Economics Bulletin

Volume 46, Issue 1

The economic impact of the 2014 oil price collapse on Rio de Janeiro: A synthetic control analysis

Caio F Ferreira
Universidade do Estado do Rio de Janeiro

Abstract

This paper evaluates the economic impact of the 2014 oil price collapse on the State of Rio de Janeiro using the Synthetic Control Method and annual data from 1990 to 2021. The synthetic control replicates Rio's pre-shock income path and serves as a counterfactual for the post-2014 period. Results show a persistent decline in GDP per capita relative to the synthetic benchmark, with a peak gap of about 14 percent in 2020. Placebo tests, an in-time falsification, and leave-one-out estimates confirm the robustness of the findings. The evidence provides the first causal estimate of a major commodity-price shock for a Brazilian state and highlights the vulnerability of resource-dependent regional economies.

Citation: Caio F Ferreira, (2026) "The economic impact of the 2014 oil price collapse on Rio de Janeiro: A synthetic control analysis", *Economics Bulletin*, Volume 46, Issue 1, pages 260-273
Contact: Caio F Ferreira - caio.ferrari@outlook.com.
Submitted: December 03, 2025. **Published:** March 30, 2026.

1. Introduction

Subnational economies that rely heavily on volatile commodity revenues are particularly vulnerable to external shocks. In resource-dependent regions, sudden declines in international prices can tighten local budget constraints, disrupt production networks linked to the commodity sector, and trigger persistent contractions in income. Although commodity traps and resource dependence are widely discussed in the literature (Guillaumont, 1999; Combes and Guillaumont, 2002; Le Billon and Good, 2016; Cárdenas et al., 2011; UNCTAD, 2019, 2021), credible causal evidence on the effects of commodity price collapses at the subnational level remains limited, especially for emerging economies.

Rio de Janeiro provides a natural setting to study this question. Oil royalties and related transfers constitute a substantial share of the state's revenues, making its fiscal capacity unusually exposed to international oil-price fluctuations. The sharp decline in global oil prices in 2014 generated a large and externally driven negative shock to Rio's revenues and to its oil-related production chain. Using annual data from 1990 to 2021 and the Synthetic Control Method (SCM), we estimate the causal effect of the 2014 oil price collapse on Rio's GDP per capita by comparing it to a data-driven counterfactual constructed from other Brazilian states. We find a persistent post-2014 divergence: Rio's GDP per capita falls to about 14 percent below its synthetic counterpart.

The paper relates to three strands of literature. First, work on commodity dependence emphasizes that the growth and governance consequences of resource reliance depend on institutions and on how resource rents are shared across levels of government, with fiscal decentralization potentially amplifying exposure to booms and busts (Perez-Sebastian and Raveh, 2016; van der Ploeg, 2011; Frankel, 2012). Second, subnational evidence highlights that commodity windfalls and busts can operate through local fiscal capacity and regional spillovers. For Brazil, influential municipal-level studies exploit cross-municipality variation in oil-related windfalls to examine public spending, living standards, and political incentives (Caselli and Michaels, 2013; Brollo et al., 2013). Internationally, related analyses document regional dynamics around resource booms and busts (Allcott and Keniston, 2018; Pellegrini et al., 2021). A third strand of the literature applies the Synthetic Control Method to quantify the macroeconomic effects of large shocks and regime changes in settings with a single treated unit. Representative applications include the growth effects of market-oriented liberalization reforms (Billmeier and Nannicini, 2013), the output consequences of catastrophic natural disasters (Cavallo et al., 2013), the income effects of institutional integration such as EU accession (Campos et al., 2019), the recovery impact of IMF programs during macroeconomic crises (Kuruc, 2022), and the output costs of major political shocks such as the Brexit referendum (Born et al., 2019).

