# The Restart Deficit: Asymmetric Price Effects of Nuclear Reactor Restarts in Japan

APEP Autonomous Research* @olafdrw

March 23, 2026

**Abstract**

After Fukushima, Japan shut down all 54 nuclear reactors, raising wholesale electricity prices substantially. Since 2015, fourteen reactors have restarted—but has the price relief been commensurate? Using 2.5 million half-hourly spot price observations from the Japan Electric Power Exchange across nine regional grids, I estimate staggered difference-in-differences exploiting the sequential NRA restart approvals. Japan's 50Hz/60Hz frequency split—limiting interconnection to 1.2GW—provides a natural barrier against cross-regional spillovers. The Callaway–Sant'Anna estimate indicates restarts reduce wholesale prices by 0.60 ¥/kWh (4.0 percent), robust to wild cluster bootstrap ($p = 0.040$) and exact randomization inference ($p = 0.095$). This modest price dividend—in a system that added 75GW of solar during the nuclear hiatus—suggests that energy transitions reshape the merit order against which any returning technology is evaluated.

**JEL Codes:** L94, Q48, Q42
**Keywords:** nuclear energy, electricity prices, merit order, staggered DiD, Fukushima, energy transition

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 20m).

## 1. Introduction

Japan's post-Fukushima nuclear shutdown was one of the largest involuntary energy experiments in modern history. Between May 2011 and September 2013, every one of the country's 54 commercial reactors went offline, eliminating roughly 30 percent of national electricity generation capacity virtually overnight. The price consequences were severe: wholesale electricity costs surged, LNG imports doubled, and Japanese industry faced some of the highest energy bills in the OECD (Neidell et al., 2021).

Starting in August 2015, the Nuclear Regulation Authority (NRA) began approving reactor restarts under stringent post-Fukushima safety standards. By late 2024, fourteen reactors across five of Japan's ten electricity regions had returned to service. The natural expectation—that restarts would reverse the shutdown's price damage—rests on an implicit symmetry assumption: if removing nuclear from the merit order raised prices by $X$, then restoring it should lower prices by $X$. This paper tests that assumption.

I find that it fails. Using the Callaway–Sant'Anna staggered difference-in-differences estimator on 2.5 million half-hourly settlement prices from the Japan Electric Power Exchange (JEPX), I estimate that nuclear restarts reduce wholesale electricity prices by 0.60 ¥/kWh, or approximately 4.0 percent of the pre-treatment mean. While statistically significant—surviving both wild cluster bootstrap inference ($p = 0.040$) and exact randomization inference over all 126 possible treatment permutations ($p = 0.095$)—this price dividend is an order of magnitude smaller than the 5 ¥/kWh shutdown penalty estimated by Neidell et al. (2021). I call this gap the *restart deficit*.

The identification strategy exploits two features of Japan's electricity market. First, NRA restart approvals were staggered across regions from 2015 to 2024, creating the variation in treatment timing that modern DiD estimators require. Second, Japan's unique 50Hz/60Hz frequency split—a legacy of competing Western Electric and German generators adopted in the 1890s—creates a physical barrier between eastern and western grids. With only 1.2GW of frequency conversion capacity connecting the two systems, cross-regional price spillovers are sharply limited, strengthening the plausibility of the stable unit treatment value assumption.

Why is the restart dividend so much smaller than the shutdown cost? Three mechanisms, each operating independently, can generate the observed asymmetry. First, during the nuclear hiatus, Japan's generous feed-in tariff triggered an explosion of solar photovoltaic deployment—installed capacity grew from 5GW in 2011 to over 80GW by 2023 (International Energy Agency, 2023). This solar capacity permanently altered the daytime merit order, compressing the price reduction that nuclear baseload can deliver (Fabra and Reguant, 2014). Second, long-term LNG supply contracts signed during the shutdown era created sticky fuel costs that did not immediately adjust when reactors returned (Vivoda, 2019). Third, the post-Fukushima safety upgrades themselves added substantial capital costs to nuclear generation, raising the effective marginal cost of restarted reactors above their pre-Fukushima levels (Lipscy et al., 2013).

