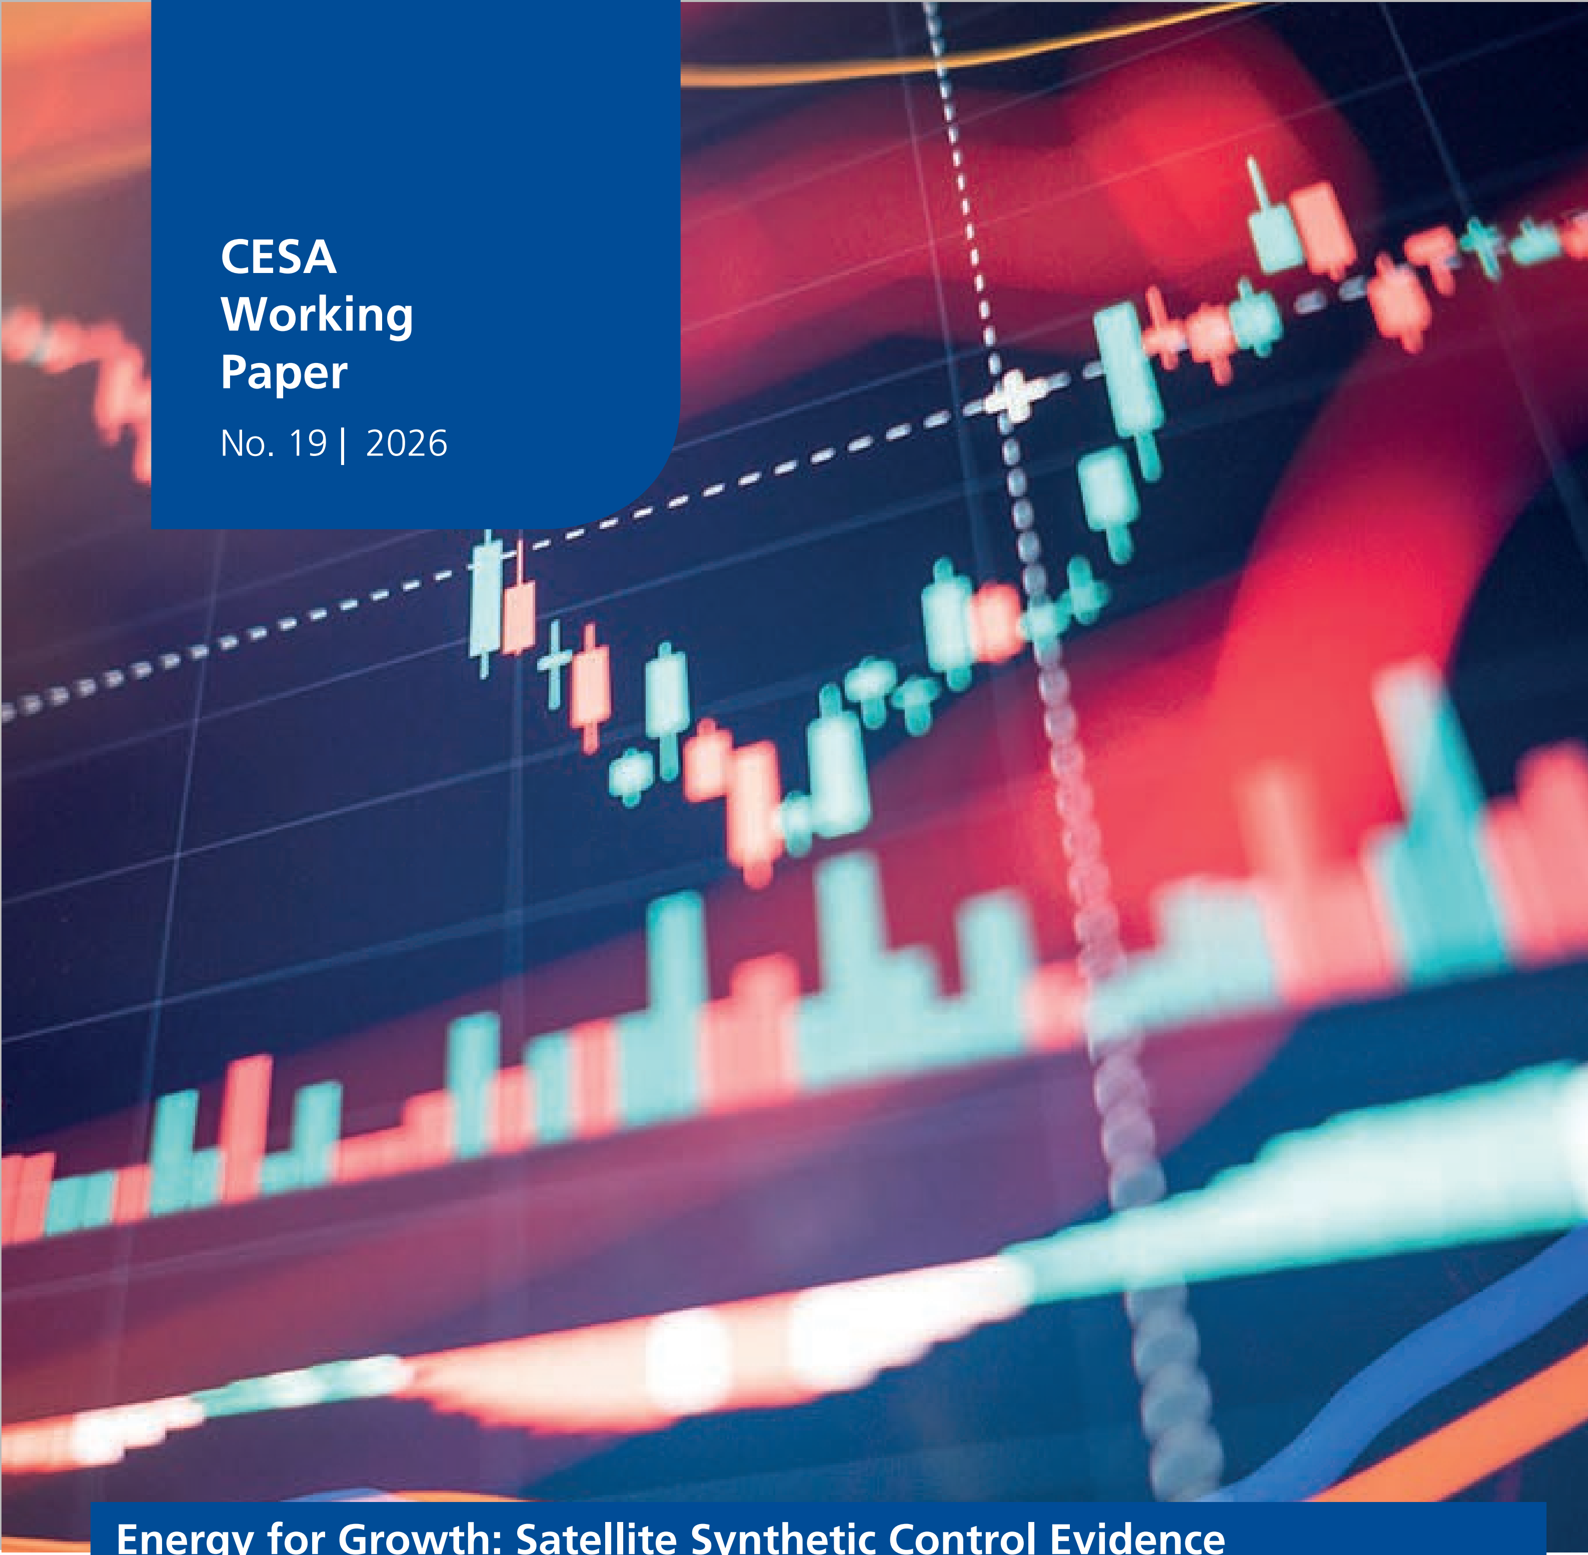

CESA Working Paper

No. 19 | 2026

Energy for Growth: Satellite Synthetic Control Evidence from Indonesia

Klaus Andresen[(a)], Ursula Müller[(a)] and Hans-Jörg Schmerer[(a)]

a FernUniversität in Hagen

FernUniversität in Hagen

Energy for Growth: Satellite Synthetic Control Evidence from Indonesia

Klaus Andresen, Ursula Müller, Hans-Jörg Schmerer *

February 22, 2026

Abstract

This study evaluates the causal impact of Indonesia's Power Transmission Development Project on regional economic development. Employing the Generalized Synthetic Control Method on district-level satellite data, we find a statistically significant increase in economic activity (p<0.001), which translates to a GDP increase of approximately 2 percent based on empirically derived nightlight-GDP elasticities. These findings provide rigorous ex-post evidence supporting the economic returns to infrastructure investments in middle-income countries, while demonstrating how satellite imagery combined with quasi-experimental methods can enable credible impact evaluation where traditional data are scarce.

JEL Classification: C23, H54, O11, O18

Keywords: Electricity infrastructure, economic development, nighttime lights, generalized synthetic control, Indonesia

*Emails: klaus.wulf-andresen@fernuni-hagen.de, ursulabmueller@gmail.com, hans-joerg.schmerer@fernuni-hagen.de.
We acknowledge the use of Claude AI (Anthropic) for assistance with programming, data preparation, and LaTeX formatting. AI tools were also employed to improve text clarity and readability. All substantive content, analysis, and conclusions remain the responsibility of the authors. An earlier version of this paper circulated under the title: Energy and Economic Growth: An Impact Evaluation of a Power Project in Indonesia.

