that manufacturing plants responded by switching products, increasing outsourcing, reallocating production within firms, relocating, or substituting toward other energy sources. We estimate a production function and show that a standard Cobb–Douglas model cannot replicate the reduced-form impacts. Allowing for limited substitutability between electricity and other inputs and for capital idleness when electricity constraints leave installed capital temporarily unused enables the model to match these reduced-form impacts. We discuss the welfare implications of electricity-market reforms that could mitigate the adverse effects of such shocks.

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

*We thank Joseph Shapiro, James Sallee, Reed Walker, and Guojun He, as well as seminar and conference participants at UC Berkeley, Institutional Barriers to Climate Adaptation (Stanford), LSE Environmental Camp, AERE Annual Conference, 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

As governments around the world pursue clean energy transitions, electricity affordability and reliability have become central policy concerns. How do adverse shocks to electricity markets affect industrial production? This paper studies a large-scale, persistent contraction in electricity supply to manufacturing.

Electricity demand in advanced economies is rising again after a long period of stagnation, a trend particularly relevant for the United States, the European Union, Japan, and South Korea. In the United States, for example, electricity demand grew by only about 0.1% per year from 2005 to 2019, but has grown by roughly 2% per year since 2020 and is projected to be about one-quarter to one-half higher by 2050 (EIA, 2026c,a). Widespread electrification in transportation, industry, and buildings, aimed at reducing reliance on fossil fuels, contributes to this trend. At the same time, data centers have emerged as a major new source of demand growth.

Electricity supply, however, remains vulnerable to shocks. Geopolitical risks and global fuel-market disruptions can raise generation costs and destabilize local electricity markets (Fabra, 2023). The transition toward weather-dependent renewable generation adds new operational challenges. Renewable resources are often located far from demand centers, while transmission infrastructure expands slowly, limiting the ability of supply to meet local demand (Davis et al., 2023; Gonzales et al., 2023).

Rising demand and supply-side vulnerabilities have made affordability and reliability central concerns in electricity markets. Heightened demand, grid constraints, and large-scale infrastructure needs can place upward pressure on electricity prices while increasing the risk of demand curtailment. These pressures are already visible in recent price trends: in the United States, average retail electricity prices rose by about 33% from 2016 to 2025 (EIA, 2026b).

Researchers, scientists, and policymakers increasingly share these concerns. The North American Electric Reliability Corporation (NERC) characterizes the current reliability outlook as one of "mounting resource adequacy challenges ... as surging demand growth continues" (NERC, 2024). Affordability concerns are similarly prominent: the International Energy Agency (IEA) identifies electricity affordability as a "key and growing concern" over the coming decades (IEA, 2026). Recent research examines the drivers of high electricity prices and proposes reforms to mitigate cost increases and improve equity (Borenstein et al., 2021; Wiser et al., 2025).

Despite growing concern, little is known about how persistent electricity supply shocks affect production in advanced economies or how firms and plants adjust to them. Existing research on electricity markets, reviewed below, has focused mainly on transitory and

relatively small price changes (Wolverton et al., 2022; von Graevenitz and Rottner, 2024) or temporary service interruptions, often in less developed economies (Fisher-Vanden et al., 2015; Allcott et al., 2016; Cole et al., 2018; Gonzales et al., 2026). A broader literature documents large economic losses from shocks to oil and gas, rather than electricity itself (Kilian, 2008; Hamilton, 2013). Moreover, much of this work relies on macro- or industry-level data, making it difficult to observe how firms and plants adjust.

This paper addresses this gap by studying electricity supply shocks to Japanese manufacturing after the 2011 Fukushima nuclear accident. Before the accident, nuclear plants generated approximately 35% of total electricity. The accident triggered widespread concerns about nuclear safety, leading to the suspension of nuclear reactors across the country and bringing nuclear generation almost completely to a halt. The decline persisted for years: nuclear generation accounted for only 2% of total electricity generation in 2013 and about 8% in 2018 (METI, 2011a). To avoid widespread blackouts, the government curtailed demand through mandatory restrictions and information campaigns. At the same time, replacing nuclear generation with higher-cost fossil-fuel generation led to sharp and persistent increases in electricity prices.

We use restricted-access plant-level data from the Census of Manufactures for 2005–2018, covering approximately 80% of manufacturing output. Because it reports inputs, output, firm identifiers, and product-level information at the plant level, the data allow us to examine both changes in production and adjustment margins such as product switching and within-firm reallocation. We further combine the census with administrative energy-use data to study fuel substitution and on-site electricity generation.

We exploit quasi-experimental regional variation in the severity of electricity supply shocks using an instrumental-variables approach. Japan's electricity system has historically been segmented into nine regional service areas. During our study period, the electricity market remained heavily regulated, and interregional electricity transfers were limited in practice, amounting to only a few percent of regional demand. This regulatory regime generated wide variation in localized shock exposure. We measure regional shock intensity based on peak day supply-demand conditions, when blackout risk is most salient. Shock intensity is defined as the gap between required and available supply, normalized by pre-accident peak demand.[1] Regions with greater losses of nuclear generation and lower initial reserve margins faced especially large shocks; in some cases, the implied shortfall approached 20%. To isolate plausibly exogenous variation, we instrument regional shocks with predicted supply shortfalls based on pre-accident nuclear reliance and initial reserve margins.

[1]Required supply equals pre-accident peak demand plus a 5% reserve margin, while available supply equals each region's maximum planned supply for that year.

Our empirical analysis proceeds in four steps. First, we show that electricity supply shocks caused persistent, economically meaningful output losses. Following the complete shutdown of nuclear power in 2013, these shocks reduced manufacturing electricity use by 6.7%, labor and materials by 1.3%, and output by 1.7%. Although our data extend through 2018, seven years after the accident, we find little evidence that these effects diminished over time. The magnitude of the output loss is striking given that electricity expenditures account for only 2% of revenue. The implied annual output loss is approximately 4.9 trillion JPY, or USD 50–60 billion, depending on the exchange rate. Moreover, the effects are particularly large among capital-intensive plants and industries, suggesting that capital intensity amplifies the output losses caused by electricity supply shocks.

Second, we show that manufacturing plants make little adjustment in response to electricity supply shocks, helping to explain the substantial output losses we document. Specifically, electricity supply shocks do not meaningfully alter several margins that could, in principle, mitigate their adverse effects: product switching, outsourcing, within-firm reallocation of production across plants, and plant relocation. The estimated impacts on all of these margins are statistically insignificant and economically small.

Additionally, we find little evidence that manufacturing plants substitute away from electricity toward other energy inputs. Our data distinguish between energy used directly in production and energy used for on-site electricity generation. Electricity supply shocks do not increase energy use in production; if anything, this input declines, suggesting that electricity and other energy inputs are not close substitutes in the production process. By contrast, plants with self-generators increase fuel use for on-site generation. This adjustment margin appears effective: plants with self-generators experience no output declines due to electricity supply shocks. However, only 6.1% of plants own self-generators, and the shocks do not significantly increase generator ownership. Thus, on-site generation is unavailable to more than 90% of manufacturing plants.

Third, we investigate what production technology can match the reduced-form responses of electricity, labor, materials, and output to electricity supply shocks. To do so, we estimate CES production functions following Gandhi et al. (2020), assuming a common technology within each industry. We use the estimated industry-level output elasticities to construct baseline cost shares and calibrate the CES distribution parameters for candidate elasticities of substitution. We then predict plants' output and input choices under the electricity price changes and rationing induced by local supply shocks, and compare these model-implied responses with the reduced-form estimates.

We find that standard Cobb–Douglas (CD) production functions fail to reproduce the reduced-form impacts of electricity supply shocks on both output and input use. First, the CD model predicts much smaller reductions in inputs and output, suggesting the

need to incorporate additional mechanisms that amplify the adverse effects of the shocks. Second, the CD model overstates the relative impact on electricity use compared with other inputs, such as materials and labor. The CD specification allows plants to substitute strongly away from electricity, generating much larger reductions in electricity use but smaller reductions in materials and labor.

To address these discrepancies, we incorporate two mechanisms: capital idleness and input complementarity. First, reduced electricity availability can leave electricity-dependent capital idle (Finn, 2000). For example, a steel plant with electricity-dependent furnaces must reduce their operating hours when electricity availability declines. The resulting capital idleness amplifies the reduction in output and lowers the marginal products of labor and materials. Our reduced-form heterogeneity results support this mechanism: the impacts of the shocks are disproportionately large among capital-intensive plants and plants in capital-intensive industries. Second, we introduce limited substitutability between electricity and other inputs by allowing the elasticity of substitution in the CES production function to differ from the Cobb–Douglas benchmark.

Incorporating these two mechanisms allows the model to match the reduced-form estimates. Using the generalized method of moments (GMM), we jointly estimate the capital-idleness elasticity and the elasticity of substitution between electricity and other inputs. The estimated capital-idleness elasticity is 0.17, meaning that a 1% reduction in electricity use leaves 0.17% of installed capital temporarily idle. Because capital's output elasticity, 0.139, is much larger than electricity's, 0.015, this channel amplifies output losses beyond electricity's direct contribution to production. The estimated elasticity of substitution is 0.59, indicating substantially greater complementarity between electricity and other inputs than under Cobb–Douglas technology. This complementarity explains the larger reductions in labor and materials and the resulting output losses.

Fourth, we use the structural model to evaluate counterfactual policies that could mitigate the output effects of electricity supply shocks. Designing these counterfactuals requires identifying which channel had larger effects on output: electricity rationing or electricity price increases. To decompose these impacts, we predict output under price-change-only and price-change-plus-rationing scenarios. The results show that price changes alone explain approximately 95% of the output decline. This result differs from the residential sector, where non-price interventions can meaningfully reduce electricity demand during scarcity (e.g., Reiss and White, 2008; He and Tanaka, 2023).

Building on the finding that higher electricity prices account for most output losses, our counterfactual analysis evaluates regional market integration, which has attracted considerable policy attention because of its potential benefits. Historically, Japan's regulatory framework limited interregional electricity transmission and trade, generating

inefficiencies. Regional segmentation can require high-cost generators to serve local demand while lower-cost capacity elsewhere remains idle. It can also require each region to maintain reserve capacity independently, increasing system-wide reserve requirements. Relative to the price-only baseline, we estimate that regional market integration reduces the shock-induced electricity-price increase by 2.8% and lowers output losses by 32.1%. These findings suggest that the economic costs of electricity supply shocks partly reflect long-standing regulations that restricted interregional electricity trade.

This paper makes several contributions. First, we provide, to our knowledge, some of the first micro-level evidence on how persistent electricity supply shocks affect manufacturing in an advanced economy. Widespread electrification has made electricity markets increasingly important in their own right, yet existing work has focused primarily on shocks to broader energy markets rather than on electricity-market-specific shocks. A related literature examines the association between electricity prices and firm performance (Kahn and Mansur, 2013; Bijnens et al., 2022). Existing quasi-experimental studies have focused on small and transitory price changes induced by regulatory network charges (von Graevenitz and Rottner, 2024) or renewable portfolio standards (Wolverton et al., 2022), finding modest or null effects on output.[2] We complement this literature by studying the Fukushima nuclear accident, which has been described as a "major turning point in Japan's energy policy" (ANRE, 2022). In contrast to the modest effects documented in prior quasi-experimental work, we find substantial output losses from the resulting large-scale and persistent electricity supply shocks.

This paper is also distinct from studies of electricity outages in developing economies, which document large production losses (Fisher-Vanden et al., 2015; Allcott et al., 2016; Cole et al., 2018; Gonzales et al., 2026). For example, Allcott et al. (2016) finds that frequent blackouts in India reduce output by about 5–10%. Although we document sizable output effects in Japan, the effects in these outage settings are several times larger than our estimates. We interpret this difference as reflecting a different margin of adjustment: plants in our setting continue to operate, albeit at higher marginal costs. In contrast, outages in developing economies often force production to stop.

We also contribute to the broader literature on oil and energy shocks by providing plant-level evidence on how manufacturers respond to disruptions in energy markets. A long-standing literature documents the classic "oil puzzle": oil shocks appear to generate large output losses despite oil's small share of expenditures (Kilian, 2008; Hamilton, 2013).[3] Existing explanations for the oil puzzle emphasize monetary policy responses,

[2] Ganapati et al. (2020); Martin et al. (2014) examine how energy prices affect the manufacturing sector, while Abeberese (2017) examines the effects of electricity prices in India.

[3] Recent evidence from the European energy crisis, however, points to more moderate aggregate output

sectoral reallocation frictions, and amplification through production networks (Bernanke et al., 1997; Davis and Haltiwanger, 2001; Baqaee and Farhi, 2019). Because much of the empirical evidence comes from time-series or industry-level analyses, less is known about how firms and plants adjust to energy shocks. We address this gap by analyzing plant-level responses to a persistent electricity supply shock. We find little evidence of adjustment along the margins we examine, including outsourcing, product switching, fuel substitution, and reallocation across firms.

This paper also contributes to research on nuclear phase-outs and their economic consequences.[4] Major nuclear accidents have repeatedly reshaped energy policy worldwide. After Chernobyl, Italy exited nuclear power; after Fukushima, Germany accelerated its nuclear phase-out, Japan rapidly shut down nuclear generation, and countries such as Switzerland and South Korea adopted or considered policies to reduce their reliance on nuclear power. A growing literature shows that nuclear shutdowns can impose sizable costs by removing low-marginal-cost generation and increasing reliance on fossil fuels. These changes can raise electricity prices (Davis and Hausman, 2016), increase emissions, and generate health damages through local air pollution (Severnini, 2017; Jarvis et al., 2022) and higher energy costs (Neidell et al., 2021; He and Tanaka, 2023). We extend this literature by studying a less-examined downstream effect: industrial production. Using plant-level data from Japan, we show that the rapid post-Fukushima shutdown of nuclear generation led to persistent electricity price increases and substantial manufacturing output losses.

2 Institutional Background

During our study period, Japan's electricity sector was heavily regulated, so the post-Fukushima nuclear phase-out generated substantial regional variation in electricity supply shocks. Japan's mainland was divided into nine regional markets, each dominated by a traditional utility company. These utilities supplied nearly all grid electricity, accounting for more than 90% of local electricity supply, while wholesale trading was negligible (JEPIC, 2019).

Japan's electricity system had two institutional features that amplified this regional variation. First, regulation limited interregional transmission and trade, preventing local

effects (Moll et al., 2023; Bachmann et al., 2024).

[4]Existing studies document the short- and long-term impacts of nuclear accidents and related disasters. Such events can disrupt supply chains (Carvalho et al., 2021), affect health (Almond et al., 2009; Danzer and Danzer, 2016), alter perceptions of nuclear risk and induce anxiety (Fu et al., 2026), and reduce housing and land prices (Kawaguchi and Yukutake, 2017).

electricity supply shortfalls from being smoothed across regions. Under the Electricity Business Act, each utility was allowed to monopolize electricity supply within its service area and was legally responsible for maintaining local supply reliability (Government of Japan, 1964). Utilities therefore had little incentive to prioritize exports to neighboring regions when doing so could threaten supply adequacy for their own customers. As a result, even when Japan experienced nationwide electricity supply shocks, interregional transfers remained extremely small, often below 1% of peak-day demand (see Section 4.3).[5]

Second, utilities operated under cost-of-service regulation, causing supply shocks to translate into asymmetric regional price changes. Retail electricity prices were set to recover realized service costs plus a regulated return. Because generation mixes differed markedly across regions, especially in their reliance on nuclear power, the post-Fukushima nuclear shutdown raised costs unevenly across utilities. Under cost-of-service regulation, these higher generation costs were largely passed through to consumers, but only after regulatory approval processes that often lasted several months to more than a year.

Prior to the Fukushima Daiichi nuclear accident, Japan's electricity market operated without major disruptions. In 2010, Japan's nine mainland regional electric utilities generated approximately 796 TWh of electricity. Non-residential users accounted for roughly 71% of electricity consumption. The generation mix was diversified: fossil-fuel thermal generation accounted for about 58% of electricity generation, while nuclear power accounted for about 35%, hydropower for about 8%, and renewables for close to 0 %. The share of nuclear power varied substantially across regions, ranging from 7% in Chugoku and 15% in Chubu to 51% in Hokkaido and 53% in Shikoku (METI, 2011a).

The Great East Japan Earthquake and subsequent Fukushima Daiichi nuclear accident in March 2011 abruptly disrupted this stability. The earthquake killed more than 15,000 people and triggered a tsunami that inundated the Fukushima Daiichi nuclear power plant. The flooding disabled cooling systems, leading to core meltdowns in multiple reactors and the release of radioactive materials. The accident immediately changed public and regulatory attitudes toward nuclear safety. 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 were offline, and most remained suspended for several years. Panel A of Figure 1 illustrates the sharp and persistent collapse in nuclear capacity utilization that followed. As a result, nuclear power accounted for only 2% of total electricity generation in 2013 and about 8% in 2018 (METI, 2011a).

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

Because nuclear power accounted for about 35% of electricity generation among the nine mainland regional utilities, and more than 50% in some utility regions, its removal precipitated a nationwide supply crisis with uneven regional exposure. The electricity shortfall was particularly severe during summer peak periods, when elevated air-conditioning demand increased the risk of blackouts. The severity of the shock varied substantially across regions. Regions with greater losses of nuclear generation and lower initial reserve margins faced especially large shocks, while others were less affected.

Policymakers responded to the electricity supply shock with a combination of demand-side and supply-side measures. On the demand side, the government prioritized avoiding blackouts through "physical rationing." In the summer of 2011, it also temporarily required large establishments in the Tokyo and Tohoku service areas to reduce peak electricity usage by 15% relative to 2010 (METI, 2011b).[6] The government also relied heavily on conservation campaigns and moral suasion through a national movement, encouraging manufacturers to improve energy efficiency, dim lighting, reduce air- conditioning use, adjust work schedules, and shift operations away from peak hours (Kakuchi, 2011). These campaigns appear to have induced electricity conservation in the residential sector (He and Tanaka, 2023), yet their effectiveness in the manufacturing sector remains unclear.

On the supply side, utilities reactivated retired and inefficient fossil-fuel plants, expanded the use of existing thermal generators, and constructed new thermal capacity. Panel B of Figure 1 shows aggregate electricity generation by energy source (METI, 2011a). Thermal generation partially replaced the lost nuclear baseload generation, but total generation remained below its pre-accident level over the medium run, while renewable power made only a limited contribution.

Replacing lost nuclear generation with thermal generation raised retail electricity prices because thermal plants had higher marginal costs. Under Japan's cost-of-service regulatory regime, these higher generation costs were largely passed through to consumers. Regions facing larger supply deficiencies experienced larger spikes in electricity prices because they had to rely more heavily on higher-cost thermal plants. Panel C shows that electricity prices rose through 2014 and declined thereafter, broadly tracking global fuel markets, but with substantial regional dispersion: regional differences in price growth reached up to 20%. Electricity supply shocks in our setting therefore also operated as input cost shocks to manufacturers.

[6]The order was issued under Article 27 of the Electricity Business Act and applied to large users with contract demand of at least 500 kW in the Tokyo and Tohoku service areas during weekday daytime peak hours. Noncompliance was subject to fines of up to one million yen; this was the first invocation of Article 27 since the 1974 oil-shock measures.

3 Data

3.1 Peak-Day Electricity Demand and Supply

We assemble regional-level data on peak-day electricity demand and supply from 2005 to 2018 to measure electricity supply shocks. Our primary data source is the Electricity Supply–Demand Verification Subcommittee (ESDV, 2012), [7] a government body that convenes annually to assess summer and winter electricity supply conditions. The subcommittee collects detailed information on regional supply capacity and demand on peak days, which the national government uses to design and implement interventions in electricity markets. The Japanese government uses peak-day level rather than the peak-hour level for the assessment.

3.2 Census of Manufactures

We use plant-level administrative data from the Census of Manufactures (CMF) for 2005–2018, compiled by the Ministry of Economy, Trade, and Industry (METI) (METI, 2012).[8] The census annually surveys all manufacturing establishments with 30 or more employees, covering approximately 80% of total manufacturing output in Japan. Japan's CMF is comparable to manufacturing census data used in studies of the United States and other major economies in scope and structure.

The CMF provides detailed plant-level information on outputs, inputs, and production characteristics. We measure output as revenue adjusted for inventory changes. We construct capital stock as the reported book value of land, buildings, and machinery, net of depreciation. We compute capital rent by multiplying each plant's deflated capital stock by an industry-year rental rate. We measure electricity purchases,[9] materials, and labor using expenditure data, and deflate all monetary values. Because the data do not report plant-level electricity prices, we deflate electricity expenditures using average regional electricity prices. This approach is appropriate in Japan's heavily regulated electricity market, where prices were largely set by regional utilities through tariff schedules (JEPIC, 2019). We describe the detailed procedures in Appendix A-1.1. Each plant is assigned a two- to four-digit industry code.[10] Throughout the paper, we use the three-digit classification as

[7] We supplement these data with information from the Agency for Natural Resources and Energy (ANRE) (ANRE, 2022) and utilities' power supply planning documents.

[8] For 2011 and 2015, corresponding manufacturing data are collected through the Economic Census for Business Activity, conducted jointly by the Ministry of Internal Affairs and Communications and METI.

[9] Throughout the main CMF analysis, electricity use refers to electricity purchased from regional utilities. The CMF does not report electricity generated on site.

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

our baseline definition of an "industry" unless otherwise noted.

We construct the analysis sample using three main restrictions. First, we exclude establishments located in municipalities severely affected by the Great East Japan Earthquake and the Fukushima nuclear accident, using an affected-area definition consistent with Carvalho et al. (2021). Second, for each outcome regression, we restrict the sample to establishments observed in both the baseline year, 2010, and the outcome year, so that the dependent variable, the change in log outcomes relative to 2010, is defined. Third, we restrict the CMF sample to establishments with nonmissing and positive values for output and primary input expenditures, including labor, materials, electricity, and capital. We winsorize these output and input variables at the 1st and 99th percentiles to mitigate the influence of outliers.

The CMF includes several additional features useful for our analysis. Firm identifiers allow us to link establishments owned by the same firm and study whether activity shifts across regions within multi-plant firms. A subset of establishments reports product-level information, enabling us to analyze product switching. Plants may respond by switching to products that require less electricity. To examine this possibility, we define a product-switching indicator and measure each plant-year's product-mix electricity intensity. We detail the construction of these product-level measures in Appendix A-1.2.

Lastly, we use publicly available aggregate plant-count data from METI in 2006, 2009, 2014, and 2016 to examine plant turnover. We rely on these aggregate data because our restricted-access microdata include only plants with 30 or more employees. The aggregate data report plant counts, new establishments, and closed establishments by prefecture, industry, and survey wave, which we map to electricity regions. Because the data do not distinguish plant entry and exit from relocation across regions, we interpret them as measures of plant turnover.