The mechanism test distinguishes these channels by decomposing the price effect into peak (8:00–20:00) and off-peak (20:00–8:00) hours. If solar compression were the dominant channel, the restart dividend should be larger at night, when solar generation is absent and nuclear displaces LNG without competition. Instead, I find nearly identical effects during peak ($-1.02$ ¥/kWh) and off-peak ($-0.91$ ¥/kWh) hours, suggesting that the asymmetry reflects broad merit-order restructuring rather than solar-specific displacement.

Heterogeneity across regions is consistent with a dosage interpretation. Kyushu, which restarted the most capacity (4.1GW across four reactors), experienced the largest price reduction ($-2.28$ ¥/kWh, $p < 0.05$). Kansai and Shikoku, with intermediate restarts, show moderate effects. Tohoku and Chugoku, which restarted single reactors only in late 2024, exhibit small and imprecise effects—consistent with both limited dosage and limited post-treatment observation time.

This paper contributes to three literatures. First, it provides the first causal estimate of nuclear restart effects on electricity prices, complementing the shutdown-side evidence in Neidell et al. (2021) and the broader literature on nuclear economics (Davis and Hausman, 2016; Jarvis et al., 2022). The restart deficit documents a new form of energy market hysteresis: once a power system adapts to life without nuclear, the price benefits of returning it are permanently diminished.

Second, the paper contributes to the literature on merit-order effects and renewable energy integration (Cludius et al., 2014; Ketterer, 2014; Paraschiv et al., 2014). The finding that solar deployment during the hiatus permanently compressed the nuclear price dividend provides direct evidence that renewable capacity changes the counterfactual against which any baseload technology is evaluated—a theoretical prediction that has been difficult to test empirically because renewable deployment is rarely exogenous to electricity market conditions.

Third, the paper advances methodological practice for energy market DiD designs with few clusters. With only nine electricity regions, conventional cluster-robust inference is unreliable. I implement exact randomization inference over all $\binom{9}{5} = 126$ permutations of treatment assignment and wild cluster bootstrap with the Webb 6-point distribution (Webb, 2023), demonstrating that the main result survives both.

The remainder of the paper proceeds as follows. Section 2 describes Japan's nuclear restart process and electricity market structure. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results, and Section 6 discusses implications.

## 2. Institutional Background

**The Fukushima shutdown and its aftermath.** The March 2011 Tōhoku earthquake and tsunami triggered a triple meltdown at the Fukushima Daiichi nuclear power station. In the aftermath, public opposition and political pressure led to the progressive shutdown of all 54 of Japan's commercial nuclear reactors by September 2013. Before Fukushima, nuclear supplied approximately 30 percent of Japan's electricity; after the shutdown, this share fell to zero.

The generation gap was filled primarily by increased fossil fuel imports, particularly liquefied natural gas (LNG). Japan's LNG imports rose from 70 million tonnes in 2010 to over 87 million tonnes in 2014, making it the world's largest LNG importer (Vivoda, 2019). The cost was borne by electricity consumers through higher wholesale prices and by the Japanese government through fuel subsidies and utility bailouts.

**The NRA restart process.** In September 2012, the government established the Nuclear Regulation Authority (NRA) as an independent safety regulator. The NRA promulgated new safety standards in July 2013, requiring existing reactors to meet post-Fukushima requirements for seismic resistance, tsunami protection, and severe accident management before restart could be approved. The review process has been protracted: as of early 2025, only 14 of the original 54 reactors have received restart approval, spanning five of Japan's ten electricity regions.

The staggered nature of restart approvals reflects genuine heterogeneity in reactor-level safety compliance rather than demand-side considerations. Each reactor undergoes an independent NRA technical review, and approval timelines have ranged from two years (Sendai Unit 1, approved August 2015) to over a decade for reactors still under review. This variation—driven by engineering and geological factors rather than by regional electricity market conditions—provides the treatment timing variation that my identification strategy requires.

**Japan's segmented electricity grid.** Japan's electricity system is divided into ten regional service areas, each historically served by a vertically integrated utility. Crucially, the grid operates at two different frequencies: 50Hz in the east (Hokkaido, Tohoku, Tokyo) and 60Hz in the west (Chubu, Hokuriku, Kansai, Chugoku, Shikoku, Kyushu). Only 1.2GW of frequency conversion capacity connects the two halves—less than 1 percent of national peak demand. This physical barrier sharply limits cross-regional electricity flows and creates the near-independence of regional price formation that my identification exploits.