1 Introduction

Reliable access to electricity is a cornerstone of economic development. Infrastructure investments reduce transaction costs, expand market access, and enable productivity gains across sectors (Straub, 2011; Servén and Calderón, 2004), and electricity infrastructure in particular has been shown to catalyze economic development by raising labor productivity and expanding employment opportunities (Dinkelman, 2011; Lipscomb et al., 2013). Yet credible causal evidence on the economic returns to specific electricity projects remains scarce, in large part because the data requirements for rigorous ex-post evaluation are difficult to meet in developing countries.

This paper evaluates the causal impact of Indonesia's Power Transmission Development Project on regional economic development. The project, financed by the World Bank with a loan commitment of 225 million US$ and completed in 2019, strengthened power transmission and distribution infrastructure in Java, Bali, and South Central Sumatra—regions that collectively account for more than 80 percent of Indonesia's GDP and 88 percent of its electricity demand.[1] The project comprised two integrated components. The Java-Bali component expanded four existing 500/150 kV substations and approximately 20–25 existing 150/20 kV substations, adding a total of 3,350 MVA of transformer capacity. The South Central Sumatra component upgraded five existing 150 kV substations and 10–15 existing 150/20 kV substations, adding 3,200 MVA.[2] These investments were designed to reduce transmission losses, improve grid stability, and support the growing electricity demand driven by Indonesia's sustained annual GDP growth of 5–6 percent since 2003. By project completion, electricity sales in the treated regions had substantially exceeded targets, reaching 200 TWh in Java and 40 TWh in Sumatra, transformer outage frequency had fallen by more than 99 percent in Java, and the total customer base in Java and Sumatra had expanded from 33.5 million in 2009 to 85.4 million in 2018. An independent evaluation commissioned by the World Bank assessed overall project performance as *moderately satisfactory*.[3] While capacity and reliability targets were met or exceeded, the evaluation acknowledged that increases in electricity sales and customer connections cannot be entirely attributed to the project, as complementary investments

[1] World Bank Indonesia Power Transmission Development (Project P 117323), https://projects.worldbank.org/en/projects-operations/project-detail/P117323.

[2] For concepts of electric power networks and energy transmission, see the International Electrotechnical Commission (IEC), https://www.electropedia.org/iev/iev.nsf/display?openform&ievref=601-01-09.

[3] World Bank Implementation Completion and Results Report, dated April 30, 2020, Report No. ICR00005098.

and broader economic trends likely played a role. This attribution challenge motivates our use of quasi-experimental methods to isolate the project's causal contribution to economic development. The World Bank's evaluation, like most ex-post project assessments by development institutions, compares realized outcomes against pre-defined targets—an approach that cannot distinguish the project's contribution from concurrent trends, complementary policies, or broader macroeconomic dynamics. Our analysis addresses this gap by constructing explicit counterfactuals through the Generalized Synthetic Control Method, estimating what economic activity in the treated regions would have been in the absence of the infrastructure investment. Figure 1 displays the geographic distribution of treated and non-treated provinces. Crucially for our identification strategy, the bulk of substation capacity was installed by 2017 (see Figure 2), which motivates our choice of treatment timing.

We employ the Generalized Synthetic Control Method (Xu, 2017) to estimate the project's effect on economic activity, measured through three complementary outcomes: provincial GDP, provincial nighttime light intensity, and district-level quarterly nighttime light data. The last of these exploits the spatial granularity of satellite imagery to achieve substantially greater statistical power than the provincial analyses. Our approach builds on a growing literature that uses satellite-derived nighttime light (NTL) data to measure economic activity where official statistics are unreliable or unavailable. Henderson et al. (2012) demonstrated that NTL provides a robust measure of true income growth, especially in countries with weak statistical capacity, and Chen and Nordhaus (2011) established the strong correlation between luminosity and economic output at subnational levels. Earlier work by Elvidge et al. (1997) and Doll et al. (2006) documented that changes in nighttime luminosity correlate with growth rates across development contexts. Donaldson and Storeygard (2016) provide a comprehensive overview of how satellite data can be applied in economics, highlighting the potential for measuring economic activity in data-scarce settings. Several studies have applied NTL specifically to infrastructure evaluation. Chakravorty et al. (2016) estimate both the welfare gains and infrastructure costs of rural electrification in the Philippines, finding that grid extension costs are typically recovered within a single year of realized benefits. Burlig and Preonas (2024), studying India's national rural electrification program, provide a more nuanced picture: while the program meaningfully expanded electricity access, economic impacts after three to five years were limited, with positive returns concentrated in larger villages. At a

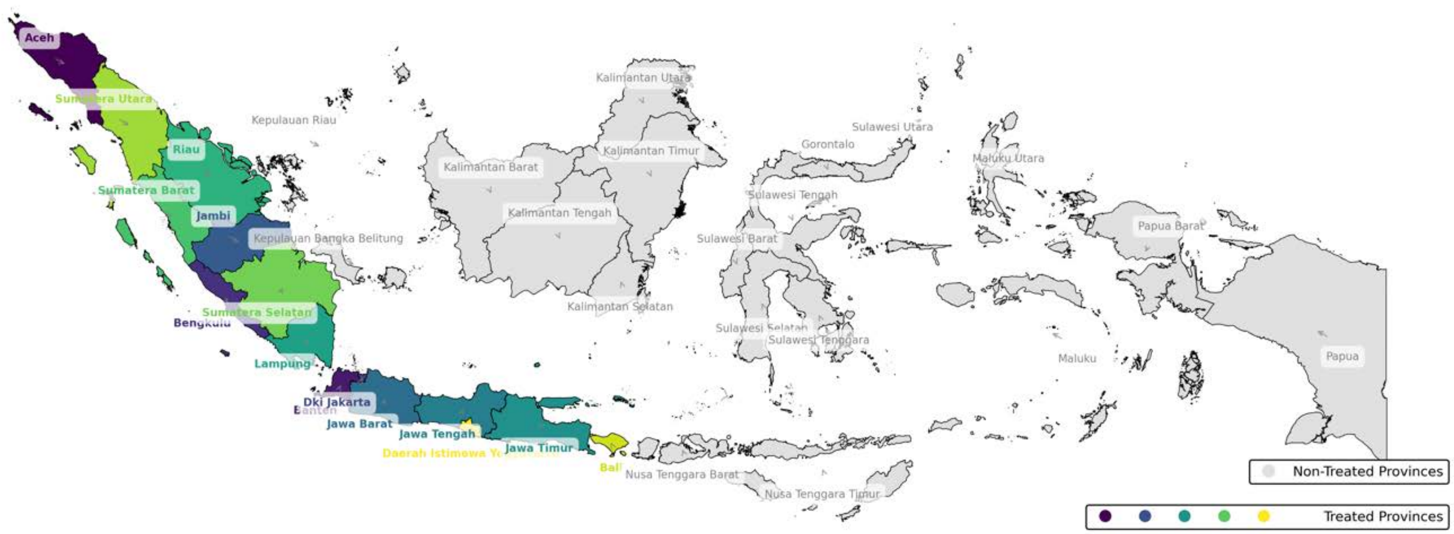

Figure 1: *Geographic Distribution of Treated and Non-Treated Provinces in Indonesia*

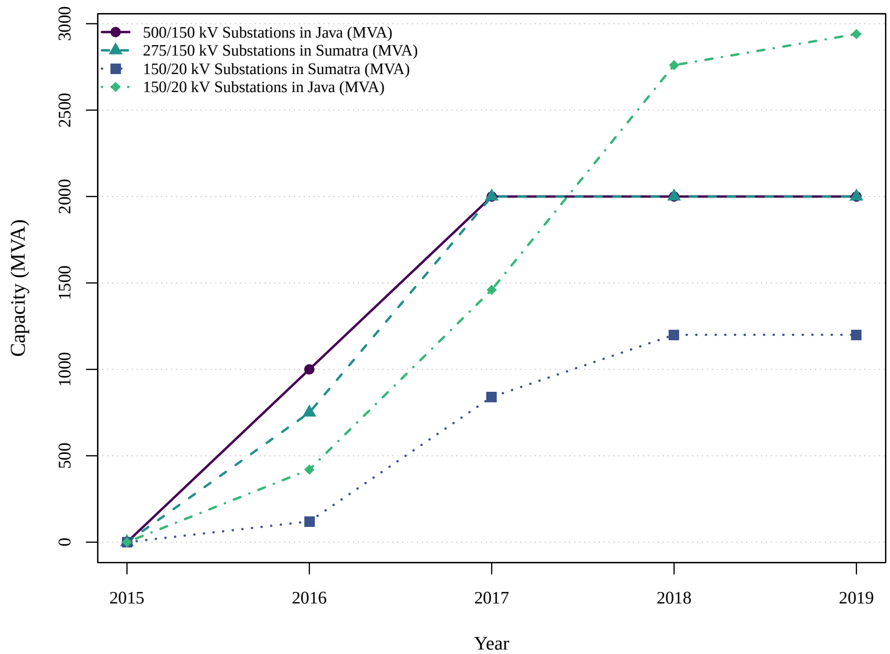

Figure 2: *Project Development (2015–2019)*
Source: Authors' calculations based on World Bank Group Implementation Status Results Reports for the Indonesia Power Transmission Development Project (P117323), Sequences 10–17.

larger scale, Bluhm et al. (2025) estimate the impact of Chinese infrastructure programs on the spatial distribution of economic activity worldwide, finding that these programs are associated with greater spatial decentralization. Hodler and Raschky (2014) use nighttime lights to document regional favoritism in the allocation of public resources across a large sample of countries.