3.3 Energy Consumption Statistics

We complement the CMF with plant-level administrative data from the Energy Consumption Statistics (ECS) provided by ANRE.[11] These surveys report physical consumption of fuel inputs, including natural gas (city gas, LPG, and LNG), gasoline, kerosene, heavy oil, diesel oil, and steam.[12] 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

[11] We also use data from the Petroleum and Other Fuels Monthly Statistics (POFMS). These two surveys collect similar variables but differ in industry coverage: POFMS focuses on a subset of energy-intensive industries, while ECS covers a broader set of industries.

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

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. Although the ECS is representative, its sampling frame differs from the CMF in industry coverage and timing.[13]

3.4 Other datasets

We supplement the plant-level data with additional regional data for robustness checks. To account for supply-chain disruptions from the earthquake and tsunami, we use prefecture-by-industry input-output tables from RIETI (RIETI, 2005). We also collect prefecture-level data on local economic conditions, including gross prefectural product, population, construction investment, consumption, and labor- market conditions (MIC, 2025). See details in Appendix A-3.

3.5 Summary Statistics

Table 1 presents summary statistics. In the CMF (N = 437,561), the average plant generates approximately 5,067 million yen in annual output (about $50–60 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 51% of revenue, respectively, while electricity represents a relatively small share at 2.0%. Among establishments covered by the ECS and POFMS (N = 381,761), approximately 8% own self-generation capacity. Across all establishments, plants purchase 20 TJ of electricity from the grid and generate 3.6 TJ on site, on average. Among generator-owning establishments, grid purchases average 68 TJ and self-generation averages 44 TJ, so self-generation accounts for roughly 39% of total electricity use.

4 Empirical Strategy

4.1 Conceptual Framework

We begin with a conceptual framework that guides our empirical analysis. Manufacturing plants produce output (Y) using capital (K), labor (L), materials (M), and electricity (E),

[13]The ECS covers fewer industries overall, but it includes plants that operate self- generators even when employment is below 30 workers.

together with Hicks-neutral productivity Ω. The production function is

$$Y = \Omega F(K, L, M, E), \tag{1}$$

where $F(\cdot)$ is differentiable and exhibits constant returns to scale. For this benchmark, suppose plants take output and input prices as given.

Differentiating $\ln Y$ with respect to electricity supply shock severity $\mathcal{S}$ yields

$$\frac{d \ln Y}{d\mathcal{S}} = \theta_E \frac{d \ln E}{d\mathcal{S}} + \theta_M \frac{d \ln M}{d\mathcal{S}} + \theta_L \frac{d \ln L}{d\mathcal{S}} + \theta_K \frac{d \ln K}{d\mathcal{S}} + \frac{d \ln \Omega}{d\mathcal{S}}, \tag{2}$$

where $\theta_X \equiv \partial \ln F(\cdot)/\partial \ln X$ denotes the output elasticity with respect to input $X \in \{K, L, M, E\}$. Appendix B-2 provides the detailed derivation.

Equation (2) shows that electricity supply shocks may affect output through several margins. They may reduce electricity use directly, but they may also affect labor, materials, capital, and productivity. Our empirical analysis therefore estimates the effects of electricity supply shocks on output, electricity use, other inputs, and productivity, except for capital, which is measured with substantial error.

This framework provides a benchmark for the direct output effect of an electricity supply shock. Consider a small, transitory, and unexpected shock, and evaluate the marginal product of electricity at the plant's pre-shock optimum. Holding labor, materials, capital, and productivity fixed at L^*, M^*, K^*, and Ω^*, respectively, the direct effect is approximated by

$$\left.\frac{d \ln Y}{d\mathcal{S}}\right|_{L=L^*, M=M^*, K=K^*, \Omega=\Omega^*} \approx \theta_E^* \frac{d \ln E}{d\mathcal{S}}. \tag{3}$$

The first-order condition links the electricity output elasticity to the electricity revenue share. Since electricity accounts for only about 2% of revenue on average, this benchmark predicts modest output losses from electricity reductions alone. This helps explain why the existing literature often finds modest, and sometimes null, output impacts of input shocks in electricity markets (von Graevenitz and Rottner, 2024; Wolverton et al., 2022).

This benchmark is not a full prediction of the plant response to the shocks we study. Electricity supply shocks after the Fukushima accident were large, persisted for years, and became increasingly predictable after the near-complete shutdown of nuclear reactors. Over such a horizon, plants can re-optimize materials, labor, and, to some extent, capital.

4.2 Econometric Specification

We estimate the effects of electricity supply shocks on the outcomes in equation (2). We begin with output and estimate the following specification:

$$\Delta \ln Y_{ijrt} = \rho_1 \mathcal{S}_{rt} \mathbb{1}\{2011 \leq t \leq 2012\} + \rho_2 \mathcal{S}_{rt} \mathbb{1}\{2013 \leq t \leq 2018\} + \mu_{jt} + \epsilon_{ijrt}, \quad (4)$$

where $\Delta \ln Y_{ijrt} = \ln Y_{ijrt} - \ln Y_{ijr,2010}$ is the change in plant-level log output relative to 2010, the year before the nuclear accident, for plant i in industry j, region r, and year t. The variable $\mathcal{S}_{rt}$ denotes the intensity of the local electricity supply shock, μ_{jt} denotes industry-by-year fixed effects, and ϵ_{ijrt} is the error term.

The coefficients ρ_1 and ρ_2 capture the effects of the shock, $d \ln Y / d\mathcal{S}$, separately for the phase-out period (2011–2012) and the post-phase-out period (2013–2018). We focus primarily on the post-phase-out period because it captures persistent and severe supply shocks. In the immediate aftermath of the accident, reactors were gradually suspended and policy uncertainty remained high, so plants may have perceived the disruption as temporary. By 2013, however, nuclear power had been fully phased out, and the resulting electricity supply shocks persisted through 2018 and beyond.

We also estimate event-study-style coefficients to examine pre-trends and dynamic effects. Because electricity supply shocks vary over time, we use $\bar{\mathcal{S}}_r$, the mean region-level shock over 2013–2018, for pre-accident years and the year-specific realized shock, $\mathcal{S}_{r\tau}$, for post-accident years. We estimate

$$\Delta \ln Y_{ijr\tau} = \sum_{t \neq 2010} \beta_t \mathbf{1}\{\tau = t\} \left[\mathbf{1}\{t < 2011\}\bar{\mathcal{S}}_r + \mathbf{1}\{t \geq 2011\}\mathcal{S}_{r\tau}\right] + \mu_{jt} + \varepsilon_{ijr\tau}, \quad (5)$$

where τ indexes calendar years in the sample. Statistically insignificant estimates of β_t in pre-accident years support the assumption that high- and low-shock regions followed similar outcome trajectories before the accident. The post-accident coefficients describe how the relationship between realized supply shocks and the outcome evolves over time.

We estimate the effects on electricity use, materials, labor, and productivity analogously by replacing output with each corresponding outcome variable. These estimates correspond to the input and productivity terms in equation (2).

We weight regressions by the log of plant-level total salary in 2010. Moreover, 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.

4.3 Constructing a Measure of Electricity Supply Shocks

We construct a measure of regional electricity supply shocks based on peak-day supply-demand conditions. Peak days are central to blackout risk because electricity supply must meet demand in real time, including on the highest-demand days.

Regional shock intensity is defined as the gap between required and available supply, normalized by pre-accident peak demand. Specifically, we define the regional shock as:

$$\mathcal{S}_{rt} = \frac{1}{D^{\text{peak}}_{r,2010}} \Big(\underbrace{1.05\, D^{\text{peak}}_{r,2010}}_{\text{required supply}} - \underbrace{S^{\text{plan,peak}}_{rt}}_{\text{available supply}} \Big) \cdot \mathbb{1}\left\{ S^{\text{plan,peak}}_{rt} < 1.05\, D^{\text{peak}}_{r,2010} \right\}. \tag{6}$$

where $D^{\text{peak}}_{r,2010}$ is realized peak-day demand in region r in 2010. Required supply equals this pre-accident peak-day demand plus a 5% reserve margin. Reserve margins absorb demand fluctuations driven by weather and economic activity; we follow the government's 5% reserve-margin threshold (ESDV, 2012). $S^{\text{plan,peak}}_{rt}$ is planned available supply capacity for the peak day in year t. Available supply declines when existing nuclear capacity is phased out and increases when supply-side measures, such as restarting retired thermal plants or constructing new thermal plants, are taken. The indicator sets the shock to zero in region-years in which available supply exceeds required supply.

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.

Appendix Table B1 reports peak-day supply from different energy sources, demand, and interregional transmission volumes for four major regions. Total supply declined relative to 2010 in all regions, largely because of the nuclear phase-out. Transmission was severely limited.[14] When Tokyo faced its most severe supply deficits, it received virtually no electricity from other regions or non-utility suppliers. 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).

[14]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.

4.4 Instrumenting for Electricity Supply Shocks

Regional electricity supply shocks may be endogenous because of post-accident policy responses. Some regions reactivated dormant thermal plants or restarted nuclear reactors. These supply-side adjustments may be correlated with unobserved local economic conditions, such as recovering industrial demand or political influence.

To address this concern, we use an instrumental-variables approach. Regions that relied more heavily on nuclear power before the accident faced larger electricity supply shocks because nuclear generation was almost completely phased out. At the same time, regions with larger pre-accident reserve margins had greater buffer capacity and were better able to absorb the loss of nuclear supply. We therefore construct the instrument vector, $\mathbf{Z}_{rt}$, from predicted nuclear-supply losses and initial reserve margins:

$$\mathbf{Z}_{rt} = \left(Sh^{\text{nuclear}}_{r,2010} \times g^{\text{nuclear}}_{t,-r},\ Sh^{RM}_{r,2010} \right). \tag{7}$$

Here, $Sh^{\text{nuclear}}_{r,2010}$ denotes the 2010 nuclear-generation share in region r, and $g^{\text{nuclear}}_{t,-r}$ is the leave-one-out national growth rate of nuclear generation. The second instrument, $Sh^{RM}_{r,2010}$, denotes region r's reserve margin in 2010. All variables are normalized by 2010 peak-day demand to match the scale of the shock measure.

4.5 Production Function Estimation

4.5.1 CES Production Function

We estimate a constant-elasticity-of-substitution (CES) production function. We assume constant returns to scale and that plants take output and input prices as given. We impose the standard timing assumptions: electricity and materials are flexible inputs chosen after plants observe productivity, whereas labor and capital are quasi-fixed inputs chosen before the current productivity innovation is realized.

The production function for plant i in year t is

$$Y_{it} = \exp(\omega_{it}) \exp(\epsilon_{it}) F(K_{it}, L_{it}, M_{it}, E_{it}), \tag{8}$$

$$F(K_{it}, L_{it}, M_{it}, E_{it}) = \left(\alpha_K K^{\rho}_{it} + \alpha_L L^{\rho}_{it} + \alpha_M M^{\rho}_{it} + \alpha_E E^{\rho}_{it} \right)^{\frac{1}{\rho}}, \qquad \sigma \equiv \frac{1}{1-\rho}.$$

Here, σ denotes the elasticity of substitution and the share parameters satisfy $\sum_j \alpha_j = 1$.

In logs, the production function can be written as

$$y_{it} = f(k_{it}, l_{it}, m_{it}, e_{it}) + \omega_{it} + \epsilon_{it}, \tag{9}$$

where lowercase variables denote logs, $f(\cdot)$ is the log of the CES aggregator, ω_{it} is plant-level log productivity, and ϵ_{it} is an idiosyncratic ex post production shock satisfying $\mathbb{E}[\epsilon_{it} \mid k_{it}, l_{it}, m_{it}, e_{it}] = 0$. Productivity follows a first-order Markov process:

$$\omega_{it} = h(\omega_{it-1}) + \eta_{it}, \tag{10}$$

where the innovation η_{it} is orthogonal to the quasi-fixed inputs, so that $\mathbb{E}[\eta_{it} \mid k_{it}, l_{it}] = 0$.

4.5.2 Estimation Strategy

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 each industry.

We recover the CES parameters in two steps. First, we estimate industry-specific output elasticities for each input. We identify the elasticities of the flexible inputs, electricity and materials, using static first-order conditions following Gandhi et al. (2020), and estimate those of the quasi-fixed inputs, labor and capital, using a control-function approach (e.g., Olley and Pakes, 1996; Ackerberg et al., 2015). Second, conditional on a candidate elasticity of substitution (σ), we recover the CES distribution parameters (α) from the estimated output elasticities. Under constant returns to scale and cost minimization, an input's normalized output elasticity equals its share of total input cost. The CES first-order conditions then map these cost shares into the distribution parameters.

This approach offers three advantages. First, it allows us to estimate a gross-output production function using static first-order conditions. As Gandhi et al. (2020) show, standard proxy-variable approaches have difficulty identifying flexible- input elasticities in gross-output specifications because they require additional variation in flexible-input demand that is often unavailable (e.g., Olley and Pakes, 1996; Ackerberg et al., 2015). Second, the first step delivers output elasticities, so the implied CES share parameters can be recovered for alternative values of σ without re-estimating the output elasticities. Third, estimating the model separately by three-digit industry accommodates technological heterogeneity across manufacturing industries. This heterogeneity is central to the counterfactual analysis because we compare the welfare impacts of physical rationing and pricing, the two main policy responses to electricity supply shocks. Without technological heterogeneity, pricing and quantity restrictions would have identical welfare consequences.

4.5.3 Estimating Output Elasticities

Flexible Inputs In the first stage, we estimate output elasticities of the flexible inputs: $\theta_E \equiv \partial f/\partial e$ (electricity) and $\theta_M \equiv \partial f/\partial m$ (materials). For a flexible input, the first-order condition implies that its log revenue share equals its log output elasticity minus a mean-zero idiosyncratic shock. For electricity,

$$\ln s_{E,it} = \ln \theta_E - \epsilon_{it}, \tag{11}$$

where $s_{E,it} \equiv \frac{P_t^E E_{it}}{P_{it} Y_{it}}$ and $P_{it} Y_{it}$ is nominal revenue. Equation (11) implies that θ_E is identified by the conditional mean of the log electricity expenditure share.[15] We obtain θ_M analogously using the materials share.

Quasi-fixed Inputs In the second stage, we estimate the output elasticities of labor and capital, θ_L and θ_K. Using the first-stage estimates, we construct the partial residual

$$\widehat{\mathcal{Y}}_{it} \equiv y_{it} - \hat{\theta}_E e_{it} - \hat{\theta}_M m_{it}, \tag{12}$$

and define implied productivity as

$$\hat{\omega}_{it}(\theta_L, \theta_K) \equiv \widehat{\mathcal{Y}}_{it} - \theta_L l_{it} - \theta_K k_{it} - \hat{\epsilon}_{it}. \tag{13}$$

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_L and θ_K by GMM using the Markov structure in (10) and moment conditions based on the quasi-fixed inputs, $\mathbb{E}[\eta_{it} \cdot l_{it}] = 0$ and $\mathbb{E}[\eta_{it} \cdot k_{it}] = 0$.

4.5.4 Recovering CES Distribution Parameters

For each industry, we first normalize the estimated output elasticities to obtain input cost shares:

$$\hat{s}_j = \frac{\hat{\theta}_j}{\sum_{k \in \{K,L,M,E\}} \hat{\theta}_k}, \qquad j \in \{K, L, M, E\}. \tag{14}$$

Under constant returns to scale and cost minimization, these normalized output elasticities correspond to input cost shares. Conditional on a candidate elasticity of substitution σ,

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

we recover the CES distribution parameters as

$$\hat{\alpha}_j(\sigma) = \frac{\hat{s}_j^{1/\sigma}}{\sum_{k\in\{K,L,M,E\}} \hat{s}_k^{1/\sigma}}, \qquad j \in \{K, L, M, E\}. \tag{15}$$

This parameterization satisfies $\sum_j \hat{\alpha}_j(\sigma) = 1$ and reproduces the estimated cost shares. Appendix B-3.2 provides the detailed derivation.

4.5.5 Production Function Results

Table 2 reports output elasticities estimated using pre-accident data from 2005–2010. Columns (1)–(3) present our estimates following Gandhi et al. (2020). We estimate the flexible-input elasticities in the first step using expenditure-share equations. The mean elasticities are 0.446 for materials and 0.015 for electricity. We then jointly estimate the quasi-fixed-input elasticities by GMM. The mean estimates are 0.354 for labor and 0.139 for capital. The electricity elasticity varies considerably relative to its mean, ranging from 0.008 at the 25th percentile to 0.017 at the 75th percentile.

Column (4) presents the OLS benchmark. OLS estimates an electricity elasticity of 0.068, more than four times our estimate of 0.015. It also produces larger estimates for materials (0.492 versus 0.446) and labor (0.423 versus 0.354), but a smaller estimate for capital (0.042 versus 0.139). The larger OLS estimates for flexible inputs are consistent with upward simultaneity bias: firms increase materials and electricity use when they experience positive productivity or demand shocks that OLS does not observe, causing OLS to overstate their output elasticities. By contrast, measurement error in capital attenuates its estimated relationship with output and causes OLS to understate the capital elasticity.

5 Electricity Supply Shocks and Plant Outcomes

5.1 Trends in Inputs and Outputs of Manufacturing Plants

We begin by describing aggregate trends in manufacturing inputs, output, and productivity among plants in high- and low-shock regions. We classify regions with mean 2013–2018 shock intensity at or above the median as high-shock regions and the remaining regions as low-shock regions. Figure 3 plots log changes from 2010 levels over the 2005–2018 period.

Figure 3 shows that, following the 2011 accident, manufacturing inputs and output declined in high-shock regions relative to low-shock regions. In Panel A, electricity use followed similar trends in the two groups before 2011. After the accident, however,

electricity use fell substantially in high-shock regions but changed only modestly in low-shock regions. This divergence widened as nuclear reactors were gradually shut down during 2011–2012 and persisted throughout the complete-shutdown period from 2013 to 2018. Panels B–D show similar, though smaller, divergences in employment, total salary, and material inputs. By contrast, Panel E shows no comparable divergence in productivity. Panel F shows a persistent relative decline in output in high-shock regions after 2011, consistent with the reductions in labor and material inputs, which together account for a large share of production costs.

We also examine the relationship between the intensity of regional electricity supply shocks and changes in electricity use and output. For Figure C1, we first residualize plant-level log changes from 2010 levels with respect to industry-by-year fixed effects and then aggregate the residuals to the nine regions. Panel A shows a negative relationship between supply shocks and residualized changes in electricity use: regions exposed to larger shocks experienced larger declines after the accident. This relationship is stronger during the complete-shutdown period (2013–2018) than during the phase-out period (2011–2012). Panel B shows no corresponding relationship before the accident, indicating that regions subsequently exposed to larger shocks did not experience systematically different pre-accident changes in electricity use. Panel C shows a similar pattern for output: larger shocks are associated with greater declines in output during the complete-shutdown period, but not during the phase-out period. Panel D likewise shows no relationship between mean supply shocks and relative changes in output before the accident.

5.2 First-stage Results: Testing the Instruments

We estimate first-stage relationships between our instruments and regional electricity supply shocks to assess the relevance. We use region-by-year data from 2011 to 2018 and regress local shock intensity on predicted changes in nuclear generation and the initial reserve margin, conditional on year fixed effects.

Table 3 shows that predicted changes in nuclear generation are negatively associated with local electricity supply shocks. For 2013–2018, when nuclear generation remained almost entirely offline, the estimated coefficient is −0.26 (Column (2)). Thus, a 1 percentage point (ppt) decline in predicted nuclear generation relative to 2010 peak-day demand corresponds to a 0.26 ppt increase in the local supply shock. During this period, the average region lost nuclear generation equivalent to approximately 17% of its 2010 peak-day demand,[16] implying a 4.4 ppt increase in the local supply shock. In the most exposed

[16]In 2010, nuclear generation accounted for an average of 18% of peak-day demand, of which 97% was lost during 2013–2018.

regions, where the predicted decline in nuclear generation reached 30% of peak-day demand, the implied increase in the supply shock was approximately 8 ppt.

The initial reserve margin ($Sh^{\text{RM}}_{r,2010}$) is negatively associated with local supply shocks. This relationship is consistent with our conjecture that pre-existing capacity buffers helped regions mitigate these shocks. The initial peak-day reserve margin averaged 10%. Given the estimated coefficient of -0.82, the model predicts that a region with a 5% reserve margin would have experienced a supply shock approximately 4.1 percentage points larger during 2013–2018.

The first stage is strong despite the fact that we have only nine independent clusters. The first-stage F-statistic is 30 for the phase-out period (2011–2012) and 26 for the post-phase-out period (2013–2018).[17]

5.3 The Effects of Supply Shocks on Manufacturing Outcomes

5.3.1 Impacts on Electricity Inputs

This section examines how electricity supply shocks affect manufacturing inputs, productivity, and output. Figure 4 presents event- study estimates based on Equation (5). We regress changes in log outcomes relative to 2010 on local electricity supply shocks. We report both OLS and 2SLS estimates. The 2SLS specification instruments local supply shocks with predicted nuclear-generation losses and initial reserve margins. The figure displays the OLS estimates as blue circles and the 2SLS estimates as red diamonds, together with 90% confidence intervals.

Panel A of Figure 4 shows that electricity supply shocks reduced electricity use in manufacturing. Before the accident, the estimated effects are statistically indistinguishable from zero. This pattern suggests similar pre-trends in regions with high and low supply shocks. Following the accident, electricity use declined, and the effect grew over time. The effect was small in 2011, when firms may have perceived the shock as temporary. It increased in 2012 and persisted through 2018. The OLS and 2SLS estimates are similar in magnitude and statistical precision.

Table 4 reports the OLS and 2SLS estimates from Equation (4). During the phase-out period (2011–2012), the 2SLS estimates in Panel B imply that a 1-percentage-point increase in shock intensity reduced electricity use by approximately 0.21%. During the post-phase-out period (2013–2018), the magnitude of this effect rose to 0.89%. Given the mean post-phase-out shock intensity of 7.5%, this estimate implies an average decline in

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

manufacturing electricity use of approximately 6.7%. In a high-shock region such as Tokyo, where the mean shock intensity reached 10.8%, the implied decline was approximately 9.6%.

5.3.2 Impacts on Labor and Materials

Figure 4 show that electricity supply shocks reduce employment (Panel B), total salary (Panel C), and materials (Panel D), and these effects grow over time. We show the event-study results in equation (5) with OLS and IV estimates using these outcomes analogously. Across these outcomes, we do not find systematic differences in pre-trends and the OLS and IV estimates are similar in magnitude.

Columns (2)–(4) of Table 4 report the effects on labor, total salary, and materials. During the post-phase-out period (2013–2018), a 1 ppt increase in the supply shock reduced each outcome by 0.17% in the IV estimates (Panel B). Given the mean supply shock of 7.5%, the implied average declines are approximately 1.3%.