Our contribution is threefold. First, we provide a transparent comparative-case causal estimate of the subnational income impact of the 2014 global oil-price collapse within a federative fiscal system, using a long pre-intervention period and a donor pool of other Brazilian states. Second, we document the magnitude and persistence of the post-2014 divergence in GDP per capita, consistent with transmission mechanisms operating through fiscal capacity and exposure to the oil-related production chain. Third, we offer a replicable empirical template that can be applied to other regions facing externally driven commodity shocks when a sufficiently long pre-shock period and a credible comparison set are available. The remainder of the paper follows a standard empirical structure. Section 2 describes the data and background. Section 3 outlines the SCM. Section 4 presents the main results. Section 5 reports robustness checks. Section 6 concludes.

2. Background and Institutional Setting

Rio de Janeiro is one of the most oil-dependent subnational economies in Brazil. A significant share of state revenues derives from offshore oil royalties, which are directly linked to international price fluctuations. This fiscal arrangement makes Rio particularly exposed to external commodity shocks. The sharp decline in global oil prices between 2014 and 2016 reduced royalty inflows and affected the state's oil-related production chain. Because other Brazilian states are much less exposed to oil royalties, the setting is well suited for constructing a synthetic counterfactual.

Although the fall in royalties contributed to fiscal stress during this period, no other state experienced a shock of similar magnitude or origin. For this reason, the 2014 oil price collapse provides a natural experiment to examine the causal effect of an external commodity shock on subnational income. Figures 1a and 1b document the evolution of the Brent price and Rio's royalties, illustrating the timing and magnitude of the shock.

Figure 1 - (a) Brent Oil Price: Historical Data (2012–2020); (b) Oil Royalty Revenues – State of Rio de Janeiro.

(a)

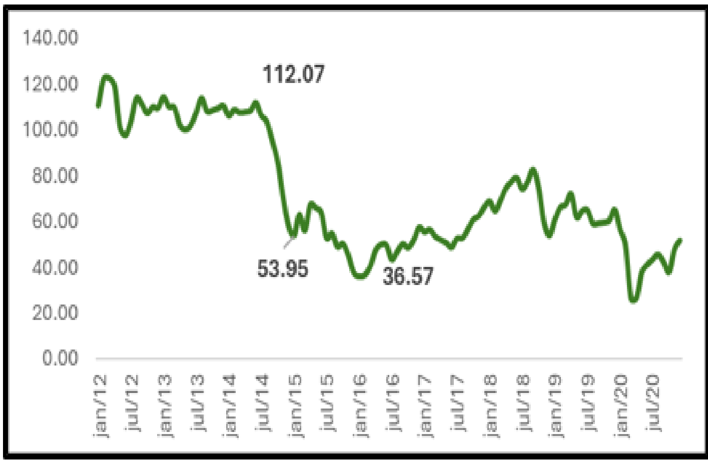

(b)

Year	Royalties (in Billions)		Annual Change (%)	Change Since 2014 (%)	Difference from 2014 (R$)
2012	R$	2.96	-	-	-
2013	R$	2.82	-5.0%	-	-
2014	R$	2.85	1.3%	-	-
2015	R$	1.85	-35.1%	-35.1%	-R$ 1.00
2016	R$	1.50	-19.1%	-47.5%	-R$ 1.35
2017	R$	1.95	29.9%	-31.7%	-R$ 0.91
2018	R$	2.97	52.6%	4.2%	R$ 0.12
2019	R$	3.07	3.3%	7.6%	R$ 0.22
2020	R$	3.10	1.2%	8.9%	R$ 0.25

2.1. Transmission channels

In federative and decentralized fiscal systems, commodity dependence can amplify vulnerability to external price shocks because revenue-sharing rules transmit global fluctuations directly to subnational budget constraints, limiting local stabilization capacity (Perez-Sebastian and Raveh, 2016). More broadly, the resource-dependence literature emphasizes that the persistence of shocks depends on fiscal and institutional arrangements and on political-economy mechanisms that can propagate downturns (van der Ploeg, 2011; Frankel, 2012). In this context, the 2014 oil price collapse can affect Rio de Janeiro's GDP per capita through three complementary channels:

- **Fiscal channel.** The first mechanism operates through public finances. When royalty-linked revenues fall, the state faces a tighter budget constraint, which can translate into reductions or delays in investment and public service provision, with demand-side and medium-run effects. This channel is consistent with evidence that oil windfalls shape local public spending and related outcomes in Brazil and can affect political incentives when revenues vary exogenously (Caselli and Michaels, 2013; Brollo et al., 2013). In Rio de Janeiro, the unusually high reliance on oil royalties implies that the 2014 collapse plausibly generated a large contractionary fiscal shock.
- **Productive channel.** A second mechanism operates through regional production and labor-market spillovers connected to the oil and gas complex. Commodity busts can depress activity along sectoral supply chains, reducing investment and employment in linked industries and propagating shocks to local demand, construction, and market services. This logic is consistent with subnational evidence on regional dynamics around resource booms and busts (Allcott and

Keniston, 2018; Pellegrini et al., 2021). For Rio de Janeiro, exposure through extraction-related services and suppliers provides an additional channel beyond the fiscal link.
- **Institutional channel.** A third mechanism works through institutional and confidence effects. Large fiscal deteriorations may heighten political uncertainty, weaken governance, and reduce private-sector confidence, amplifying the contraction and contributing to persistence. This mechanism is in line with the broader resource-curse literature that highlights institutional and political-economy channels behind medium-run underperformance (Ploeg, 2011; Frankel, 2012). In Rio de Janeiro, the combination of fiscal stress and institutional fragility may therefore reinforce persistence in the post-2014 divergence.

While the empirical strategy estimates the overall causal effect on GDP per capita, the channels above motivate the hypotheses stated in Section 2.2 and guide the interpretation of the post-2014 dynamics in Section 4.

2.2. Economic hypotheses

Building on the channels of section 2.1, we focus on two hypotheses that summarize the main propagation mechanisms of the shock:

- **H1 (fiscal transmission):** the 2014 oil-price collapse reduced Rio de Janeiro's GDP per capita by tightening the state's budget constraint through lower royalty-linked revenues, contributing to cuts or delays in public investment and weaker local demand.
- **H2 (persistence):** the effect is persistent because subnational governments have limited capacity to smooth large revenue shocks under rigid expenditures and constrained borrowing, slowing the recovery even after prices stabilize.

These hypotheses are not separately identified, but they provide an organizing interpretation for the magnitude and persistence of the estimated divergence after 2014.

3. Methodology

This study employs the Synthetic Control Method (SCM) introduced by Abadie and Gardeazabal (2003) and formalized by Abadie et al. (2010). SCM constructs a weighted combination of untreated units that best reproduces the pre-treatment trajectory of the treated unit. This weighted combination—referred to as the synthetic control—approximates the counterfactual path that the treated unit would have followed in the absence of the intervention.

Formally, the synthetic control is defined as:

$$\hat{Y}_{0t} = \sum_{j=1}^{J} w_j Y_{jt} \qquad (1)$$

where Y_{jt} denotes the outcome for donor state j at time t, and the weights w_j are non-negative and sum to one. These weights are chosen to minimize the distance between Rio de Janeiro and the synthetic combination of donor states in terms of pre-treatment outcomes and socioeconomic predictors. The long pre-treatment period (1990–2013) provides substantial information for achieving a close fit.

The identifying assumption is that, conditional on matching pre-treatment characteristics and outcome dynamics, the synthetic control represents the path Rio would have followed had the 2014 oil price collapse not occurred. The estimated treatment effect corresponds to the post-2014 difference between Rio's observed GDP per capita and that of its synthetic counterpart.

Inference follows standard practice in the SCM literature. As Abadie (2021) emphasizes, the credibility of SCM hinges on achieving a close replication of the treated unit's outcome path during a sufficiently long pre-intervention period. First, placebo-in-space tests reassign the treatment to each donor state to assess whether effects of comparable magnitude arise under artificial interventions. Second, an in-time placebo assigns a false treatment year (2011) to verify the absence of anticipatory effects. Third, leave-one-out exercises iteratively remove each positively weighted donor state to assess the sensitivity of results to donor composition. Together, these procedures evaluate the robustness of the estimated treatment effect.