An important methodological consideration concerns the choice of satellite data source. Gibson et al. (2021) demonstrate—specifically in the Indonesian context—that the Visible Infrared Imaging Radiometer Suite (VIIRS) offers substantial improvements over the older Defense Meteorological Satellite Program (DMSP) in capturing economic variation. The remote sensing literature broadly confirms the superiority of VIIRS (Elvidge et al., 2017; Chen and Nordhaus, 2019), though both data sources face limitations in low-density rural areas, where household lights are unlikely to be illuminated at the satellite's overpass time of approximately 1:30 a.m.[4] The use of NTL as a proxy for economic activity has been validated across diverse contexts, including China (Glawe

[4]Night lights have also been used as a proxy for electricity access; see Elvidge et al. (2011, 2017, 2020); Min et al. (2013); Min and Gaba (2014); Falchetta et al. (2019). Alternative remote sensing sources, such as daytime Landsat imagery, may serve as superior cross-sectional predictors (Goldblatt et al., 2020).

and Mendez, 2024; Glawe, 2025), and performs better in urbanized economies with moderate agricultural shares (Keola et al., 2015; Bickenbach et al., 2016).

Across all three specifications, we find consistent evidence of positive treatment effects. The district-level quarterly analysis—which provides the greatest statistical power with 169 treated units—yields a statistically significant average treatment effect ($p < 0.001$), which translates to a GDP increase of approximately 2 percent based on empirically derived NTL-GDP elasticities. This estimate is consistent in direction and order of magnitude with the direct provincial GRDP estimate of 1.08 percent, though the latter fails to reach statistical significance due to limited observations. The spatial distribution of district-level effects reveals substantial heterogeneity, with the largest positive effects concentrated near substation locations.

This study makes two contributions. First, it provides rigorous ex-post causal evidence on the economic returns to a major electricity infrastructure project in a middle-income country. While the existing literature has documented the benefits of electrification in general terms, credible project-level evaluations using quasi-experimental methods remain rare, particularly for transmission and distribution infrastructure as opposed to generation capacity or last-mile connections. Second, our methodological approach—combining the Generalized Synthetic Control Method with district-level satellite data and empirically derived NTL-GDP elasticities—offers a replicable framework for evaluating infrastructure projects in data-scarce environments. Since both nighttime light data and the global gridded GDP estimates we use to derive elasticities (Rossi-Hansberg and Zhang, 2025) are available worldwide, this approach is in principle applicable to infrastructure evaluation globally.

2 Dataset

Indonesia stands out for its remarkably comprehensive provincial-level data, allowing for detailed subnational analyses. Table 1 provides summary statistics for all variables used in our empirical exercise.

Our primary outcome variables are Gross Regional Domestic Product (GRDP), sourced from Badan Pusat Statistik (BPS), the official government statistical agency of Indonesia (Badan Pusat Statistik, 2025), and nighttime light (NTL) intensity from satellite imagery. GRDP is measured

Table 1: *Descriptive Statistics*

Variable	Mean	SD	Min	Max	N	T	Obs
GRDP (Provincial)	296679	423073	17120	2050466	34	12	408
RH-GDP (District)	1.8762	4.1061	0.0000	57.0539	522	11	5742
NTL (Provincial)	55517	80169	785	516011	34	12	408
NTL (District)	4998	6843	0	107939	522	151	78822
Labor Force	3931940	5603406	209971	25399990	34	12	408
Foreign Hotel Guests	59351	260371	0	2468625	34	12	408
Disaster Affected	130961	306057	0	2451635	34	12	408
Elec. Distributed	7190	12023	178	58564	32	11	351

Note: GRDP is in 2010 Constant Prices, Billion IDR. District GDP is in Billion USD. Nightlight Intensity is measured in nWatts$\cdot$cm$^{-2}\cdot$sr^{-1}. N is the number of cross-sectional units and T is the number of time periods. See text for details on data sources and methodology.

in 2010 constant prices at the provincial level. The nighttime light data stem from the Earth Observation Group, Payne Institute for Public Policy's Monthly and Yearly Cloud-free DNB Composite product (Elvidge et al., 2021). This dataset is available starting from 2012 when the VIIRS satellite became operational. For each province and district, we calculated the sum of light intensity values across all pixels within the respective administrative boundaries, using boundary information from the Humanitarian Data Exchange (Humanitarian Data Exchange, 2024). We addressed the administrative changes that occurred in 2023, when the original provinces of Papua and Papua Barat were further subdivided into six provinces, by aggregating data from the new provinces back to their parent provinces. We later identified data inconsistencies associated with this administrative change and excluded the Papua and Papua Barat provinces from the analysis. As a complementary measure of local economic activity, we use district-level GDP estimates derived from the global gridded GDP dataset compiled by Rossi-Hansberg and Zhang (2025). This dataset employs Random Forest models trained on high-resolution spatial data to predict local GDP distribution globally. The predictors include population density, nighttime light emissions, land use characteristics (urban areas, cropland, forest), CO_2 emissions from manufacturing, heavy industry, and transportation, as well as vegetation indices and terrain ruggedness. The model is trained on subnational GDP data from countries across North America, South America, Europe, Africa, and Asia, and then applied globally to generate consistent local GDP estimates, achieving an out-of-sample R^2 above 0.92 for GDP levels and above 0.62 for annual changes. We aggregated the gridded estimates to the district level by calculating the sum of GDP values across all grid cells within each district's administrative boundaries, using boundary information from the

Humanitarian Data Exchange (Humanitarian Data Exchange, 2024). The resulting dataset provides annual GDP estimates at the district level in billion USD, covering the period 2012–2022.

The provincial-level specifications include three additional covariates. Labor force size and total foreign hotel guests are sourced from BPS (Badan Pusat Statistik, 2025). The disaster-affected population is obtained from the *Data Informasi Bencana Indonesia* (DIBI) database (Badan Nasional Penanggulangan Bencana, 2024), provided by the Indonesian National Disaster Management Authority.[5] This variable captures the aggregate number of individuals affected by natural disasters, including fatalities, missing persons, injuries, and displacement due to disaster events.

3 Methodology

We employ the Generalized Synthetic Control Method (GSC) developed by Xu (2017).[6] The GSC method relaxes the parallel trends assumption required by standard difference-in-differences by explicitly modeling unobserved heterogeneity through a factor structure, which is permitted to be correlated with treatment assignment. For our application, an additional advantage is that it accommodates settings where observed covariates are unavailable—as is the case with granular satellite data—by capturing unobserved heterogeneity through latent factors. Furthermore, the unique geographic characteristic of Indonesia, where treated areas (Java and Sumatra) and control areas (other islands) are separated by ocean, minimizes the problem of spillover effects that may cause serious estimation bias.

The assumed functional form in the generalized synthetic control approach is given by:

$$Y_{it} = \delta_{it} D_{it} + X_{it}' \beta + \lambda_i' f_t + \varepsilon_{it} \tag{1}$$

where Y_{it} is the outcome variable for unit i at time t, δ_{it} represents the heterogeneous treatment effect, D_{it} is the treatment indicator, X_{it}' is a vector of observed covariates, β is a vector of unknown parameters, f_t denotes unobserved common factors, λ_i represents unit-specific factor loadings,

[5] Badan Nasional Penanggulangan Bencana (BNPB).

[6] Despite its name, the GSC method differs substantially from the original synthetic control method of Abadie et al. (2010). Rather than constructing explicit weights on control units to form a synthetic counterfactual, GSC uses control units only to estimate a common factor structure; counterfactuals for treated units are then constructed using unit-specific factor loadings estimated from pre-treatment treated outcomes alone.

and ε_{it} is the error term. The method requires strict exogeneity of the idiosyncratic error term:

$$\varepsilon_{it} \perp D_{js}, X_{js}, \lambda_j, f_s \quad \forall\, i, j, t, s. \tag{2}$$

In our setting, selection into the electricity infrastructure program was based on provincial development priorities and economic potential. Such selection criteria reflect differential exposure to common economic trends—heterogeneity that the factor structure $\lambda_i' f_t$ can capture. For instance, if more developed areas were selected for early electrification, and development levels determine how strongly a region responds to national economic shocks, then this selection mechanism operates through the factor loadings λ_i rather than the idiosyncratic error ε_{it}. The strict exogeneity assumption thus requires only that treatment assignment is uncorrelated with unmodeled, unit-specific shocks—a weaker requirement than the parallel trends assumption imposed by standard difference-in-differences.

The core intuition of the GSC method is as follows. First, a common factor structure is estimated using only the control units, which by assumption share the same underlying factors as the treated units. Second, factor loadings for each treated unit are estimated by regressing the unit's pre-treatment outcomes onto the estimated factors via OLS. These loadings capture how strongly each treated unit responds to the common factors. Finally, the estimated factors and treated-unit loadings are combined to construct counterfactual outcomes for treated units in post-treatment periods. Importantly, control units serve only to identify the factor space—their individual factor loadings are not used in constructing treated counterfactuals.

