HSR and Regional Economies: Evidence from the Shinkansen in Japan

Andrew Ilisoi

Independent Researcher

August 2025

Abstract

This paper evaluates the long-run economic effects of Shinkansen high-speed rail (HSR) station openings across Japan. Using newly constructed panel datasets at the prefectural (1950–2024) and municipal (1980–2023) levels, I apply modern difference-in-differences estimators that address biases in two-way fixed effects under staggered adoption.

At the prefectural scale, Shinkansen access generates persistent population and migration inflows, modest increases in housing activity, and sharp commercial land value capitalization (≈12–13% within a decade). Unemployment counts also rise, reflecting transitional labor market frictions. By contrast, GDP, taxable income, and prefectural revenues remain flat, suggesting no broad productivity or fiscal expansion.

Municipal results reveal the mechanisms underlying these redistributional patterns. Host municipalities experience immediate and lasting increases in fiscal revenue (+5–9%) and taxable income (+3–5%), while municipalities within 10–30 km exhibit lagged housing construction booms (+18–21%) and medium-run increases in live births (+3–7%). Hospital capacity and municipal financial need also rise near stations, indicating greater service demands.

Taken together, the evidence supports a redistributional interpretation: the Shinkansen reallocates population, housing demand, and fiscal capacity toward connected municipalities, without generating aggregate growth in income or output. This result contrasts with large GDP gains documented for China and aligns more closely with European evidence of spatial reallocation. For policymakers, the findings imply that justifying high-speed rail investment requires emphasizing redistributional benefits—local revitalization, spatial equity, and cohesion—rather than aggregate productivity gains.

Keywords: High-speed rail, regional development, difference-in-differences, Japan, migration, land markets.

JEL Codes: R11, R12, O18, C23

1. Introduction

Since its debut, the Shinkansen has been Japan's most visible symbol of high-speed rail (HSR). Today the network connects nearly all major regional centers. Beyond reducing travel times, policymakers have long debated whether HSR drives regional growth or merely redistributes economic activity. This paper provides new empirical evidence on the long-run effects of Shinkansen station openings on local and regional economies.

The central research question is whether access to the Shinkansen stimulates population and income growth, or instead reshuffles people and resources without expanding the overall economic pie. Research on HSR effects in China shows large GDP effects; Japan evidence is more muted. Early work emphasized tourism and short-run commerce, while more recent studies point to limited productivity gains. This ambiguity makes Japan a particularly valuable case: it was the world's first HSR system, has the longest operating history, and exhibits staggered expansion across both prefectures and municipalities.

The empirical challenge is that station openings occurred at different times across regions, and treatment effects may vary across cohorts and over time. I employ modern DiD estimators that address biases in TWFE under staggered adoption. The primary specification uses the Callaway & Sant'Anna (2021) group-time estimator, which recovers cohort-specific dynamic treatment effects without negative weighting. Robustness checks include the Sun & Abraham (2021) interaction-weighted TWFE correction and a matched-sample C&S DiD that balances units on baseline levels and pre-treatment outcome slopes. This design ensures that results are not artifacts of estimator choice or non-parallel counterfactual trends.

The analysis combines newly constructed prefectural and municipal panel datasets covering 1950–2024 and 1980–2023, respectively, with outcomes spanning demographics, housing, land markets, fiscal accounts, and social indicators. At the prefectural level, Shinkansen access is associated with persistent population and migration inflows, rising land values, and transitional unemployment, while GDP and income remain flat. At the municipal level, effects are sharper and more localized, including fiscal gains in host cities and housing/demographic adjustments in nearby areas

This paper contributes to the discussion by integrating prefectural and municipal evidence with modern DiD estimators, offering a unified view of how the Shinkansen reshaped regional and local economies.

This paper proceeds as follows:

- Section 2 reviews the institutional background of the Shinkansen's expansion.
- Section 3 describes the data construction.
- Section 4 presents the empirical strategy.
- Section 5 reports results for prefectures and municipalities.
- Section 6 discusses robustness and diagnostics.
- Section 7 interprets mechanisms and situates the findings in an international context.
- Section 8 concludes.
- References.
- Appendix

2. Background and Related Literature

2.1 Institutional Background

Japan inaugurated the Shinkansen in 1964 with the opening of the Tōkaidō line between Tokyo and Osaka. Over the following six decades, the network expanded in stages, linking most major regional centers across Honshū, Kyūshū, and Hokkaidō. By 2022, the Shinkansen exceeded 3,000 kilometers of track, making it one of the most extensive high-speed rail systems in the world.

The staggered expansion of the Shinkansen generates substantial variation in the timing and geography of station openings. New lines typically opened in 100–300 kilometer segments, with stations distributed across both major urban centers and smaller regional municipalities. This sequencing provides the empirical leverage for this paper: different prefectures and municipalities received treatment at different times, allowing identification of dynamic effects in a difference-in-differences framework.

National planners promoted the Shinkansen not only as a transportation innovation but also as a regional development policy. Initially intended to reduce travel times, integrate regional markets, and mitigate the dominance of Tokyo and Osaka, the Shinkansen later became a vehicle for regional revitalization in declining or peripheral areas. This debate underpins the Shinkansen's policy rationale.

2.2 High-Speed Rail in Comparative Perspective

International research on high-speed rail offers mixed evidence. In China, Banerjee, Duflo, and Qian (2020) document large positive GDP effects from HSR, while Faber (2014) shows similar productivity gains from China's highway expansion. In contrast, evidence from Europe points to more modest or redistributional impacts: Ahlfeldt and Feddersen (2010) find market access gains around German transit stations; Chen and de Abreu e Silva (2014) link Spain's AVE to employment and education outcomes; Bouf and Desmaris (2020) and Di Ruocco (2024) highlight redistributional effects in France; and Blanquart (2017) reviews theoretical mechanisms emphasizing reallocation rather than net growth. Cheng (2022) provides a comprehensive survey underscoring the heterogeneity of HSR effects across contexts.

2.3 The Shinkansen in the Literature

Compared to China and Europe, the Japanese experience remains less studied. Early work emphasized tourism and short-run commerce effects. More recent contributions have documented population redistribution but limited productivity gains: Nickelsburg, Ahluwalia, and Yang (2018) analyze housing affordability pressures, while Wang et al. (2024) show demographic relocations toward HSR-connected areas. These studies suggest that Japan's Shinkansen may have functioned more as a redistributive force than as a productivity engine.

2.4 Contribution of This Paper

Building on this literature, my contribution is twofold. Substantively, I extend beyond prior Japan studies (Nickelsburg et al. 2018; Wang et al. 2024) by integrating both prefecture-level and municipal-level evidence into a single framework. Methodologically, I apply modern DiD estimators (Callaway & Sant'Anna 2021; Sun & Abraham 2021) that correct for biases in traditional two-way fixed effects under staggered adoption. This combination allows me to revisit the longstanding question of whether the Shinkansen generated regional growth or merely redistributed population and land values.

3. Data

I construct two panel datasets at the prefectural and municipal levels to examine the effects of Shinkansen station openings. The prefectural panel covers 47 prefectures from 1950 to 2024, while the municipal panel spans 1980 to 2023 and incorporates all municipalities exposed to Shinkansen stations. Together, these datasets provide consistent coverage of demographic, economic, land, fiscal, and social outcomes across multiple geographic scales.

3.1 Sources

The majority of variables were obtained through API queries from *e-Stat*, Japan's official government statistics portal, which compiles census, finance, land, labor, and social indicators from national ministries. Additional measures were drawn from the Statistics Bureau and Japanese Statistical Yearbooks.

3.2 Treatment Definition

Treatment years are defined as the opening year of the **first Shinkansen station** within a prefecture or municipality. For municipalities without stations, indirect exposure is assigned if the jurisdiction lies within 10, 20, or 30 kilometers of a station site. This staggered rollout provides the temporal and spatial variation needed for difference-in-differences analysis. (Formal specification is presented in Section 4.)

3.3 Outcomes

The outcomes fall into five domains:

1. **Demographics**: total population, in-migration, out-migration, and net migration.
2. **Economic aggregates**: GDP, taxable income, labor force participation, unemployment counts.
3. **Land and housing markets**: commercial, residential, and industrial land values; new housing starts and residential construction.
4. **Fiscal accounts**: prefectural revenues, municipal financial revenue and need, municipal tax revenues.
5. **Social indicators**: live births, kindergarten enrollment, hospital capacity.

3.4 Data Handling

Population series were linearly interpolated to fill intercensal years, and variables were transformed into logarithms; where zeros occur, log(1+x) transformations were applied. Summary statistics and descriptive plots are provided in Appendix B. Detailed notes on dataset harmonization, including municipality mergers, historical yearbook digitization, and naming crosswalks, are also reported in Appendix B.

To provide a sense of the data, Table 1 reports data coverage statistics for the main outcomes measured in the prefectural and municipal panels. Prefectural coverage spans long-run demographics, GDP, migration, land markets, and unemployment, while municipal coverage captures fiscal accounts, housing activity, demographic indicators, and service provision. The table highlights the breadth of economic and social outcomes observed across both levels of analysis.

Outcome	Level	Years Covered
Population (log)	Prefecture; Municipality	1950–2023; 1980-2020
Taxable Income (log)	Prefecture; Municipality	1985–2023; 1985–2023
GDP (log, real)	Prefecture	2001–2020
Labor Force (log)	Prefecture; Municipality	1975–2020; 1980–2020
Unemployed (log)	Prefecture; Muncipality	1975–2020; 1980–2020
In-migrants (log)	Prefecture; Municipality	1975–2024; 1996–2023
Net Migration (log)	Prefecture	1975–2024
Commercial Land Value (log)	Prefecture	1976–2023
New Residential Buildings (log)	Prefecture	1997–2023
Newly Constructed Dwellings (log)	Prefecture; Municipality	1975–2023; 2000–2023
Financial Revenue (log)	Prefecture; Municipality	1975–2021; 2000–2021
Financial Need (log)	Prefecture; Municipality	1975–2021; 2000–2021
Tax Revenue (log)	Prefecture; Municipality	1975–2021; 2000–2021
Live Births (log)	Prefecture; Municipality	1980–2022; 1995–2022
Hospitals (log1p)	Prefecture; Municipality	1975–2022; 1997–2022
Kindergarten Pupils (log)	Municipality	1980–2023
Kindergartens (log1p)	Municipality	1980–2023

Table 1. Description of Prefecture- and Municipal-Level Panel Data

4. Empirical Strategy

The identification challenge in this setting is that Shinkansen stations opened at different times across regions, and treatment effects may vary by cohort and evolve dynamically. A naïve two-way fixed effects (TWFE) regression would pool all units and periods, but this approach is now known to produce biased estimates when adoption is staggered and effects are heterogeneous (Goodman-Bacon 2021). To avoid these pitfalls, we rely on estimators explicitly designed for staggered adoption.

4.1 Baseline: Callaway & Sant'Anna (2021)

Our primary estimator is the group-time average treatment effects (ATT) framework of Callaway & Sant'Anna (2021), implemented via the att_gt() and aggte() functions in the **did** package for R. This design recovers dynamic treatment effects for each cohort g at each event time e = t − g, and aggregates them across cohorts to form event-study profiles.

Formally,

$$Y_{it} = \alpha_i + \lambda_t + \sum_g ATT(g,t) D_{it}^{g} + \varepsilon_{it},$$

where Y_{it} is the outcome for unit i in year t, α_i and λ_t are unit and year fixed effects, and D_{it}^{g} is an indicator equal to one if unit i belongs to cohort g and t ≥ g.

- **Control group:** not-yet-treated units, ensuring already-treated units are never used as counterfactuals.
- **Estimator:** doubly robust (DR) influence-function estimator.
- **Inference:** wild bootstrap with 999 replications (bstrap=TRUE, biters=999), clustered at the prefecture level for prefecture data and at the unit level for municipalities.
- **Outputs:** dynamic ATTs are aggregated into short-run (0–5 years) and medium-run (6–10 years) horizon averages using aggte(type="dynamic").

4.2 Robustness: Sun & Abraham (2021)

As a complementary design, we implement the interaction-weighted estimator of Sun & Abraham (2021) using the sunab() syntax in the **fixest** package for R. This estimator corrects the negative weighting problem in TWFE by estimating cohort-specific event-time interactions and then reweighting them consistently.