All variables used to construct the synthetic control (GDP per capita, industry share, population growth, education indicators, tax revenues, and energy consumption) were obtained from publicly available state-level sources (IPEADATA, IBGE) and harmonized to annual frequency.

Our baseline donor pool includes all other Brazilian states and the Federal District to ensure comparability in the national macroeconomic and institutional environment. Appendix Table A1 reports the pre-2014 distribution of state-level oil royalties: Rio de Janeiro accounts for a disproportionate share, while most states receive negligible amounts, underscoring the asymmetry of exposure that motivates the comparative-case design. Nonetheless, because some donors may have non-negligible oil and gas exposure, we assess sensitivity to donor composition by re-estimating the SCM after excluding oil-exposed states based on pre-2014 royalty shares (see Section 5). We also report leave-one-out estimates excluding each positively weighted donor in turn, showing that results are not driven by any single donor unit.

The SCM predictor set is designed to capture slow-moving determinants of income and to reproduce Rio de Janeiro’s pre-2014 outcome dynamics. We include the covariates reported in Table 2 in the next section, constructed using pre-treatment information only: each predictor is measured as an average over 2003–2013 (or at the closest available pre-treatment benchmark years when only decennial data exist, e.g., census-based education in 1991, 2000, and 2010). To discipline the fit in the pre-treatment path, we also include pre-2014 GDP per capita summarized using averages over subperiods (as listed in Table 3), which reduces noise and prevents overfitting to year-specific shocks. All predictors are standardized in the SCM optimization in the usual way.

Finally, as a complementary robustness check, we also implement an augmented synthetic control estimator as an alternative weighting scheme (reported in Section 5).

4. Results

Figure 2 displays the evolution of real GDP per capita for Rio de Janeiro and its synthetic counterpart. The synthetic control reproduces the broad pre-treatment trajectory of Rio’s income, though the match is not exact in all years. This level of fit is typical in applications with long historical periods and heterogeneous regional economies and remains adequate for constructing a credible counterfactual.

After the 2014 oil price collapse, Rio’s GDP per capita diverges sharply from the synthetic benchmark. The estimated gaps indicate a shortfall of approximately 3% in 2015, rising to 9% in both 2016 and 2017. The divergence remains sizable in 2018 (5%) and 2019 (9%), and peaks at about 14% in 2020. Although the gap narrows slightly in 2021, it remains substantial.