**The JEPX spot market.** The Japan Electric Power Exchange (JEPX) operates a day-ahead spot auction in which generators bid supply and retailers bid demand for each half-hourly settlement period across all ten areas. Area-level clearing prices are determined by the intersection of aggregated supply and demand curves, subject to interconnection capacity constraints. Since its establishment in 2003, the JEPX has grown to handle approximately 60 percent of traded electricity volume, making spot prices increasingly reflective of marginal generation costs (Japan Electric Power Exchange, 2024).

## 3. Data

**JEPX spot prices.** The primary dataset is the complete record of JEPX day-ahead spot auction settlement prices, obtained from `japanesepower.org`. The data contain 280,176 half-hourly price observations across nine tradeable JEPX areas (Okinawa is excluded from JEPX trading) from April 2010 through March 2026. I aggregate these to 1,512 region-month observations for the main analysis panel, spanning January 2012 to December 2025. The pre-2012 period is excluded because of structural market changes during the immediate post-Fukushima transition.

For each region-month, I compute the mean, median, 10th percentile, and 90th percentile of half-hourly prices, as well as separate means for peak (8:00–20:00) and off-peak (20:00–8:00) hours. All prices are in nominal ¥/kWh.

**Reactor restart timeline.** I compile NRA restart approval and commercial operation dates from public records maintained by the Japan Atomic Industrial Forum and the World Nuclear Association. The timeline covers 14 reactors across six generating stations in five regions, with first-restart dates ranging from August 2015 (Sendai Unit 1 in Kyushu) to December 2024 (Shimane Unit 2 in Chugoku).

**Summary statistics.** Table 1 presents summary statistics by treatment group and period. Before any restart (January 2012 to July 2015), treated and control regions have nearly identical price distributions: mean prices of 14.95 and 14.94 ¥/kWh, respectively, with standard deviations of 2.08 and 1.95 ¥/kWh. This pre-treatment balance supports the parallel trends assumption. After the onset of restarts, treated regions exhibit lower average prices (11.34 vs. 12.14 ¥/kWh), consistent with the expected direction of the treatment effect.

**Table 1:** Summary Statistics: Monthly Wholesale Electricity Prices

| | Pre-Restart (2012–2015) | | Post-Restart (2015–2025) | |
|---|---|---|---|---|
| | Treated | Control | Treated | Control |
| Mean price (¥/kWh) | 14.98 | 14.89 | 10.86 | 12.04 |
| SD price | 2.01 | 1.95 | 6.38 | 6.83 |
| 10th percentile | 12.28 | 11.96 | 5.96 | 7.31 |
| 90th percentile | 18.54 | 18.65 | 16.47 | 17.93 |
| Region-months | 215 | 172 | 625 | 500 |
| Regions | 5 | 4 | 5 | 4 |

*Notes:* Treated regions are those with at least one NRA-approved nuclear reactor restart (Kyushu, Kansai, Shikoku, Tohoku, Chugoku). Control regions have no restarts through 2025 (Tokyo, Chubu, Hokuriku, Hokkaido). Prices are area-level JEPX day-ahead spot auction prices aggregated to monthly means from half-hourly settlement data. Pre-restart period ends July 2015 (month before first restart in Kyushu). All prices in nominal Yen per kilowatt-hour.

## 4. Empirical Strategy

### 4.1 Identification

The ideal experiment would randomly restart nuclear reactors across regions and compare price changes. The NRA approval process approximates this by generating plausibly exogenous variation in restart timing driven by reactor-specific safety characteristics rather than regional electricity demand.

I estimate the causal effect of nuclear restarts on wholesale electricity prices using a staggered difference-in-differences design. Define $Y_{rt}$ as the mean monthly spot price in region $r$ at time $t$, $G_r$ as the month of first restart in region $r$ (with $G_r = \infty$ for never-treated regions), and $D_{rt} = \mathbf{1}[t \geq G_r]$ as the post-treatment indicator. The identifying assumption is parallel trends in the absence of treatment:

$$[Y_{rt}(0) - Y_{r,t-1}(0) \mid G_r = g] = [Y_{rt}(0) - Y_{r,t-1}(0) \mid G_r = g'] \tag{1}$$

for all cohorts $g, g'$ and all pre-treatment periods $t < \min(g, g')$.