Implementation:

$$Y_{it} = \alpha_i + \lambda_t + \sum_e \beta_e \cdot 1\{t - G_i = e\} + \varepsilon_{it},$$

with feols(y ~ sunab(g, Year) | unit_id + Year, data) estimated by OLS. Cluster-robust standard errors are computed at the prefecture or municipality level, and dynamic profiles are aggregated into horizon averages (0–5 and 6–10 years).

4.3 Robustness: Matched-Sample DiD

As a further robustness check, we implement a matched-sample DiD. Units are first matched on baseline outcome levels and pre-treatment outcome slopes using nearest-neighbor Mahalanobis distance. We then re-estimate dynamic ATTs within the matched sample using the Callaway & Sant'Anna (2021) estimator. This design ensures balance in both levels and pre-trends, reducing concerns about non-parallel counterfactual trajectories.

Dynamic treatment effects are then re-estimated within the matched sample using the Callaway & Sant'Anna estimator:

```
att_matched <- att_gt(
  yname        = outcome,
  tname        = "Year",
  idname       = "unit_id",
  gname        = "g",
  data         = df_matched,
  control_group = "notyettreated",
  est_method    = "dr",
  bstrap       = TRUE,
  biters       = 999,
  clustervars   = "unit_id"
)
```

This design ensures balance in both outcome levels and slopes, thereby reducing concerns that counterfactual trends differ systematically between treated and control units.

4.4 Prefecture Specification

At the prefectural level, treatment is defined by the year of the first Shinkansen station opening within a prefecture's borders. We estimate dynamic ATTs separately for three complementary samples:

(i) **one-station** prefectures, where only a single station was ever opened;
(ii) **mult-10** prefectures, where multiple stations opened but analysis is restricted to a 10-year event window; and
(iii) **no-overlap** prefectures, where prefectures exposed to closely spaced multiple openings are excluded.

These sample definitions provide alternative views of staggered treatment timing and mitigate concerns that overlapping station cohorts contaminate estimates.

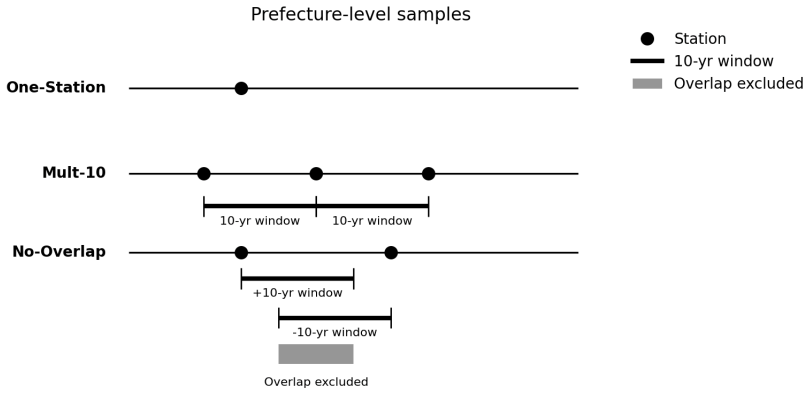


Figure D1-1: Prefecture-level samples

For a visualization of the sample data represented by a map of Japan and its prefectures, please see Figures E1-1, E1-2, and E1-3.

4.5 Municipal Specification

At the municipal level, host municipalities are coded as directly treated in their station-opening year. To capture spatial spillovers, we also define indirect treatment for municipalities located within 10, 20, and 30 kilometers of a new station. Dynamic ATTs are estimated separately for each exposure band, allowing us to trace the spatial extent of local impacts. This design captures both the immediate effects on host municipalities and the broader adjustments in nearby catchments.

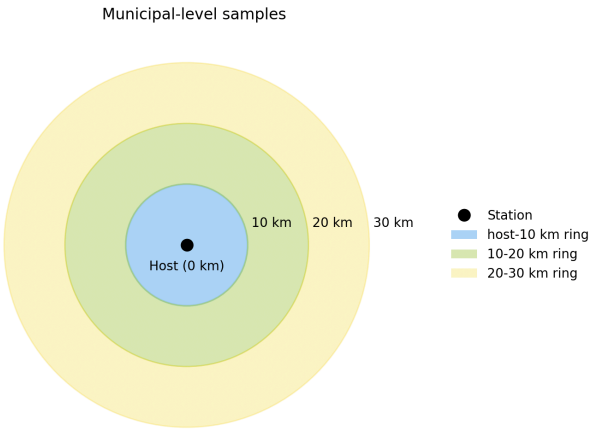


Figure D1-2: Municipal-level samples

For a visualization of the sample data represented by a map of Japan, its prefectures, and its municipalities, please see Figure E2-1.

4.7 Event-Study Implementation

Event-study plots report dynamic ATTs from ten years before to ten years after treatment. We adopt −1 as the reference period for our Sun-Abraham estimator, normalizing coefficients relative to the year immediately preceding station opening. For the DiD estimators (both done with the C&S model), no such reference period is used due. Pre-trend checks compare treated units to never-treated or not-yet-treated units with a varying base period. Horizon averages are computed for short-run (0–5 years) and medium-run (6–10 years) post-treatment windows, complementing the visual evidence with summary magnitudes.

4.8 Diagnostics and Pre-Trend Checks

A central identifying assumption is that, absent treatment, treated and control units would have followed parallel trends. We evaluate this assumption using several complementary checks:

1. **Event-time leads.** Dynamic ATTs are estimated up to ten years before treatment. We screen for significant pre-treatment coefficients (pointwise 95% confidence intervals). For the main outcomes highlighted in the results section, no pre-treatment leads are flagged.
2. **Visual inspection.** Event-study plots (Figures C1-1 to C2-23) display full dynamic paths. For population, migration, land values, and municipal fiscal revenue, pre-treatment trajectories are flat and parallel, lending visual credibility to the design.
3. **Matched-sample balance.** The matched-sample DiD robustness design explicitly balances treated and control units on both baseline levels and pre-treatment outcome slopes. This reduces concerns that residual differences in pre-trends drive the results.

5. Results

5.1 Prefecture-Level Effects

Population.
Shinkansen access is associated with robust and persistent population increases. In the one-station specification, the interpolated population series (log_Pop_lin) passes parallel-trend checks and shows statistically significant post-treatment gains. Callaway–Sant'Anna estimates (Table A1) indicate a short-run increase of about 4% (0–5 years) and a medium-run increase of 8.9% (6–10 years). The growth path is monotonic, with no evidence of pre-treatment divergence.

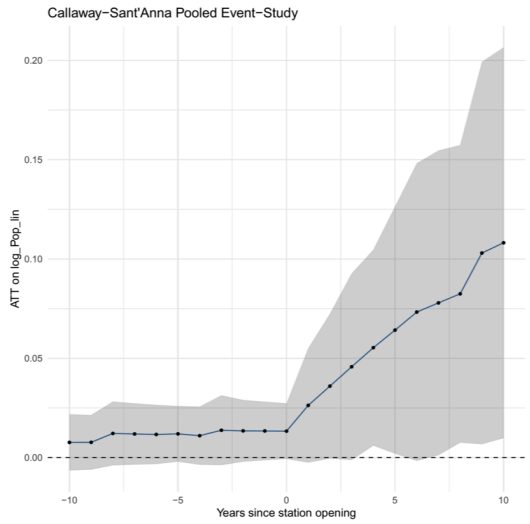


Figure C1-1: log Pop Lin - C&S DiD - 1 station sample

Migration.
Population gains appear partially driven by in-migration. In the Mult-10 and No-Overlap samples, in-migration rises by roughly 2.7% in the short run and 4.9% in the medium run, while out-migration remains flat. Consequently, net migration increases by about 2.5% over both horizons. Robustness checks using matched-sample DiD yield nearly identical effects (≈3% short run; 3–5% medium run).

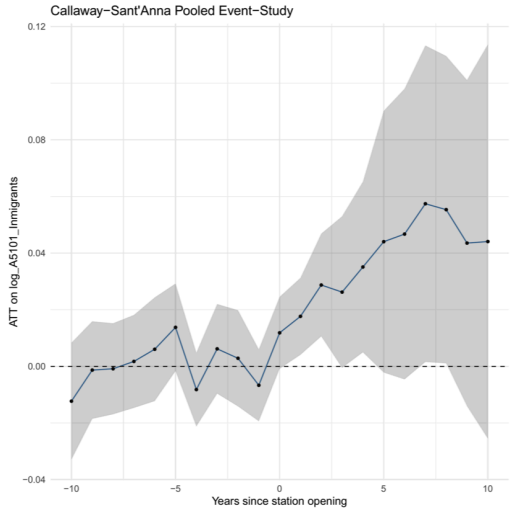


Figure C1-2: ln-migrants - C&S DiD - mult 10 sample

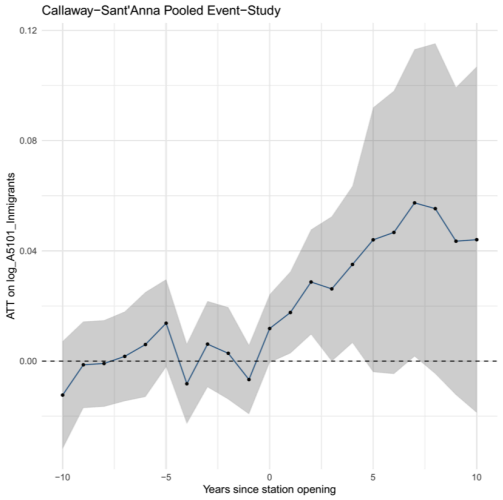


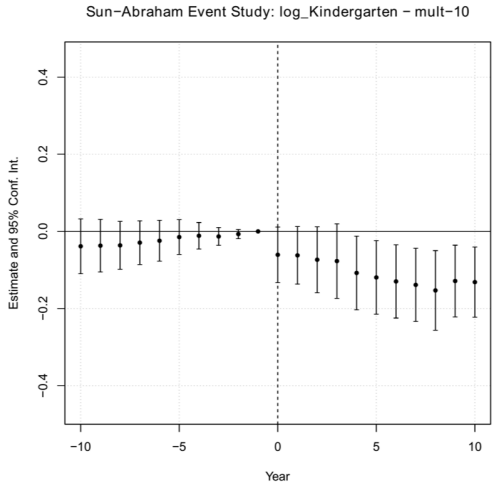
Figure *C1-7: In-migrants - C&S DiD - no overlap sample*

Housing and Land Markets.
Housing activity responds strongly to station openings. New residential building starts increase by ~2.2% in the short run and ~5.2% in the medium run under CS-DiD (Figures C1-6 and C1-11). These estimates are somewhat larger in matched-sample designs (~4–8%) (Figures C1-15 and C1-20) and under Sun–Abraham (~5–7%) (Figures C1-23 and C1-26). Land markets display the clearest capitalization effects: commercial land values rise by ~3% within five years and by 12–13% after a decade (Figures C1-3, C1-8, and C1-17). The magnitude and consistency across estimators (CS-DiD, matched, Sun–Abraham) make this one of the most robust findings.

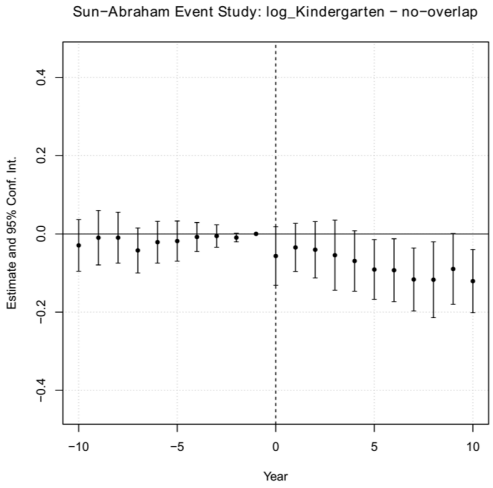
Labor Market.
Unemployment counts rise in treated prefectures, consistent with transitional frictions. CS-DiD estimates indicate increases of 3–4% in the short run and nearly 10% in the medium run. Matched-sample and Sun–Abraham robustness checks yield similar magnitudes, reinforcing the interpretation that inflows of job seekers and reallocation pressures, rather than net job creation, drive this result. Relevant figures are C1-4, C1-9, C1-13, C1-18, and C1-22.