Figure 2. Trends in GDP per capita: Rio de Janeiro vs. synthetic Rio de Janeiro

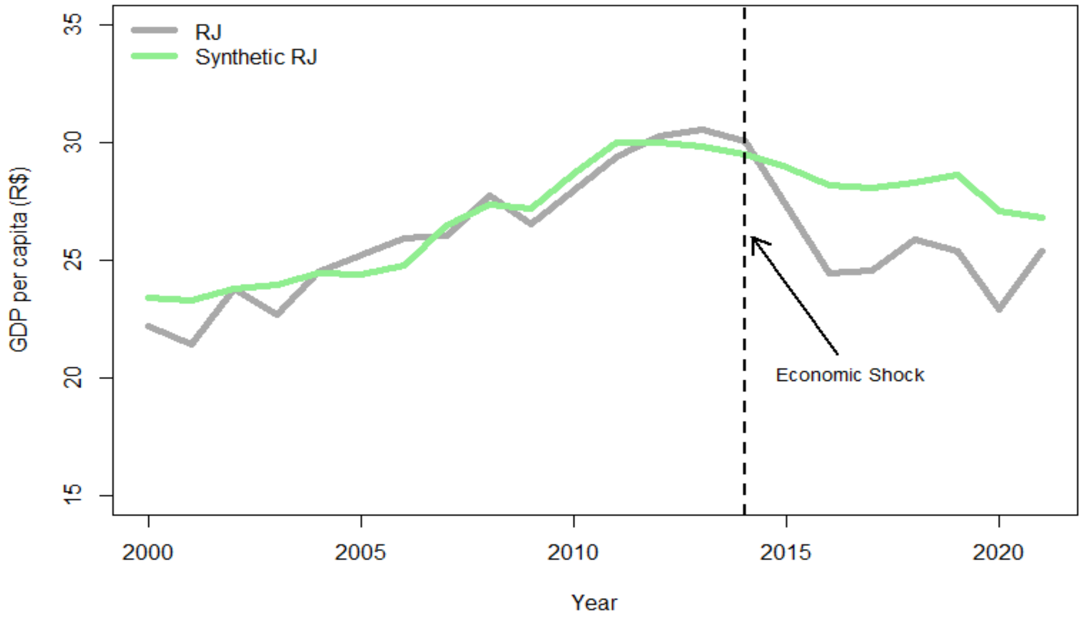

Table 1 reports on the donor weights. The synthetic control is primarily composed of Paraná, followed by São Paulo, the Federal District, Santa Catarina, and Espírito Santo. These states jointly reproduce Rio's structural characteristics in the pre-intervention period and contribute to the overall balance achieved in the optimization procedure.

Table 1. State weights in the synthetic RJ.

State	Weights
Paraná (PR)	0.563
Distrito Federal (DF)	0.172
São Paulo (SP)	0.162
Santa Catarina (SC)	0.053
Espirito Santo (ES)	0.049

Note: Table reports only states with wj > 0.

Figure 3 plots the estimated treatment effect $(Y_{0t} - \hat{Y}_{0t})$over time. The effects remain close to zero throughout the pre-treatment period, suggesting the absence of systematic pre-trends. After the shock, the effect becomes negative, persistent, and economically large, consistent with a substantial decline relative to the synthetic benchmark.

The magnitude and persistence of the post-2014 gap are consistent with a combination of fiscal and productive transmission mechanisms. The initial divergence emerges after the collapse in oil prices and deepens as the shock translates into tighter public budgets and weaker activity in the oil-related production network. The persistence of the gap is consistent with limited subnational smoothing capacity and slower recovery in investment and local demand. Additionally, Table 2 reports the root mean squared prediction error (RMSPE) in the pre- and post-treatment periods. The low pre-treatment RMSPE indicates a close fit between Rio de Janeiro and its synthetic counterpart before 2014, while the post-treatment RMSPE increases substantially after 2014, consistent with the estimated divergence.

Figure 3. Gap in GDP per capita: Rio de Janeiro vs. synthetic Rio de Janeiro.