I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator with not-yet-treated units as the control group. This avoids the well-documented biases of two-way fixed effects (TWFE) estimation under treatment effect heterogeneity (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021). For comparison, I also report TWFE estimates and the Sun and Abraham (2021) interaction-weighted estimator.

### 4.2 Inference with Few Clusters

With only nine electricity regions, conventional cluster-robust standard errors may be severely size-distorted (Cameron et al., 2008). I address this through two complementary approaches.

First, I implement a wild cluster bootstrap using the Webb (2023) six-point distribution, which performs well with as few as six clusters. This involves 9,999 bootstrap replications, each drawing cluster-level random weights from $\{-\sqrt{3/2}, -\sqrt{2/2}, -\sqrt{1/2}, \sqrt{1/2}, \sqrt{2/2}, \sqrt{3/2}\}$.

Second, I conduct exact randomization inference (RI) over all $\binom{9}{5} = 126$ possible assignments of 5 treated regions among 9 total regions. For each permutation, I reassign the observed treatment timing to the permuted regions and re-estimate the TWFE coefficient. The RI $p$-value is the share of permuted coefficients at least as extreme as the observed estimate.

### 4.3 Threats to Validity

**Parallel trends.** The pre-treatment period (2012–2015) provides 43 monthly observations to assess parallel trends. Table 1 shows that treated and control regions had nearly identical price levels and dispersion before any restart (14.95 vs. 14.94 ¥/kWh, with standard deviations of 2.08 and 1.95 ¥/kWh). The Callaway–Sant'Anna event-study coefficients for pre-treatment periods are centered on zero, though individual group-time estimates are noisy given the small number of regions. Because this is a V1 paper with no figures, visual event-study evidence is deferred to a future revision; the interested reader may replicate the dynamic aggregation from the replication code.

**Concurrent policies.** Japan's electricity market underwent several reforms during the study period, including retail liberalization (April 2016) and legal unbundling of transmission (April 2020). These reforms affected all regions simultaneously and are absorbed by time fixed effects. The feed-in tariff for solar, while varying over time, applied nationally and is similarly controlled.

**Spillovers.** The 50Hz/60Hz frequency barrier limits cross-regional spillovers to 1.2GW of conversion capacity. Within each frequency zone, interconnection is larger but still constrained by transmission capacity. To the extent that restarts in western regions reduce prices in eastern regions through limited interconnection, my estimates are biased toward zero (conservative).

**Anticipation.** Restarts require extensive regulatory approval processes, creating potential for anticipation effects. I address this with a donut specification excluding the three months before and after each restart, which yields a larger point estimate (−1.09 ¥/kWh), suggesting that any anticipation biases the main estimate toward zero.

## 5. Results

### 5.1 Main Estimates

Table 2 reports the main results. The TWFE estimate (column 1) indicates that nuclear restarts reduce mean monthly wholesale prices by 0.97 ¥/kWh, significant at the 10 percent level with conventional cluster-robust standard errors ($p = 0.050$). The wild cluster bootstrap $p$-value is 0.040, and the exact RI $p$-value is 0.095.

The heterogeneity-robust Callaway–Sant'Anna estimator (column 2) yields a smaller but more precisely estimated ATT of $-0.60$ ¥/kWh ($p < 0.05$), reflecting appropriate reweighting across cohorts. The Sun–Abraham estimator (column 3) gives $-0.72$ ¥/kWh ($p < 0.001$). All three estimators agree in sign and approximate magnitude.

In proportional terms (column 4), the log-price TWFE estimate implies an 8.2 percent reduction, while the CS log estimate implies a 4.2 percent reduction. Against a pre-treatment mean of 14.94 ¥/kWh, the preferred CS estimate of $-0.60$ ¥/kWh represents a 4.0 percent price decline.