We further examine effects across worker types and find that non-regular workers experienced disproportionate employment losses. Appendix Table C1 shows little effect on wages per worker. The estimated coefficient for the post-phase-out period is -0.01 (Panel B, Column (1)), suggesting that firms adjusted employment rather than wages. The estimated employment effects are similar for male and female workers: -0.17 in Column (2) and -0.20 in Column (3), respectively. However, the effects differ substantially between regular and non-regular workers. The coefficient for regular workers is small and statistically insignificant at -0.07 (Panel B, Column (4)). By contrast, the coefficient for non-regular workers is -0.90 (Panel B, Column (5)), more than ten times as large in magnitude. These results indicate that firms adjusted employment primarily by reducing non-regular positions. Although non-regular workers accounted for only 10% of manufacturing employment, they accounted for nearly 60% of the implied employment decline.

We use a back-of-the-envelope calculation to quantify the annual employment loss. In 2010, manufacturing plants employed 7.7 million workers. Applying our post-phase-out estimate to this baseline implies that supply shocks reduced manufacturing employment by approximately 97,000 workers per year during 2013–2018.[18]

[18]The employment total covers manufacturing plants with at least four employees (METI, 2012). To obtain this figure, we assume that the post-phase-out estimate applies to the 2010 employment aggregate.

5.3.3 Impacts on Productivity

We estimate the effects of electricity supply shocks on TFP. Theory offers competing predictions. TFP may remain unchanged under a neoclassical model in which productivity is exogenous. Alternatively, it may decline if electricity shocks distort the input mix or lead to underutilization of quasi-fixed inputs, consistent with evidence from past oil shocks (Kilian, 2008). Productivity may also rise if unrealized efficiency gains exist and shocks induce manufacturing plants to improve efficiency by adopting electricity-saving technologies or better managerial practices. Anecdotal evidence suggests that Japanese manufacturers undertook extensive energy-conservation efforts, potentially resulting in efficiency gains (Kimura and Nishio, 2016). We therefore estimate the total effect of electricity supply shocks on TFP, which captures the net contribution of these competing mechanisms

Panel E in Figure 4 provides no evidence that productivity improved in response to the shocks. The OLS and IV estimates are close to zero and statistically insignificant throughout the study period. Column (5) of Table 4 likewise reports very small point estimates, ranging from -0.02 to 0.01 across specifications and time periods.

Appendix C-2 examines whether electricity supply shocks induced electricity-biased technical change. Because the electricity–labor ratio depends on both relative factor prices and relative productivity, we remove the variation explained by relative prices to construct a descriptive measure of technical change. This measure evolves similarly in high- and low-shock regions, providing little evidence that more severe supply shocks induced electricity-biased technical change.

5.3.4 Impacts on Output

Electricity supply shocks significantly reduce output, with effects that grow over time through 2015 and then persist thereafter as described in Panel F in Figure 4. As is similar with results for labor and material, we find no significant pre-trends, and the results of the OLS and IV are similar.

Column (6) of Table 4 shows that the IV estimate during the post-phase-out period is -0.23 (Panel B), while the phase-out-period estimate is small and statistically insignificant. Given the mean shock of 7.5%, the estimate implies that output in the average region fell by 1.7%.

A back-of-the-envelope calculation illustrates the economic magnitude. In 2010, manufacturing generated 289 trillion JPY in annual output (METI, 2012). Applying our post-phase-out estimate implies that supply shocks reduced annual manufacturing output by approximately 4.9 trillion JPY, on the order of $40–60 billion, depending on the exchange

rate.

5.4 Heterogeneous Impacts

The heterogeneity analysis sheds light on the channels through which electricity supply shocks affect manufacturing outcomes. We examine whether the effects vary with plant- and industry-level electricity and capital intensity, measured in 2010.[19] For each measure, we split the sample at the median, estimate Equation (4) separately for the two groups, and report the 2SLS estimates for the complete-shutdown period.

The effect of electricity supply shocks on electricity use varies little with baseline electricity or capital intensity, measured at either the plant or industry level. Panel A of Figure 5 shows that the subgroup estimates generally lie close to the full-sample estimate of −0.89. For each sample split, we fail to reject equality between the estimates for the high- and low-intensity groups. This limited heterogeneity is consistent with Japanese policy: plants within the same region generally faced common rationing requirements and electricity price increases, regardless of their characteristics.

By contrast, the output impacts vary sharply with capital intensity. Panel B of Figure 5 reports a coefficient of −0.25 for capital-intensive plants, compared with 0.00 for plants with low capital intensity. The corresponding coefficient is −0.38 for plants in capital-intensive industries, compared with 0.05 for those in less capital-intensive industries. We find no comparable heterogeneity by electricity intensity.

Electricity use declines similarly across subgroups, yet output falls more at capital-intensive plants and in capital-intensive industries. This contrast is consistent with electricity supply shocks amplifying output losses through capital utilization rather than through larger reductions in electricity use. The energy-shock literature emphasizes this mechanism: electricity shortages can reduce the operating hours of installed machinery, lowering effective capital services and output (Berndt and Wood, 1975; Finn, 2000). Motivated by this evidence, Section 7.3.1 incorporates electricity-dependent capital utilization into the production technology.

5.5 Alternative Explanations and Robustness Checks

5.5.1 Supply-Chain Disruptions and Local Demand Shocks

We examine whether supply-chain disruptions provide an alternative explanation for our finding that electricity supply shocks led to large output contractions. Previous research

[19] Electricity intensity is electricity expenditure divided by revenue, and capital intensity is the capital stock divided by revenue.

documents that damage from the Great East Japan Earthquake propagated through input–output linkages (Carvalho et al., 2021). Such disruptions could therefore have contributed to the estimated output declines.

We assess this alternative explanation by controlling for two measures of exposure to supply-chain disruptions (see Appendix A-3 for details on their data and construction). The first captures direct exposure as the share of intermediate inputs sourced from the three most severely affected prefectures: Iwate, Miyagi, and Fukushima. The second captures indirect exposure through upstream linkages, calculated using the Leontief inverse. We interact these pre-disaster, time-invariant exposure measures with year indicators, allowing the effects of supply-chain disruptions to vary flexibly over time.

Table C3 shows that controlling for supply-chain disruptions does not materially alter our estimates. Column (1) reproduces the main 2SLS coefficients during the post-phase-out period: -0.89 for electricity use and -0.23 for output. Controlling separately for direct or indirect exposure yields an electricity coefficient of -0.87 and an output coefficient of -0.20 in both specifications (Columns (2) and (3)). Including both exposure measures produces estimates of -0.88 for electricity use and -0.20 for output (Column (4)). Thus, the supply-chain specifications leave the electricity response essentially unchanged and produce economically meaningful output effects of similar magnitude.

We further investigate whether the observed contractions reflect localized demand shocks rather than supply-side constraints. We control for time-varying prefecture-level economic conditions, including log gross prefectural product in non-manufacturing sectors, total and working-age populations, construction investment, aggregate consumption, and the job-openings-to-applicants ratio. Adding these controls yields 2SLS coefficients of -0.75 for electricity use and -0.23 for output (Column (5)). Thus, the electricity response is somewhat attenuated, whereas the output point estimate is nearly identical to the baseline, though less precisely estimated.

5.5.2 Robustness Check

Our results are robust to plant-level controls, alternative specifications, and weighting schemes. Column (1) of Appendix Table C4 adds controls for pre-shock (2010) log revenue, electricity, capital, and labor cost shares, export intensity, and multi-plant firm status. The estimated effects remain stable after including these controls. Column (2) replaces the first-difference specification with a two-way fixed effects specification that includes plant and year fixed effects. Column (3) reports unweighted estimates, which are similar to those from our baseline specification weighted by log total salary.

6 Plants' Responses to Electricity Supply Shocks

6.1 Impacts on Observed Plant Behaviors

Why did electricity supply shocks lead to such large output contractions? To help answer this question, we examine how and to what extent manufacturing plants adapted to these shocks. If plants had limited capacity to adapt, this could partly explain the large observed impacts.

We begin by testing whether plants switched products or increased outsourcing. A natural hypothesis is that plants shifted toward less electricity-intensive products (von Graevenitz et al., 2025) or increased outsourcing to reduce their reliance on electricity (Fisher-Vanden et al., 2015). We use measures of product switching, product-level electricity intensity, outsourcing inputs, and the outsourcing share of material inputs to investigate these adjustment margins (see Appendix A-1.2 for details on data construction). We regress these outcomes on electricity supply shock intensity using both OLS (Panel A) and IV (Panel B), following the specification in equation (4).

We find no evidence that electricity supply shocks induced product switching or increased outsourcing. Table 5 shows that the IV estimates for the post-phase-out period (2013–2018) are small and statistically insignificant, with coefficients of 0.07 for product switching (Column (1)) and 0.03 for product-level electricity intensity (Column (2)), suggesting limited adjustment along the product-choice margin. Columns (3) and (4) likewise provide no evidence that the shocks increased either outsourcing inputs or the outsourcing share of material inputs. If anything, the point estimates are negative.

We further find no evidence that multi-plant firms reallocated activity across establishments, a margin of adjustment to localized input shocks suggested by previous studies (Chen et al., 2025). In our data, 42% of plants belong to multi-plant firms, accounting for 66% of output. Manufacturing firms may respond to localized shocks by shifting production from plants in high-shock regions to plants in low-shock regions, thereby partially maintaining firm-level production. To test this margin, we aggregate the data to the firm-by-region level and examine whether shocks in one region increase the same firm's activity in other regions. For example, shocks in Tokyo may induce the same firm to increase production in Osaka and Hokkaido (see Appendix A-1.2 for details on data construction). Appendix Table D1 provides limited evidence of such reshuffling. In the IV estimates for the post-phase-out period, shocks are associated with an increase in electricity use at the firm's other plants, with a coefficient of 0.84 (Panel B, Column (1)). However, we detect no corresponding response in output, with a coefficient of -0.04 (Column (2)). This absence of reshuffling in output, despite some increase in electricity use at other plants, suggests that firms faced substantial difficulty shifting production across plants.

Additionally, Appendix Table D2 shows no statistically significant effects of supply shocks on plant turnover. In the matched pre-accident sample, approximately 1.2% and 5.5% of manufacturing plants entered and exited the sector, respectively. The IV estimates for entry and exit are −0.10 and 0.18, respectively. At the mean shock of 7.5%, these point estimates correspond to approximately 0.1 fewer entrants and 0.7 additional exits per 1,000 plants ($1.17\% \times (-0.10) \times 7.5\% \times 1{,}000 \approx -0.1$ for entry and $5.46\% \times 0.18 \times 7.5\% \times 1{,}000 \approx 0.7$ for exit). These results suggest limited responses in plant turnover.

6.2 Substitution Between Electricity and Other Fuels

We examine whether plants substituted away from grid electricity and toward alternative energy sources, as electricity and other types of energy are often considered close substitutes. In theory, there are two types of substitution. First, manufacturing plants may switch to other fuels in production. Some electricity-using processes, such as those relying on motor-driven machinery or electric furnaces, may in principle be substituted with fuel-using processes, such as boilers, furnaces, or other combustion-based equipment, though such substitution may be costly. Second, plants may substitute indirectly by using alternative fuels for on-site electricity self-generation. Existing studies document increases in on-site electricity generation in response to electricity price increases (von Graevenitz and Rottner, 2024) and outages (Allcott et al., 2016).

To investigate whether electricity supply shocks induced greater fuel use for production and on-site generation, we construct two separate outcome variables: fuel used directly for production and fuel used for self-generation. We then estimate the main specification in equation (4) using these variables as outcomes, with both OLS and IV.

We find little evidence of substitution away from electricity toward fuel used in production. If these inputs were strong substitutes, electricity supply shocks should cause fuel use to decline less than other inputs, or even to increase. However, we find no such pattern. In Column (1) of Table 6, the post-phase- out IV estimate is −0.35 and statistically imprecise. This estimate is more negative than the corresponding estimates for labor and materials, both −0.17, which are generally considered complements to electricity (Section 5.3). These results provide no evidence that plants substituted fuels used directly in production for electricity.

In contrast, Columns (2) and (3) of Table 6 show that supply shocks increased on-site self-generation. In 2010, 1,872 of 30,776 manufacturing plants (6.1%) owned self-generators. These plants increased fuel use for on-site generation. Column (2) shows that the post-phase-out 2SLS estimate for fuel used for on-site generation is 4.50, implying about a 34% increase for a region exposed to the mean shock (4.50×0.075). The corre-

sponding estimate for self-generated electricity is 2.33 (Column (3)), implying an increase of approximately 18% (2.33×0.075).

Our heterogeneity results suggest that on-site generation could mitigate the adverse effects of electricity supply shocks. To examine this possibility, we estimate the specification in equation (4) separately for plants with and without generators, using grid electricity purchases and output as outcomes. In the CMF, electricity inputs measure purchases from the grid and exclude electricity generated on site. Appendix Figure D1 shows that grid electricity purchases declined in both groups. The post-phase-out IV coefficients are -1.03 for plants with generators and -0.86 for plants without generators, implying reductions of approximately 7.8% and 6.5%, respectively, at the mean shock. Despite the decline in grid purchases, output among plants with generators did not fall: the coefficient is 0.27 and is not statistically significant. In contrast, output among plants without generators declined, with a coefficient of -0.28. This pattern is consistent with on-site generation helping plants offset reductions in grid electricity purchases and mitigate the adverse effects on output.

We find little evidence of an extensive-margin response in generator ownership. Column (4) of Table 6 reports a post-phase-out IV estimate of 0.06. For a region exposed to the mean shock, this estimate implies an increase of approximately 0.0045, or 0.45 percentage points ($0.06 \times 0.075 \times 100$), equivalent to about 4.5 additional generator-owning plants per 1,000 plants. Thus, more than 90% of manufacturing plants would remain without on-site generation.

Overall, this section highlights the difficulty manufacturing plants faced in adapting to large-scale and persistent electricity supply shocks. Even using data through 2018, by which time the shocks were broadly expected to persist, we find limited evidence of adjustment through product switching, outsourcing, production reshuffling, plant entry and exit, or direct substitution toward other fuels. Self- generation is an important exception: plants with access to this margin increased on-site generation and experienced no detectable output contraction. Nevertheless, because only a small minority of plants generated electricity on site, most manufacturing plants remained unable to offset the supply shocks through this channel. This limited scope for adjustment could partly explain why electricity supply shocks generated large output losses.

7 From Reduced-Form Effects to Production Mechanisms

This section uses the production function estimated in Section 4.5 to interpret the reduced-form effects of electricity supply shocks. We predict the impacts of electricity supply shocks using this production function and compare these predictions with the reduced-

form estimates. This comparison sheds light on the mechanisms that could reproduce the reduced-form responses

7.1 Simulation Framework

We use the production function to predict electricity use, materials, labor, and output with and without the observed electricity supply shocks. The resulting differences correspond to the reduced-form effects estimated using Equation (2). We construct the simulation sample from the plants in the main electricity regression and restrict it to plants that can be matched to industry production-function estimates and regional price inputs. We use this resulting common sample for all simulated outcomes. We use 2010 as the baseline year and average conditions over 2013–2018 to represent the post-phase-out period.

We hold technologies fixed across the baseline and shock scenarios. Each plant i is characterized by its 2010 productivity, $\omega_{i,2010}$, output, $Y_{i,2010}$, and generator-ownership status, $G_{i,2010}$. For each industry j, we normalize the estimated output elasticities to construct the baseline cost shares, s^0_{hj}, for $h \in \{K, L, M, E\}$, where the superscript 0 denotes the baseline. We hold these shares fixed as the elasticity of substitution, σ, varies and recover the corresponding CES distribution parameters as described in Section 4.5.4.

We assume that each plant faces isoelastic demand,

$$Y_i = B_i P_i^{-\eta}, \tag{16}$$

where B_i is a plant-specific demand shifter and $\eta > 1$ is the absolute value of the price elasticity of demand. We calibrate B_i so that demand at the baseline price reproduces the plant's observed output in 2010. We set $\eta = 5$, which lies within the range of sector-specific elasticities of substitution estimated for U.S. manufacturing by Shapiro and Walker (2018).[20]

Electricity supply shocks affected manufacturing plants through two channels: higher electricity prices and quantity rationing. The prediction, therefore requires exogenous changes in these parameters. For the rationing channel, we require each plant to reduce its peak-season electricity use by its region's average supply shock over 2013–2018, $\bar{\mathcal{S}}_r$. Because the peak season lasts 90 days, this restriction applies for approximately one-quarter of the year. At the sample mean, it reduces peak-season electricity use by 7.5%. For the price channel, we estimate how regional supply shocks affected electricity prices. Electricity prices also rose in 2014 and fell in 2015–2016 as global energy prices changed. Our regression separates the price response to regional supply shocks from these common

[20] We hold plant-specific productivity ω_i and the demand shifter B_i fixed across scenarios.

movements. As shown in Table E2, a 1-percentage-point increase in shock intensity raises the log electricity price by 0.0112, approximately 1.12%. During the post-phase-out period, the weighted mean predicted log-price change is 8.4 log points, corresponding to an average price increase of approximately 9.0% in levels.

For each plant, we use its production technology and demand to predict profit-maximizing input use and output with and without the supply shocks. We hold technology and demand fixed across the two scenarios and vary only electricity prices and the rationing constraint. Because rationing applies only during the peak season, we solve the plant's optimization problem separately for the peak and off-peak seasons and aggregate the predictions to obtain annual outcomes. Table E1 summarizes the simulation inputs, and Appendix E-3 describes the procedure.

The simulation also allows responses to differ by access to self-generation. In Section 6.2, we show that plants with generators reduce their purchases from grid utilities but exhibit no meaningful declines in other inputs or output. Based on this evidence, we assume that these plants offset reductions in grid purchases through self-generation and maintain their baseline levels of non-electricity inputs and output.

7.2 Cobb-Douglas Technology

We first consider a standard Cobb–Douglas benchmark, which imposes a unit elasticity of substitution. We use the model to predict the effects of electricity supply shocks on electricity use, labor, materials, and output and compare these predictions with the reduced-form effects evaluated at the mean shock intensity in Section 5.3.

The Cobb–Douglas benchmark fails to reproduce the impacts of electricity supply shocks. As Figure 6 shows, the model predicts that these shocks reduce labor, materials, and output by only about 0.4% (Panel B), substantially less than the reduced-form estimates of 1.3–1.7% (Panel A). At the same time, it predicts an 8.9% decline in electricity use (Panel B), exceeding the reduced-form estimate of 6.7% (Panel A) by 2.2 percentage points. Thus, the model understates the contractions in other inputs and output while overstating the decline in electricity use. The Cobb-Douglas model allows plants to substitute away from electricity relatively easily, preventing the electricity shock from generating sufficiently large contractions in other inputs and output.

7.3 Amplification Mechanisms

7.3.1 Capital Idleness

The discrepancies between the reduced-form estimates and the Cobb–Douglas predictions reveal two limitations. First, the Cobb–Douglas model substantially underpredicts the output losses from electricity supply shocks. We address this limitation by allowing electricity shortages to idle installed capital. Consider a steel plant that operates electric furnaces. When electricity is unavailable, the plant must reduce furnace operating hours, leaving part of its capacity idle. The resulting decline in effective capital services lowers output and reduces the marginal products of labor and materials, prompting the plant to use less of both inputs. Capital idleness thus amplifies the output loss beyond the direct effect of reduced electricity use.

We introduce capital idleness for two reasons. First, the literature on energy shocks establishes capital utilization as an important mechanism underlying large fluctuations in aggregate output (see, e.g., Berndt and Wood, 1975; Finn, 2000). Electricity unavailability can similarly amplify output losses by reducing capital utilization. Second, the heterogeneity results in Section 5.4 support this mechanism. Electricity supply shocks cause disproportionately large output losses at more capital-intensive plants and in more capital-intensive industries, whereas their effects do not vary systematically with electricity intensity. This pattern is consistent with capital utilization as the underlying mechanism.

We therefore modify the model so that production depends on effective capital services, $\widetilde{K}$, rather than directly on the installed capital stock, K. For example, production at a steel plant depends on the services provided by its electric furnaces rather than simply on its installed furnace capacity. Electricity constraints can reduce furnace utilization and, consequently, the effective capital services available for production without changing the installed capital stock.

We introduce the effective-capital-utilization elasticity, κ, which measures the percentage reduction in effective capital services caused by a 1% reduction in electricity use, evaluated at the 2010 baseline. Specifically, we define

$$\widetilde{K}_{it} = \left[(1-\kappa)K_{it}^{\rho} + \kappa\,(v_i E_{it})^{\rho}\right]^{1/\rho}, \qquad v_i \equiv \frac{K_{i,2010}}{E_{i,2010}}, \tag{17}$$

where K_{it} is the installed capital stock and $v_i E_{it}$ is the amount of capital that plant i can operate with its electricity input. We assume full capital utilization in 2010, so v_i measures the amount of capital that the plant can operate per unit of electricity. Because

$v_i E_{i,2010} = K_{i,2010}$ by construction,

$$\left.\frac{\partial \ln \widetilde{K}_{it}}{\partial \ln E_{it}}\right|_{t=2010} = \kappa. \tag{18}$$

Using effective capital services instead of installed capital changes the baseline input cost shares. The Appendix E-4 describes how we incorporate capital idleness into the simulation. After this adjustment, the baseline capital and electricity cost shares are

$$\widetilde{s}^0_{Kj}(\kappa) = (1-\kappa)s^0_{Kj}, \qquad \widetilde{s}^0_{Ej}(\kappa) = s^0_{Ej} + \kappa s^0_{Kj}. \tag{19}$$

where s^0_{Kj} and s^0_{Ej} are the baseline capital and electricity cost shares in industry j, and $\widetilde{s}^0_{Kj}(\kappa)$ and $\widetilde{s}^0_{Ej}(\kappa)$ are the corresponding shares after accounting for electricity-dependent capital utilization.

A positive κ therefore shifts κs^0_{Kj} from the capital cost share to the electricity cost share. Electricity affects output both as a direct production input and by determining how much installed capital can operate. Consequently, higher electricity prices or tighter rationing have stronger effects on unit costs, input demand, and output when capital idleness is more important.

The magnitude of κ is ambiguous a priori. Some machinery may require electricity in nearly fixed proportions, leading electricity shortages to idle substantial capital. However, retrofitting machinery or changing operating practices may limit capital idleness. We therefore estimate κ from the data.

7.3.2 Input Complementarity

The second discrepancy is that the Cobb–Douglas model predicts too large a decline in electricity use and too small declines in labor and materials. We address this discrepancy by allowing the elasticity of substitution, σ, to differ from one. When $\sigma < 1$, electricity is complementary to other inputs, so electricity supply shocks induce plants to reduce labor and materials rather than substitute these inputs for electricity. This stronger complementarity produces a smaller decline in electricity use and larger declines in labor, materials, and output, bringing the model predictions closer to the reduced-form estimates.