Income and Fiscal Aggregates.
By contrast, GDP, taxable income, and prefectural revenues show no robust measurable responses. Sun–Abraham estimates of taxable income are flat at ≈1% in the short-term (Figure C1-25). Prefectural fiscal revenues and intergovernmental grants exhibit no systematic increases.

Social Indicators.
Some social outcomes adjust in ways consistent with consolidation. Kindergarten enrollment declines by 6–14% after treatment, particularly in the medium run, as populations concentrate around station areas.



C1-21: Kindergartens - Sunab - mult 10 sample



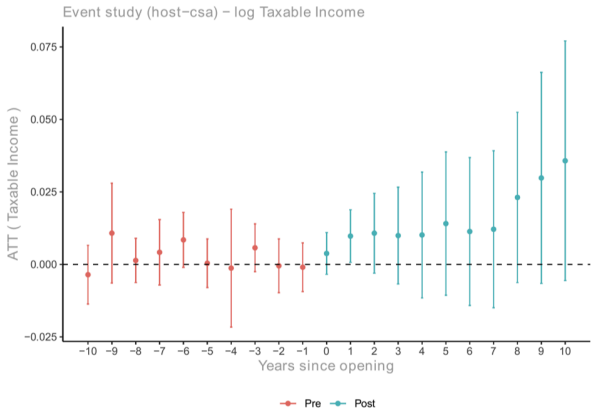
C1-24: Kindergartens - Sunab - no overlap sample

Synthesis.
Taken together, the prefectural results highlight a redistributional pattern: station openings attract population inflows (~4–9%), increase housing demand (~2–5%), and substantially raise commercial land

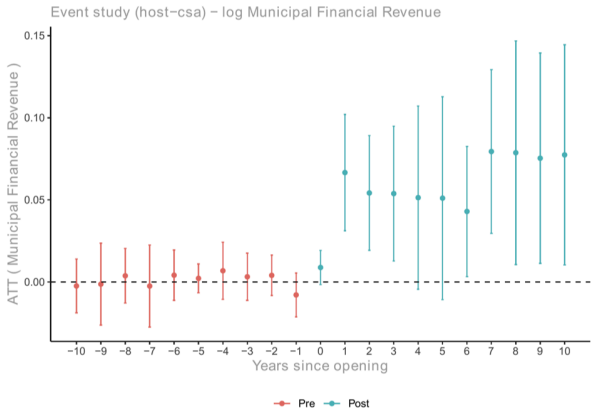
values (~12–13%), but do not stimulate GDP, taxable income, or fiscal expansion. Rising unemployment counts reinforce the interpretation of transitional labor market pressures rather than productivity growth.

5.2 Municipal-Level Effects

Municipal Finances.
The strongest and most immediate municipal effect of Shinkansen access is fiscal. Host municipalities show large and persistent increases in *financial revenue* (+5% in the short run, +9% in the medium run), significant from the opening year (e = 0) and with no pre-trend violations. In addition, *taxable income* rises in host municipalities by about 3–5% in the medium run. Together, these results indicate that station openings expand the municipal fiscal base directly, through higher property values, increased business activity, and a larger local tax base.



C2-8: Taxable Income - C&S DiD - host csa sample



C2-9: Municipal Fin. Rev. - C&S DiD - host csa sample

Housing Construction.
Housing markets adjust in both host and nearby municipalities. Host municipalities show positive effects under Sun–Abraham (+5% short run; +7% medium run). In the 30 km band, new residential dwellings increase strongly in the medium run (~+20%). These lagged responses, appearing after event years 7–8, are consistent with planning and construction cycles. Relevant Figures are C2-6, C2-11, C2-13, C2-14, C2-19, and C2-23)

Demographic Responses.
Live births increase in municipalities within 10–30 km. Medium-run gains are ~+6.5% at 10 km and ~+3% at 30 km, with significance beginning at event year 3. These increases plausibly reflect in-migration into station-adjacent areas followed by family formation and fertility timing effects. Relevant Figures are C2-1, C2-3, and C2-5.

Hospitals and Financial Need.
Hospital capacity also expands in 20–30 km catchments: hospital counts rise significantly in the medium run. Municipal financial need increases by ~3–4% within 10 km, consistent with greater demand for services and infrastructure as populations concentrate around stations. Relevant Figures are: Hospitals - C2-4, C2-7, C2-15, C2-20, and C2-21. Financial Need - C2-2 and C2-18.

Other Outcomes.
Other outcomes are less consistent. Municipal tax revenue, kindergarten pupils, and unemployment show some post-treatment significance under certain estimators but fail robustness or pre-trend screens, so we do not treat them as causal.

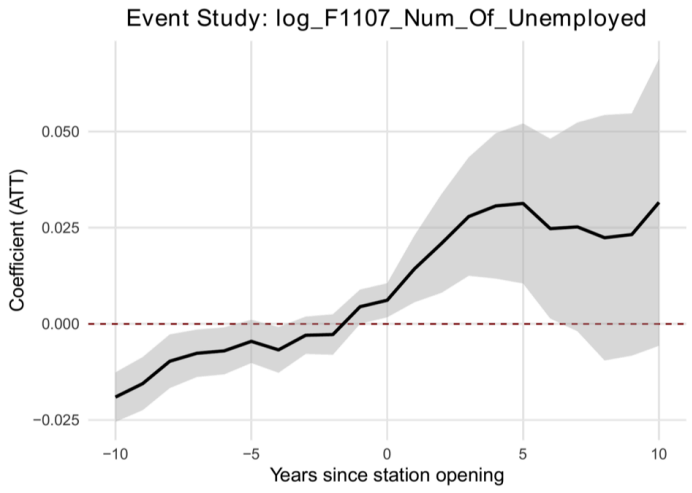


Figure C2-10: Unemployed - DiD - 10km csa sample

Synthesis.
At the municipal level, the clearest effects operate through *fiscal revenues and taxable income in host municipalities*, lagged *housing construction* in broader catchments, *increases in live births* in nearby municipalities, and *hospital expansion* in the 20–30 km band. Municipal financial need also rises close to stations, reflecting service demands. These results confirm that the most tangible Shinkansen impacts occur at the local scale, where fiscal, housing, demographic, and service adjustments directly reshape municipal economies.

5.3 Comparative Synthesis

Prefectural and municipal results tell a consistent redistributional story. At the prefecture level, Shinkansen station openings produce clear demographic and market responses: population increases of 4–9%, driven by in-migration (+3–5%) and net migration gains (~+3%), alongside new residential construction (+2–7%), sharp commercial land capitalization (~+13%), and higher unemployment counts (+3–10%). In contrast, GDP, taxable income, and prefectural revenues remain flat, indicating no broad productivity or fiscal expansion.

Municipal analysis clarifies the channels behind these patterns. Host municipalities record immediate and persistent fiscal revenue gains (+5–9%), expanding their tax bases from the year of opening onward. Neighboring municipalities experience large, lagged housing construction booms (+18–21% by years 6–10) and medium-run increases in live births (+3–7%), consistent with migration inflows followed by household formation. Service pressures rise as well: hospital counts expand modestly in 20–30 km bands, and municipal financial need increases by 3–4% within 10 km, reflecting greater infrastructure demand.

Taken together, these findings indicate that the Shinkansen reallocated population, housing demand, and fiscal capacity toward connected municipalities without raising aggregate GDP or prefectural revenues. Prefecture-level outcomes capture the broad redistribution, while municipal results expose its mechanisms: fiscal strengthening in host cities, construction surges and demographic inflows in surrounding areas, and rising local service burdens. The Shinkansen thus reshaped the geography of economic activity through redistribution and capitalization, rather than functioning as a catalyst for aggregate productivity growth.

Outcome	Level	Sample	Estimator(s)	Short-run (0–5y)	Medium-run (6–10y)	Robust?
Population	Prefecture	1-station	CSDID	+4.0%	+8.9%	Yes
Net Migration	Prefecture	Mult-10 / No-overlap	CSDID, DID	+2.6–2.7%	+2.8–2.9%	Yes
Immigrants	Prefecture	Mult-10 / No-overlap	CSDID, DID	+2.6–2.8%	+4.5–5.1%	Yes
New Residential Buildings	Prefecture	Mult-10 / No-overlap	CSDID, DID, SunAb	+3.7–4.4%	+6.5–6.9%	Yes
Commercial Land Values	Prefecture	No-overlap	CSDID, DID	+3.1%	+12.7%	Yes
Unemployed	Prefecture	Mult-10 / No-overlap	CSDID, DID, SunAb	+3.5–3.8%	+8.6–9.5%	Yes
Municipal Financial Revenue	Municipality	Host-csa	CSDID, DID, SunAb	+5.0%	+7.9%	Yes
Municipal Financial Need	Municipality	10km-csa	CSDID, SunAb	+0.4–3.1%	+1.3–4.4%	Yes
New Residential Dwellings	Municipality	10–30 km	CSDID, DID, SunAb	+3.5–6.9%	+18.7–21.1 %	Yes

Live Births	Municipality	10–30 km	CSDID only	+1.2–1.4%	+3.1–6.5%	Suggestive
Hospitals	Municipality	30 km	CSDID, DID, SunAb	~0%	+2% (small)	Suggestive

Table A9. Comparative Summary of Robust Outcomes: Prefecture vs. Municipality

6. Robustness and Diagnostics

Our main results rely on the Callaway & Sant'Anna (2021) estimator, which is designed for staggered adoption and avoids the negative-weighting problems documented in Goodman-Bacon (2021). To verify that findings are not estimator-specific or driven by pre-trend imbalances, we replicate the analysis using two complementary designs.

6.1 Matched-Sample DiD

As a further robustness check, we implement a matched-sample DiD. Units are matched on baseline outcome levels and pre-treatment outcome slopes using nearest-neighbor Mahalanobis distance, after which dynamic ATTs are re-estimated with the Callaway & Sant'Anna estimator. This design directly balances pre-trends, reducing concerns about non-parallel trajectories. Results closely resemble the baseline: in-migration gains of 3–5%, residential construction increases of 4–8%, and commercial land value gains of about 13%. At the municipal level, matched estimates confirm fiscal revenue gains and medium-run housing responses.

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Builds	mult-10	0.037 (~3.7%)	0.083 (~8.3%)
	no-overlap	0.035 (~3.5%)	0.081 (~8.1%)
In-migrants	mult-10	0.025 (~2.5%)	0.040 (~4.0%)
	no-overlap	0.029 (~2.9%)	0.053 (~5.3%)
Commercial Land Price	no-overlap	0.032 (~3.2%)	0.128 (~12.8%)
Unemployed	mult-10	0.034 (~3.4%)	0.087 (~8.7%)
	no-overlap	0.034 (~3.4%)	0.095 (~9.5%)
Net Migration	mult-10	0.027 (~2.7%)	0.031 (~3.1%)
	no-overlap	0.029 (~2.9%)	0.033 (~3.3%)

Table A2. Prefecture Pre-trend Match DiD Horizon Averages

6.2 Sun & Abraham (2021)

The Sun–Abraham interaction-weighted TWFE correction provides an alternative aggregation of cohort-specific effects. Estimates broadly confirm our baseline findings. For prefectural outcomes, Sun–Abraham detects medium-run increases in residential construction (+5–7%) and unemployment (+4–7%), while taxable income remains flat (≈0%). At the municipal level, Sun–Abraham corroborates immediate fiscal revenue gains in host municipalities (+5% short run; +7% medium run) and medium-run housing responses in nearby municipalities (+18% at 10 km; +17% at 30 km). These results reinforce the interpretation that Shinkansen openings reallocate population and housing demand without stimulating broad income growth.

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Builds	mult-10	0.075 (~7.5%)	0.073 (~7.3%)
	no-overlap	0.054 (~5.4%)	0.062 (~6.2%)
Kindergartens	mult-10	-0.083 (~-8.3%)	-0.136 (~-13.6%)
	no-overlap	-0.058 (~-5.8%)	-0.107 (~-10.7%)
Taxable Income	no-overlap	-0.001 (~-0.1%)	-0.001 (~-0.1%)
Unemployed	mult-10	0.043 (~4.3%)	0.074 (~7.4%)

Table A3. Prefecture Sun-Abraham Horizon Averages

6.3 Diagnostics

Across specifications, event-study figures show no statistically significant pre-treatment leads for the headline outcomes (population, migration, land, housing, municipal revenue). Pointwise pre-trend checks therefore support the identifying assumption. In addition, the matched-sample design ensures balance in both outcome levels and slopes. Finally, restricting the estimation window to 1975–2015 — where covariate coverage is more complete — yields results consistent with the full-sample estimates, suggesting that early decades do not drive the findings.