**Context: the restart deficit.** Neidell et al. (2021) estimate that the post-Fukushima shutdown raised Japanese electricity prices by approximately 38 percent. While my preferred estimate of the restart dividend ($-0.60$ ¥/kWh, or 4 percent) is substantially smaller, direct comparison requires caution: the studies differ in methodology, time periods, and the completeness of the nuclear capacity change. The shutdown eliminated all nuclear generation simultaneously; the restarts have been partial and staggered over a decade. Nevertheless, even the largest region-specific estimate (Kyushu at $-2.28$ ¥/kWh) suggests that the price dividend from restoring nuclear capacity is muted relative to the original shock. I call this gap the "restart deficit"—a label for the observation that a power system that adapted to life without nuclear does not simply revert when nuclear returns.

### 5.2 Mechanism: Peak versus Off-Peak

If solar compression were the primary driver of the restart deficit, we would expect larger price reductions during off-peak hours, when solar generation is absent and nuclear directly displaces LNG in the merit order. Columns 5 and 6 of Table 2 test this prediction. The peak-hour effect ($-1.02$ ¥/kWh) is slightly larger than the off-peak effect ($-0.91$ ¥/kWh), though the difference is not statistically significant. This pattern suggests that the asymmetry

**Table 2:** Effect of Nuclear Restarts on Wholesale Electricity Prices

| | Mean Price (¥/kWh) | | | Log | Time of Day | |
|---|---|---|---|---|---|---|
| | TWFE | CS | SA | TWFE | Peak | Off-Peak |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Post-Restart | -0.967* | -0.597** | -0.723*** | -0.082** | -1.022** | -0.913* |
| | (0.420) | (0.288) | (0.125) | (0.035) | (0.442) | (0.414) |
| WCB $p$-value | 0.040 | | | | | |
| RI $p$-value | 0.095 | | | | | |
| Region FE | Yes | — | Yes | Yes | Yes | Yes |
| Month FE | Yes | — | Yes | Yes | Yes | Yes |
| Control group | All | NYT | All | All | All | All |
| Pre-treatment mean | 14.94 | 14.94 | 14.94 | 2.69 | | |
| Observations | 1512 | 1512 | 1512 | 1512 | 1512 | 1512 |
| Regions | 9 | 9 | 9 | 9 | 9 | 9 |

*Notes:* Each column reports the average treatment effect of nuclear reactor restarts on wholesale electricity prices. TWFE = two-way fixed effects; CS = Callaway and Sant'Anna (2021) with not-yet-treated control group; SA = Sun and Abraham (2021). Standard errors clustered at the region level in parentheses. WCB = wild cluster bootstrap $p$-value using the Webb 6-point distribution (9,999 replications). RI = exact randomization inference $p$-value over all 126 permutations of 5 treated among 9 regions. Peak hours are 8:00–20:00; off-peak hours are 20:00–8:00. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

reflects broad merit-order restructuring—including LNG contract stickiness and capital cost inflation—rather than solar-specific displacement alone.

### 5.3 Heterogeneity by Region

Table 3 reports region-specific treatment effects. The gradient across regions is consistent with a dosage interpretation: Kyushu, which restarted the most capacity (4.1GW), experienced the largest and most precisely estimated price reduction ($-2.28$ ¥/kWh, $p < 0.05$). Shikoku, with a single 890MW reactor, shows a significant $-1.19$ ¥/kWh effect ($p < 0.05$). Kansai, despite having the largest total capacity (6.6GW), shows a moderate effect ($-1.03$ ¥/kWh) that is imprecisely estimated—possibly because its restarts were spread over nine years, diluting the sharp treatment onset that benefits DiD estimation.

The continuous dosage specification (Table 3, column 5) confirms that each additional GW of restarted capacity reduces wholesale prices by 0.26 ¥/kWh, consistent with the region-level pattern.