7.3.3 Calibration

We jointly estimate the elasticity of substitution, σ, and the utilization elasticity, κ, using minimum-distance GMM. The empirical targets are the post-phase-out 2SLS impacts of electricity supply shocks on electricity use, labor, materials, and output. For each

candidate pair (σ, κ), we calculate the model-implied responses using the simulation procedure described above. We then construct four moment discrepancies as the differences between the model-implied and reduced-form responses. We select the parameter pair that minimizes the sum of the squared discrepancies. Appendix E-4 presents the detailed procedure.

We obtain $\widehat{\sigma} = 0.59$ and $\widehat{\kappa} = 0.17$, which minimize the discrepancy between the model-implied and reduced-form impacts of electricity supply shocks. At the mean post-phase-out supply shock, the reduced-form estimates imply declines of 6.7% in electricity use, 1.3% in labor, 1.3% in materials, and 1.7% in output. The calibrated model predicts corresponding declines of 6.8%, 1.4%, 1.4%, and 1.6%. The model therefore closely reproduces the input and output responses.

The estimate $\widehat{\sigma} = 0.59$ indicates limited substitutability between electricity and other inputs. Direct empirical estimates of substitution between electricity and other inputs remain scarce. [21] The broader literature on substitution between energy and other inputs generally reports elasticities below one, typically ranging from approximately 0.2 to 0.8.[22] Our estimate falls within this range and implies greater input complementarity than a Cobb–Douglas technology.

Panel D of Figure 6 varies σ while holding $\kappa = 0.17$ and the baseline cost shares fixed. Lowering σ from its calibrated value of 0.59 to 0.2 strengthens input complementarity. The predicted decline in electricity use falls from 6.8% to 4.4%, while the declines in labor and materials increase from 1.4% to 2.0% and the output loss increases from 1.6% to 2.1%. Conversely, raising σ to 1.5 increases input substitutability. The predicted decline in electricity use rises to 13.8%, while the declines in labor and materials fall to 1.0% and the output loss falls to 1.4%.

The estimate $\widehat{\kappa} = 0.17$ implies that, at the baseline, a 1% reduction in electricity use reduces effective capital services by 0.17%. This utilization channel is quantitatively important because electricity's direct output elasticity is only 0.015, approximately one-tenth of capital's output elasticity of 0.139. Panel E of Figure 6 illustrates the extent to which capital idleness amplifies the effects of electricity supply shocks. Eliminating the utilization channel by setting $\kappa = 0$ reduces the predicted output loss from 1.6% to 0.6%, a reduction of approximately 61%. The corresponding declines in labor and materials fall

[21] Aggarwal (2026) estimates labor–electricity elasticities of approximately 0.1 in India's formal manufacturing sector and 0.2–0.5 in its informal manufacturing sector.

[22] Using U.S. manufacturing-plant data and variation in energy prices, Hawkins-Pierot and Wagner (2026) estimates a common elasticity of substitution among capital, labor, and energy of 0.26 across industries and 0.24 among energy-intensive industries. van der Werf (2008) reports energy substitution elasticities of approximately 0.16–0.62 across industries, while Bretschger and Jo (2024) estimates labor–energy elasticities of 0.61–0.80 using French manufacturing-firm data.

from 1.4% to 0.5%. Conversely, setting $\kappa = 0.5$ increases the predicted output loss from 1.6% to 3.3%, more than doubling the effect. The corresponding declines in labor and materials increase to 2.9%.

8 Policy Counterfactuals

In this section, we use the structural model to evaluate counterfactual policies that could mitigate the output effects of electricity supply shocks.

8.1 Decomposing the Effects of Pricing and Rationing

Electricity supply shocks affected manufacturing plants through higher electricity prices and peak-season rationing. We use the structural model to decompose the effects of these two channels during the post-phase-out period (2013–2018). Their relative importance indicates whether policies should focus on reducing electricity costs or relaxing physical supply constraints.

To separate the price and rationing channels, we compare plant-level predictions for electricity use, labor, materials, and output across three scenarios: the 2010 no-shock baseline, a price-only scenario, and a combined price-and-rationing scenario. The decomposition is

$$\underbrace{\ln Z_i^{\text{sim}}(\overline{p}_r^E, \overline{\mathcal{S}}_r) - \ln Z_i^{\text{sim}}(1,0)}_{\text{Total impact}} = \underbrace{\ln Z_i^{\text{sim}}(\overline{p}_r^E, \overline{\mathcal{S}}_r) - \ln Z_i^{\text{sim}}(\overline{p}_r^E, 0)}_{\text{Rationing contribution}} + \underbrace{\ln Z_i^{\text{sim}}(\overline{p}_r^E, 0) - \ln Z_i^{\text{sim}}(1,0)}_{\text{Price contribution}}. \tag{20}$$

where $Z_i^{\text{sim}}(\cdot,\cdot)$ denotes plant i's simulated annual outcome, with the first argument representing the relative electricity price and the second representing peak-season rationing intensity. We normalize the 2010 baseline to $(1,0)$. For region r, $\overline{p}_r^E$ is the predicted post-phase-out relative price, and $\overline{\mathcal{S}}_r$ is the mean supply shock used as the rationing intensity.

The first term on the right-hand side measures the rationing contribution by comparing the combined and price-only scenarios. The second term measures the price contribution by comparing the price-only scenario with the no-shock baseline. By construction, the two contributions sum to the total impact. This ordering assigns their interaction to the rationing contribution.

Table 7 shows that pricing accounts for most of the model-implied decline. At the mean post-phase-out supply shock, the structural model predicts a 6.76% decline in electricity use. The price contribution is 6.39 percentage points, or 94.5% of the total decline, while the rationing contribution is 0.37 percentage points, or 5.5%. The model also predicts declines of 1.39% in labor, 1.39% in materials, and 1.57% in output. Pricing accounts for 1.32, 1.32, and 1.49 percentage points of these declines, respectively, corresponding to 94.9% of each total effect. Rationing accounts for only the remaining 5.1%. Although the physical-rationing constraint binds for 98.1% of plants during summer, it applies for only 90 days per year and is generally small in magnitude. Consequently, rationing accounts for only about 5% of the annualized effects on electricity use, labor, materials, and output.

These results indicate that the nuclear shutdown primarily affected manufacturing plants through higher electricity prices. Utilities replaced nuclear generation with more expensive fossil generation, raising generation costs. Binding quantity constraints played a much smaller role. This pattern differs from the residential sector, where non-price interventions can meaningfully reduce electricity demand during scarcity (e.g., Reiss and White, 2008; He and Tanaka, 2023).

8.1.1 Elasticity Benchmarks

We report elasticities with respect to electricity-price changes using both the reduced-form estimates and the structural model. These elasticities provide benchmarks for the counterfactual analysis.

The reduced-form analysis implies an own-price elasticity of -0.80 for electricity, cross-price elasticities of -0.15 for both labor and materials, and an output elasticity of -0.20 (Table E3, Column (1)). We calculate each elasticity as the ratio of the main reduced-form effect of the supply shock on the outcome to its estimated effect on the log electricity price. At the mean post-phase-out supply shock, the 2010-total-salary-weighted mean predicted log-price change is 8.4 log points, corresponding to an average electricity-price increase of approximately 9.0% in levels. This estimated price response captures only the change in the market price and does not include the shadow price associated with physical rationing.

The structural model yields corresponding price-only elasticities of -0.76, -0.16, -0.16, and -0.18 (Column (2)). We calculate the plant-level responses using the model's closed-form expressions. For example, the own-price elasticity of electricity demand is

$$\epsilon_{E,i} \equiv \frac{\partial \ln E_i}{\partial \ln p^E} = -\sigma + (\sigma - \eta) s^C_{E,ij},$$

where $s^C_{E,ij}$ is electricity's share of the plant's unit cost. We then aggregate the plant-level

elasticities using 2010 total-salary weights. Appendix E-3.2 describes the derivation.[23]

8.2 Counterfactual: Regional Market Integration

We use the estimated production framework to evaluate regional market integration as a supply-side intervention that can mitigate the economic costs of electricity supply disruptions. Section 6.1 shows that manufacturing plants have limited capacity to adapt to supply shocks. We therefore focus on a policy that improves electricity-market functioning and lowers electricity prices.

Japan's electricity market has historically been segmented because regulations limited interregional electricity transmission. This segmentation can create dispatch inefficiencies: high-marginal-cost generators may serve local demand while lower-cost capacity elsewhere remains idle (Hausman, 2025). It also limits reserve sharing, requiring regions to maintain redundant backup capacity and raising system-wide costs. The potential gains from greater market integration have therefore attracted substantial policy attention. We quantify the benefits of relaxing these constraints through interregional trade and more efficient reserve sharing.

We quantify the resulting changes in output and producer surplus attributable to electricity supply shocks. In particular, we approximate the policy-induced change in aggregate producer surplus as

$$\Delta\Pi \approx \sum_i \Pi_i\left(p^E_{\text{base}}\right)\,(1-\eta)\,s^C_{E,i}\,\Delta \ln p^E, \tag{21}$$

where $s^C_{E,i}$ is plant i's electricity cost share and $\Delta \ln p^E$ is the policy-induced change in the electricity price. This approximation follows from the model's electricity-price elasticities; Appendix E-3 describes the counterfactual simulation procedure.

Estimating these effects requires counterfactual electricity prices under market integration (Appendix F-0.1). We estimate regional supply curves and solve for integrated-grid equilibrium prices, allowing lower-cost generators to serve a wider market and enabling reserve sharing. Full integration reduces the national average log electricity price by 2.8 log points relative to the observed scenario.[24] The mean log-price increase relative to the

[23]Our own-price elasticity estimates are consistent with the quasi-experimental evidence and are closest to von Graevenitz and Rottner (2024), who report estimates of approximately −0.4 to −0.6. Previous studies report industrial own-price elasticities ranging from −0.09 to −0.20 (Gerster and Lamp, 2024) and values as large as −1.5 in settings with substantial price shocks (Martin et al., 2014).

[24]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

no-shock baseline falls from 8.4 to 5.7 log points, corresponding to level increases of approximately 9.0% and 5.8%, respectively. Price changes remain heterogeneous, with large declines in high-cost demand centers such as Tokyo and offsetting increases in low-cost exporting regions (Table F1, Column (2)).

Table 7 summarizes the simulation results. Under the price-only structural baseline (the price-contribution row of Panel B), electricity use declines by 6.3%, while labor and materials each decline by 1.3%. Real output and producer surplus fall by 1.5% and 1.2%, respectively. Evaluated at 2010 levels, the output decline corresponds to an annual loss of 4,320 billion JPY, compared with the reduced-form estimate of 4,920 billion JPY.[25] Producer surplus falls by 691 billion JPY.

Market integration with reserve sharing substantially mitigates these costs (Panel C). It reduces the decline in electricity use from 6.3% to 4.3% and the output decline from 1.5% to 1.0%. Annual output losses fall from 4,320 billion JPY to 2,935 billion JPY, while producer-surplus losses fall from 691 billion JPY to 470 billion JPY. Both losses decline by approximately 32.1% relative to the baseline.

These results indicate that limited interregional coordination amplified the costs of the post-Fukushima electricity supply shock. Greater market integration recovers nearly one-third of the price-induced losses by improving system-wide dispatch and reserve sharing.

9 Conclusion

This paper shows that the persistent electricity supply shock following the Fukushima accident imposed sizable costs on Japanese manufacturing. A 6% decline in electricity use reduced revenue by 1.7%, implying annual losses of roughly 4.9 trillion JPY despite electricity's 2% expenditure share. Plants showed little adjustment through technology, product mix, fuel substitution, or production reallocation; self-generation protected only a small subset of producers.

Our structural model explains these losses through strong input complementarity and capital idleness and closely matches the reduced-form estimates. The counterfactuals show that higher electricity prices drove most of the output decline and that greater grid integration could recover nearly one-third of the baseline losses. The results highlight the importance of affordable, efficiently allocated electricity for industrial resilience.

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.

[25] These correspond to approximately 45–55 billion USD, using an exchange-rate range of 80–95 JPY per USD (around 2010).

References

Abeberese, A. B. (2017). Electricity cost and firm performance: Evidence from india. *Review of Economics and Statistics*, 99.

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

Aggarwal, R. (2026). Input substitution in electricity generation and industrial production: Evidence from india. *Journal of Development Economics*, 179:103614.

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

Almond, D., Edlund, L., and Palme, M. (2009). Chernobyl's subclinical legacy: Prenatal exposure to radioactive fallout and school outcomes in Sweden. *The Quarterly Journal of Economics*, 124(4):1729–1772.

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

Bachmann, R., Baqaee, D., Bayer, C., Kuhn, M., Löschel, A., Moll, B., Peichl, A., Pittel, K., and Schularick, M. (2024). What if? the macroeconomic and distributional effects for germany of a stop of energy imports from russia. *Economica*, 91(364):1157–1200.

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.

Berndt, E. R. and Wood, D. O. (1975). Technology, prices, and the derived demand for energy. *The Review of Economics and Statistics*, 57.

Bijnens, G., Konings, J., and Vanormelingen, S. (2022). The impact of electricity prices on european manufacturing jobs. *Applied Economics*, 54(1):38–56.

Borenstein, S., Fowlie, M., and Sallee, J. (2021). Designing electricity rates for an equitable energy transition. Working Paper 314, Energy Institute at Haas.

Bretschger, L. and Jo, A. (2024). Complementarity between labor and energy: A firm-level analysis. *Journal of Environmental Economics and Management*, 124:102934.

Cabinet Office (2023). National accounts of japan: Capital stock and capital cost data. Economic and Social Research Institute, Cabinet Office, Government of Japan. Used to construct industry-year capital rental rates as nominal capital cost divided by real net capital stock.

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.

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.

Danzer, A. M. and Danzer, N. (2016). The long-run consequences of Chernobyl: Evidence on subjective well-being, mental health and welfare. *Journal of Public Economics*, 135:47–60.

Davis, L. and Hausman, C. (2016). Market impacts of a nuclear power plant closure. *American Economic Journal: Applied Economics*, 8(2):92–122.

Davis, L. W., Hausman, C., and Rose, N. L. (2023). Transmission impossible? prospects for decarbonizing the US grid. *Journal of Economic Perspectives*, 37(4):155–180.

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

EIA (2026a). Annual energy outlook 2026. U.S. Energy Information Administration, Annual Energy Outlook. Accessed June 12, 2026.

EIA (2026b). Electric power monthly: Table 5.3. average price of electricity to ultimate customers: Total by end-use sector, 2016–march 2026. U.S. Energy Information Administration, Electric Power Monthly. Accessed June 12, 2026.

EIA (2026c). Fossil generation could rise with faster-than-expected growth in data center power demand. U.S. Energy Information Administration, Today in Energy. Accessed June 12, 2026.

ESDV (2012). Meeting materials and reports on electricity supply and demand verification. Electricity Supply–Demand Verification Subcommittee, Advisory Committee for Natural Resources and Energy, Ministry of Economy, Trade and Industry. In Japanese; materials used for regional peak-day electricity demand and supply.

Fabra, N. (2023). Reforming European electricity markets: Lessons from the energy crisis. *Energy Economics*, 126:106963.

Finn, M. G. (2000). Perfect competition and the effects of energy price increases on economic activity. *Journal of Money, Credit and Banking*, 32:400–416.

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.

Fu, R., Sohn, Y., Shen, Y., and Noguchi, H. (2026). Invisible threat, tangible harm: Radiation anxiety and birth outcomes after Fukushima. *Journal of Health Economics*, 107:103125.

Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in us manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12.

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.

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.

Gonzales, L. E., Ito, K., and Reguant, M. (2026). Zero energy day: How nationwide blackouts affect the economy. NBER Working Paper 35066, National Bureau of Economic Research.

Government of Japan (1964). Electricity business act. Act No. 170 of 1964, e-Gov Law Search. Article 18 governed the duty to supply under the pre-reform regime.

Hamilton, J. D. (2013). Historical oil shocks. In *Routledge handbook of major events in economic history*, pages 239–265. Routledge.

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

Hawkins-Pierot, J. T. and Wagner, K. R. H. (2026). Technology lock-in and costs of delayed climate policy. *Journal of Political Economy Microeconomics*. Forthcoming.

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 (2026). Electricity 2026. Technical report, International Energy Agency, Paris. Licence: CC BY 4.0. Accessed June 12, 2026.

Jarvis, S., Deschenes, O., and Jha, A. (2022). The private and external costs of germany's nuclear phase-out. *Journal of the European Economic Association*, 20(3):1311–1346.

JEPIC (2019). The Electric Power Industry in Japan 2019. Technical report, Japan Electric Power Information Center, Tokyo.

Kahn, M. E. and Mansur, E. T. (2013). Do local energy prices and regulation affect the geographic concentration of employment? *Journal of Public Economics*, 101.

Kakuchi, S. (2011). Energy-saving 'setsuden' campaign sweeps japan after fukushima. *The Guardian*.

Kawaguchi, D. and Yukutake, N. (2017). Estimating the residential land damage of the Fukushima nuclear accident. *Journal of Urban Economics*, 99:148–160.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of economic literature*,

46(4):871–909.

Kimura, O. and Nishio, K.-I. (2016). Responding to electricity shortfalls: Electricity-saving activities of households and firms in japan after fukushima. *Economics of Energy & Environmental Policy*, 5(1):51–72.

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.

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 (2011a). Electric power survey statistics, fy2010. e-Stat, Portal Site of Official Statistics of Japan. Tables 2-(1), 2-(5), 3-(1), and 3-(2); in Japanese.

METI (2011b). Restriction on electricity use under article 27 of the electricity business act. Ministry of Economy, Trade and Industry. Order for large-lot users in the Tokyo Electric Power Company and Tohoku Electric Power Company service areas to reduce summer peak electricity use by 15 percent; in Japanese.

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.

Moll, B., Schularick, M., and Zachmann, G. (2023). The power of substitution: The great german gas debate in retrospect. *Brookings Papers on Economic Activity*, 2023(2):395–455.

Neidell, M., Uchida, S., and Veronesi, M. (2021). The unintended effects from halting nuclear power production: Evidence from fukushima daiichi accident. *Journal of Health Economics*, 79:102507.

NERC (2024). 2024 long-term reliability assessment. Technical report, North American Electric Reliability Corporation. Corrected July 2025. Accessed June 12, 2026.

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

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., Nielsen, M. O., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *Stata Journal*, 19.

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

Severnini, E. (2017). Impacts of nuclear plant shutdown on coal-fired power generation and infant health in the tennessee valley in the 1980s. *Nature Energy*, 2(4):17051.

Shapiro, J. S. and Walker, R. (2018). Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade. *American Economic Review*, 108.

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.

van der Werf, E. (2008). Production functions for climate policy modeling: An empirical analysis. *Energy Economics*, 30(6):2964–2979.

von Graevenitz, K., Krug, J., and Rottner, E. (2025). Rising energy prices without falling consumption? the role of energy price dispersion in a multi-product world. Technical report, ZEW Discussion Papers.

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.

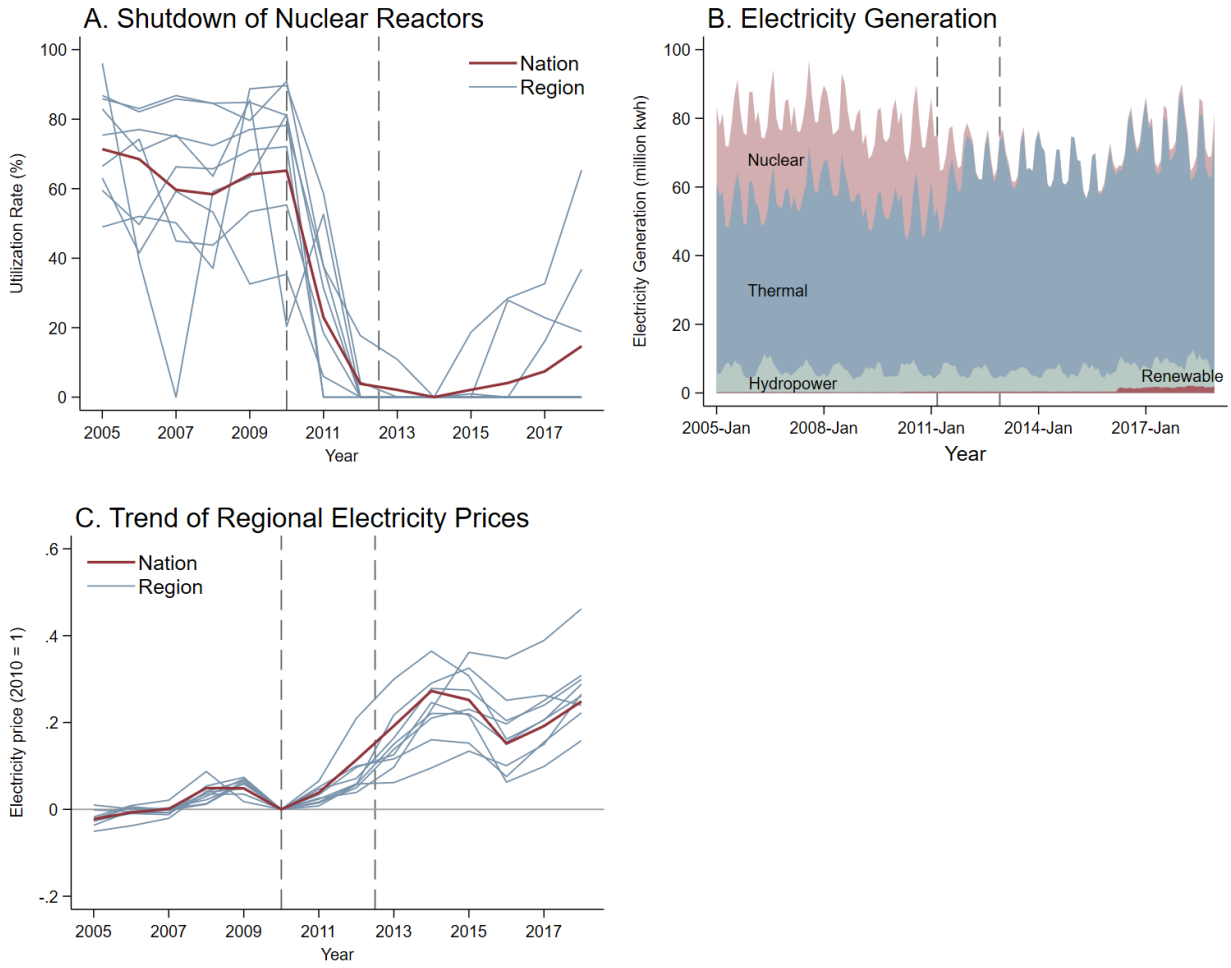
Wiser, R. H., O'Shaughnessy, E., Barbose, G. L., Cappers, P., and Gorman, W. (2025). Factors influencing recent trends in retail electricity prices in the United States. *The Electricity Journal*, 38(4):107516.