6.4 Summary

Taken together, the robustness checks show that our findings are not artifacts of estimator choice, pre-trend imbalances, or sample composition. Callaway & Sant'Anna, Sun–Abraham, and matched-sample DiD converge on the same qualitative pattern: modest but persistent population and migration inflows, substantial capitalization into land and housing markets, transitional increases in unemployment, and null effects on GDP, taxable income, and prefectural revenues. At the municipal level, fiscal and housing effects are robust across methods, underscoring the localized redistributional nature of Shinkansen access.

7. Mechanisms and Discussion

The results suggest that Shinkansen access reshaped regional economies through population reallocation and land market capitalization, rather than through productivity or income growth. Several mechanisms emerge from the prefectural and municipal evidence.

7.1 Migration as the Primary Channel

Population gains in treated prefectures are driven by in-migration, with out-migration unchanged. This pattern indicates that the Shinkansen facilitated relocation into connected areas rather than reducing outflows. At the municipal level, increases in live births in 10–30 km catchments provide further evidence that migration brought in younger households who subsequently formed families. The timing, significance emerging after several years, is consistent with household relocation and adjustment dynamics.

7.2 Land and Housing Markets

Housing and land markets provide the clearest capitalization channel. Residential construction rose by 2–5% at the prefectural level and by 17–27% at the municipal level in the medium run, reflecting demand for housing stock in accessible locations. Commercial land values increased by more than 12% within a decade, suggesting that firms and households were willing to pay substantial premiums for proximity to high-speed rail. These findings align with urban economics models of accessibility capitalization, where transportation improvements are absorbed into property markets rather than directly expanding output.

7.3 Fiscal Adjustments

Municipal finances also respond to station openings. Host municipalities record immediate and persistent increases in financial revenue, consistent with an enlarged tax base via higher property values, increased business activity, or greater eligibility for intergovernmental transfers tied to economic activity. At the same time, nearby municipalities register increases in financial need, reflecting rising demand for infrastructure and public services as populations concentrate in catchment areas. Together, these patterns suggest that fiscal adjustments are a key channel of redistribution, reinforcing the localized rather than aggregate nature of the gains.

7.4 Labor Market Dynamics

Unemployment counts rise in treated prefectures, even as population and housing expand. This pattern is most plausibly explained by transitional frictions: new migrants may not find immediate employment, or smaller prefectures may lack sufficient absorptive capacity in labor markets. These transitional effects underscore that migration-induced population growth does not automatically translate into net employment gains.

7.5 International Comparison

The Japanese experience contrasts sharply with findings from China. Banerjee, Duflo, and Qian (2020) document large positive GDP effects of high-speed rail through enhanced market access, while Faber (2014) reports similar productivity gains from China's national highway system. In contrast, Japan's Shinkansen shows redistribution without aggregate growth: population relocates, land values rise, but GDP and taxable income remain flat.

The results align more closely with European evidence. Studies of France and Spain point to modest or redistributional effects: Bouf and Desmaris (2020) and Di Ruocco (2024) show that high-speed rail in France reshaped the geography of activity without raising regional GDP, while Chen and de Abreu e Silva (2014) document educational and employment impacts of Spain's AVE but limited broad productivity gains. Blanquart (2017) synthesizes this literature, emphasizing that high-speed rail often acts as a catalyst for spatial reallocation rather than a driver of aggregate economic growth.

Taken together, the contrast highlights the importance of national context. China's rapid industrialization provided large market-access gains, whereas Japan's mature, urbanized economy channeled high-speed rail into redistributional rather than productivity-enhancing outcomes.

7.6 Interpretation

Taken together, these mechanisms support a redistributional interpretation of the Shinkansen. The network attracts population inflows, capitalizes into land and housing values, and strengthens municipal fiscal bases, but these gains do not scale into broad productivity or income growth. For policymakers, the implication is clear: the Shinkansen should be viewed less as an engine of aggregate growth and more as a tool of spatial redistribution and regional cohesion.

8. Conclusion

This paper has examined the long-run effects of Shinkansen station openings on Japanese prefectures and municipalities between 1950 and 2024. Using modern difference-in-differences estimators designed for staggered adoption, we find consistent evidence that station openings increase population, stimulate migration inflows, raise land values, and expand housing construction. At the municipal level, station openings generate immediate fiscal revenue gains in host municipalities and catalyze housing construction in nearby catchments.

At the same time, we find no evidence of broad productivity or fiscal expansion. Prefectural GDP, taxable income, and aggregate fiscal revenues remain flat, while unemployment counts rise modestly in the aftermath of station openings. These results suggest that the Shinkansen functions primarily as a redistributional amenity: it reallocates people and economic activity toward connected areas and capitalizes into land markets, but does not generate sustained aggregate growth.

The policy implications are twofold. First, high-speed rail investments should not be justified on the basis of aggregate GDP growth alone. Their principal contribution lies in reshaping the geography of economic activity, enhancing accessibility, and redistributing population and fiscal capacity across space. Second, if redistribution is the main channel, then the welfare consequences depend on whether these reallocations align with broader policy goals. For example, revitalizing peripheral regions or promoting national cohesion.

Future work should explore heterogeneity across opening cohorts, examine longer-run persistence, and integrate tourism and industrial policy channels. Cross-national comparisons will also be valuable, as the Japanese experience contrasts with the large GDP effects reported in China and the more localized gains observed in Europe. By combining prefectural and municipal perspectives, this study contributes to a clearer understanding of how high-speed rail reshapes, but does not necessarily grow, regional economies.

References

Core Econometric Methods

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

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

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

Japan-Specific HSR Studies

Nickelsburg, John, Manmohan Ahluwalia, and Yang. 2018. *High-Speed Rail Economics, Urbanization and Housing Affordability in Japan (1964–2010).* Working Paper, UCLA Anderson School of Management.

Wang, Jingyuan, Shintaro Terabe, Hideki Yaginuma, Haruka Uno, and Yu Suzuki. 2024. "Do High-Speed Railways Have an Impact on Population Change? Evidence from Japan." *Transport Policy* 149: 12–25.

Europe and International HSR Evidence

Ahlfeldt, Gabriel, and Andreas Feddersen. 2010. "Market Access and GDP: Transit Station Effects in Germany." *Journal of Economic Geography* 10(1): 1–28.

Blanquart, Camille. 2017. "The Local Economic Impacts of High-Speed Railways: A Theoretical Review." *European Transport Research Review* 9(4): 1–12.

Bouf, Jean, and Nicolas Desmaris. 2020. *Economic Impacts of High-Speed Rail in France: Regional Per Capita GDP Effects.* Research Report.

Chen, Lisa, and Tiago de Abreu e Silva. 2014. "High-Speed Rail, Education, Employment, and GDP: A Structural Model from Spain." *Transport Research Part A: Policy and Practice* 61: 187–202.

Di Ruocco, Ilaria. 2024. "Economic Impact of High-Speed Rail on Rural Areas: Evidence from France." *Science Publishing Group.*

Cheng, J. 2022. "Socioeconomic Impacts of High-Speed Rail: A Review of Studies." *Transportation Reviews* 42(3): 345–372.

Other Transport Infrastructure Papers

Banerjee, Abhijit, Esther Duflo, and Nancy Qian. 2020. “On the Road: Access to Transportation Infrastructure and Economic Growth in China.” *Journal of Development Economics* 145: 102442.

Donaldson, Dave, and Richard Hornbeck. 2016. “Railroads and American Economic Growth: A ‘Market Access’ Approach.” *Quarterly Journal of Economics* 131(2): 799–858.

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

Appendix

Appendix A. Estimator Horizon Tables

Table A1. Prefecture CS-DiD Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Builds	mult-10	0.022 (~2.2%)	0.052 (~5.2%)
	no-overlap	0.022 (~2.2%)	0.052 (~5.2%)
In-migrants	mult-10	0.027 (~2.7%)	0.049 (~4.9%)
	no-overlap	0.027 (~2.7%)	0.049 (~4.9%)
Commercial Land Price	mult-10	0.030 (~3.0%)	0.125 (~12.5%)
	no-overlap	0.030 (~3.0%)	0.125 (~12.5%)
Unemployed	mult-10	0.036 (~3.6%)	0.096 (~9.6%)
	no-overlap	0.036 (~3.6%)	0.096 (~9.6%)
Population	1-station	0.040 (~4.0%)	0.089 (~8.9%)
Net Migration	mult-10	0.025 (~2.5%)	0.025 (~2.5%)
	no-overlap	0.025 (~2.5%)	0.025 (~2.5%)

Table A2. Prefecture Pre-trend Match DiD Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Builds	mult-10	0.037 (~3.7%)	0.083 (~8.3%)
	no-overlap	0.035 (~3.5%)	0.081 (~8.1%)
In-migrants	mult-10	0.025 (~2.5%)	0.040 (~4.0%)
	no-overlap	0.029 (~2.9%)	0.053 (~5.3%)
Commercial Land Price	no-overlap	0.032 (~3.2%)	0.128 (~12.8%)
Unemployed	mult-10	0.034 (~3.4%)	0.087 (~8.7%)
	no-overlap	0.034 (~3.4%)	0.095 (~9.5%)
Net Migration	mult-10	0.027 (~2.7%)	0.031 (~3.1%)
	no-overlap	0.029 (~2.9%)	0.033 (~3.3%)

Table A3. Prefecture Sun-Abraham Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Builds	mult-10	0.075 (~7.5%)	0.073 (~7.3%)
	no-overlap	0.054 (~5.4%)	0.062 (~6.2%)
Kindergartens	mult-10	-0.083 (~-8.3%)	-0.136 (~-13.6%)
	no-overlap	-0.058 (~-5.8%)	-0.107 (~-10.7%)
Taxable Income	no-overlap	-0.001 (~-0.1%)	-0.001 (~-0.1%)
Unemployed	mult-10	0.043 (~4.3%)	0.074 (~7.4%)

Table A4. Prefecture-level summary table (consolidated).

Outcome	Sample	Estimator(s) with significance	Short-run (0–5y)	Medium-run (6–10y)	Robust?
Commercial Land Values	mult-10	CSDID	0.030 (~3.0%)	0.125 (~12.5%)	No
Population	1-station	CSDID	0.040 (~4.0%)	0.089 (~8.9%)	No
Commercial Land Values	no-overlap	CSDID, DID	0.031 (~3.1%)	0.127 (~12.7%)	Yes
Immigrants	mult-10	CSDID, DID	0.026 (~2.6%)	0.045 (~4.5%)	Yes
Immigrants	no-overlap	CSDID, DID	0.028 (~2.8%)	0.051 (~5.1%)	Yes
Net Migration	mult-10	CSDID, DID	0.026 (~2.6%)	0.028 (~2.8%)	Yes
Net Migration	no-overlap	CSDID, DID	0.027 (~2.7%)	0.029 (~2.9%)	Yes
Unemployed	no-overlap	CSDID, DID	0.035 (~3.5%)	0.095 (~9.5%)	Yes
New Residential Buildings	mult-10	CSDID, DID, SUNAB	0.044 (~4.4%)	0.069 (~6.9%)	Yes
New Residential Buildings	no-overlap	CSDID, DID, SUNAB	0.037 (~3.7%)	0.065 (~6.5%)	Yes
Unemployed	mult-10	CSDID, DID, SUNAB	0.038 (~3.8%)	0.086 (~8.6%)	Yes
Kindergarten Pupils	mult-10	SUNAB	-0.083 (~-8.3%)	-0.136 (~-13.6%)	No
Kindergarten Pupils	no-overlap	SUNAB	-0.058 (~-5.8%)	-0.107 (~-10.7%)	No