Formally, the GSC estimator proceeds in three steps. In the first step, an interactive fixed effects model is estimated using only control group data, minimizing the sum of squared residuals to obtain $\hat{\beta}$, $\hat{F}$, and factor loadings for control units $\hat{\Lambda}_{co}$:

$$(\hat{\beta}, \hat{F}, \hat{\Lambda}_{co}) = \arg\min_{\tilde{\beta}, \tilde{F}, \tilde{\Lambda}_{co}} \sum_{i \in \mathcal{C}} (Y_i - X_i\tilde{\beta} - \tilde{F}\tilde{\lambda}_i)'(Y_i - X_i\tilde{\beta} - \tilde{F}\tilde{\lambda}_i) \tag{3}$$

This model is estimated via an iterative procedure closely related to principal components analysis: conditional on β, the factors and loadings are extracted as the principal components of the residual matrix that explain maximal variance. While the control units' factor loadings $\hat{\Lambda}_{co}$ are essential for identifying the factor space $\hat{F}$, they are not used in subsequent steps. In the second step, factor loadings for each treated unit are estimated via OLS:

$$\hat{\lambda}_i = (\hat{F}_0'\hat{F}_0)^{-1}\hat{F}_0'(Y_i^0 - X_i^0\hat{\beta}), \quad i \in \mathcal{T} \tag{4}$$

where the superscript 0 denotes pre-treatment periods. In the third step, counterfactual outcomes for treated units in post-treatment periods are constructed as $\hat{Y}_{it}(0) = X_{it}'\hat{\beta} + \hat{\lambda}_i'\hat{f}_t$. The treatment effect for each treated unit and time period is estimated as:

$$\hat{\delta}_{it} = Y_{it}(1) - \hat{Y}_{it}(0) \tag{5}$$

where $Y_{it}(1)$ is the observed outcome and $\hat{Y}_{it}(0)$ is the estimated counterfactual. The Average Treatment Effect on the Treated (ATT) is obtained by averaging these unit-specific effects across treated units and post-treatment periods.

The number of factors is determined through cross-validation. For a given number of factors r, the method first estimates the interactive fixed effects model using only control unit data. The algorithm then iterates through each pre-treatment period s: treated unit observations at time s are held out, factor loadings for treated units are estimated via OLS using the remaining pre-treatment data, and predictions are generated for the held-out treated observations. The optimal number of factors r^* is selected by minimizing the mean squared prediction error across all held-out predictions.

Uncertainty estimates are obtained through a parametric bootstrap procedure. The procedure first constructs an empirical distribution of prediction errors by iteratively leaving out each control unit, resampling the remaining controls with replacement, applying the GSC method to this resampled dataset, and computing the prediction error for the left-out unit. Bootstrapped samples are then generated by adding resampled residuals to fitted values, drawing from two distinct empirical distributions: in-sample residuals of the IFE model for control units, and the constructed prediction errors for treated units. Because the IFE model is estimated using only control group data, it fits control units better than treated units, and the prediction error distribution is therefore wider than the in-sample residual distribution. The GSC method is applied to each bootstrapped sample to generate a distribution of ATT estimates, from which standard errors and confidence intervals are computed.

We extend the standard GSC procedure with an additional screening step, not part of Xu (2017), to ensure the validity of the common factor assumption for each treated unit. In a first stage, we estimate the GSC model using all treated units. For each treated unit and each pre-treatment period, the `gsynth` package reports a period-specific effect estimate $\hat{\delta}_{it}$ along with a bootstrap standard error $\widehat{\mathrm{SE}}_{it}^{\mathrm{boot}}$ obtained from the parametric bootstrap procedure. A two-sided p-value is then computed under a normal approximation as $p_{it} = 2[1 - \Phi(|\hat{\delta}_{it} / \widehat{\mathrm{SE}}_{it}^{\mathrm{boot}}|)]$, where Φ denotes the standard normal CDF. Treated units for which any pre-treatment period exhibits a significant deviation at the 5% level are excluded, as this indicates that the common factor structure does not adequately capture these units' outcome dynamics. In a second stage, we re-estimate the GSC model using only the treated units with adequate pre-treatment fit, while maintaining the identical control group. This ensures that the ATT is estimated using only units for which the counterfactual is credibly constructed.

4 Analysis

The causal relationships between the variables used in our study can be represented by the following graph.

$$
\begin{tikzcd}
 & \text{Quantity} \arrow[r] \arrow[ddr] & \text{GDP} \\
\text{Electricity Infrastructure} \arrow[ur] \arrow[dr] & & \\
 & \text{Stability} \arrow[r] \arrow[uur] & \text{NTL}
\end{tikzcd}
$$

The main aim of the project is to increase the potential quantity (distributed GWh) and stability (less power outages) of electricity supply, which should be visible in systematic changes in the night light graphs used in our analysis. Output and nighttime lights do not have a direct causal link between each other but share the same mediators[7], which leads to the strong correlations reported in the literature.

[7] Confounders are controlled for by inclusion of covariates and the linear factors in our model

4.1 Impact on the Gross Regional Domestic Product

To estimate the causal effect of the electricity infrastructure program on economic development, we employ the generalized synthetic control method (Xu, 2017), setting the treatment date to 2017 when most of the project was completed (Gregan, 2017). The outcome variable is Gross Regional Domestic Product (GRDP) measured in 2010 constant market prices (billion rupiahs). As control variables, we include the size of the labor force, the number of foreign hotel guests, and the population affected by disasters. These covariates are included to capture province-specific shocks that may not be fully absorbed by the common factor structure. An optimal number of $r^* = 2$ linear factors was determined via cross-validation.

To ensure the validity of our causal estimates, we assess the quality of pre-treatment fit for each treated province by examining whether the estimated treatment effects in the pre-intervention period are statistically indistinguishable from zero. Three provinces—Jawa Tengah, Jawa Barat, and Riau—exhibit poor pre-treatment fit, with one or more pre-treatment periods showing statistically significant deviations from their synthetic counterparts ($p < 0.05$). Following best practices in synthetic control methods, we exclude these provinces from our analysis, focusing on the 12 provinces with good pre-treatment fit. This approach maintains the integrity of the control group while improving the validity of the parallel trends assumption.

Figure 3 presents the covariate coefficient estimates, the estimated treatment effects over time, and the average treatment effect. The coefficient associated with the labor force variable indicates a positive significant effect of labor force size on GRDP ($p < 0.001$), while the number of hotel guests and the population affected by disasters are not statistically significant predictors.

The middle panel of Figure 3 compares the treatment effect estimates between the full sample (dashed line with circles) and the filtered sample with good pre-treatment fit (solid line with triangles). The results reveal a mildly positive effect of the electricity infrastructure program on GRDP in the initial post-treatment period (2017-2018), with point estimates suggesting modest increases in provincial economic output immediately following project completion. However, this effect is accompanied by stark increasing uncertainty over time, as evidenced by the widening confidence intervals in later periods (2019-2023). The average treatment effect across all post-treatment periods is 5,194.36 billion rupiahs (approximately 1.08% of mean counterfactual GRDP),

Figure 3: *Estimated Effect on GRDP*

Covariate	Estimate	S.E.	95% CI		p-value
Labor force	0.0216***	0.0015	0.0187	0.0245	<0.001
Total hotel guests	0.0149	0.0147	−0.0138	0.0437	0.309
Disaster affected	−0.0014	0.0011	−0.0036	0.0008	0.218

Notes: Significance levels: $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. Sample restricted to 12 treated provinces with good pre-treatment fit.

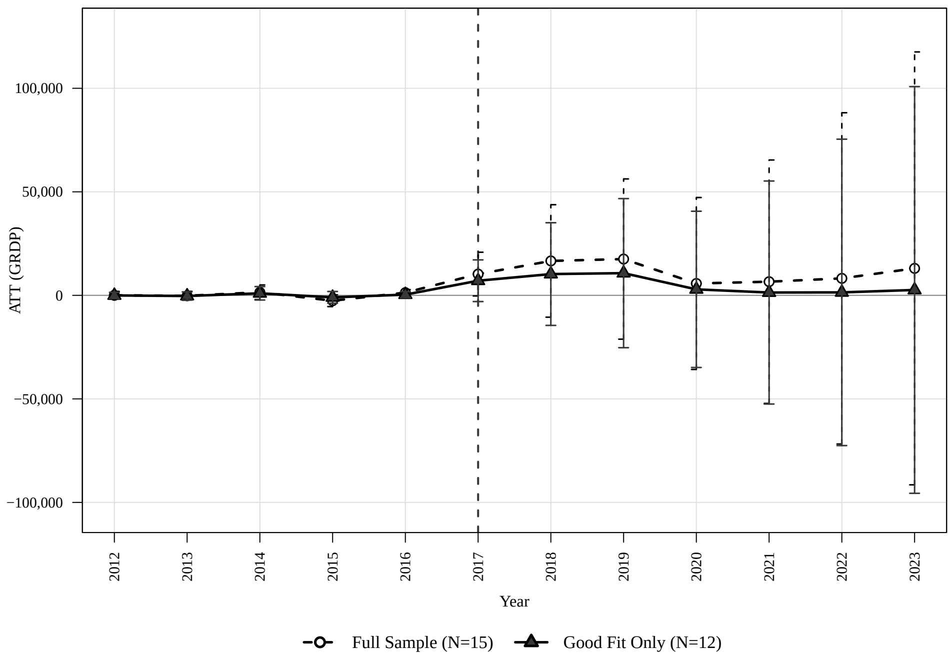

Average Treatment Effect on GRDP				
Estimate	S.E.	95% CI		p-value
5,194.36	23,986.99	-41,819.27	52,207.99	0.829

Note: Comparison of estimated treatment effects between the full sample (N=15, dashed line) and the filtered sample with good pre-treatment fit (N=12, solid line). Vertical line marks the treatment year (2017). Error bars represent 95% confidence intervals. Three provinces (Jawa Tengah, Jawa Barat, Riau) were excluded due to poor pre-treatment fit. Average treatment effect across all post-treatment periods (2017-2023). Sample restricted to 12 treated provinces with good pre-treatment fit.

but this estimate is not statistically significant (p = 0.829, 95% CI: [-41,819.27, 52,207.99]).