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.

Wolverton, A., Shadbegian, R., and Gray, W. B. (2022). The U.S. manufacturing sector's response to higher electricity prices: Evidence from state-level renewable portfolio standards. NBER Working Paper 30502, National Bureau of Economic Research.

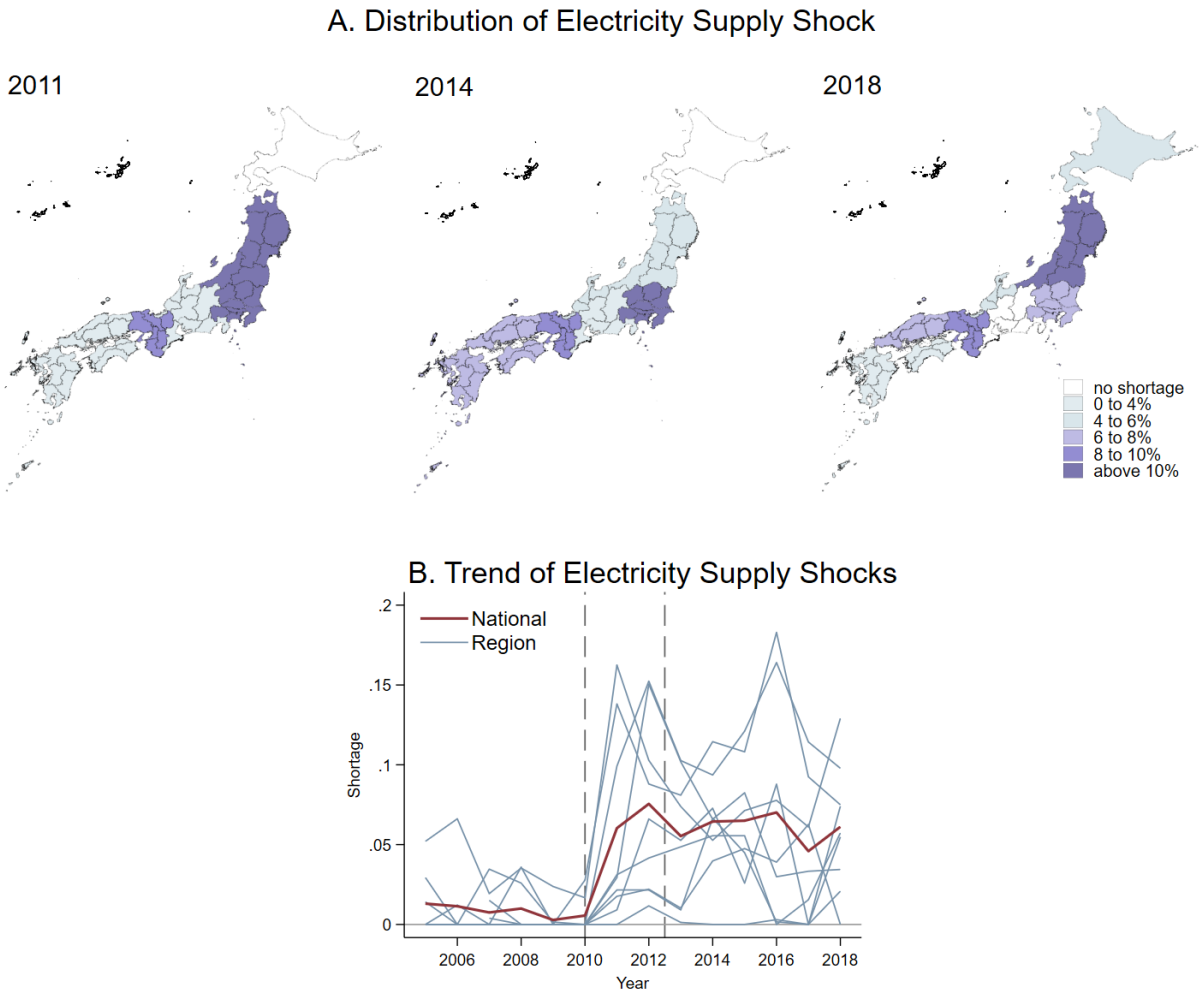
Figures and Tables

Figure 1: Nuclear shutdown and electricity supply shocks



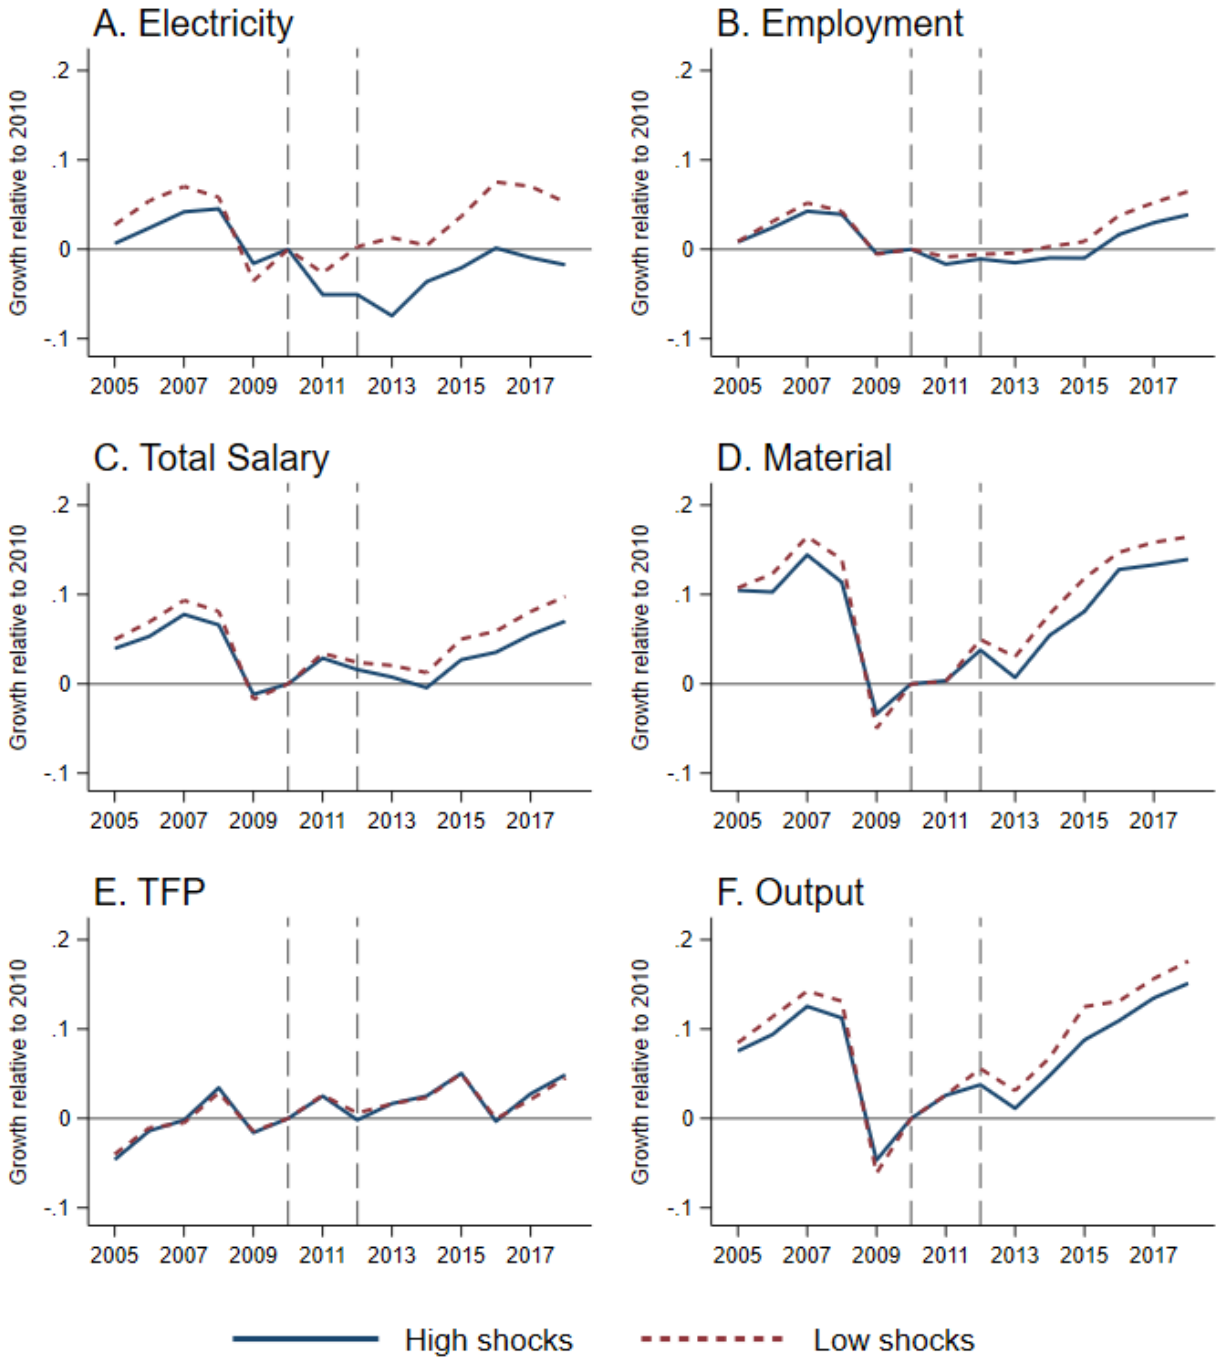
Notes: Panel A shows the utilization of nuclear reactors, defined as power generation divided by supply capacity. Panel B depicts the trend in monthly national electricity demand by energy source. Panel C presents the trend in average electricity prices. In all panels, the vertical dotted line indicates 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



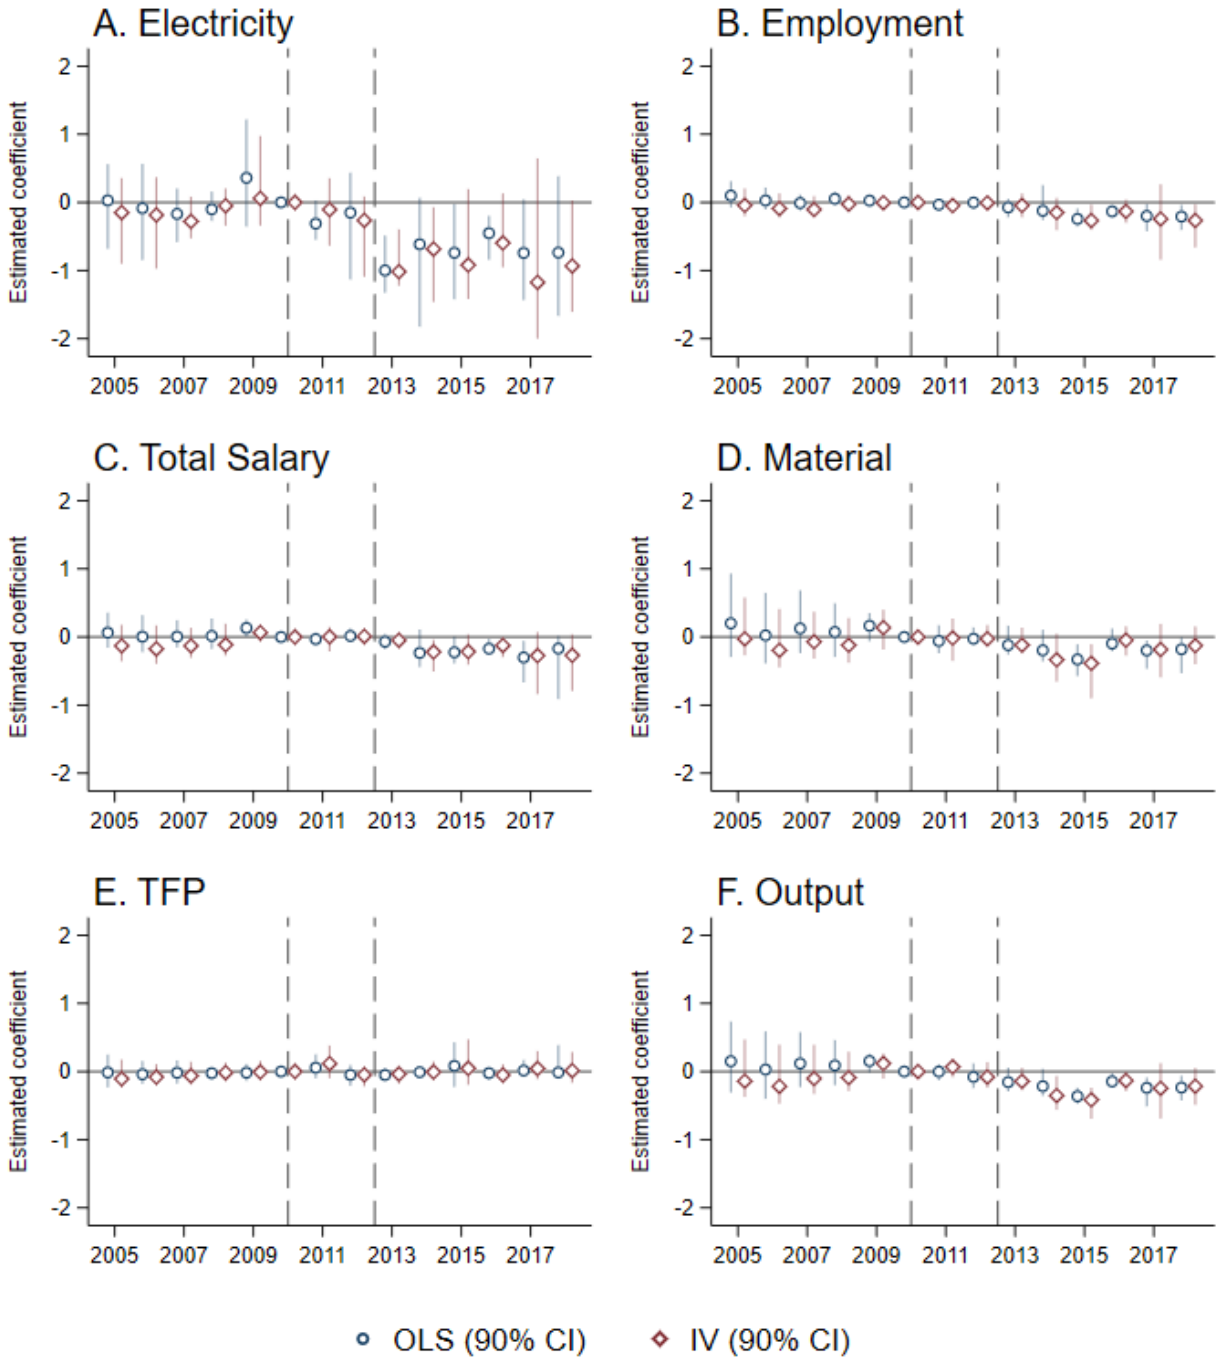
Notes: Panel A presents the geographic distribution of our measure of electricity supply shocks from 2011 to 2018. Panel B shows the corresponding trends, where the bold line represents the national average and the thin gray lines represent individual regions. See the main text for the definition of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

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



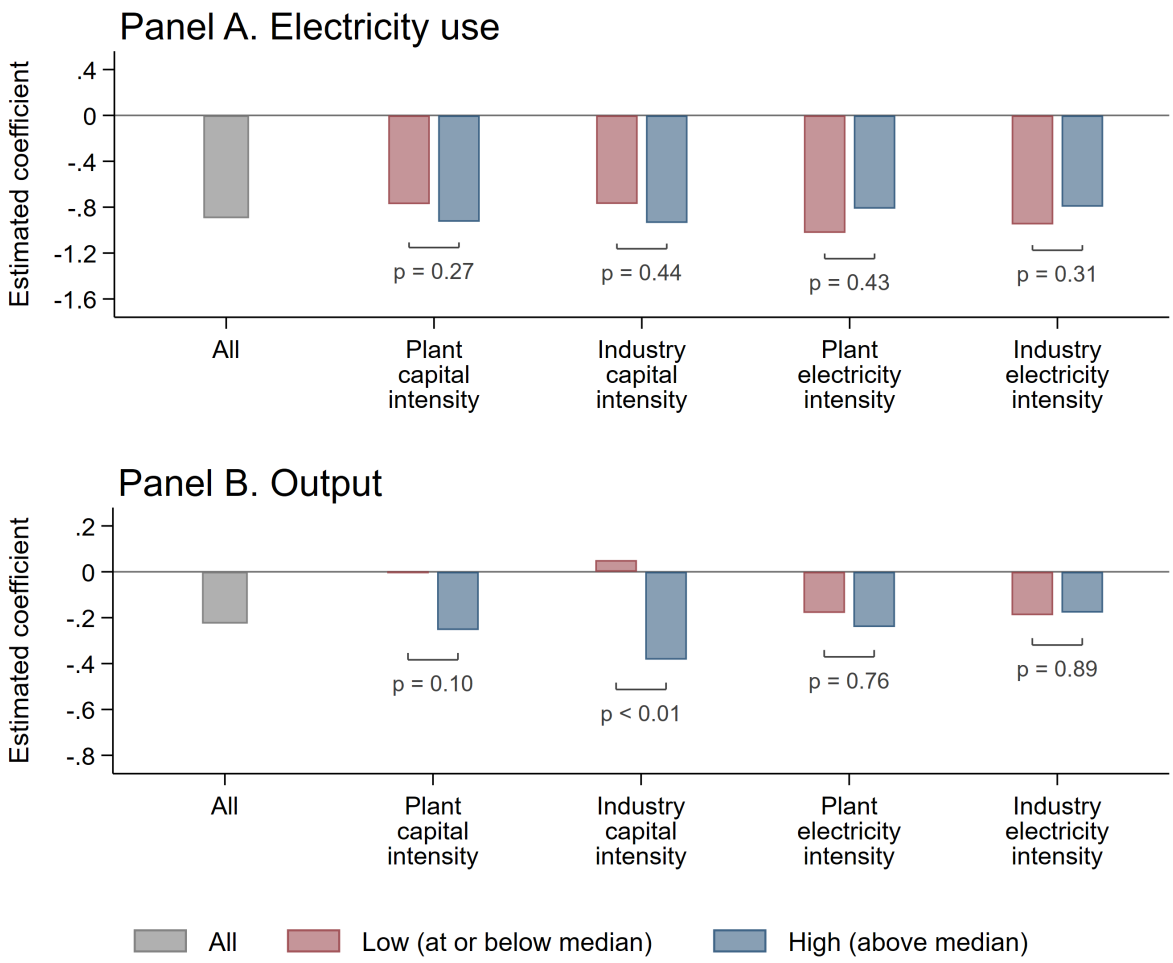
Notes: We classify regions with mean supply shocks during 2013–2018 at or above the median across the nine regions as high-shock regions and the remaining regions as low-shock regions. Outcomes are normalized by their 2010 values, so the y-axis represents changes relative to 2010.

Figure 4: Event study results of effects of supply shock on various manufacture's activities



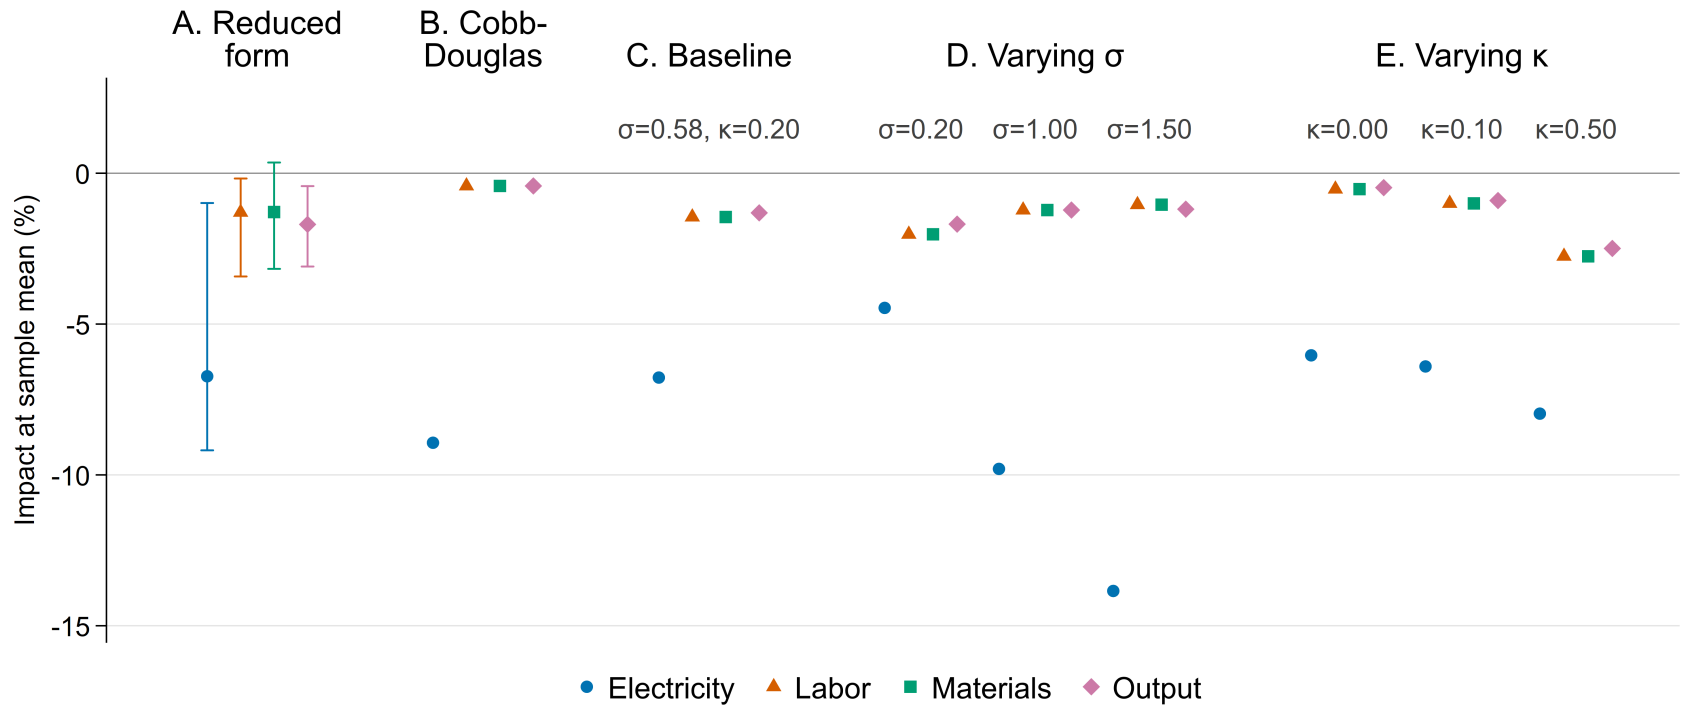
Notes: The figure reports event-study estimates from equation (5). Blue circles denote OLS estimates, and red diamonds denote 2SLS estimates. Vertical bars show 90% confidence intervals obtained using a wild cluster bootstrap with clustering by region. Confidence intervals extending below −2 are truncated at −2. For pre-accident years, we use each region's mean supply shock during the post-phase-out period. All regressions include three-digit-industry-by-year fixed effects and use 2010 log-salary weights.