Taxable Income	no-overlap	SUNAB	-0.001 (~-0.1%)	-0.001 (~-0.1%)	No

Table A5. Municipal CS-DiD Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
Taxable Income	host-csa	0.010 (~1.0%)	0.022 (~2.2%)
New Residential Dwellings	30km-csa	0.040 (~4.0%)	0.174 (~17.4%)
Municipal Financial Revenue	host-csa	0.048 (~4.8%)	0.071 (~7.1%)
	10km-csa	0.005 (~0.5%)	0.014 (~1.4%)
Live Births	10km-csa	0.014 (~1.4%)	0.065 (~6.5%)
	20km-csa	0.012 (~1.2%)	0.038 (~3.8%)
	30km-csa	0.014 (~1.4%)	0.031 (~3.1%)
Hospitals	20km-csa	0.008 (~0.8%)	0.012 (~1.2%)
	30km-csa	0.001 (~0.1%)	-0.005 (~-0.5%)

Table A6. Municipal Pre-trend Match DiD Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
Unemployed	10km-csa	0.022 (~2.2%)	0.025 (~2.5%)
New Residential Dwellings	10km-csa	0.051 (~5.1%)	0.269 (~26.9%)
	20km-csa	-0.018 (~-1.8%)	0.145 (~14.5%)
	30km-csa	0.029 (~2.9%)	0.200 (~20.0%)
Municipal Tax Revenue	host-csa	0.039 (~3.9%)	0.034 (~3.4%)
Municipal Financial Revenue	host-csa	0.054 (~5.4%)	0.096 (~9.6%)
Kindergarten Pupils	20km-csa	0.099 (~9.9%)	0.217 (~21.7%)
Hospitals	30km-csa	-0.001 (~-0.1%)	-0.009 (~-0.9%)

Table A7. Municipal Sun-Abraham Horizon Averages

Outcome	Sample	Short-run (0–5y)	Medium-run (6–10y)
New Residential Dwellings	10km-csa	0.087 (~8.7%)	0.153 (~15.3%)
	host-csa	-0.068 (~-6.8%)	-0.077 (~-7.7%)
Municipal Financial Revenue	host-csa	0.047 (~4.7%)	0.070 (~7.0%)
	10km-csa	0.003 (~0.3%)	0.011 (~1.1%)
Hospitals	10km-csa	0.008 (~0.8%)	0.012 (~1.2%)
	30km-csa	-0.002 (~-0.2%)	-0.014 (~-1.4%)

Table A8. Municipal-level summary table (consolidated)

Outcome	Sample	Estimator(s) with significance	Short-run (0–5y)	Medium-run (6–10y)	Robust?
Taxable Income	host-csa	CSDID	0.010 (~1.0%)	0.022 (~2.2%)	No
Live Births	10km-csa	CSDID	0.014 (~1.4%)	0.065 (~6.5%)	No
Live Births	20km-csa	CSDID	0.012 (~1.2%)	0.038 (~3.8%)	No
Live Births	30km-csa	CSDID	0.014 (~1.4%)	0.031 (~3.1%)	No
Hospitals	20km-csa	CSDID	0.008 (~0.8%)	0.012 (~1.2%)	No
New Residential Dwellings	30km-csa	CSDID, DID	0.035 (~3.5%)	0.187 (~18.7%)	Yes
Municipal Financial Revenue	host-csa	CSDID, DID, SUNAB	0.050 (~5.0%)	0.079 (~7.9%)	Yes
Hospitals	30km-csa	CSDID, DID, SUNAB	-0.001 (~-0.1%)	-0.009 (~-0.9%)	Yes
Municipal Financial Need	10km-csa	CSDID, SUNAB	0.004 (~0.4%)	0.013 (~1.3%)	Yes
Unemployed	10km-csa	DID	0.022 (~2.2%)	0.025 (~2.5%)	No
New Residential Dwellings	20km-csa	DID	-0.018 (~-1.8%)	0.145 (~14.5%)	No
Municipal Tax Revenue	host-csa	DID	0.039 (~3.9%)	0.034 (~3.4%)	No
Kindergarten	20km-csa	DID	0.099 (~9.9%)	0.217 (~21.7%)	No

Pupils					
New Residential Dwellings	10km-csa	DID, SUNAB	0.069 (~6.9%)	0.211 (~21.1%)	Yes
New Residential Dwellings	host-csa	SUNAB	-0.068 (~-6.8%)	-0.077 (~-7.7%)	No
Hospitals	10km-csa	SUNAB	0.008 (~0.8%)	0.012 (~1.2%)	No

Table A9. Comparative Summary of Robust Outcomes: Prefecture vs. Municipality

Outcome	Level	Sample	Estimator(s)	Short-run (0–5y)	Medium-run (6–10y)	Robust?
Population	Prefecture	1-station	CSDID	+4.0%	+8.9%	Yes
Net Migration	Prefecture	Mult-10 / No-overlap	CSDID, DID	+2.6–2.7%	+2.8–2.9%	Yes
Immigrants	Prefecture	Mult-10 / No-overlap	CSDID, DID	+2.6–2.8%	+4.5–5.1%	Yes
New Residential Buildings	Prefecture	Mult-10 / No-overlap	CSDID, DID, SunAb	+3.7–4.4%	+6.5–6.9%	Yes
Commercial Land Values	Prefecture	No-overlap	CSDID, DID	+3.1%	+12.7%	Yes
Unemployed	Prefecture	Mult-10 / No-overlap	CSDID, DID, SunAb	+3.5–3.8%	+8.6–9.5%	Yes
Municipal Financial Revenue	Municipality	Host-csa	CSDID, DID, SunAb	+5.0%	+7.9%	Yes
Municipal Financial Need	Municipality	10km-csa	CSDID, SunAb	+0.4–3.1%	+1.3–4.4%	Yes
New Residential Dwellings	Municipality	10–30 km	CSDID, DID, SunAb	+3.5–6.9%	+18.7–21.1 %	Yes
Live Births	Municipality	10–30 km	CSDID only	+1.2–1.4%	+3.1–6.5%	Suggestive

Hospitals	Municipality	30 km	CSDID, DID, SunAb	~0%	+2% (small)	Suggestive

Appendix B. Data Sources and Construction

B.1 e-Stat: Japan's National Data Portal

The majority of data in this study is sourced from **e-Stat**, Japan's official government statistics portal (https://www.e-stat.go.jp/). e-Stat is administered by the **Statistics Bureau of Japan** and serves as a central access point for official statistics compiled by national ministries and agencies. It contains both historical and contemporary data across the following domains:

- **Population and Demographics** (Population Census, Basic Resident Register).
- **Economic Activity** (Prefectural Accounts, Establishment and Enterprise Census, Labor Force Survey).
- **Finance and Public Accounts** (Local Government Finance Statistics, municipal and prefectural fiscal data).
- **Housing and Land** (Housing and Land Survey, Building Starts Statistics, Land Price Survey).
- **Social Indicators** (School Basic Survey, Vital Statistics).

e-Stat provides programmatic access through an **API**, which allows users to query and download series in machine-readable formats (CSV, JSON, XML). For this project, nearly all prefectural and municipal variables — including GDP, taxable income, labor force indicators, housing starts, municipal revenues, and service indicators such as hospitals and schools — were retrieved directly from e-Stat using custom API scripts.

B.2 Population and Yearbook Data

While e-Stat provides population counts from the modern censuses, its historical coverage is incomplete at the prefectural level before the 1970s. To extend the population series consistently back to 1950 (the start of our prefectural panel), we collected figures from **prefectural statistical yearbooks**. These volumes, published annually by each prefecture, contain historical census tabulations and intercensal population estimates. We digitized and manually reconciled these figures with the e-Stat population data, ensuring consistency in definitions and coverage. Where necessary, we interpolated between census benchmarks to create a continuous annual panel.

The resulting population series is a **synthesized panel**:

- **1950–1970:** derived primarily from prefectural yearbooks.
- **1970 onward:** taken from e-Stat Population Census and Basic Resident Register data.
- **All years:** harmonized into a consistent series (log and log-linear versions), with adjustments for boundary changes.

B.3 Harmonization and Transformations

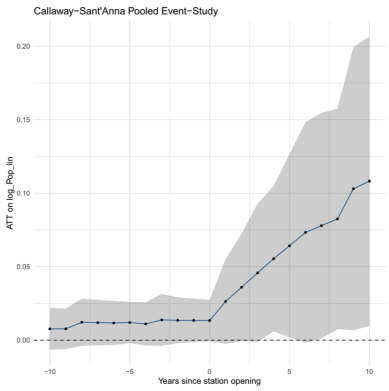
To address municipal mergers and naming inconsistencies (particularly during the Great Heisei Consolidation of the 2000s), we constructed a harmonized crosswalk linking our raw municipal panel to official municipality codes. We began from a GIS boundary shapefile containing municipality codes (adm_code) and names, and joined these to our master panel by prefecture and municipality name. Because municipality names vary across sources and over time (e.g., “Happou Cho” vs. “Happo Cho”), we manually standardized names and corrected common variants prior to joining. The procedure successfully matched all municipalities except for a handful of non-municipal aggregates (e.g., “Ku Area” in Tokyo), which were dropped.

The result is a panel where each observation is associated with a stable municipal identifier (Municipality_Code), ensuring consistency across the 1980–2023 period despite boundary changes and mergers.

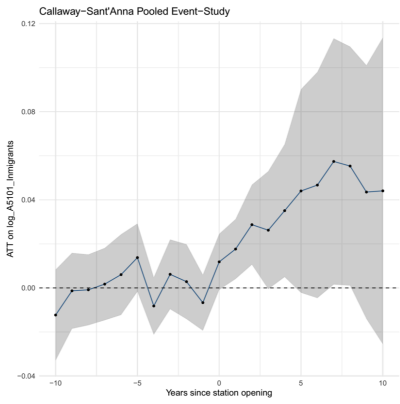
- **Monetary values:** deflated to constant prices using national GDP deflators.
- **Zeros in count data:** transformed using log(1+x) for outcomes such as hospitals, births, and kindergarten pupils.

Appendix C. Event-Study Figures

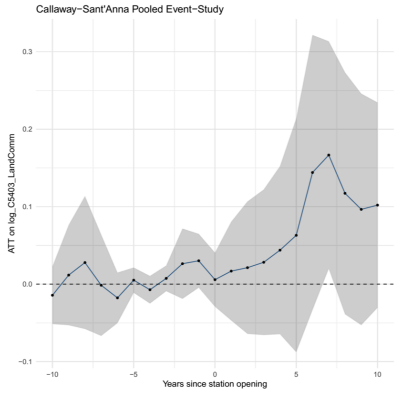
C.1 Prefecture Dynamic ATT Plots (Outcome, Sample, Estimator combination)



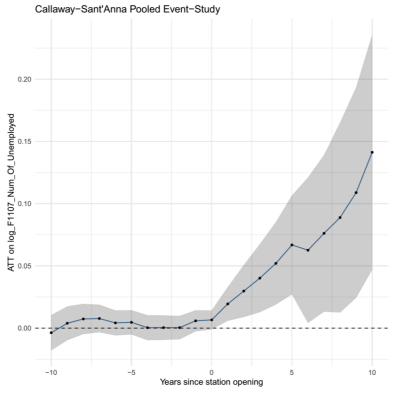
C1-1: log Pop Lin - C&S DiD - 1 station sample



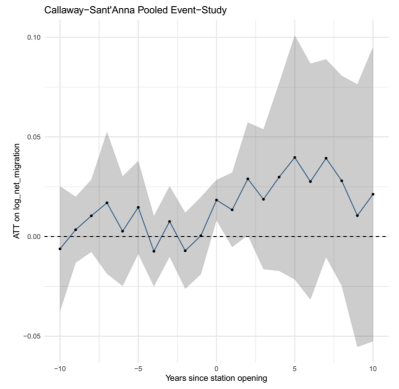
C1-2: ln-migrants - C&S DiD - mult 10 sample



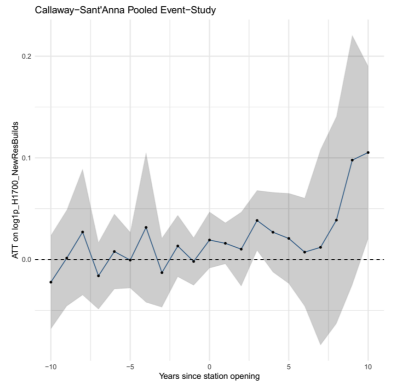
C1-3: Commercial Land Price - C&S DiD - mult 10 sample