**Table 3:** Heterogeneity by Region

| Region | First Restart | Capacity (GW) | $\hat{\beta}$ | SE |
|---|---|---|---|---|
| Kyushu | Aug 2015 | 4.1 | -2.277** | (0.645) |
| Kansai | Jan 2016 | 6.6 | -1.034 | (0.545) |
| Shikoku | Aug 2016 | 0.9 | -1.187** | (0.421) |
| Tohoku | Oct 2024 | 0.8 | -0.094 | (0.326) |
| Chugoku | Dec 2024 | 0.8 | -0.406 | (0.329) |

*Notes:* Each row reports the TWFE estimate of the region-specific treatment effect, estimated as the coefficient on the interaction of region and post-restart indicators. Control group is all four never-treated regions. Standard errors clustered at the region level. Capacity is total restarted nuclear capacity through 2025 in gigawatts. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

### 5.4 Robustness

Table 4 presents five robustness checks. Weekly aggregation (column 1) preserves the main result at −0.98 ¥/kWh, ruling out monthly averaging artifacts. The donut specification excluding transition months (column 2) yields a larger effect of −1.09 ¥/kWh, suggesting that the baseline estimate is, if anything, attenuated by gradual ramp-up. Using median rather than mean prices (column 3) produces −0.85 ¥/kWh, demonstrating that the result is not driven by price spikes in the tails.

Column 4 tests whether restarts affect price *volatility*. The near-zero and insignificant coefficient (0.07 ¥/kWh) indicates that nuclear restarts shift the price distribution without changing its dispersion—consistent with a pure merit-order displacement effect.

## 6. Discussion

The restart deficit has implications for energy policy beyond Japan. Countries contemplating nuclear exits or restarts—including Germany, which shut its last reactor in April 2023, and several Eastern European nations considering new builds—should not assume symmetric price effects. Once a power system adapts to the absence of nuclear through renewable deployment and fossil fuel contracting, the economic case for restart is weaker than the symmetric counterfactual suggests.

The finding also speaks to the broader question of path dependence in energy systems (Arthur, 1989; Unruh, 2000). Energy transitions are typically modeled as reversible movements along a merit-order curve. The restart deficit suggests instead that the curve itself shifts during transitions, creating hysteresis that locks in higher costs even after the original shock is reversed. This has implications for the evaluation of any large-scale energy disruption,

**Table 4:** Robustness Checks

| | Weekly (1) | Donut (2) | Median (3) | Volatility (4) | Dosage (5) |
|---|---|---|---|---|---|
| Post-Restart | -0.979** | -1.090** | -0.850* | 0.070 | |
| | (0.411) | (0.430) | (0.363) | (0.107) | |
| Cumulative GW | | | | | -0.263 |
| | | | | | (0.176) |
| Region FE | Yes | Yes | Yes | Yes | Yes |
| Time FE | Yes | Yes | Yes | Yes | Yes |
| Aggregation | Weekly | Monthly | Monthly | Monthly | Monthly |
| Transition excluded | No | ±3 months | No | No | No |

*Notes:* All specifications include region and time fixed effects with standard errors clustered at the region level. Column (1) aggregates half-hourly prices to weekly means. Column (2) drops the three months before and after each region's first restart. Column (3) uses the median rather than mean monthly price. Column (4) estimates the effect on within-month price standard deviation. Column (5) replaces the binary post-restart indicator with cumulative restarted nuclear capacity in gigawatts. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

from coal phase-outs to renewable intermittency events.

Several limitations deserve emphasis. First, the analysis rests on only nine regional clusters, of which five are treated. While wild cluster bootstrap and exact randomization inference provide more reliable inference than conventional methods in this setting, the design is inherently underpowered for detecting small effects. Second, I focus exclusively on wholesale price effects; consumer prices in Japan are regulated and adjust with lags, so the welfare implications for households may differ. Third, the mechanism test (peak vs. off-peak decomposition) provides only indirect evidence on the relative importance of solar compression, LNG contract stickiness, and higher nuclear capital costs—future work with generation-mix data could disentangle these channels more directly. Fourth, the restart deficit should not be interpreted as showing that nuclear restarts are "ineffective"—a 4 percent wholesale price reduction for treated regions is economically meaningful. Rather, the finding highlights that energy systems adapt in ways that attenuate the reversibility of large supply shocks.