Figure 5: Heterogeneous impacts of electricity supply shocks



Notes: The figure reports 2SLS coefficients from equation (4) for electricity use (Panel A) and output (Panel B). Gray bars show the full-sample estimates. Red and blue bars show estimates for plants at or below and above the median, respectively, of the indicated plant- or three-digit-industry characteristic measured in 2010. Brackets report wild-cluster-bootstrap p-values for tests of equality between the two subgroup coefficients, clustering by region. All regressions include year-by-three-digit-industry fixed effects and use 2010 log-salary weights.

Figure 6: Reduced-Form Estimates and Structural Simulations



Notes: The figure compares reduced-form impacts at the weighted mean post-phase-out supply shock with structural predictions. Reduced-form whiskers show 90% confidence intervals. Panel B reports the Cobb–Douglas benchmark, and Panel C reports the calibrated baseline with $\sigma = 0.59$ and $\kappa = 0.17$. Panel D varies $\sigma \in \{0.20, 1.00, 1.50\}$ while holding $\kappa = 0.17$ and baseline cost shares fixed. Panel E varies $\kappa \in \{0.00, 0.10, 0.50\}$ while holding $\sigma = 0.59$. Within every specification, markers appear in the order electricity, labor, materials, and output. Generator owners retain their simulated grid-electricity response, while their structural non-electricity and output responses are set to zero.

Table 1: Summary Statistics

	Mean (1)	S.D. (2)
Panel A. Plant Characteristics		
Revenue (million yen)	5,067	11,494
Employment (headcount)	127	166
Capital stock (million yen)	1,003	2,405
Investment (million yen)	185	532
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	0.020	0.023
Panel C. Plant Energy Usage		
Fuel input for production (TJ)	30	114
Fuel input for self-generation (TJ)	4.2	31
Self-generated electricity (TJ)	3.6	24
Grid electricity purchased (TJ)	20	55
Self-generation ownership (yes = 1)	0.08	0.27

Notes: Panels A and B include all plants in our sample constructed from the Census of Manufactures (N = 437,561). Expenditure shares in Panel B are calculated as the ratio of input expenditure to plant revenue. Panel C uses data from the Energy Consumption Statistics and the Petroleum and Other Fuels Monthly Statistics (N = 381,761). Monetary values are deflated to 2010 prices.

Table 2: Distribution of Estimated Output Elasticities

	Structural Estimates			Benchmark
	Mean	p25	p75	OLS
	(1)	(2)	(3)	(4)
Materials (θ_M)	0.446	0.386	0.506	0.492
Electricity (θ_E)	0.015	0.008	0.017	0.068
Labor (θ_L)	0.354	0.259	0.432	0.423
Capital (θ_K)	0.139	0.083	0.170	0.042

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 with converged, strictly positive elasticities. Column (4) provides the mean benchmark across the same industries based on industry-specific OLS regressions that regress log revenue on log inputs. Materials and electricity are treated as flexible inputs and estimated in the first step from expenditure-share equations. Labor and capital are treated as quasi-fixed inputs and estimated jointly in the second step by GMM.

Table 3: First-Stage Regression: Impact of Supply Composition on Supply Shocks

	Dep Var: Supply Shocks	
	Phase-out	Post-phase-out
	(1)	(2)
Predicted nuclear supply growth	-0.45	-0.26
	(0.00)***	(0.00)***
Reserve margin in 2010	-1.05	-0.82
	(0.00)***	(0.00)***
Year FE	Yes	Yes
Observations	18	54

Notes: The table reports first-stage regression estimates. The dependent variable is the electricity supply shock. Predicted nuclear supply growth is defined as $Sh^{Nuc}_{r,2010} \times g^{Nuc}_{t,-r}$, where $Sh^{Nuc}_{r,2010}$ is region r's nuclear generation share in 2010 and $g^{Nuc}_{t,-r}$ is the leave-one-out national growth rate in nuclear generation. Column (1) covers the nuclear phase-out period (2011–2012), and Column (2) covers the post-phase-out period (2013–2018). P-values are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Effects of Supply Shocks on Inputs and Outputs

	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$					
	Electricity (1)	Labor (2)	Total Salary (3)	Materials (4)	TFP (5)	Output (6)
Panel A. The effects of supply shocks using OLS						
Supply Shocks	-0.24	-0.02	-0.01	-0.04	0.01	-0.04
(Phase-out)	[0.11]	[0.42]	[0.72]	[0.27]	[0.91]	[0.23]
Supply Shocks	-0.63**	-0.15**	-0.18**	-0.15	-0.01	-0.19***
(Post-phase-out)	[0.03]	[0.02]	[0.04]	[0.13]	[0.73]	[0.00]
Panel B. The effects of supply shocks using 2SLS						
Supply Shocks	-0.21	-0.03	0.01	-0.01	0.00	-0.02
(Phase-out)	[0.36]	[0.11]	[0.91]	[0.82]	[0.95]	[0.61]
Supply Shocks	-0.89*	-0.17	-0.17*	-0.17	-0.02	-0.23**
(Post-phase-out)	[0.06]	[0.12]	[0.06]	[0.18]	[0.80]	[0.04]
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. The dependent variables are log differences from 2010. The phase-out 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 fixed effects, where industries are defined at the three-digit level, and use 2010 log labor weights. P-values from wild cluster bootstrap using regions as clusters are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: Effects of Supply Shocks on Plant Behavior

	Product-level		Outsourcing	
	Switching (1)	Intensity (2)	Outsourcing (3)	Share in Mat. (4)
Panel A. Effects of supply shocks using OLS				
Supply shocks (Phase-out)	0.09 [0.23]	0.02 [0.50]	-0.02 [0.88]	0.00 [0.96]
Supply shocks (Post-phase-out)	0.03 [0.72]	0.03 [0.27]	-0.06 [0.86]	-0.01 [0.76]
Panel B. Effects of supply shocks using 2SLS				
Supply shocks (Phase-out)	0.07 [0.33]	0.02 [0.32]	-0.02 [0.95]	-0.05 [0.51]
Supply shocks (Post-phase-out)	0.07 [0.52]	0.03 [0.44]	-0.36 [0.37]	-0.08 [0.46]
Year-by-industry FE	Yes	Yes	Yes	Yes

Notes: The table reports regression estimates. Column (1) uses a dummy variable for product switching. Column (2) measures product-level electricity intensity. Column (3) uses log differences in outsourcing relative to 2010. Column (4) uses log differences in the outsourcing share in materials relative to 2010. The phase-out 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 fixed effects, where industries are defined at the three-digit level, and use 2010 log-labor weights. P-values from wild cluster bootstrap using regions as clusters are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 6: Effects of Supply Shocks on Energy Usage

	Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$			
	Fuel Consumption for:			
	Production (1)	Self-generation (2)	Electricity Generated (3)	Generator ownership (4)
Panel A. Effects of supply shocks using OLS				
Supply shocks (Phase-out)	-0.12 [0.39]	0.14 [0.87]	0.59 [0.28]	0.04** [0.02]
Supply shocks (Post-phase-out)	-0.26 [0.19]	3.45** [0.04]	1.51* [0.09]	0.01 [0.87]
Panel B. Effects of supply shocks using 2SLS				
Supply shocks (Phase-out)	-0.12 [0.35]	0.67 [0.51]	0.95 [0.23]	0.06*** [0.00]
Supply shocks (Post-phase-out)	-0.35 [0.53]	4.50* [0.05]	2.33 [0.12]	0.06 [0.54]
Year-by-industry FE	Yes	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 self-generated electricity. Column (4) reports generator ownership, defined by self-generation equal to at least 5% of grid purchases. The phase-out row uses data from fiscal years 2007–2010 and 2011–2012. The post-phase-out row uses data from fiscal years 2007–2010 and 2013–2018. All regressions include year-by-industry fixed effects, where industries are defined at the two-digit level, and use 2010 employment weights. P-values from wild cluster bootstrap using regions as clusters are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 7: Structural Simulation Impacts and Counterfactuals

	Impacts at mean shock				Monetary losses (JPY bn)	
	Electricity (1)	Labor (2)	Materials (3)	Output (4)	Output (5)	Profit (6)
Panel A. Reduced-form impacts at the mean shock						
Reduced-form estimate	-6.7%	-1.3%	-1.3%	-1.7%	-4,920	–
Panel B. Structural estimation						
Total impact	-6.8%	-1.4%	-1.4%	-1.6%	-4,552	-728
Price contribution	-6.4%	-1.3%	-1.3%	-1.5%	-4,320	-691
(Share of total impact)	(94.5%)	(94.9%)	(94.9%)	(94.9%)	(94.9%)	(94.9%)
Rationing contribution	-0.4%	-0.1%	-0.1%	-0.1%	-232	-37
(Share of total impact)	(5.5%)	(5.1%)	(5.1%)	(5.1%)	(5.1%)	(5.1%)
Panel C. Counterfactual analysis						
Market integration	-4.3%	-0.9%	-0.9%	-1.0%	-2,935	-470

Notes: Columns (1)–(4) report annualized log impacts at the mean post-phase-out supply shock, expressed in percent. Columns (5) and (6) report annual aggregate output and profit losses in billion constant 2010 Japanese yen. Panel A reports the main reduced-form IV estimates multiplied by the mean shock. Panel B reports predictions from the preferred dynamic-labor structural model. The total impact compares predicted electricity prices and seasonal rationing with the 2010 baseline. The price contribution changes predicted prices while holding rationing at the no-rationing baseline. The rationing contribution is the difference between the total and price-only impacts and therefore includes their interaction. Panel C compares the price-only baseline reported in the price-contribution row of Panel B with full market integration and reserve sharing. Percentage responses are weighted by 2010 total salary. Aggregate losses apply these weighted responses to aggregate 2010 output. Contribution shares are calculated from underlying unrounded estimates. Generator owners retain their simulated grid-electricity response, while their labor, materials, output, and profit responses are set to zero based on the heterogeneous IV evidence.

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 Data Appendix

A-1 Census of Manufactures

A-1.1 Deflators and Capital Rent

We deflate inputs and outputs using official price indexes. We obtain industry-specific output and input deflators from the Bank of Japan and use them to deflate output and intermediate inputs, including materials (Bank of Japan, 2026). We deflate electricity expenditures using regional average electricity prices from FEPC (FEPC, 2026). We deflate total wages and salaries using the Consumer Price Index from the Statistics Bureau of Japan (Statistics Bureau of Japan, 2026). We deflate capital stock using capital deflators constructed from national capital-stock series published by the Cabinet Office (Cabinet Office, 2026). We compute capital rental costs by multiplying each plant's deflated capital stock by an industry-year rental rate, measured as nominal capital cost divided by real net capital stock from the JIP database (Cabinet Office, 2023). Industry classifications vary across datasets but mainly follow two-digit classifications.

A-1.2 Multi-plant Firms and Product Switching

We use firm identifiers in the CMF to examine whether firms reallocate activity across establishments in response to localized electricity supply shocks (Section 6.1). For each firm-region-year, we aggregate electricity use and output across the firm's establishments located outside the focal region and compare these outcomes with their 2010 levels. For example, we construct each firm's output in Tokyo and non-Tokyo regions in each year to investigate how shocks in Tokyo affect production in non-Tokyo regions.

We also use product-level records in the CMF to measure product switching and changes in product mix (Section 6.1). We define product switching as an indicator equal to one if a plant's product set differs from its 2010 product set because the plant added or dropped at least one product. To measure changes in product-mix electricity intensity, we first calculate electricity expenditure per unit of output for single-product plants in 2010. We then construct a plant-year measure as the sales-weighted average electricity intensity of the products produced by the plant in that year.

A-2 Energy Consumption Statistics (ECS)

We combine the ECS with plant-level data from the Petroleum and Other Fuels Monthly Statistics (POFMS). The ECS is an annual survey covering a broad range of industries,

whereas POFMS is a monthly census of establishments in nine energy-intensive manufacturing industries: pulp, paper, and paperboard; chemicals; chemical fibers; petroleum products; ceramic, stone, and clay products; glass products; iron and steel; non-ferrous metal smelting and refining; and machinery. To avoid duplicate reporting, the ECS excludes establishments covered by POFMS. We aggregate the monthly POFMS observations to the fiscal-year level and harmonize the two datasets before pooling them.

The CMF, ECS, and POFMS all provide plant-level data but use different sampling frames. Our CMF sample covers manufacturing establishments with at least 30 employees. POFMS is a census of all establishments meeting its industry-specific eligibility criteria, which generally target large and energy-intensive plants. The ECS fully enumerates establishments on the self-generation roster, designated energy-management plants, and large establishments exceeding energy-use thresholds, while randomly sampling the remaining establishments by industry and size. Consequently, relative to the CMF, the pooled ECS–POFMS sample is more concentrated among large, energy-intensive, and self-generating plants. The surveys also have different response rates: in 2018, the response rates were 69.6% for ECS and 97.6% for POFMS, whereas the 2019 CMF achieved a response rate of 95.2%.

We construct total fuel consumption in common energy units in two steps since the surveys report a broad range of energy inputs, including electricity, gaseous fuels, petroleum products, coal and coke, and steam. First, we standardize each fuel's reported quantity. Most establishments report physical quantities, but some report monetary values; in these cases, we recover physical quantities by dividing expenditures by the corresponding regional fiscal-year fuel price. Because fuels are reported in different units, we express physical quantities in liters, kilograms, or cubic meters, as appropriate for each fuel. Second, we convert the resulting quantities into a common energy unit using conversion factors based primarily on the 2013 standard gross calorific values published by the ANRE (Agency for Natural Resources and Energy, 2022). We aggregate energy content across fuel types and divide by 1,000 to express total fuel use in GJ.

We construct fuel consumption for on-site generation using the reported purpose of fuel consumption. We define fuel used for self-generation as fuel consumed by boilers, cogeneration systems, and diesel generators. For each fuel type, we calculate fuel used directly in production as total fuel consumption minus fuel used by these self-generation facilities. Observations with negative constructed values are excluded, and the resulting measures are winsorized at the 1st and 99th percentiles.

The ECS does not directly report equipment ownership, so we infer self-generator ownership using self-generated electricity and grid electricity purchases. We classify a plant as owning a self-generator once its self-generated electricity reaches at least 5% of

its grid electricity purchases. A plant with positive self-generation and no grid purchases is also classified as an owner. Ownership begins in the first year this condition is satisfied and persists thereafter. For the heterogeneity analysis by pre-accident ownership, plants are classified as owners if they satisfy this condition by 2010.

A-3 Robustness-Control Variables

This appendix describes the regional variables used in the robustness checks in Section 5.5. We construct supply-chain exposure measures using prefecture-by-industry input-output tables from RIETI (RIETI, 2005). We define disaster-affected areas as Iwate, Miyagi, and Fukushima prefectures, which accounted for more than 99% of deaths and missing persons from the earthquake and tsunami. The direct exposure measure captures the share of a prefecture-industry's intermediate inputs sourced from disaster-affected areas. For example, if an iron factory sources 50% of its materials from disaster-affected areas, its direct exposure is 50%. The indirect exposure measure incorporates upstream input-output linkages to capture exposure through suppliers' suppliers. For example, a machinery plant that procures iron from suppliers using materials from disaster-affected areas would have indirect exposure. We use the 2005 tables, the latest available wave before the Fukushima accident. Because these exposure measures are time-invariant, the robustness specifications interact the pre-disaster measures with year indicators.

We construct local demand controls from prefecture-level economic indicators published by the Ministry of Internal Affairs and Communications (MIC, 2025). These controls include gross prefectural product in non-manufacturing sectors, total population, working-age population, construction investment, consumption, and the job-to-applicant ratio. Construction investment may be relevant because the earthquake increased reconstruction demand. The job-to-applicant ratio measures labor market tightness, which is especially relevant in the short run.

B Empirical Strategy

B-1 Measuring Electricity Supply Shocks

Table B1: Peak-Day Electricity Supply-Demand Conditions

	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-2 Derivation of the Small-Share Benchmark

Suppose P denotes the output price and P^K, P^L, P^M, and P^E denote input prices. A plant solves

$$\max_{K,L,M,E} \; P\Omega F(K,L,M,E) - P^K K - P^L L - P^M M - P^E E. \tag{B1}$$

Taking logs gives

$$\ln Y = \ln \Omega + \ln F(K,L,M,E). \tag{B2}$$

Differentiating with respect to shock severity $\mathcal{S}$ yields

$$\frac{d \ln Y}{d\mathcal{S}} = \frac{d \ln \Omega}{d\mathcal{S}} + \frac{1}{F}\frac{dF}{d\mathcal{S}}. \tag{B3}$$

By the chain rule,

$$\begin{aligned} \frac{1}{F}\frac{dF}{d\mathcal{S}} &= \frac{F_K}{F}\frac{dK}{d\mathcal{S}} + \frac{F_L}{F}\frac{dL}{d\mathcal{S}} + \frac{F_M}{F}\frac{dM}{d\mathcal{S}} + \frac{F_E}{F}\frac{dE}{d\mathcal{S}} \\ &= \frac{F_K K}{F}\frac{d \ln K}{d\mathcal{S}} + \frac{F_L L}{F}\frac{d \ln L}{d\mathcal{S}} + \frac{F_M M}{F}\frac{d \ln M}{d\mathcal{S}} + \frac{F_E E}{F}\frac{d \ln E}{d\mathcal{S}}. \end{aligned} \tag{B4}$$

Define the output elasticity with respect to each input as

$$\theta_K \equiv \frac{F_K K}{F}, \qquad \theta_L \equiv \frac{F_L L}{F}, \qquad \theta_M \equiv \frac{F_M M}{F}, \qquad \theta_E \equiv \frac{F_E E}{F}. \tag{B5}$$

Substituting these expressions into the derivative above yields

$$\frac{d \ln Y}{d\mathcal{S}} = \theta_E \frac{d \ln E}{d\mathcal{S}} + \theta_M \frac{d \ln M}{d\mathcal{S}} + \theta_L \frac{d \ln L}{d\mathcal{S}} + \theta_K \frac{d \ln K}{d\mathcal{S}} + \frac{d \ln \Omega}{d\mathcal{S}}. \tag{B6}$$

Since electricity is a flexible input, under price-taking behavior, the first-order condition for electricity is

$$P\Omega F_E = P^E. \tag{B7}$$

Multiplying this condition by electricity use and dividing by nominal revenue $PY = P\Omega F$ gives

$$\theta_E = \frac{F_E E}{F} = \frac{P^E E}{PY} = s_E. \tag{B8}$$

For the small-share benchmark, evaluate this condition at the pre-shock optimum, denoted by stars. Holding labor, materials, capital, and productivity fixed at L^*, M^*, K^*, and Ω^*, respectively, sets the non-electricity terms in equation (2) to zero. The direct

first-order effect of the shock is therefore

$$\left.\frac{d\ln Y}{d\mathcal{S}}\right|_{L=L^*,\,M=M^*,\,K=K^*,\,\Omega=\Omega^*} \approx \theta_E^* \frac{d\ln E}{d\mathcal{S}} = s_E^* \frac{d\ln E}{d\mathcal{S}}. \tag{B9}$$

The approximation is local and abstracts from induced adjustments in non-electricity inputs and productivity. This establishes equation (3).

B-3 Production Function Estimation

B-3.1 Derivation of the Flexible-Input Share Equation

The plant's static problem with respect to electricity input is

$$\max_{E_{it}} \; P_{it}\,\mathbb{E}\Big[\exp\big(f(k_{it},l_{it},m_{it},e_{it})+\omega_{it}+\epsilon_{it}\big)\mid k_{it},l_{it},m_{it},e_{it},\omega_{it}\Big] - P_t^E E_{it}, \tag{B10}$$

where $e_{it} \equiv \ln E_{it}$.