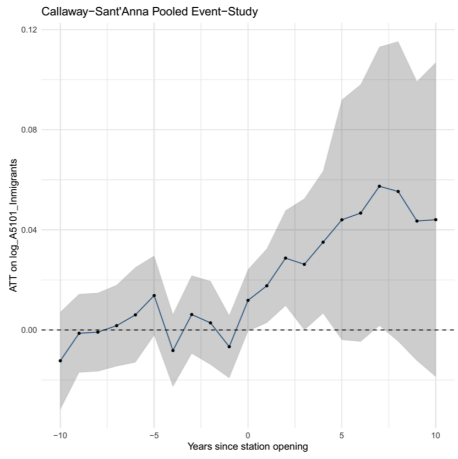
C1-4: Unemployed - C&S DiD - mult 10 sample



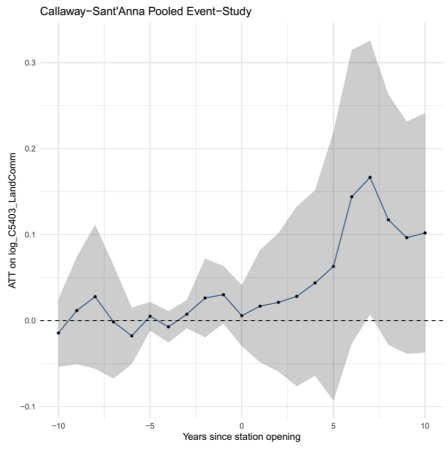
C1-5: Net Migration - C&S DiD - mult 10 sample



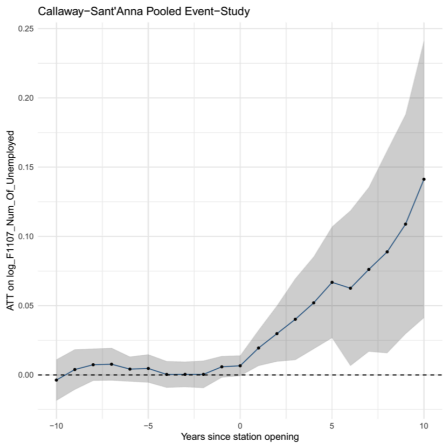
C1-6: log1p New Residential Builds - C&S DiD - mult 10 sample



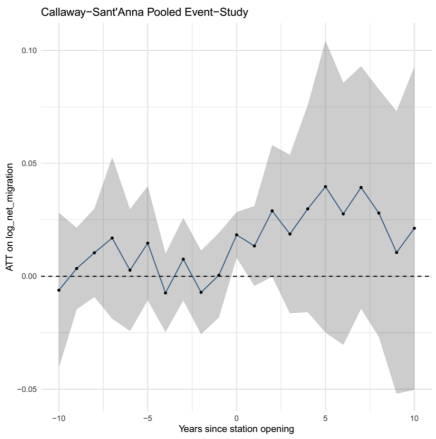
C1-7: In-migrants - C&S DiD - no overlap sample



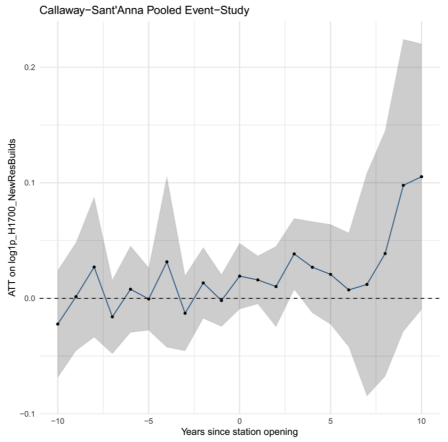
C1-8: Commercial Land Price - C&S DiD - no overlap sample



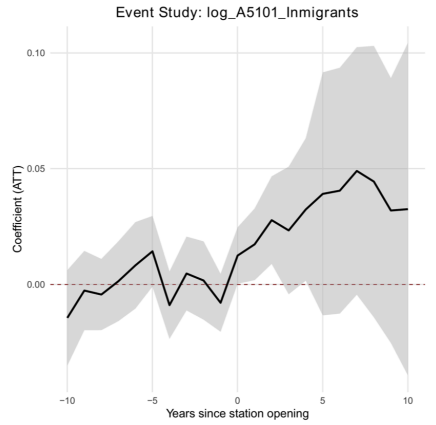
C1-9: Unemployed - C&S DiD - no overlap sample



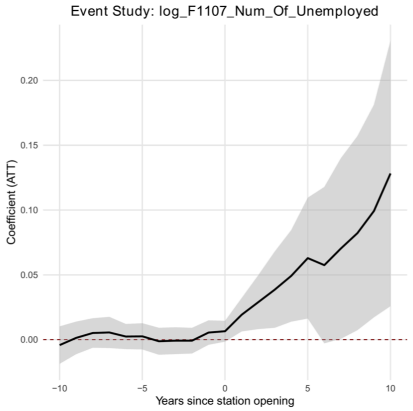
C1-10: Net Migration - C&S DiD - no overlap sample



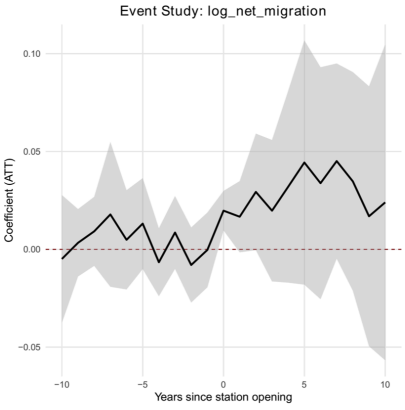
C1-11: log1p New Residential Builds - C&S DiD - no overlap sample



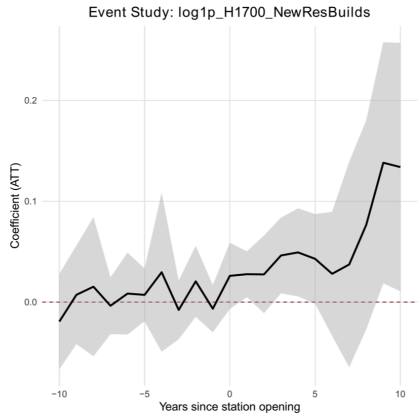
C1-12: In-migrants - DiD - mult 10 sample



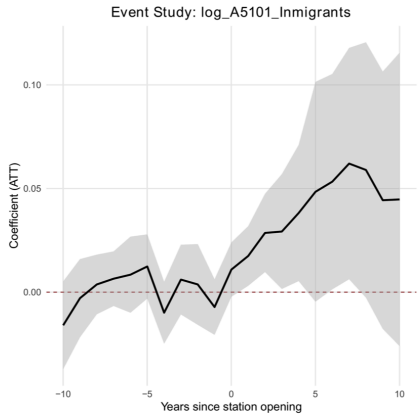
C1-13: Unemployed - DiD - mult 10 sample



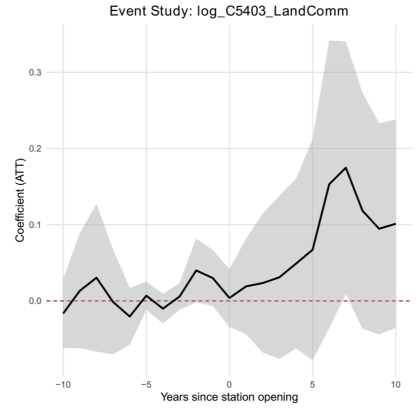
C1-14: Net Migration - DiD - mult 10 sample



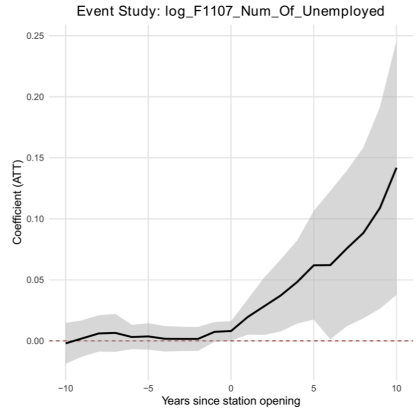
C1-15: log1p New Residential Builds - DiD - mult 10 sample



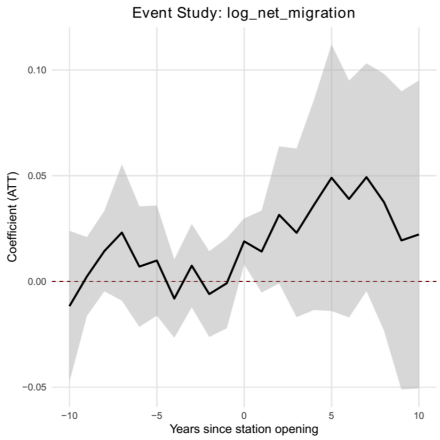
C1-16: In-migrants - DiD - no overlap sample



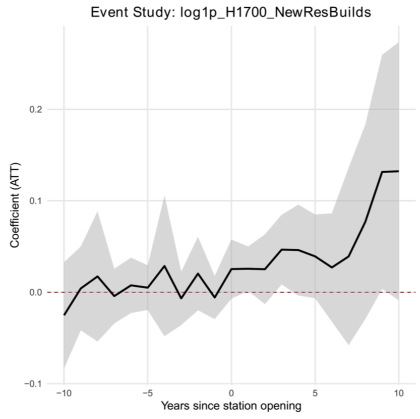
C1-17: Commercial Land Price - DiD - no overlap sample



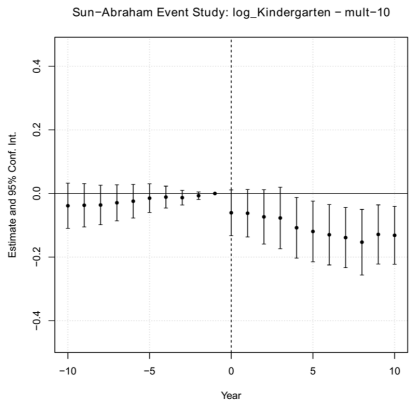
C1-18: Unemployed - DiD - no overlap sample



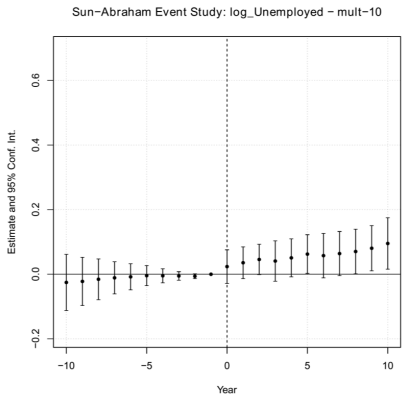
C1-19: Net Migration - DiD - no overlap sample



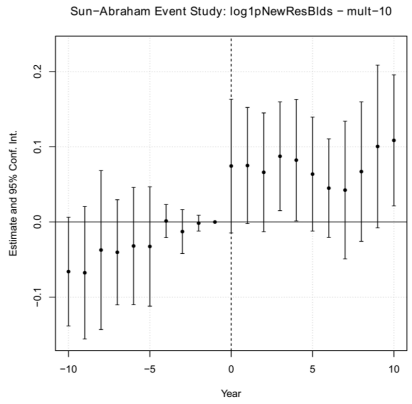
C1-20: log1p New Residential Builds - DiD - no overlap sample



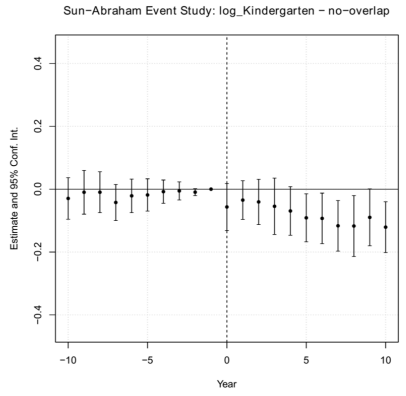
C1-21: Kindergartens - Sunab - mult 10 sample



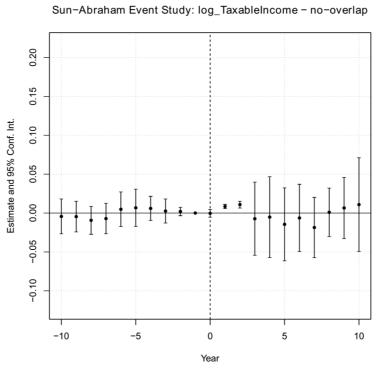
C1-22: Unemployed - Sunab - mult 10 sample