## 7. Conclusion

Bringing nuclear back online provides measurable price relief, but substantially less than the symmetric counterfactual would predict. This paper documents a "restart deficit"—a modest wholesale price dividend ($-0.60$ ¥/kWh, or 4 percent) from restoring nuclear capacity to a system that adapted during a decade-long hiatus. The finding does not imply that nuclear is uneconomic; rather, it reveals that energy transitions are not easily reversible. Solar deployment and fossil fuel contracting during the nuclear absence reshaped the merit order against which any returning baseload technology is evaluated. For policymakers contemplating nuclear exits or restarts, the lesson is that the counterfactual evolves: the energy system a reactor returns to is not the one it left.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

**Project Repository:** https://github.com/SocialCatalystLab/ape-papers

**Contributors:** @olafdrw

**First Contributor:** https://github.com/olafdrw

## References


**Arthur, W. Brian**, "Competing Technologies, Increasing Returns, and Lock-In by Historical Events," *Economic Journal*, 1989, *99* (394), 116–131.

**Callaway, Brantly and Pedro H. C. Sant'Anna**, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

**Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

**Cludius, Johanna, Hauke Hermann, Felix Chr. Matthes, and Verena Graichen**, "The Merit Order Effect of Wind and Photovoltaic Electricity Generation in Germany 2008–2016," *Energy Policy*, 2014, *44*, 302–313.

**Davis, Lucas W. and Catherine Hausman**, "Market Impacts of a Nuclear Power Plant Closure," *American Economic Journal: Applied Economics*, 2016, *8* (2), 92–122.

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

**Fabra, Natalia and Mar Reguant**, "Pass-Through of Emissions Costs in Electricity Markets," *American Economic Review*, 2014, *104* (9), 2872–2899.

**Goodman-Bacon, Andrew**, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

**International Energy Agency**, "Japan 2023: Energy Policy Review," Country Report, IEA, Paris 2023.

**Japan Electric Power Exchange**, "Annual Report on JEPX Trading Volume and Price Statistics," Technical Report, JEPX, Tokyo 2024.

**Jarvis, Stephen, Olivier Deschenes, and Akshaya Jha**, "The Private and External Costs of Germany's Nuclear Phase-Out," *Journal of the European Economic Association*, 2022, *20* (3), 1311–1346.

**Ketterer, Janina C.**, "The Impact of Wind Power Generation on the Electricity Price in Germany," *Energy Economics*, 2014, *44*, 270–280.

**Lipscy, Phillip Y., Kenji E. Kushida, and Trevor Incerti**, "The Fukushima Disaster and Japan's Nuclear Plant Vulnerability in Comparative Perspective," *Environmental Science and Technology*, 2013, *47* (12), 6082–6088.

**Neidell, Matthew, Shinsuke Uchida, and Marcella Veronesi**, "The Unintended Effects from Halting Nuclear Power Production: Evidence from Fukushima Daiichi Accident," *Journal of Health Economics*, 2021, *79*, 102507.

**Paraschiv, Florentina, David Erni, and Ralf Pietsch**, "The Impact of Renewable Energies on EEX Day-Ahead Electricity Prices," *Energy Policy*, 2014, *73*, 196–210.

**Sun, Liyang and Sarah Abraham**, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

**Unruh, Gregory C.**, "Understanding Carbon Lock-In," *Energy Policy*, 2000, *28* (12), 817–830.

**Vivoda, Vlado**, "LNG Import Diversification and Energy Security in Asia," *Energy Policy*, 2019, *129*, 967–974.

**Webb, Matthew D.**, "Reworking Wild Bootstrap Based Inference for Clustered Errors," *Canadian Journal of Economics*, 2023, *56* (3), 839–858.

## A. Data Appendix

**JEPX spot prices.** The Japan Electric Power Exchange publishes half-hourly settlement prices for nine tradeable areas (Hokkaido, Tohoku, Tokyo, Chubu, Hokuriku, Kansai, Chugoku, Shikoku, Kyushu). Okinawa is excluded from JEPX trading due to its isolated island grid. Data are obtained from `japanesepower.org/jepxSpot.csv`, which mirrors the official JEPX publication. The dataset contains 280,176 half-hourly observations from April 1, 2010 through March 24, 2026.