The first-order condition is

$$P_{it}\,\mathbb{E}\left[\frac{\partial Y_{it}}{\partial E_{it}} \mid k_{it},l_{it},m_{it},e_{it},\omega_{it}\right] = P_t^E. \tag{B11}$$

Using

$$\frac{\partial Y_{it}}{\partial E_{it}} = \frac{\partial f(k_{it},l_{it},m_{it},e_{it})}{\partial e_{it}}\frac{Y_{it}}{E_{it}}, \tag{B12}$$

we obtain

$$\frac{P_t^E E_{it}}{P_{it}Y_{it}} = \theta_{E,it}\,\mathbb{E}\big[\exp(\epsilon_{it})\mid k_{it},l_{it},m_{it},e_{it}\big]\exp(-\epsilon_{it}), \tag{B13}$$

where

$$\theta_{E,it} \equiv \frac{\partial f(k_{it},l_{it},m_{it},e_{it})}{\partial e_{it}}. \tag{B14}$$

Imposing the normalization $\mathbb{E}\big[\exp(\epsilon_{it})\mid k_{it},l_{it},m_{it},e_{it}\big]=1$ yields

$$\ln s_{E,it} = \ln\theta_{E,it} - \epsilon_{it}. \tag{B15}$$

B-3.2 Recovering CES Distribution Parameters

Suppose the production aggregator within an industry is CES:

$$F(K,L,M,E) = \left(\sum_{j\in\{K,L,M,E\}} \alpha_j X_j^{\rho}\right)^{1/\rho}, \qquad \rho \equiv \frac{\sigma-1}{\sigma}, \tag{B16}$$

where $\sum_j \alpha_j = 1$ and σ is the elasticity of substitution across inputs. For $\sigma \neq 1$, the associated unit-cost function is proportional to

$$c(p) = \left[\sum_{j \in \{K,L,M,E\}} \alpha_j^{\sigma} p_j^{1-\sigma} \right]^{1/(1-\sigma)}. \tag{B17}$$

By Shephard's lemma, input j's share of total cost is

$$s_j^C(p) = \frac{\alpha_j^{\sigma} p_j^{1-\sigma}}{\sum_{k \in \{K,L,M,E\}} \alpha_k^{\sigma} p_k^{1-\sigma}}. \tag{B18}$$

Under constant returns to scale and cost minimization, an input's normalized output elasticity equals its cost share. We therefore construct the industry-specific cost-share target as

$$\hat{s}_j = \frac{\hat{\theta}_j}{\sum_{k \in \{K,L,M,E\}} \hat{\theta}_k}. \tag{B19}$$

Expressing input prices relative to their reference levels gives $p_j^0 = 1$ for every input. Equation (B18) then implies

$$\hat{s}_j = \frac{\hat{\alpha}_j(\sigma)^{\sigma}}{\sum_{k \in \{K,L,M,E\}} \hat{\alpha}_k(\sigma)^{\sigma}}. \tag{B20}$$

Conditional on a candidate σ and imposing $\sum_j \hat{\alpha}_j(\sigma) = 1$, the CES distribution parameters are therefore

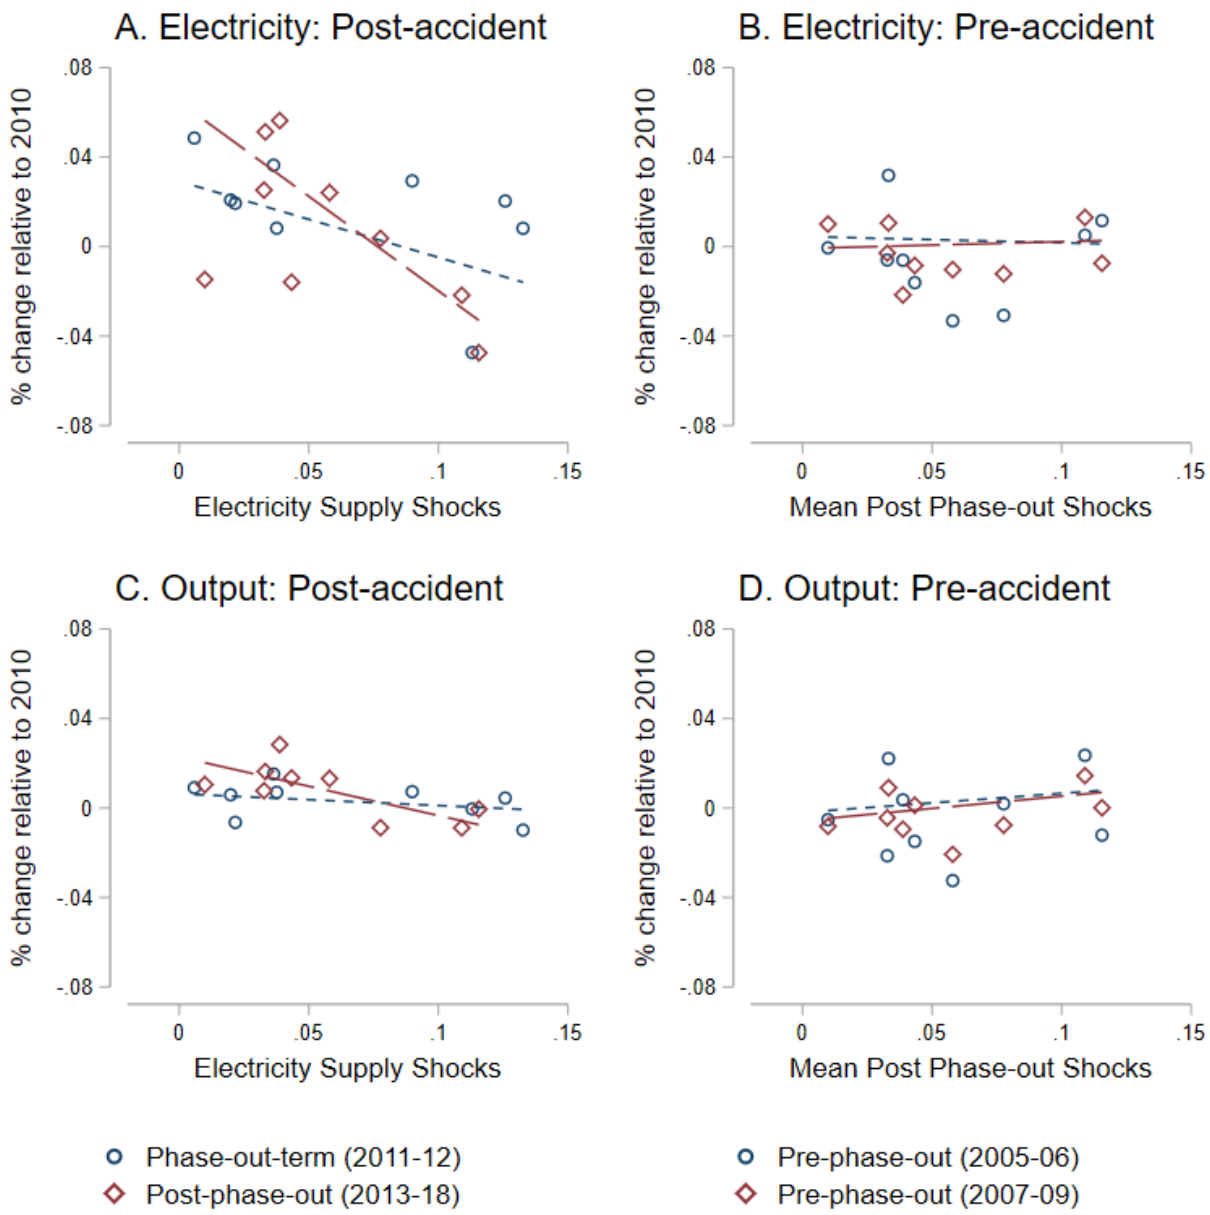
$$\hat{\alpha}_j(\sigma) = \frac{\hat{s}_j^{1/\sigma}}{\sum_{k \in \{K,L,M,E\}} \hat{s}_k^{1/\sigma}}, \qquad j \in \{K, L, M, E\}. \tag{B21}$$

Substituting Equation (B21) into the cost-share equation reproduces $\hat{s}_j$ for every candidate value of σ.

C Main Results

C-1 Additional Results

Figure C1: 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.

Table C1: 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	0.01	-0.01	-0.05	-0.03	-0.21
(Phase-out Terms)	[0.71]	[0.77]	[0.46]	[0.64]	[0.26]
Supply Shocks	-0.04	-0.19**	-0.18**	-0.19	-0.43
(Post-Phase-out)	[0.51]	[0.01]	[0.02]	[0.11]	[0.16]
Panel B: The effects of Supply Shocks using 2SLS					
Supply Shocks	0.03	-0.02	-0.07*	-0.01	-0.40
(Phase-out Terms)	[0.34]	[0.61]	[0.08]	[0.79]	[0.10]
Supply Shocks	-0.01	-0.17**	-0.20**	-0.07	-0.90*
(Post-Phase-out)	[0.79]	[0.04]	[0.04]	[0.46]	[0.08]
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 regressions include year-by-industry (3-digit) fixed effects and use 2010 log-salary weights. We use wild cluster bootstrap with regions as clusters. P-values are reported in brackets. $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$.

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.

Recent manufacturing estimates are approximately 0.1 for labor–electricity substitution in India (Aggarwal, 2026), 0.24–0.26 in the United States (Hawkins-Pierot and Wagner, 2026), and 0.5–0.55 for broader energy inputs in Indonesia (Abe et al., 2026).

We assess whether relative price movements account for the observed evolution of E/L after Fukushima. Rather than estimate σ from potentially endogenous relative prices, we consider $\sigma \in \{0.2, 0.5, 0.7\}$. For each value, we recover the change in electricity-biased technology as

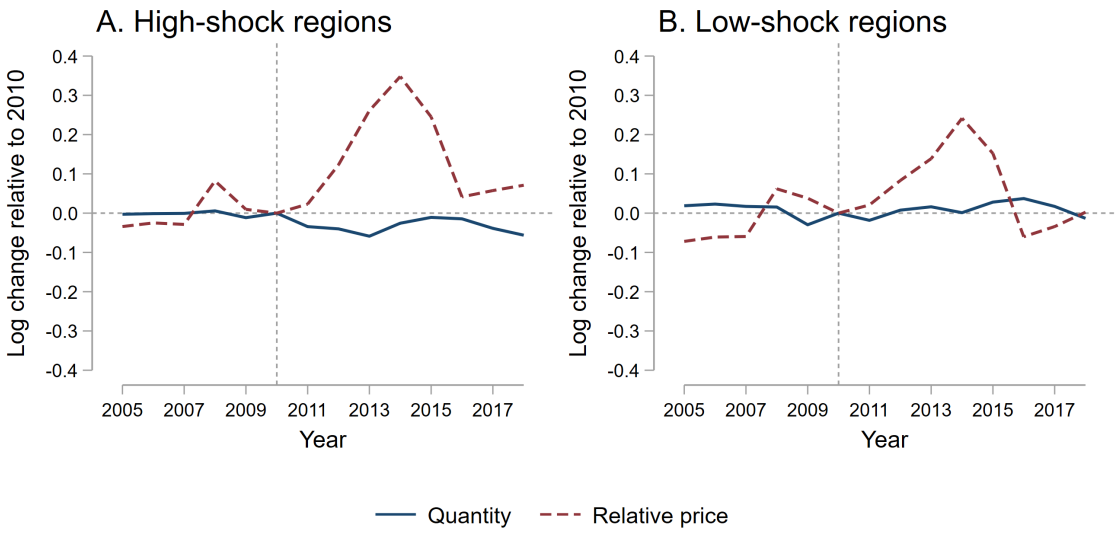
$$\Delta\omega^E_{it}(\sigma) = \frac{\Delta \ln(E_{it}/L_{it}) + \sigma \Delta \ln(p^E_{it}/p^L_{it})}{1-\sigma}. \tag{C2}$$

This measure captures changes in electricity productivity relative to labor productivity under the CES framework.

Appendix Figure C2 plots the relative quantity and price series for high- and low-shock regions. Relative prices rose more and the electricity-to-labor ratio fell more in high-shock regions (Panel A) than in low-shock regions (Panel B), consistent with substitution away from electricity. Appendix Table C2 reports OLS and 2SLS estimates of the effect of supply shocks on implied electricity-biased productivity. The post-phase-out estimates are statistically insignificant across all assumed elasticities ($\sigma \in \{0.2, 0.5, 0.7\}$). Depending on the assumed elasticity, the coefficient estimates change sign and range from -0.6 to 0.3. Overall, we find little evidence that supply shocks induced electricity-biased technical change.

[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 Electricity Inputs and Prices



Notes: The figure plots changes in the log electricity-to-labor ratio and the log relative price of electricity to labor. Panel A displays regions with supply shocks above the median ("High-shock"), while Panel B displays those below the median ("Low-shock"). All variables are log-differences relative to 2010.

Table C2: Effects of Supply Shocks on Electricity-Biased Technical Change

	Assumed Elasticity of Substitution		
	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 0.7$
	(1)	(2)	(3)
Panel A. Effects of supply shocks using OLS			
Supply shocks	-0.26**	-0.41**	-0.66**
(Phase-out Terms)	[0.05]	[0.03]	[0.03]
Supply shocks	-0.38*	-0.13	0.33
(Post-Phase-out)	[0.06]	[0.69]	[0.54]
Panel B. Effects of supply shocks using 2SLS			
Supply shocks	-0.23	-0.38**	-0.65*
(Phase-out Terms)	[0.21]	[0.05]	[0.08]
Supply shocks	-0.60	-0.27	0.30
(Post-Phase-out)	[0.31]	[0.85]	[0.77]
Year-by-industry FE	Yes	Yes	Yes

Notes: The dependent variable is the implied change in electricity-biased technology, $\Delta\omega^E(\sigma) = [\Delta \ln(E/L) + \sigma\Delta \ln(p^E/p^L)]/(1-\sigma)$. Columns (1)–(3) assume elasticities of substitution of 0.2, 0.5, and 0.7, respectively. The phase-out rows use data from 2005–2010 and 2011–2012. The post-phase-out rows use data from 2005–2010 and 2013–2018. All regressions include year-by-industry fixed effects, where industries are defined at the three-digit level, and use 2010 log-labor weights. The 2SLS specifications instrument supply shocks with predicted changes in nuclear generation and the initial reserve margin. P-values from a wild cluster bootstrap using regions as clusters are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C-3 Alternative Explanations and Robustness Checks

Table C3: Supply-Chain Exposure and Local-Demand Controls

	(1)	(2)	(3)	(4)	(5)
	Dep Var: $\Delta \ln(E) = \ln(E) - \ln(E_{2010})$				
Panel A: The effects of Supply Shocks using OLS					
Supply Shocks	-0.24	-0.24*	-0.24*	-0.24	-0.22
(Phase-out Terms)	[0.11]	[0.10]	[0.10]	[0.11]	[0.19]
Supply Shocks	-0.63**	-0.68**	-0.69***	-0.69***	-0.55***
(Post-Phase-out)	[0.03]	[0.01]	[0.01]	[0.01]	[0.01]
Panel B: The effects of Supply Shocks using 2SLS					
Supply Shocks	-0.21	-0.29	-0.28	-0.28	-0.28
(Phase-out Terms)	[0.36]	[0.16]	[0.16]	[0.16]	[0.28]
Supply Shocks	-0.89*	-0.87**	-0.87**	-0.88**	-0.75**
(Post-Phase-out)	[0.06]	[0.02]	[0.01]	[0.01]	[0.02]
	Dep Var: $\Delta \ln(R) = \ln(R) - \ln(R_{2010})$				
Panel C: The effects of Supply Shocks using OLS					
Supply Shocks	-0.04	-0.01	-0.01	-0.01	-0.02
(Phase-out Terms)	[0.23]	[0.72]	[0.75]	[0.81]	[0.54]
Supply Shocks	-0.19***	-0.18**	-0.18**	-0.18*	-0.18
(Post-Phase-out)	[0.00]	[0.05]	[0.05]	[0.06]	[0.11]
Panel D: The effects of Supply Shocks using 2SLS					
Supply Shocks	-0.02	0.00	0.00	0.00	0.00
(Phase-out Terms)	[0.61]	[0.95]	[0.96]	[0.95]	[0.98]
Supply Shocks	-0.23**	-0.20**	-0.20**	-0.20*	-0.23
(Post-Phase-out)	[0.04]	[0.05]	[0.05]	[0.06]	[0.12]
Direct supply-chain	No	Yes	No	Yes	No
Indirect supply-chain	No	No	Yes	Yes	No
Local-demand	No	No	No	No	Yes
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes

Notes: Column (1) reproduces the baseline estimates in Table 4 using the full analysis sample. Columns (2)–(5) exclude plants in Iwate, Miyagi, and Fukushima and use a common complete-case sample. Direct exposure measures the share of intermediate inputs sourced from these prefectures. Indirect exposure incorporates upstream linkages through the Leontief inverse. Supply-chain exposure measures are interacted with year indicators. Local-demand controls include gross prefectural product in non-manufacturing sectors, total and working-age populations, construction investment, consumption, and the job-to-applicant ratio. See Appendix A-3 for variable construction. All regressions include three-digit-industry-by-year fixed effects and use 2010 log-labor weights. P-values from a wild cluster bootstrap using regions as clusters are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table C4: Robustness Checks

	Plant Controls (1)	Two-way FE (2)	No Weights (3)
	Dep Var: $\Delta \ln(E) = \ln(E) - \ln(E_{2010})$		
Panel A: The effects of Supply Shocks using OLS			
Supply Shocks	-0.26*	-0.17	-0.23
(Phase-out Terms)	[0.08]	[0.23]	[0.12]
Supply Shocks	-0.64**	-0.48**	-0.62**
(Post-Phase-out)	[0.03]	[0.02]	[0.03]
Panel B: The effects of Supply Shocks using 2SLS			
Supply Shocks	-0.32	-0.21	-0.29
(Phase-out Terms)	[0.22]	[0.47]	[0.21]
Supply Shocks	-0.82**	-0.79***	-0.78**
(Post-Phase-out)	[0.02]	[0.01]	[0.02]
	Dep Var: $\Delta \ln(R) = \ln(R) - \ln(R_{2010})$		
Panel C: The effects of Supply Shocks using OLS			
Supply Shocks	-0.06	-0.03	-0.04
(Phase-out Terms)	[0.36]	[0.44]	[0.19]
Supply Shocks	-0.16	-0.14	-0.19***
(Post-Phase-out)	[0.24]	[0.17]	[0.00]
Panel D: The effects of Supply Shocks using 2SLS			
Supply Shocks	-0.06	0.01	-0.04
(Phase-out Terms)	[0.39]	[0.94]	[0.33]
Supply Shocks	-0.20	-0.17	-0.24***
(Post-Phase-out)	[0.19]	[0.35]	[0.00]
Year-by-Industry FE	Yes	Yes	Yes

Notes: Columns (1)–(3) add plant controls, use two-way fixed effects, and omit the baseline weights, respectively. All regressions include three-digit-industry-by-year fixed effects. See the main text for details on each robustness check. We use a wild cluster bootstrap to obtain p-values, clustering by region. P-values are reported in brackets. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

D Plants' Responses to Supply Shocks

Table D1: Spillover Effects: Production Reallocation Across Regions

Dep Var: $\Delta \ln(Y) = \ln(Y) - \ln(Y_{2010})$	Electricity (1)	Output (2)
Panel A: The effects of Supply Shocks using OLS		
Supply Shocks (Phase-out Terms)	0.24 [0.16]	-0.02 [0.70]
Supply Shocks (Post-Phase-out)	0.48** [0.04]	-0.06 [0.59]
Panel B: The effects of Supply Shocks using 2SLS		
Supply Shocks (Phase-out Terms)	0.19 [0.44]	-0.03 [0.74]
Supply Shocks (Post-Phase-out)	0.84** [0.04]	-0.04 [0.80]
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 D2: 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	0.09	-0.10	0.18
	[0.82]	[0.94]	[0.75]
Year-by-Industry FE	Yes	Yes	Yes

Notes: The table reports estimates of equation (4) 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. $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$.

D-1 Electricity Supply Shocks and On-Site Generation

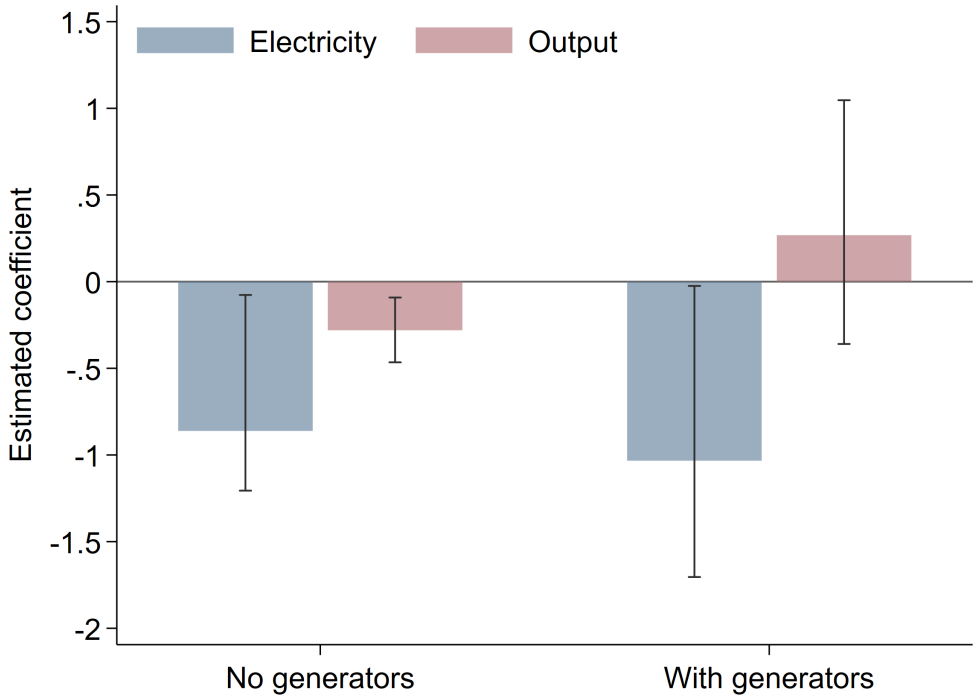


Figure D1: Effects of Supply Shocks by Generator Ownership

Notes: The figure reports post-phase-out 2SLS estimates for grid electricity purchases and output, estimated separately for plants with and without on-site generators. Generator ownership is defined as self-generated electricity reaching at least 5% of grid purchases by 2010; the non-generator group consists of the remaining matched plants. The regressions use data from 2005–2010 and 2013–2018, include three-digit-industry-by-year fixed effects, and use 2010 log-salary weights. The vertical bars show 90% confidence intervals obtained from a wild cluster bootstrap using regions as clusters.

E From Reduced-Form Effects to Production Mechanisms

E-1 Simulation Inputs

Table E1 summarizes the technology, demand, and policy inputs used in the simulation. Three features of the scenario construction are worth noting. First, we hold technology and demand parameters fixed between the no-shock baseline and the post-phase-out scenario. Second, we express all input prices relative to their 2010 levels. Relative non-electricity prices remain one, while the two scenarios differ in the relative electricity price and shock-season rationing.

Table E1: Simulation Inputs

Object	Level	Description or source	Baseline	Post-phase-out
Panel A: Technology				
σ	Common	Elasticity of substitution	$\widehat{\sigma}$	$\widehat{\sigma}$
$\widehat{\boldsymbol{\theta}}_j$	Industry	Estimated output elasticities	$\widehat{\boldsymbol{\theta}}_j$	$\widehat{\boldsymbol{\theta}}_j$
Ω_i	Plant	Baseline productivity	Ω_i	Ω_i
κ	Common	Capital-utilization elasticity	$\widehat{\kappa}$	$\widehat{\kappa}$
v_i	Plant	Baseline capital–electricity ratio; $\frac{K_{2010}}{E_{2010}}$	v_i	v_i
G_i	Plant	Generator Ownership	G_i	G_i
Panel B: Demand				
η	Common	Demand elasticity; 5	5	5
B_i	Plant	Demand shifter; $Y_{i,2010}$	B_i	B_i
Panel C: Policies				
$\mathcal{S}_r$	Region	Shock-season electricity rationing	0	$\overline{\mathcal{S}}_r$
ϕ	Common	Shock-season share of the year	-	90/365
p_r^E	Region	Regional electricity price	1	$\overline{p}_r^E$
p^K, p^L, p^M	Common	Non-electricity input prices	1	1

E-2 Predicted Electricity-Price Changes

To construct the shock-induced electricity-price input used in the simulation, we estimate how regional electricity supply shocks affected the prices paid by manufacturing plants. Observed electricity prices increased through 2014 and declined during 2015–2016, largely reflecting nationwide movements in global energy costs. We therefore control for nationwide electricity-price trends to isolate the component attributable to regional supply shocks. We estimate the following specification separately for the phase-out period

(2011–2012) and the post-phase-out period (2013–2018):

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

where p^E_{rt} denotes the average electricity price in region r and year t, which we assign to plant i; $\mathcal{S}_{rt}$ denotes the region–year supply shock; and γ_{jt} denotes three-digit-industry-by-year fixed effects. These fixed effects absorb nationwide price movements, including common changes in fossil-fuel costs. We estimate Equation (E1) by OLS and 2SLS, using the instruments defined in Equation (7) for the latter. All specifications use the same 2010 total-salary weights as the main analysis, and standard errors are clustered by region.