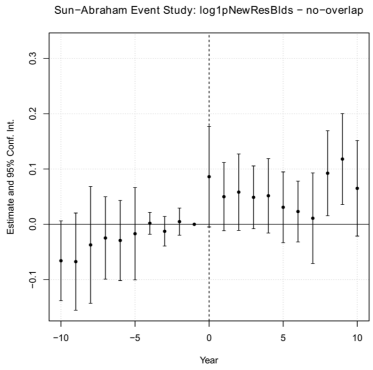
C1-23: log1p New Residential Builds - Sunab - mult 10 sample



C1-24: Kindergartens - Sunab - no overlap sample

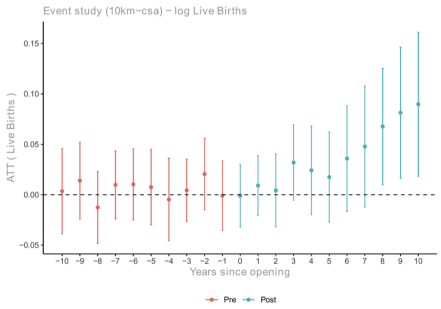


C1-25: Taxable Income - Sunab - no overlap sample

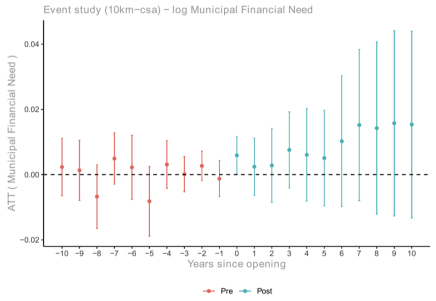


C1-26: log1p New Residential Builds - Sunab - no overlap sample

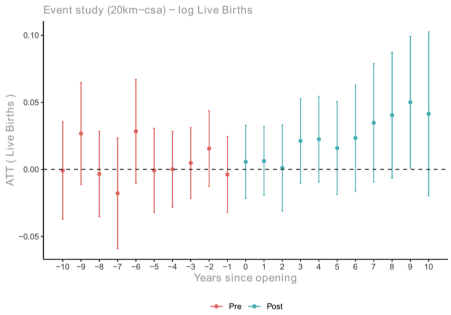
C.2 Municipal Dynamic ATT Plots (Outcome, Sample, Estimator combination)



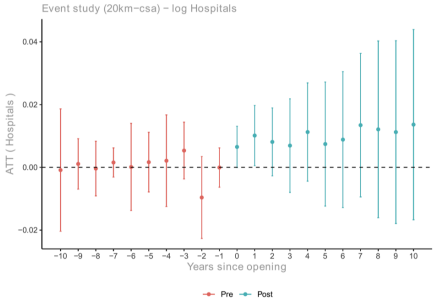
C2-1: Live Births - C&S DiD - 10km csa sample



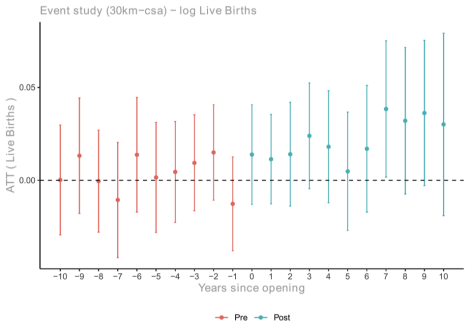
CS2-2: Municipal Financial Need - C&S DiD - 10km csa sample



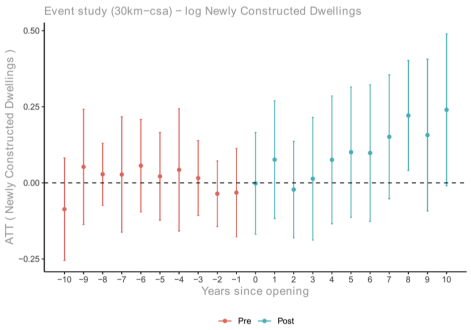
C2-3: Live Births - C&S DiD - 20km csa sample



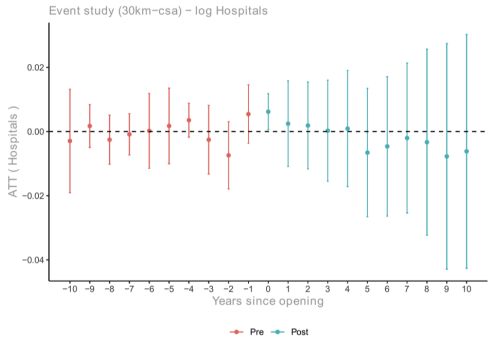
C2-4: log1p Hospitals - C&S DiD - 20km csa sample



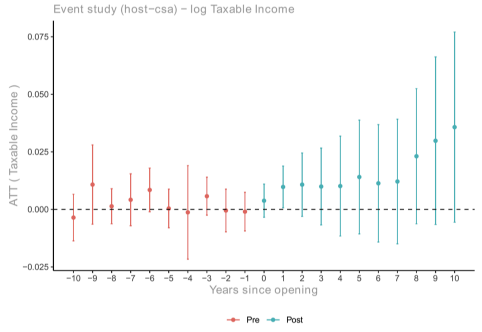
C2-5: Live Births - C&S DiD - 30km csa sample



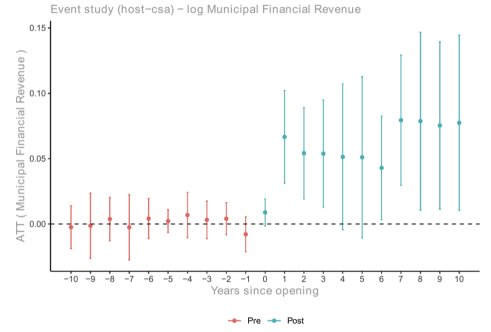
C2-6: log1p Newly Constructed Dwellings - C&S DiD - 30km csa sample



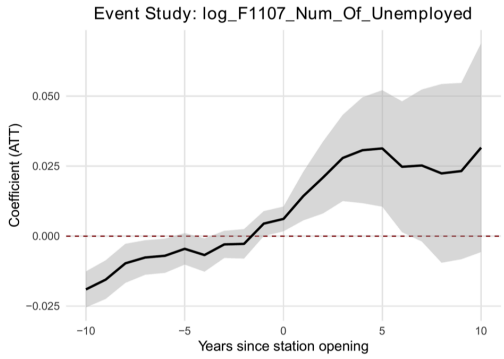
C2-7: log1p Hospitals - C&S DiD - 30km csa sample



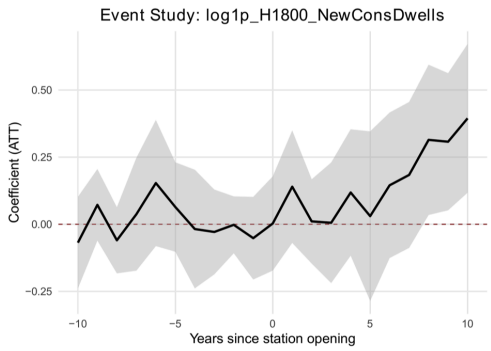
C2-8: Taxable Income - C&S DiD - host csa sample



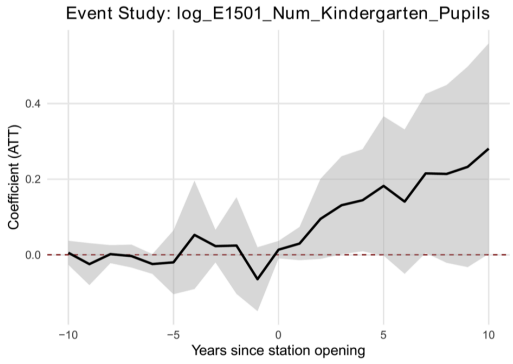
C2-9: Municipal Financial Revenue - C&S DiD - host csa sample



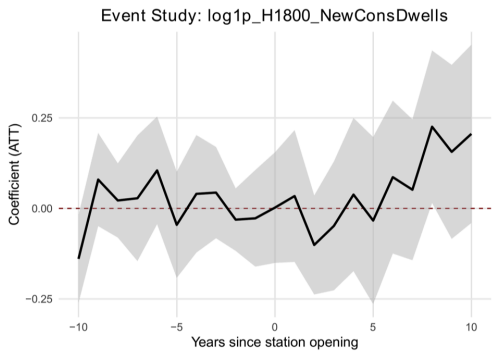
C2-10: Unemployed - DiD - 10km csa sample



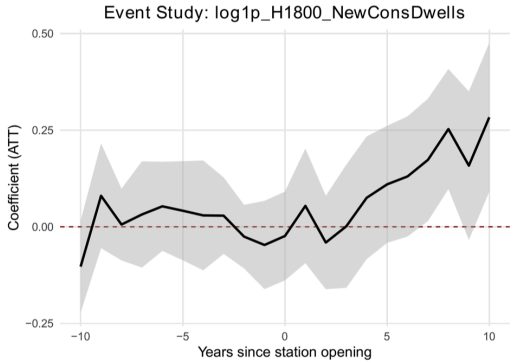
C2-11: log1p Newly Constructed Dwellings - DiD - 10km csa sample



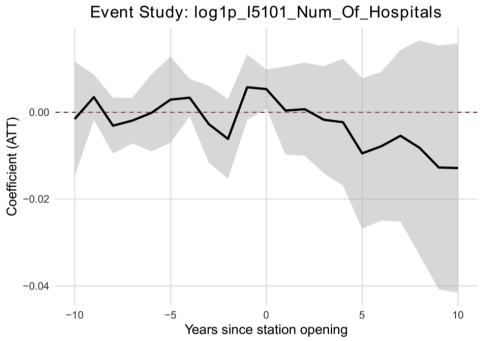
C2-12: Kindergarten Pupils - DiD - 20km csa sample



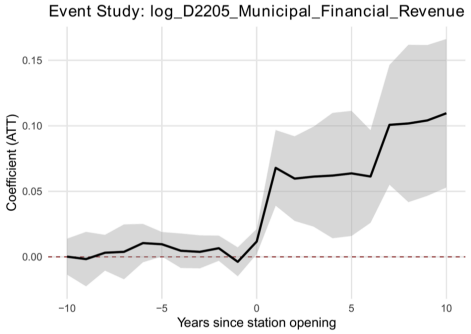
C2-13: log1p Newly Constructed Dwellings - DiD - 20km csa sample



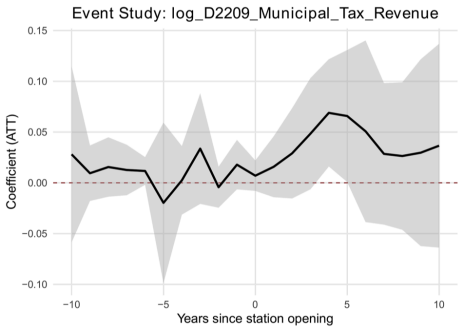
C2-14: log1p Newly Constructed Dwellings - DiD - 30km csa sample



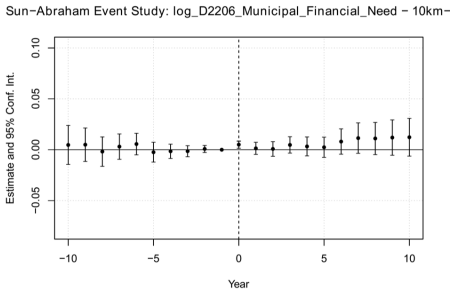
C2-15: log1p Hospitals - DiD - 30km csa sample



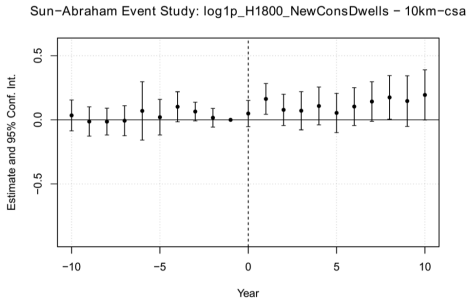
C2-16: Municipal Financial Revenue - DiD - host csa sample



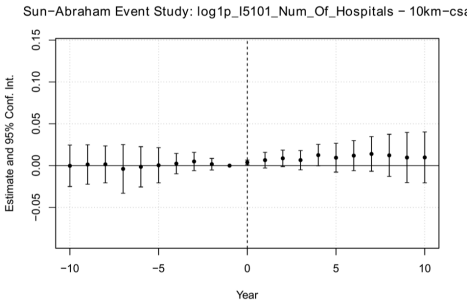
C2-17: Municipal Tax Revenue - DiD - host csa sample