The increasing uncertainty over time can be attributed to several factors. First, as we move further from the treatment period, the influence of unobserved confounders and secular trends becomes more pronounced, making it challenging to isolate the causal effect of the infrastructure program. Second, the economic disruptions caused by the COVID-19 pandemic starting in 2020 introduce additional volatility that complicates the estimation of infrastructure effects. Third, the diverging post-treatment trajectories across treated provinces may reflect heterogeneous treatment effects driven by local economic conditions, complementary infrastructure investments, or differential exposure to external shocks.

4.2 Impact on Nighttime Light (Yearly - Provincial)

As a complementary measure of economic activity, we examine the impact of the electricity infrastructure program on nighttime light (NTL) intensity using satellite imagery data. Following the same methodological approach as for GRDP, we employ the generalized synthetic control method (Xu, 2017) with the treatment date set to 2017. The outcome variable is the sum of nighttime light intensity across each province, which serves as a proxy for economic development and electrification (Henderson et al., 2012). We include the same control variables: labor force size, total foreign hotel guests, and disaster-affected population. Cross-validation selected an optimal number of $r^* = 1$ latent factors.

Following our pre-treatment fit assessment procedure, six provinces exhibited significant deviations from their synthetic counterparts during the pre-intervention period ($p < 0.05$), indicating poor counterfactual estimation. These provinces—Jawa Barat, Jawa Tengah, Jawa Timur, Sumatera Barat, Sumatera Selatan, and Sumatera Utara—were excluded from the analysis, leaving 9 treated provinces with good pre-treatment fit. Notably, the excluded provinces include Java's three largest provinces by population and economic output, suggesting that the synthetic control method faces particular challenges in constructing valid counterfactuals for these dominant economic centers using NTL data. This more stringent filtering (compared to the GRDP analysis, which excluded only three provinces) reflects the greater variability in nighttime light measurements across provinces.

Figure 4: *Estimated Effect on Nighttime Light*

Covariate	Estimate	S.E.	95% CI		p-value
Labor force	0.0176^{***}	0.0047	0.0084	0.0268	0.0002
Total hotel guests	-0.0181	0.0181	-0.0537	0.0174	0.317
Disaster affected	0.0006	0.0016	-0.0025	0.0038	0.682

Notes: Significance levels: $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. Sample restricted to 9 treated provinces with good pre-treatment fit.

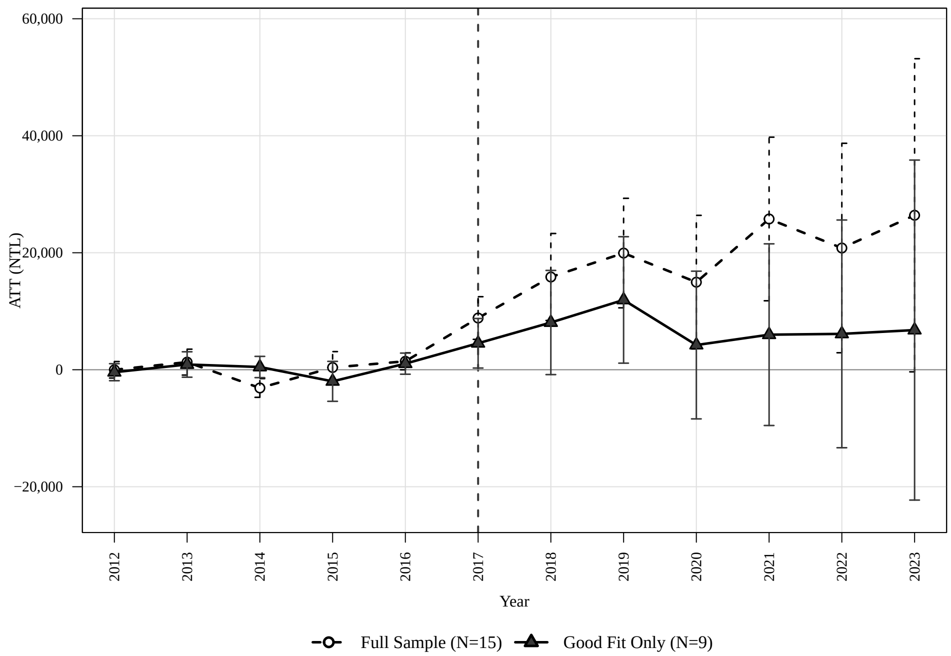

Average Treatment Effect on NTL				
Estimate	S.E.	95% CI		p-value
6,808.96	6,755.75	-6,432.06	20,049.98	0.314

Note: Comparison of estimated treatment effects between the full sample (N=15, dashed line) and the filtered sample with good pre-treatment fit (N=9, solid line). Vertical line marks the treatment year (2017). Error bars represent 95% confidence intervals. Six provinces (Jawa Barat, Jawa Tengah, Jawa Timur, Sumatera Barat, Sumatera Selatan, Sumatera Utara) were excluded due to poor pre-treatment fit. Average treatment effect across all post-treatment periods (2017-2023). Sample restricted to 9 treated provinces with good pre-treatment fit.

Figure 4 presents the covariate coefficient estimates, treatment effect trajectories, and average treatment effect. Consistent with the GRDP results, labor force size exhibits a statistically significant positive relationship with nighttime light intensity ($p < 0.001$), while total hotel guests and disaster-affected population are not significant predictors.

The middle panel of Figure 4 displays the estimated treatment effects over time, comparing the full sample (dashed line) with the filtered sample of provinces with good pre-treatment fit (solid line). The results indicate a positive but modest effect of the infrastructure program on nighttime light intensity in the initial post-treatment years, with point estimates suggesting increased luminosity following project completion. However, similar to the GRDP analysis, uncertainty increases substantially in later periods, particularly after 2020. The average treatment effect across all post-treatment periods is 6,808.96 units of light intensity, but this estimate is not statistically significant (p = 0.314, 95% CI: [-6,432.06, 20,049.98]).

The convergence of findings between NTL and GRDP analyses—both showing positive but statistically insignificant treatment effects with increasing uncertainty over time—strengthens the robustness of our conclusions. The lack of statistical significance may reflect several factors: the relatively short post-treatment observation window, the confounding influence of the COVID-19 pandemic on economic activity after 2020, heterogeneous treatment effects across provinces with varying baseline infrastructure and economic conditions, and the possibility that the economic benefits of electricity infrastructure may manifest more gradually than our observation period allows. Additionally, NTL as a proxy measure may not fully capture the distributional benefits of rural electrification, which could improve household welfare without substantially increasing aggregate light emissions visible from satellite imagery.

4.3 Impact on Nighttime Light (Quarterly - District)

To exploit the spatial granularity of satellite imagery and increase statistical power, we extend our analysis to the district level using quarterly nighttime light data. This approach yields a substantially larger sample of treated units compared to the provincial analyses, allowing for more precise estimation of treatment effects. We employ the generalized synthetic control method (Xu, 2017) with quarterly observations from 2012 to 2023 and the treatment date set to 2017Q1. The outcome variable is the sum of nighttime light intensity within each district,