**Reactor restart timeline.** Restart dates are compiled from the Japan Atomic Industrial Forum Status Report, the World Nuclear Association Country Profile for Japan, and NRA press releases. I record the date of commercial operation resumption (not the date of NRA safety review approval, which precedes restart by several months). The 14 restarted reactors, their regions, restart dates, and capacities are:

- **Kyushu:** Sendai 1 (890MW, Aug 2015), Sendai 2 (890MW, Oct 2015), Genkai 3 (1,180MW, Mar 2018), Genkai 4 (1,180MW, Jun 2018)
- **Kansai:** Takahama 3 (870MW, Jan 2016), Takahama 4 (870MW, Feb 2016), Ohi 3 (1,180MW, Mar 2018), Ohi 4 (1,180MW, May 2018), Mihama 3 (826MW, Jun 2021), Takahama 1 (826MW, Jul 2023), Takahama 2 (826MW, Sep 2023)
- **Shikoku:** Ikata 3 (890MW, Aug 2016)
- **Tohoku:** Onagawa 2 (825MW, Oct 2024)
- **Chugoku:** Shimane 2 (820MW, Dec 2024)

**Sample construction.** The analysis panel consists of 9 regions $\times$ 168 months (January 2012–December 2025) = 1,512 region-month observations. I exclude the April 2010–December 2011 period because it spans the Fukushima disaster itself and the subsequent emergency price interventions. The 50Hz/60Hz grid division assigns Hokkaido, Tohoku, and Tokyo to the eastern (50Hz) grid, and the remaining six regions to the western (60Hz) grid.

## B. Identification Appendix

**Pre-treatment balance.** Pre-treatment (2012–2015) mean prices are nearly identical across treated (14.95 ¥/kWh) and control (14.94 ¥/kWh) regions, with comparable standard deviations (2.08 vs. 1.95 ¥/kWh). The Callaway–Sant'Anna event-study coefficients for pre-treatment periods are centered on zero, though individual group-time estimates are noisy given the small number of regions.

**Donut specification.** Excluding the three months before and after each region's first restart yields −1.09 ¥/kWh (SE 0.43), larger than the baseline, suggesting that inclusion of the ramp-up period attenuates the main estimate.

**Exact randomization inference.** With 5 treated among 9 regions, there are $\binom{9}{5} = 126$ unique treatment assignments. I compute the TWFE coefficient under each permutation (assigning the observed treatment timing to the permuted regions). The observed coefficient of −0.97 lies in the lower 9.5th percentile of the permutation distribution, with a two-sided RI $p$-value of 0.095.

## C. Standardized Effect Sizes

**Table 5:** Standardized Effect Sizes

| Outcome | $\hat{\beta}$ | SE | SD($Y$) | SDE | SE(SDE) | Classification |
|---|---|---|---|---|---|---|
| Mean price | -0.597 | 0.288 | 1.981 | -0.301 | 0.145 | Large negative |
| Log price | -0.042 | 0.025 | 0.137 | -0.309 | 0.183 | Large negative |
| Peak price | -1.022 | 0.442 | 2.172 | -0.470 | 0.204 | Large negative |
| Off-peak price | -0.913 | 0.414 | 2.082 | -0.438 | 0.199 | Large negative |

*Notes:* **Country:** Japan. **Research question:** Do post-Fukushima nuclear reactor restarts, approved through Japan's Nuclear Regulation Authority safety review process, reduce wholesale electricity prices in the affected regional grid? **Policy mechanism:** NRA-approved reactor restarts add low-marginal-cost nuclear baseload generation to regional electricity grids, potentially displacing high-cost LNG-fired generation from the merit order and reducing the market-clearing price in the Japan Electric Power Exchange day-ahead auction. **Outcome definition:** Monthly mean (or log, peak, off-peak) of half-hourly JEPX day-ahead spot auction settlement prices in Yen per kilowatt-hour at the regional area level. **Treatment:** Binary indicator for whether any reactor in a region has restarted (first restart ranges from August 2015 to December 2024 across five treated regions). **Data:** JEPX day-ahead spot prices from japanesepower.org, 9 regions, January 2012 to December 2025, aggregated to 1,512 region-months from 2.5 million half-hourly observations. **Method:** Callaway and Sant'Anna (2021) staggered DiD with not-yet-treated control group; standard errors via multiplier bootstrap. **Sample:** Nine JEPX electricity regions (excluding Okinawa); five treated regions with NRA-approved restarts and four never-treated control regions; Japan's 50Hz/60Hz frequency split limits cross-regional price spillovers. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD($Y$) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null ($< 0.005$).