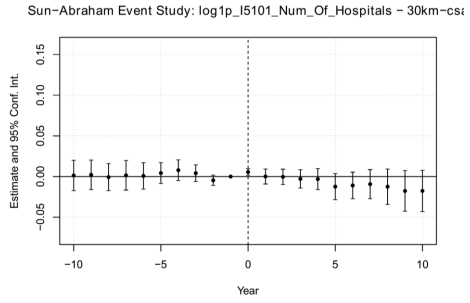
C2-18: Municipal Financial Need - Sunab - 10km csa sample



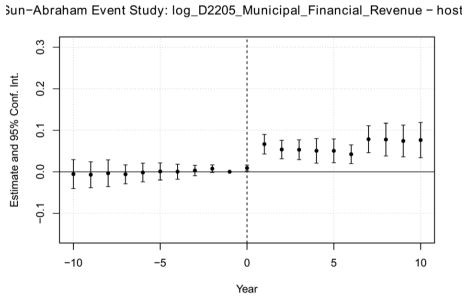
C2-19: log1p Newly Constructed Dwellings - Sunab - 10km csa sample



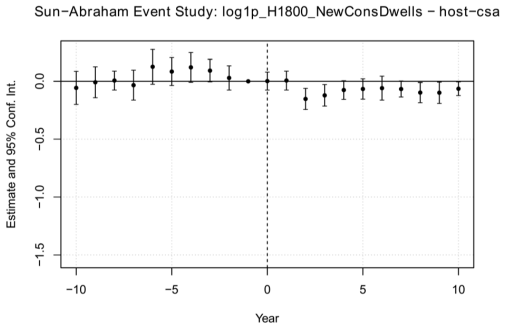
C2-20: log1p Hospitals - Sunab - 10km csa sample



C2-21: log1p Hospitals - Sunab - 30km csa sample



C2-22: Municipal Financial Revenue - Sunab - host csa sample



C2-23: log1p Newly Constructed Dwellings - Sunab - host csa sample

Appendix D. Sample Visualizations Simple

D.1 Plots

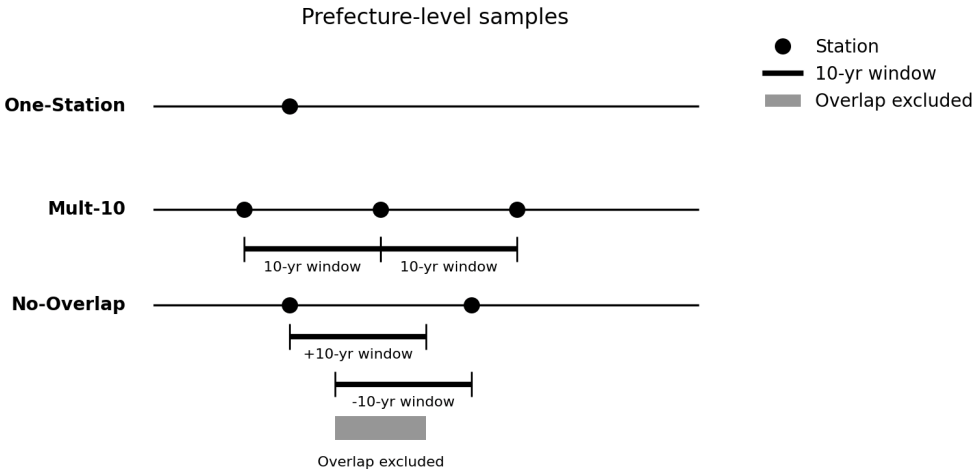


Figure D1-1: Prefecture-level samples

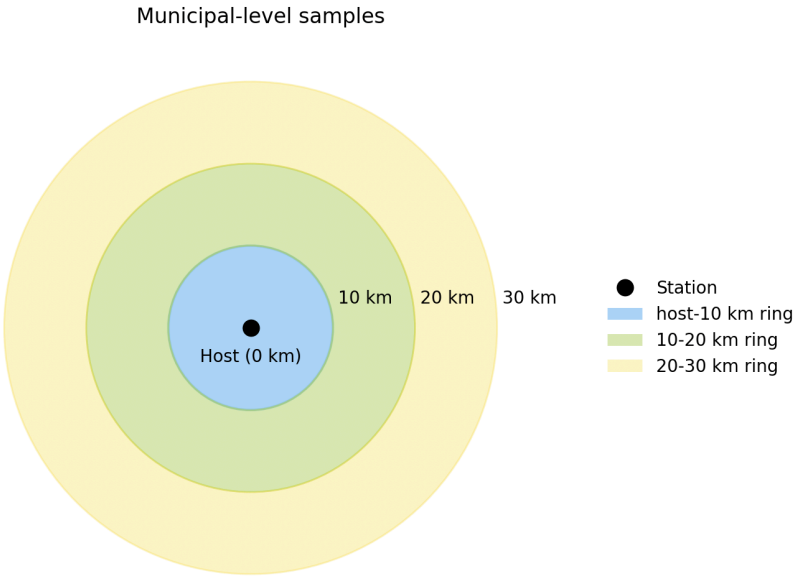


Figure D1-2: Municipal-level samples

Appendix E. Sample Visualizations with Japan Overlay

E.1 Prefecture Plots

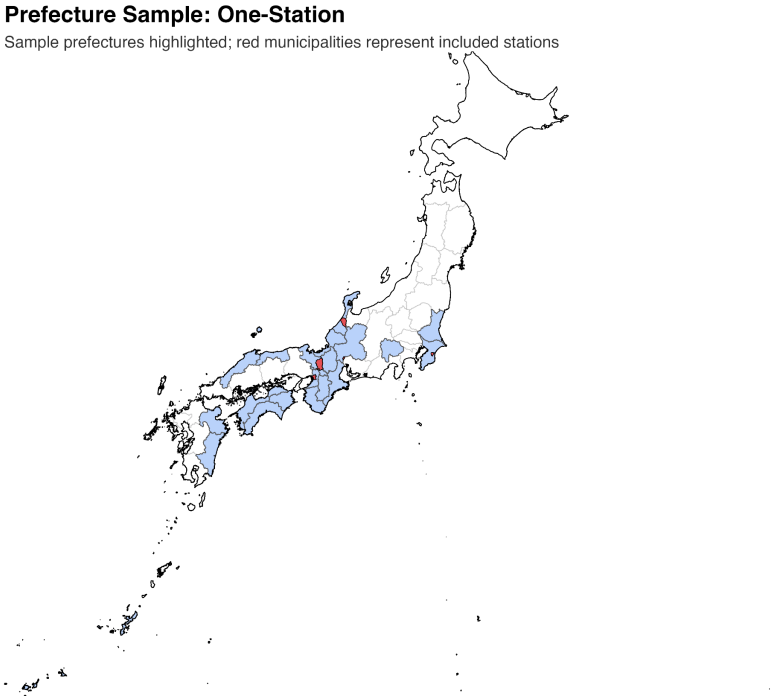


Figure E1-1: Prefecture one-station sample visualization

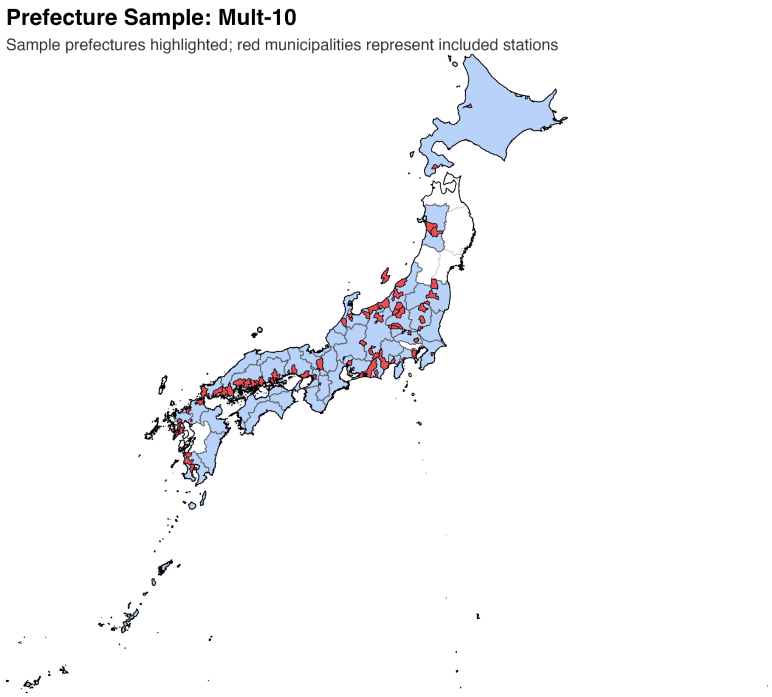


Figure E1-2: Prefecture mult-10 sample visualization

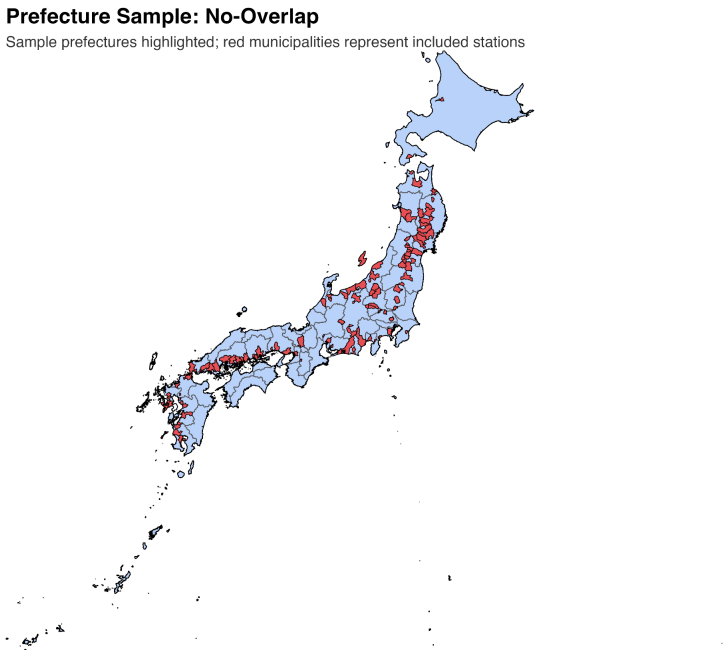


Figure E1-3: Prefecture no-overlap sample visualization

E.2 Municipality Plot

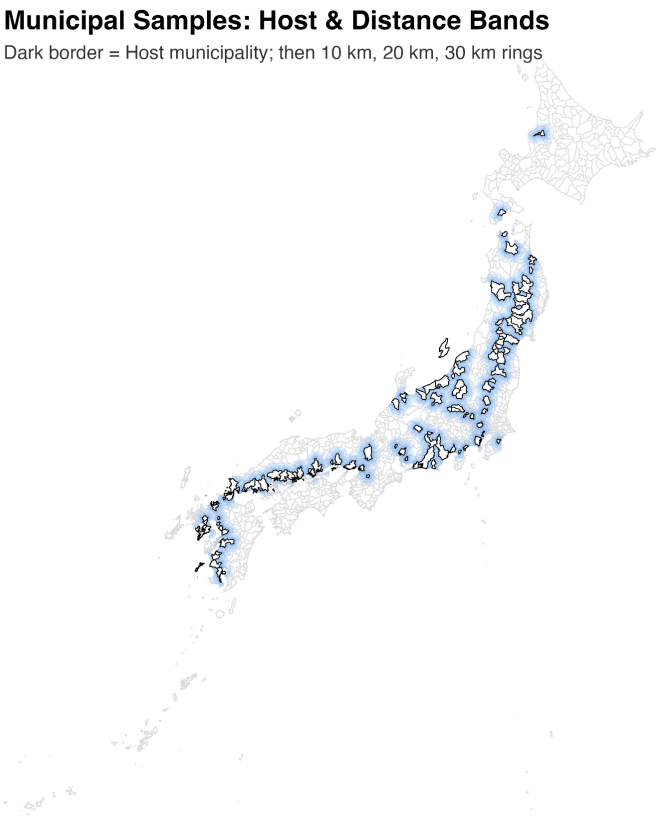


Figure E2-1: Municipality all samples visualization