Figure 5: *Estimated Effect on Nighttime Light (District-Level)*

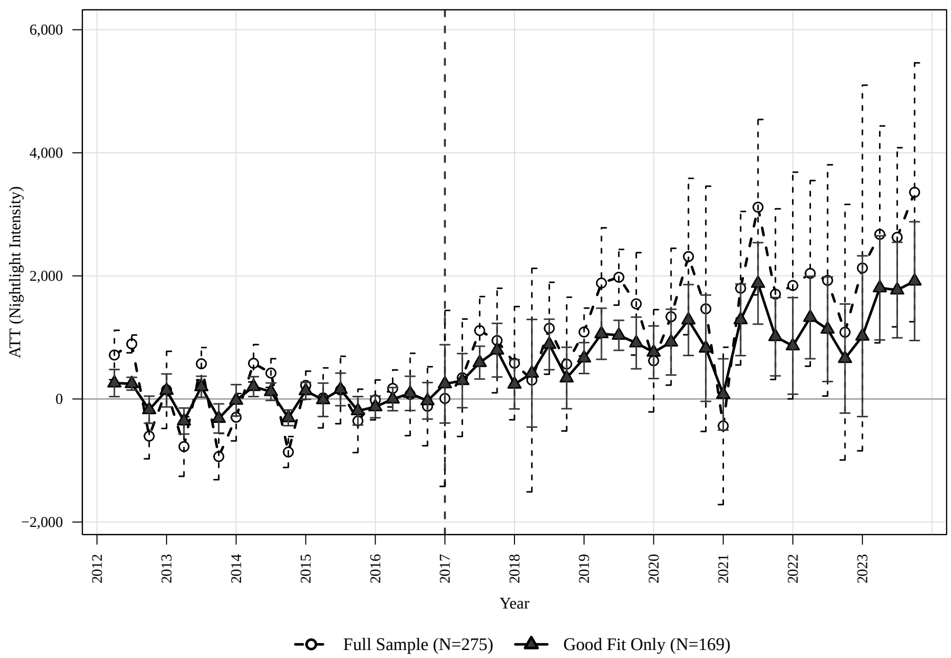

Average Treatment Effect on NTL				
Estimate	S.E.	95% CI		p-value
928.92	273.22	393.41	1,464.43	0.0007***

Note: Comparison of estimated treatment effects between the full sample (N=275, dashed line) and the filtered sample with good pre-treatment fit (N=169, solid line). Vertical line marks the treatment year (2017). Error bars represent 95% confidence intervals. 106 districts were excluded due to poor pre-treatment fit. Significance levels: $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. Average treatment effect across all post-treatment periods (2017Q1–2023Q4). Sample restricted to 169 treated districts with good pre-treatment fit.

aggregated as quarterly averages from monthly observations. We include unit fixed effects, since district-level panel data for socioeconomic covariates are not available at this temporal frequency. Cross-validation selected an optimal number of $r^* = 2$ latent factors.

Following our pre-treatment fit assessment procedure, 106 of the 275 treated districts exhibited significant deviations from their synthetic counterparts during the pre-intervention period ($p < 0.05$), indicating poor counterfactual estimation. These districts were excluded from the analysis, leaving 169 treated districts with good pre-treatment fit across 15 provinces. The exclusion rate is lower compared to the provincial analyses (39% vs. 40% for provincial NTL).

The average treatment effect across all post-treatment quarters is 928.92 units of light intensity, and this estimate is statistically significant ($p < 0.001$, 95% CI: [393.41, 1464.43]). This finding is consistent with the direction of effects observed in both the GRDP and provincial NTL analyses, which showed positive but statistically insignificant point estimates. The increased precision at the quarterly-district level suggests that the yearly-provincial analyses may have been underpowered to detect what is a modest but real effect of the infrastructure program on local economic activity.

Figure 6: *Spatial Distribution of District-Level Treatment Effects*

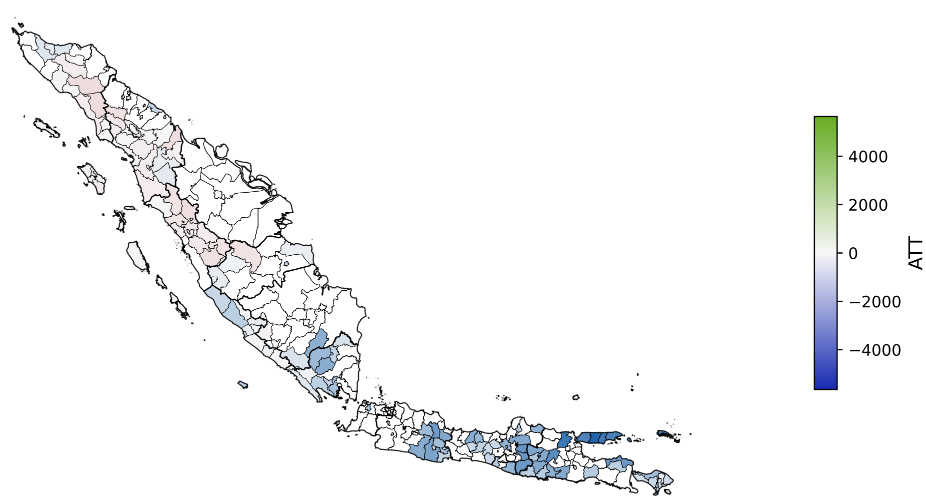

Note: Average treatment effects on nighttime light intensity for districts with good pre-treatment fit (N=169). Colored districts show estimated effects ranging from -666 (red) to $+5,642$ (blue). Districts with boundaries only were excluded due to poor pre-treatment fit.

Figure 6 displays the spatial distribution of treatment effects across districts. The 106 districts shown with boundaries only were excluded from the analysis due to poor pre-treatment fit, highlighting areas where the synthetic control method could not construct valid counterfactuals. The map reveals substantial heterogeneity in the program's impact, with estimated effects ranging from -666 to $+5,642$ units of light intensity. The largest positive effects are concentrated in districts where substation expansions were physically located, suggesting that the direct infrastructure investments drove measurable increases in local economic activity. Our treatment definition classifies all districts on Java and Sumatra as treated, on the assumption that grid-level improvements generate spillover effects across the interconnected transmission network — for instance, capacity additions at one node may relieve congestion elsewhere in the system. The spatial pattern of results, however, suggests that these indirect spillovers may be limited or even negative in some districts. One possible interpretation is that improved infrastructure in directly treated areas attracted economic activity away from neighboring districts, though we can only speculate about the mechanisms underlying the negative estimates.

All three specifications—GRDP, yearly-provincial NTL, and quartely-district NTL—point toward a positive effect of the electricity infrastructure program on economic activity. The district-level analysis, by exploiting the spatial granularity inherent in nighttime light data, provides sufficient statistical power to detect what the provincial analyses could only suggest: a modest but statistically significant positive impact of Indonesia's electricity infrastructure program on local economic development.

4.4 Interpreting the Results

To express the nighttime light treatment effects in terms of GDP, we estimate the elasticity between NTL and GDP at the district level using local GDP estimates from Rossi-Hansberg and Zhang (2025).[8] Panel A of Table 2 presents both pooled OLS and two-way fixed effects specifications. We use the two-way fixed effects estimate of 0.057, which captures within-district variation over time while controlling for common shocks.

Panel B reports the overall treatment effects. The district-level NTL analysis yields an ATT of 41.23% relative to the counterfactual. To convert this to an implied GDP effect, we apply the exact

[8]For calculating the ATT, we use NTL rather than high-resolution GDP estimates directly because the latter are only available through 2022, whereas NTL data are up to date.

formula for percentage changes derived from the log-log elasticity specification:

$$\%\Delta GDP = (1 + \%\Delta NTL)^{\varepsilon} - 1 = (1.4123)^{0.057} - 1 \approx 0.0197, \tag{6}$$

which implies a GDP increase of approximately 2%. For comparison, the direct province-level GRDP estimates yield an overall effect of 1.08%. The slightly lower province-level estimate likely reflects downward bias from the asymmetric COVID-19 shock, as evidenced by the sharp decline in GRDP ATT estimates beginning in period 4 (2020). The district-level NTL analysis, with its larger sample of 169 treated units compared to 12 provinces, provides greater statistical power and is less susceptible to this bias.

Table 2: *Elasticity Estimates and Treatment Effects*

Panel A: Nightlight Intensity and District GDP (Rossi-Hansberg and Zhang, 2025)

	Dependent variable: Log(GDP)	
	Pooled OLS	Two-way FE
Log(Nightlight Intensity)	0.592***	0.057**
	(0.031)	(0.023)
District Fixed Effects	No	Yes
Year Fixed Effects	No	Yes
Observations	4,232	4,232
R^2	0.510	0.011

Panel B: Overall Treatment Effects

	ATT	Counterfactual	ATT/Counterfactual
NTL (district-level)	928.92	2,253.26	0.4123
GRDP (province-level)	5,194.36	479,517.11	0.0108
		Implied GDP Effect	
From NTL via elasticity		2.0%	
Direct GRDP estimate		1.08%	

Note: Panel A shows regression results relating district-level nighttime light intensity to district-level GDP using estimates from Rossi-Hansberg and Zhang (2025). Robust standard errors (HC1) for pooled OLS; standard errors clustered by district for two-way FE. Panel B reports overall ATT as the average ATT divided by the mean post-treatment counterfactual. The NTL sample includes 169 treated districts with good pre-treatment fit. The GRDP sample includes 12 treated provinces with good pre-treatment fit (excluding Jawa Tengah, Jawa Barat, Riau). The implied GDP effect from NTL is calculated using $(1 + \%\Delta NTL)^{0.057} - 1$. *p<0.1; **p<0.05; ***p<0.01.

5 Discussion

This study employed the generalized synthetic control method to evaluate the causal impact of Indonesia's electricity infrastructure program on regional economic development. Our findings across three complementary specifications—provincial GRDP, provincial nighttime light intensity, and district-level nighttime light intensity—converge on a consistent narrative: the infrastructure program generated positive effects on economic activity, though the magnitude and statistical precision of these estimates vary across analytical approaches.