Table E2 shows no detectable price response during 2011–2012: both the OLS and 2SLS estimates are close to zero. This pattern is consistent with adjustment lags in regulated tariffs under Japan's cost-of-service regulation. During 2013–2018, the OLS and 2SLS estimates are 0.78 and 1.12, respectively. The 2SLS estimate implies that a 1 ppt increase in the supply shock raises the log electricity price by 0.0112, corresponding to an approximately 1.12% increase.

For the simulation, we normalize the 2010 electricity price to one and use the post-phase-out 2SLS estimate to construct each region's predicted price:

$$\overline{p}^E_r \equiv \frac{p^E_{r,\text{post}}}{p^E_{r,2010}} = \exp\left(\widehat{\beta}_{\text{post}} \, \overline{\mathcal{S}}_r\right), \tag{E2}$$

where $\overline{\mathcal{S}}_r$ is the region's mean supply shock over 2013–2018. The sample's weighted mean shock is 7.5%. The corresponding weighted mean predicted log-price change is 8.4 log points, while the weighted mean of predicted relative prices implies an approximately 9.0% increase in levels. In Tokyo, where the mean shock is 10.8%, the predicted log-price change is 12.9 log points, corresponding to an approximately 13.8% increase in levels. Because regulated tariffs do not vary seasonally within a year, we apply the same predicted price during the shock and off-shock seasons.

Table E2: Effects of Electricity Supply Shocks on Electricity Prices

	Dep Var: $\Delta \ln p^E = \ln p^E_{rt} - \ln p^E_{r,2010}$	
	Phase-out (1)	Post-phase-out (2)
Panel A. The effects of supply shocks using OLS		
Supply shock	0.03	0.78***
	(0.19)	(0.20)
Panel B. The effects of supply shocks using 2SLS		
Supply shock	0.00	1.12**
	(0.32)	(0.35)
Year-by-Industry FE	Yes	Yes

Notes: Column (1) uses data from 2005–2010 and 2011–2012. Column (2) uses data from 2005–2010 and 2013–2018. Panel B instruments the supply shock with predicted nuclear-supply losses and the 2010 reserve margin. All regressions include three-digit-industry-by-year fixed effects and use 2010 total-salary weights. Standard errors clustered by region are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

E-3 Predictions under the Standard CES Technology

We first predict inputs and output under the electricity-supply-shock scenario using the standard CES technology. In this subsection, we abstract from electricity-induced capital idleness by setting $\kappa = 0$. We introduce electricity-dependent capital utilization in Appendix E-4.

E-3.1 Baseline and price-only equilibria

The simulation considers alternative values of the elasticity of substitution, σ. For each candidate value, we recover the CES distribution parameters α using Equation (15), while holding the baseline cost shares in Equation (14) fixed.

Plant i in industry j chooses inputs, output, and its output price to maximize profits:

$$\begin{aligned} \max_{K,L,M,E,P,Y} \quad & \Pi_i = P_i Y_i - p^K K_i - p^L L_i - p^M M_i - p^E E_i, \\ \text{subject to} \quad & Y_i = \Omega_i \left(\alpha_{Kj} K_i^\rho + \alpha_{Lj} L_i^\rho + \alpha_{Mj} M_i^\rho + \alpha_{Ej} E_i^\rho \right)^{1/\rho}, \\ & Y_i = B_i P_i^{-\eta}, \qquad \rho = \frac{\sigma - 1}{\sigma}. \end{aligned} \tag{E3}$$

For $\sigma \neq 1$, substitution of Equation (15) into the dual cost function gives

$$c_{ij}(p) = \frac{1}{\Omega_i}\left[\sum_{h\in\{K,L,M,E\}} s_{hj}^0 (p^h)^{1-\sigma}\right]^{1/(1-\sigma)}, \tag{E4}$$

where each p^h is an input-price index relative to its 2010 level. Hence $p^{h,0} = 1$ and $c_{ij}(p^0) = 1/\Omega_i$ for every value of σ. When only the electricity price changes, the relative unit-cost index is

$$\frac{c_{ij}(p^1)}{c_{ij}(p^0)} = \left[1 - s_{Ej}^0 + s_{Ej}^0 (p^{E,1})^{1-\sigma}\right]^{1/(1-\sigma)}. \tag{E5}$$

This expression makes clear that comparisons across σ hold the baseline electricity share fixed. Constant-elasticity demand implies the markup $\mu = \eta/(\eta-1)$, so the optimal output price and output are

$$P_i^*(p) = \mu c_{ij}(p), \qquad Y_i^*(p) = B_i[\mu c_{ij}(p)]^{-\eta}. \tag{E6}$$

The corresponding conditional factor demands are

$$X_{ih}^*(p) = \frac{Y_i^*(p)}{\Omega_i}\left(\frac{c_{ij}(p)\Omega_i \alpha_{hj}}{p^h}\right)^{\sigma}, \qquad h \in \{K, L, M, E\}. \tag{E7}$$

We express each simulated outcome as the same log difference from the no-shock baseline used in the empirical analysis. Let $p^0 = (1, 1, 1, 1)$ denote the vector of 2010-relative price indices, and let p^1 denote the post-phase-out relative-price vector. Define

$$\Delta \ln c_{ij} \equiv \ln c_{ij}(p^1) - \ln c_{ij}(p^0), \qquad \Delta \ln p^h \equiv \ln p^{h,1} - \ln p^{h,0}. \tag{E8}$$

Because productivity, demand, and the markup are held fixed across simulations, they cancel from the log differences. Taking logs of Equations (E6) and (E7) gives

$$\begin{aligned}
\Delta \ln Y_i &= -\eta\, \Delta \ln c_{ij}, \\
\Delta \ln R_i &= (1-\eta)\, \Delta \ln c_{ij}, \\
\Delta \ln X_{ih} &= (\sigma-\eta)\, \Delta \ln c_{ij} - \sigma\, \Delta \ln p^h, \qquad h \in \{K, L, M, E\}.
\end{aligned} \tag{E9}$$

Here, $R_i \equiv P_i Y_i$ denotes revenue and $\Delta \ln Z_i \equiv \ln Z_i^1 - \ln Z_i^0$ for any outcome Z_i. Thus, we predict the plant-level log changes used in the empirical analysis without using observed plant-level baseline inputs.

For the Cobb–Douglas benchmark, $\sigma = 1$, we evaluate these expressions at the contin-

uous limit of the CES cost function. In this case,

$$\Delta \ln c_{ij} = \sum_{h \in \{K,L,M,E\}} s^0_{hj} \Delta \ln p^h. \tag{E10}$$

E-3.2 Electricity-price elasticities

Define electricity's share in unit cost as

$$s^C_{E,ij}(p) = \frac{s^0_{Ej}(p^E)^{1-\sigma}}{\sum_{h \in \{K,L,M,E\}} s^0_{hj}(p^h)^{1-\sigma}}. \tag{E11}$$

At the baseline relative-price vector $p^0 = (1, 1, 1, 1)$, this expression equals s^0_{Ej} for every σ. By Shephard's Lemma, $\partial \ln c_{ij}(p)/\partial \ln p^E = s^C_{E,ij}(p)$. Combining this result with demand and the conditional factor demands gives

$$\begin{aligned}
\frac{d \ln Y_i}{d \ln p^E} &= -\eta s^C_{E,ij}, \\
\frac{d \ln X_{ih}}{d \ln p^E} &= (\sigma - \eta) s^C_{E,ij}, \qquad h \in \{K, L, M\}, \\
\frac{d \ln E_i}{d \ln p^E} &= -\sigma + (\sigma - \eta) s^C_{E,ij}.
\end{aligned} \tag{E12}$$

Producer surplus is $\Pi_i = (P_i - c_{ij})Y_i = (\mu - 1)c_{ij}Y_i$, implying

$$\frac{d \ln \Pi_i}{d \ln p^E} = (1 - \eta) s^C_{E,ij}. \tag{E13}$$

E-3.3 Seasonal rationing

Rationing applies only during the shock season S. For a plant in region r, shock-season electricity use must satisfy

$$E^S_i \le \overline{E}^S_i \equiv (1 - \overline{\mathcal{S}}_r) E^{S,0}_i, \tag{E14}$$

where $E^{S,0}_i$ denotes baseline shock-season electricity use. We first solve the unconstrained equilibrium under post-phase-out prices and denote the resulting price-only electricity demand by $E^{S,\text{ref}}_i$.

If $E^{S,\text{ref}}_i \le \overline{E}^S_i$, rationing is slack and the price-only solution applies. If $E^{S,\text{ref}}_i > \overline{E}^S_i$, we

represent the binding constraint with an incremental shadow price. Define

$$\Delta \ln \widetilde{p}_i^E \equiv \ln\left(\frac{\widetilde{p}_i^E}{p_r^{E,1}}\right), \qquad \Delta \ln E_i^{S,R} \equiv \ln\left(\frac{\overline{E}_i^S}{E_i^{S,\text{ref}}}\right). \tag{E15}$$

Using a first-order approximation around the unconstrained post-price equilibrium, the shadow-price adjustment is

$$\Delta \ln \widetilde{p}_i^E = \frac{\Delta \ln E_i^{S,R}}{-\sigma + (\sigma - \eta) s_{E,ij}^C}, \tag{E16}$$

where the electricity cost share is evaluated at the unconstrained post-price equilibrium. The corresponding additional rationing-induced changes are

$$\begin{aligned} \Delta \ln Y_i^{S,R} &= -\eta s_{E,ij}^C \Delta \ln \widetilde{p}_i^E, \\ \Delta \ln X_{ih}^{S,R} &= (\sigma - \eta) s_{E,ij}^C \Delta \ln \widetilde{p}_i^E, \qquad h \in \{K, L, M\}, \\ \Delta \ln \Pi_i^{S,R} &= (1 - \eta) s_{E,ij}^C \Delta \ln \widetilde{p}_i^E. \end{aligned} \tag{E17}$$

The total shock-season response equals the price-only response plus this rationing adjustment. During the off-season O, plants face post-phase-out prices but no rationing.

Annual aggregation. We combine the seasonal predictions using the shock-season share $\phi = 90/365$:

$$\Delta \ln Z_i \equiv \phi \, \Delta \ln Z_i^S + (1 - \phi) \, \Delta \ln Z_i^O. \tag{E18}$$

This expression is a first-order approximation to the log change in an annual outcome constructed by aggregating the two seasons in levels.

E-3.4 Plants with generators

For plants that own generators, we retain the model-simulated response of grid-purchased electricity but set the structural responses of capital, labor, materials, output, and revenue to zero. This treatment is motivated by the heterogeneous reduced-form evidence in Section 6.2: generator owners reduce grid purchases, consistent with substitution toward self-generation, but exhibit no meaningful declines in non-electricity inputs or output. The ownership adjustment is imposed after the CES solution and is not implied by the first-order conditions above.

E-4 Incorporating Capital Idleness Induced by Electricity Supply Shocks

We use effective capital services, rather than installed capital, to predict inputs and output under electricity supply shocks. Effective capital services are defined in Equation (17) as

$$\widetilde{K}_{it} = \left[(1-\kappa)K_{it}^{\rho} + \kappa(v_i E_{it})^{\rho}\right]^{1/\rho}, \qquad v_i = \frac{K_{i,2010}}{E_{i,2010}}.$$

Because $v_i E_{i,2010} = K_{i,2010}$, effective and installed capital coincide at the 2010 baseline: $\widetilde{K}_{i,2010} = K_{i,2010}$. Moreover, the baseline elasticities of effective capital services with respect to installed capital and electricity are

$$\left.\frac{\partial \ln \widetilde{K}_{it}}{\partial \ln K_{it}}\right|_{t=2010} = 1-\kappa, \qquad \left.\frac{\partial \ln \widetilde{K}_{it}}{\partial \ln E_{it}}\right|_{t=2010} = \kappa. \tag{E19}$$

To incorporate capital idleness, we adjust the estimated capital and electricity output elasticities. Let $\widehat{\theta}_{Kj}$ denote the estimated capital elasticity. Equation (E19) shows that, at the 2010 baseline, a 1% change in installed capital changes effective capital services by $(1-\kappa)$%. Using the chain rule, the adjusted capital elasticity ($\widetilde{\theta}_{Kj}(\kappa)$) is

$$\begin{aligned}
\widetilde{\theta}_{Kj}(\kappa) &\equiv \left.\frac{\partial \ln Y_{it}}{\partial \ln K_{it}}\right|_{t=2010} \\
&= \left.\frac{\partial \ln Y_{it}}{\partial \ln \widetilde{K}_{it}}\frac{\partial \ln \widetilde{K}_{it}}{\partial \ln K_{it}}\right|_{t=2010} \\
&= (1-\kappa)\widehat{\theta}_{Kj}.
\end{aligned} \tag{E20}$$

Electricity affects output directly and through capital idleness. At the 2010 baseline, a 1% change in electricity changes effective capital services by κ%. The adjusted electricity elasticity ($\widetilde{\theta}_{Ej}(\kappa)$) therefore equals the original electricity elasticity plus κ times the capital elasticity:

$$\begin{aligned}
\widetilde{\theta}_{Ej}(\kappa) &\equiv \left.\frac{d \ln Y_{it}}{d \ln E_{it}}\right|_{t=2010} \\
&= \widehat{\theta}_{Ej} + \left.\frac{\partial \ln Y_{it}}{\partial \ln \widetilde{K}_{it}}\frac{\partial \ln \widetilde{K}_{it}}{\partial \ln E_{it}}\right|_{t=2010} \\
&= \widehat{\theta}_{Ej} + \kappa\widehat{\theta}_{Kj}.
\end{aligned} \tag{E21}$$

The output elasticities of labor and materials remain unchanged:

$$\widetilde{\theta}_{Lj}(\kappa) = \widehat{\theta}_{Lj}, \qquad \widetilde{\theta}_{Mj}(\kappa) = \widehat{\theta}_{Mj}. \tag{E22}$$

This adjustment reallocates part of the original capital elasticity to electricity but does not change the sum of the output elasticities.

Using these adjusted output elasticities, we recalculate the industry-specific baseline cost shares as

$$\widetilde{s}^0_{hj}(\kappa) \equiv \frac{\widetilde{\theta}_{hj}(\kappa)}{\sum_{g\in\{K,L,M,E\}} \widetilde{\theta}_{gj}(\kappa)}, \qquad h \in \{K, L, M, E\}. \tag{E23}$$

Because the sum of the output elasticities is unchanged, these shares can equivalently be written as

$$\begin{aligned} \widetilde{s}^0_{Kj}(\kappa) &= (1-\kappa)s^0_{Kj}, & \widetilde{s}^0_{Lj}(\kappa) &= s^0_{Lj}, \\ \widetilde{s}^0_{Mj}(\kappa) &= s^0_{Mj}, & \widetilde{s}^0_{Ej}(\kappa) &= s^0_{Ej} + \kappa s^0_{Kj}. \end{aligned} \tag{E24}$$

Finally, for each candidate value of σ, we recover the CES distribution parameters from these adjusted baseline shares following Equation (15). Thus, κ changes the baseline allocation of the capital and electricity shares.

E-4.1 Minimum-Distance Estimation

We jointly estimate the elasticity of substitution, σ, and the capital-utilization elasticity, κ, by minimizing the distance between the model-implied changes and their reduced-form counterparts. We first construct the normalized baseline shares in Equation (14). For each candidate κ, we form the effective shares in Equation (E23); for each candidate σ, we map those fixed shares into CES coefficients using Equation (15). We then obtain annual plant-level changes in electricity, labor, materials, and output using the simulation procedure described in Appendix E-3. We aggregate these predictions using the same baseline salary weights as in the main reduced-form regressions. For each outcome, the empirical target is the post-phase-out 2SLS coefficient multiplied by the weighted mean supply shock. The moment discrepancies are

$$g_q(\sigma, \kappa) \equiv \frac{\sum_i w_i \Delta \ln q_i^{\text{sim}}(\sigma, \kappa)}{\sum_i w_i} - \widehat{\beta}^{\text{RF}}_q \overline{\mathcal{S}}_w, \qquad q \in \{E, L, M, Y\}, \tag{E25}$$

where w_i denotes the corresponding regression weight and $\overline{\mathcal{S}}_w \equiv \sum_i w_i \mathcal{S}_i / \sum_i w_i$ is the weighted mean supply shock.

Using an identity weighting matrix, we define the minimum-distance criterion as

$$Q(\sigma, \kappa) = \sum_{q \in \{E,L,M,Y\}} g_q(\sigma, \kappa)^2. \tag{E26}$$

The estimator is

$$(\widehat{\sigma}, \widehat{\kappa}) = \underset{\sigma>0,\, 0\leq\kappa\leq 1}{\arg\min}\; Q(\sigma, \kappa). \tag{E27}$$

We first evaluate $Q(\sigma, \kappa)$ over a coarsely spaced grid of candidate parameter values to identify the region containing the minimum and then evaluate it over a more finely spaced grid within that region to obtain $(\widehat{\sigma}, \widehat{\kappa})$.

E-5 Results for Elasticities

Table E3: Comparison of Elasticity Estimates

	Estimation method	
	Reduced-form (1)	Structural model (2)
Own-price elasticity (electricity)	-0.80	-0.76
Cross-price elasticity (labor)	-0.15	-0.16
Cross-price elasticity (materials)	-0.15	-0.16
Output elasticity	-0.20	-0.18
Preferred parameters	$\sigma = 0.59$, $\kappa = 0.17$	

Notes: The table reports elasticities of inputs and output with respect to electricity-price changes. Column (1) divides each main reduced-form IV coefficient by the estimated IV response of log electricity prices to the supply shock. This reduced-form ratio does not explicitly incorporate the model-implied shadow price of a binding rationing constraint: the outcome coefficient may reflect both price and rationing responses, whereas the denominator contains only the market-price response. Column (2) reports finite-change structural elasticities from the price-only simulation under $\widehat{\sigma} = 0.59$ and $\widehat{\kappa} = 0.17$. For each plant, we divide its predicted log outcome response by its predicted log market-price change and then take the 2010-total-salary-weighted mean; rationing and its shadow price are inactive by construction. Generator owners retain their simulated grid-electricity response, while their structural labor, materials, and output responses are set to zero.

F Counterfactual Analysis

F-0.1 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:

$$p_{rt} = \left(\frac{Q_{rt}}{c_{rt}}\right)^{\beta}, \tag{F1}$$

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 $d \ln Q / d \ln p = 1/\beta$).[2]

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

$$Q_{rt}^* = c_{rt} \left(p_t^*\right)^{1/\beta}. \tag{F2}$$

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

$$\bar{Q}_t = \sum_r Q_{rt}^* = \left(p_t^*\right)^{1/\beta} \sum_r c_{rt}, \tag{F3}$$

which yields

$$p_t^* = \left(\frac{\bar{Q}_t}{\sum_r c_{rt}}\right)^{\beta}. \tag{F4}$$

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

$$p_t^* = \left(\frac{\bar{Q}_t}{(1+\kappa)\sum_r c_{rt}}\right)^{\beta}. \tag{F5}$$

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

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

Estimation. We set $\beta = 5$, corresponding to a supply elasticity $d \ln Q / d \ln p = 0.2$.[3] Rewrite (F1) in logs:

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

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

$$A_{rt} = \gamma_1 Sh_{rt}^{\text{Nuclear}} + \gamma_2 Sh_{rt}^{\text{Hydro}} + \gamma_3 Sh_{rt}^{\text{Thermal}} + \gamma_4 Sh_{rt}^{\text{Others}} + \tau_t + \delta_r + \varepsilon_{rt}. \tag{F6}$$

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

$$\widehat{c}_{rt} = \exp\left(-\widehat{A}_{rt}/\beta\right), \tag{F7}$$

and substitute $\widehat{c}_{rt}$ into (F5) to obtain the integrated-market counterfactual price p_t^*.

Results. Table F1, Column (2), 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 2010-total-salary-weighted national average log-price change falls from 8.4 log points in the baseline to 5.7 log points under integration, corresponding to level increases of approximately 9.0% and 5.8%, respectively.

[3] 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.

Table F1: Regional Price Changes: Baseline and Market Integration

		$100 \times \Delta \ln p^E$	
#	Region	Baseline (1)	Market integration (2)
1	Hokkaido	1.1%	-12.6%
2	Tohoku	8.7%	1.7%
3	Tokyo	12.2%	3.1%
4	Hokuriku	4.3%	18.7%
5	Chubu	3.6%	5.0%
6	Kansai	12.9%	8.0%
7	Chugoku	3.6%	12.8%
8	Shikoku	5.0%	2.3%
9	Kyushu	6.4%	10.1%
	National	8.4%	5.7%

Notes: Entries report $100 \times \Delta \ln p^E$. Column (1) reports the shock-induced price increase (baseline), and Column (2) reports the counterfactual price change under market integration. National values are weighted averages over the current plant sample using 2010 total salary.

References for Appendix

Abe, Tatsuya, Arlan Brucal, and Yuta Toyama. 2026. "Energy-Augmenting Productivity and Carbon Pricing: Evidence from Production Microdata." June, Working Paper.

Agency for Natural Resources and Energy. 2022. "Standard Calorific Values and Carbon Emission Factors by Energy Source." Ministry of Economy, Trade and Industry, Revised January 2020 and additionally revised November 2022; in Japanese.

Aggarwal, Raavi. 2026. "Input Substitution in Electricity Generation and Industrial Production: Evidence from India." *Journal of Development Economics* 179 103614. 10.1016/j.jdeveco.2025.103614.

Bank of Japan. 2026. "Corporate Goods Price Index (CGPI)." Bank of Japan Statistics, https://www.boj.or.jp/en/statistics/pi/index.htm.

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.

Cabinet Office. 2026. "Quarterly Estimates of Net Capital Stocks of Fixed Assets." Economic and Social Research Institute, National Accounts of Japan, https://www.esri.cao.go.jp/en/sna/stock/stock_top.html.

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

FEPC. 2026. "Electricity Review Japan." FEPC, https://www.fepc.or.jp/english/library/electricity_eview_japan/.

Hawkins-Pierot, Jonathan T., and Katherine R. H. Wagner. 2026. "Technology Lock-In and Costs of Delayed Climate Policy." *Journal of Political Economy Microeconomics*. 10.1086/740232, Forthcoming.

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.

Statistics Bureau of Japan. 2026. "Consumer Price Index." Ministry of Internal Affairs and Communications, https://www.stat.go.jp/english/data/cpi/index.html.

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.