5.1 Interpretation

The district-level quarterly analysis provides the most precise and statistically robust evidence of the program's impact. With an average treatment effect of 928.92 units of nighttime light intensity ($p < 0.001$), this specification demonstrates that the infrastructure investment produced measurable increases in local economic activity. Translating this effect through the estimated NTL-GDP elasticity of 0.057 implies a GDP increase of approximately 2.0 percent, while the direct provincial GRDP estimate of 1.08 percent points in the same direction at a similar order of magnitude. The convergence of these independently derived estimates—one from satellite imagery and one from official economic statistics—strengthens confidence in our findings. The provincial-level analyses, while pointing in the same positive direction, failed to achieve statistical significance. The GRDP analysis (ATT = 5,194.36 billion rupiahs, $p = 0.829$) and provincial NTL analysis (ATT = 6,808.96, $p = 0.314$) both suffered from limited statistical power due to the small number of treated units (12 and 9 provinces, respectively, after excluding units with poor pre-treatment fit). The district-level analysis, by contrast, leveraged 169 treated units, providing substantially greater precision for detecting what appears to be a real but modest treatment effect. The spatial distribution of district-level treatment effects reveals substantial heterogeneity in the program's impact across regions. Estimated effects range from -666 to $+5,642$ units of nighttime light intensity, with positive effects more pronounced in Java and southern Sumatra. This heterogeneity likely reflects variation in local conditions that mediate the economic returns to infrastructure investment, including baseline electrification levels, complementary infrastructure, and economic structure.

5.2 Limitations

Several limitations temper the conclusions that can be drawn from this analysis. First, a notable challenge was the high rate of pre-treatment fit failures, particularly in the nighttime light specifications. Six provinces were excluded from the provincial NTL analysis, and 106 of 275 districts were excluded from the district-level analysis due to statistically significant deviations from their synthetic counterparts during the pre-intervention period. The excluded provinces include Java's three largest provinces by population and economic output. Consequently, our estimates may better capture the treatment effect for mid-sized provincial economies and semi-urban districts, while the program's impact on Indonesia's largest economic centers remains uncertain. Secondly, the COVID-19 pandemic, which disrupted economic activity globally beginning in early 2020, presents a significant challenge for interpretation. The pandemic's effects are visible in the widening confidence intervals observed in later post-treatment periods across all specifications. Indonesia implemented substantial mobility restrictions and experienced significant economic contraction in 2020. Since treated provinces are more urbanized on average than the synthetic controls, the pandemic induced asymmetric negative shocks that confound treatment effect estimation, explaining the decline in GRDP treatment effect estimates beginning in 2020. This asymmetry is less pronounced using nighttime lights. Thirdly, Nighttime light intensity is an imperfect proxy for economic activity. While the correlation between NTL and GDP is well-documented (Henderson et al., 2012), rural electrification may improve household welfare through non-luminous channels—refrigeration, productive equipment, communication devices—that satellite imagery cannot capture. Additionally, our analysis cannot distinguish between the intensive margin (improved reliability for connected users) and the extensive margin (new connections), which likely have different economic implications. Finally, the treatment date of 2017 is an approximation based on when most project components became operational. Infrastructure was completed on a rolling basis, with some components operational before 2017 and others continuing through 2019. This temporal imprecision may attenuate estimated treatment effects.

5.3 Methodological Contribution

Our analysis demonstrates how satellite imagery can unlock the full potential of modern causal inference methods for infrastructure evaluation. The generalized synthetic control method requires both many cross-sectional units and many time periods to reliably estimate the latent factor structure that underpins valid counterfactual construction. Official economic statistics rarely satisfy both requirements: Indonesia's GRDP data are annual and exist only at the provincial level, yielding 12 treated provinces observed over 5 pre-treatment years. Nighttime light data, by contrast, are available at high spatial and temporal resolution. By moving to district-level quarterly observations, we increase the sample from 12 provinces to 169 districts and from 5 annual to 20 quarterly pre-treatment periods, which enables us to detect a modest but real treatment effect that aggregate analyses could only suggest. To translate nighttime light effects into GDP terms, we estimate elasticities using the global gridded GDP dataset compiled by Rossi-Hansberg and Zhang (2025). Since both nighttime light data and the Rossi-Hansberg GDP estimates are available globally, our approach is in principle applicable to infrastructure evaluation anywhere in the world. The key advantage of using nighttime lights as the primary outcome—rather than gridded GDP estimates directly—is that NTL data are published continuously, whereas gridded GDP products are released only at intervals. This allows researchers to evaluate recent projects without waiting for updated GDP estimates.

6 Conclusion

This study evaluates the causal impact of Indonesia's electricity infrastructure program on regional economic development using the generalized synthetic control method. Across three specifications—provincial GRDP, provincial nighttime light intensity, and district-level quarterly nighttime light intensity—we find consistent evidence of positive treatment effects. The district-level analysis, which provides the greatest statistical power with 169 treated units, yields a statistically significant average treatment effect ($p < 0.001$) implying a GDP increase of approximately 2.0%. This estimate is consistent in direction with the direct provincial GDP estimate of 1.08%, though the latter fails to reach statistical significance due to the limited number of observations. Our analysis highlights methodological challenges for synthetic control applications in infrastructure evaluation. The high rate of pre-treatment fit failures—particularly for Indonesia's most populous and economically dominant provinces—reflects the difficulty of constructing valid counterfactuals when treated units are structurally dissimilar from the available donor pool. The spatial distribution of district-level effects reveals substantial heterogeneity, with positive effects more pronounced in Java and southern Sumatra, while northern Sumatran districts show more mixed results. These findings provide evidence that Indonesia's electricity infrastructure program contributed positively to regional economic development in the years following its implementation. Future research with longer post-treatment windows will be valuable for assessing whether these effects persist or grow over time.

References

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, 105(490):493–505.

Badan Nasional Penanggulangan Bencana (2024). Data informasi bencana indonesia (dibi). National Disaster Management Authority of Indonesia, Bidang Pengelolaan Data dan Sistem Informasi (PDSI), Pusdatinkom. Available at: https://dibi.bnpb.go.id/baru.

Badan Pusat Statistik (2025). Statistics indonesia official website. Official Government Statistical Agency of Indonesia. Available at: https://www.bps.go.id/id.

Bickenbach, F., Bode, E., Nunnenkamp, P., and Söder, M. (2016). Night lights and regional gdp. *Review of World Economics*, 152:425–447.

Bluhm, R., Dreher, A., Fuchs, A., Parks, B. C., Strange, A. M., and Tierney, M. J. (2025). Connective financing: Chinese infrastructure projects and the diffusion of economic activity in developing countries. *Journal of Urban Economics*, 145:103730.

Burlig, F. and Preonas, L. (2024). Out of the darkness and into the light? development effects of rural electrification. *Journal of Political Economy*, 132(9):2937–2971.

Chakravorty, U., Emerick, K., and Ravago, M.-L. (2016). Lighting up the last mile: The benefits and costs of extending electricity to the rural poor. *Resources for the Future Discussion paper*, pages 16–22.

Chen, X. and Nordhaus, W. D. (2011). Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences*, 108(21):8589–8594.

Chen, X. and Nordhaus, W. D. (2019). VIIRS nighttime lights in the estimation of cross-sectional and time-series gdp. *Remote Sensing*, 11(9):1057.

Dinkelman, T. (2011). The effects of rural electrification on employment: New evidence from south africa. *American Economic Review*, 101(7):3078–3108.

Doll, C. N., Muller, J.-P., and Morley, J. G. (2006). Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics*, 57(1):75–92.

Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4):171–198.

Elvidge, C. D., Baugh, K., Zhizhin, M., Hsu, F. C., and and, T. G. (2017). VIIRS night-time lights. *International Journal of Remote Sensing*, 38(21):5860–5879.

Elvidge, C. D., Baugh, K. E., Kihn, E. A., Kroehl, H. W., Davis, E. R., and Davis, C. W. (1997). Relation between satellite observed visible-near infrared emissions, population, economic activity and electric power consumption. *International Journal of Remote Sensing*, 18(6):1373–1379.

Elvidge, C. D., Baugh, K. E., Sutton, P. C., Bhaduri, B., Tuttle, B. T., Ghosh, T., Ziskin, D., and Erwin, E. H. (2011). Who's in the dark-satellite based estimates of electrification rates. In *Urban remote sensing: Monitoring, synthesis and modeling in the urban environment*, pages 211–224. Wiley Blackwell.

Elvidge, C. D., Hsu, F.-C., Zhizhin, M., Ghosh, T., Taneja, J., and Bazilian, M. (2020). Indicators of electric power instability from satellite observed nighttime lights. *Remote Sensing*, 12(19):3194.

Elvidge, C. D., Zhizhin, M., Ghosh, T., Hsu, F.-C., and Taneja, J. (2021). Annual time series of global VIIRS nighttime lights derived from monthly averages: 2012 to 2019. *Remote Sensing*, 13(5):922.

Falchetta, G., Pachauri, S., Parkinson, S., and Byers, E. (2019). A high-resolution gridded dataset to assess electrification in sub-saharan africa. *Scientific data*, 6(1):110.

Gibson, J., Olivia, S., Boe-Gibson, G., and Li, C. (2021). Which night lights data should we use in economics, and where? *Journal of Development Economics*, 149:102602.

Glawe, L. (2025). Nonlinearities in the institutions-growth relationship in a dynamic panel data framework: Evidence from china's provinces. *China Economic Review*, 94:102583.

Glawe, L. and Mendez, C. (2024). Harmonized luminosity and economic activity across provinces in china: cross-sectional differences, regional time series, and inequality dynamics. *Applied Economics*, 0(0):1–17.

Goldblatt, R., Heilmann, K., and Vaizman, Y. (2020). Can medium-resolution satellite imagery measure economic activity at small geographies? evidence from landsat in vietnam. *The World Bank Economic Review*, 34(3):635–653.

Gregan, T. (2017). Indonesia - indonesia power transmission development project : P117323 - implementation status results report : Sequence 12. Implementation status results report, World Bank Group, Washington, D.C.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. *American economic review*, 102(2):994–1028.

Hodler, R. and Raschky, P. A. (2014). Regional favoritism. *The Quarterly Journal of Economics*, 129(2):995–1033.

Humanitarian Data Exchange (2024). Indonesia administrative level 1 boundaries (provinces). Available at: https://data.humdata.org/dataset/cod-ab-idn. Accessed: December 2024.

Keola, S., Andersson, M., and Hall, O. (2015). Monitoring economic development from space: using nighttime light and land cover data to measure economic growth. *World Development*, 66:322–334.

Lipscomb, M., Mobarak, A. M., and Barham, T. (2013). Development effects of electrification: Evidence from the topographic placement of hydropower plants in brazil. *American Economic Journal: Applied Economics*, 5(2):200–231.

Min, B. and Gaba, K. M. (2014). Tracking electrification in vietnam using nighttime lights. *Remote Sensing*, 6(10):9511–9529.

Min, B., Gaba, K. M., Sarr, O. F., and Agalassou, A. (2013). Detection of rural electrification in africa using dmsp-ols night lights imagery. *International journal of remote sensing*, 34(22):8118–8141.

Rossi-Hansberg, E. and Zhang, J. (2025). Local gdp estimates around the world. Technical report, National Bureau of Economic Research.

Servén, L. and Calderón, C. (2004). The effects of infrastructure development on growth and income distribution. *Documentos de Trabajo (Banco Central de Chile)*, (270):1–47.

Straub, S. (2011). Infrastructure and development: A critical appraisal of the macro-level literature. *The Journal of Development Studies*, 47:683–708.

Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1):57–76.

A Impact on Electricity Distribution

As a supplementary analysis, we examine whether the electricity infrastructure program increased the amount of electricity distributed across treated provinces. Using the same generalized synthetic control framework, we estimate the causal effect on electricity distribution measured in gigawatt-hours (GWh). The model specification includes the same control variables: labor force size, foreign hotel guests, and disaster-affected population. Cross-validation selected an optimal number of $r^* = 0$ latent factors.

The pre-treatment fit assessment reveals substantially more problematic synthetic matches for electricity distribution compared to the GRDP analysis. Nine provinces—Jawa Timur, DKI Jakarta, Jawa Barat, Jawa Tengah, Riau, Sumatera Utara, Lampung, Sumatera Selatan, and Banten—exhibit statistically significant deviations from their synthetic counterparts in one or more pre-treatment periods ($p < 0.05$). Notably, these excluded provinces include all four major Java provinces and several large Sumatran provinces that represent the bulk of Indonesia's electricity consumption. This extensive exclusion reduces our analytical sample to only 6 treated provinces with good pre-treatment fit, limiting the generalizability of our findings to the smaller and less industrialized provinces in the treatment group.

Figure 7 presents the results. The covariate coefficient for labor force is positive and statistically significant ($p < 0.001$), indicating that provinces with larger labor forces have higher electricity distribution, consistent with the relationship between economic activity and energy consumption. Neither total hotel guests nor disaster-affected population are significant predictors of electricity distribution. The middle panel compares treatment effect estimates between the full sample (dashed line) and the filtered sample (solid line). The estimates show considerable divergence between the two samples, particularly in later post-treatment periods. For the filtered sample, the average treatment effect across all post-treatment periods is 125.75 GWh, but this estimate is not statistically significant (p = 0.453, 95% CI: [-202.38, 453.88]).

Figure 7: *Estimated Effect on Electricity Distribution*

Covariate	Estimate	S.E.	95% CI		p-value
Labor force	0.0029***	0.0003	0.0022	0.0035	<0.001
Total hotel guests	−0.0001	0.0015	−0.0030	0.0028	0.966
Disaster affected	−0.0002	0.0002	−0.0005	0.0002	0.343

Notes: Significance levels: $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. Sample restricted to 6 treated provinces with good pre-treatment fit.

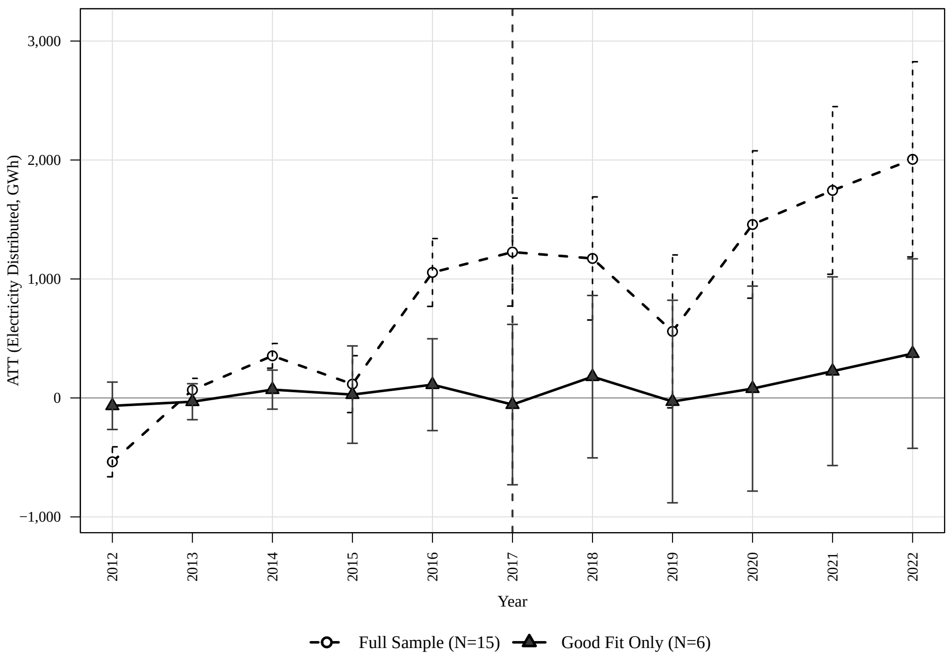

Average Treatment Effect on Electricity Distribution				
Estimate	S.E.	95% CI		p-value
125.75	167.42	-202.38	453.88	0.453

Note: Comparison of estimated treatment effects between the full sample (N=15, dashed line) and the filtered sample with good pre-treatment fit (N=6, solid line). Vertical line marks the treatment year (2017). Error bars represent 95% confidence intervals. Nine provinces were excluded due to poor pre-treatment fit: Jawa Timur, DKI Jakarta, Jawa Barat, Jawa Tengah, Riau, Sumatera Utara, Lampung, Sumatera Selatan, and Banten. Average treatment effect across all post-treatment periods (2017–2022). Sample restricted to 6 treated provinces with good pre-treatment fit. Units: GWh.

CESA Working Paper Series

CESA Working Paper

Series

CESA Working Paper No. 9 | 2024
Marco Kerkemeier
(Co-)Explosiveness of corporate credit spreads

CESA Working Paper No. 10 | 2024
Christoph Wegener, Robinson Kruse-Becher and Tony Klein
EU ETS Market Expectations and Rational Bubbles

CESA Working Paper No. 11 | 2025
Anna Krumme and Matthias Westphal
Monetary returns to upper secondary schooling, the evolution of unobserved heterogeneity, and implications for employer learning

CESA Working Paper No. 12 | 2025
Robinson Kruse-Becher and Philip Letixerant
Oil price expectations in explosive phases

CESA Working Paper No. 13 | 2025
Sarah Rezaei, Hendrik Sonnabend and Bastian Westbrock
Overconfidence and Goal-Setting

CESA Working Paper No. 14 | 2025
Rainer Baule, Sebastian Schlie and Xiaozhou Zhou
The Term Structure of Intraday Return Autocorrelations

CESA Working Paper No. 15 | 2025
Mario Lackner and Hendrik Sonnabend
Positive Feedback Shapes Gender Gaps in Adolescent Risk-Taking –
Causal Evidence from Real-Risk Competitions

FernUniversität in Hagen

CESA
Working
Paper
Series

CESA Working Paper No. 16 | 2025
Hendrik Sonnabend and Matthias Westphal
Family First: The causal effect of family size on cultural participation

CESA Working Paper No. 17 | 2025
Malte Petersen, Hendrik Sonnabend and Jürgen Weibler
Do Fields of Study Shape Social Preferences? Selection and Curricular Imprinting in Economically Oriented versus Prosocial Programs

CESA Working Paper No. 18 | 2026
Anna Krumme and Matthias Westphal
Is Academic-Track Schooling Worthwhile Without College? Decomposing Monetary Returns to Education

CESA Working Paper No. 19 | 2026
Klaus Andresen, Ursula Müller and Hans-Jörg Schmerer
Energy for Growth: Satellite Synthetic Control Evidence from Indonesia

FernUniversität in Hagen