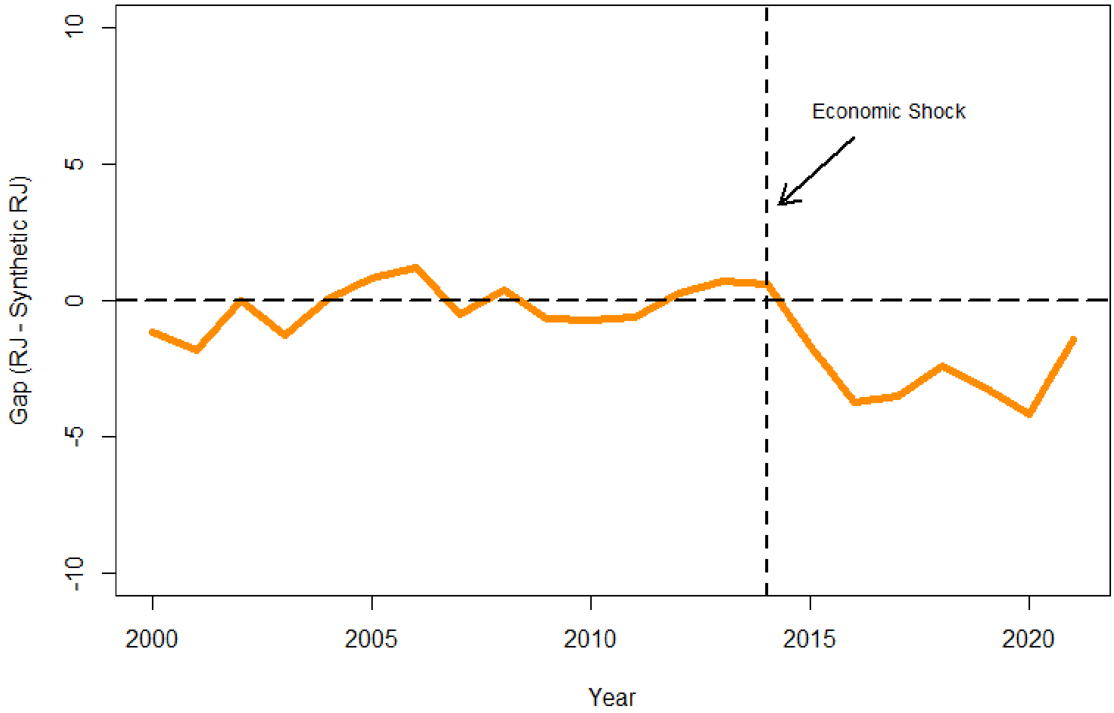

Table 2. Pre- and post-treatment RMSPE.

Unit (treated)	Pre-treatment RMSPE	Post-treatment RMSPE	Post/Pre RMSPE ratio
Rio de Janeiro	0.8752	2.8603	3.2683

Table 3 summarizes the pre-treatment predictor balance. Synthetic Rio matches the treated unit more closely than the simple average of the donor pool, supporting the credibility of the counterfactual. We also re-estimate the SCM using only pre-treatment outcome lags (GDP per capita) to assess sensitivity to the predictor set; results remain qualitatively unchanged (Appendix Figure A1).

Table 3. Averages of predictors of the GDP per capita before the crisis.

Predictors Mean	Rio de Janeiro	Synthetic RJ	Brazil Mean
Industry Share (% GDP)	26.27	20.81	19.85
Population (% growth)	2.78	2.27	1.45
IDHM - Education (1991)	0.39	0.33	0.24
IDHM - Education (2000)	0.53	0.54	0.41
IDHM - Education (2010)	0.68	0.69	0.61
State Tax Revenue (ICMS)	7.19	7.75	3.04
log Energy Consumption	16.04	16.02	14.51
GDP per capita (1990 - 2004)	22.085	22.487	12.49
GDP per capita (2005 - 2009)	26.290	26.033	14.92
GDP per capita (2010 - 2013)	29.548	29.623	17.31

Notes: All variables are averaged for the period 2003-2013, except the IDHM – Education Human Development Index, which was available only in 1991, 2000 and 2010 (due to the decennial census). The last column shows the donor pool average (26 states and Federal District)

A potential concern is that other macro shocks, such as the COVID-19 pandemic, overlap with the post-treatment period. The SCM design compares Rio de Janeiro to a synthetic unit built from other

Brazilian states, so nationwide shocks that affect all states (e.g., the 2020 recession) are largely differenced out and do not mechanically generate a treated–control gap. Concurrent shocks would bias the estimate only if they affected Rio de Janeiro disproportionately relative to the donor pool. While differential pandemic severity across states could contribute to the magnitude of the 2020–2021 gap, the divergence emerges well before the pandemic (from 2015 onward) and remains negative throughout 2015–2019. We therefore interpret the post-2014 pattern as primarily reflecting the oil-price shock and its propagation, and view 2020–2021 as consistent with persistence but subject to this caveat.

Overall, the results indicate a pronounced and persistent negative effect of the 2014 oil price collapse on Rio de Janeiro's GDP per capita. In the absence of the shock, the state's income trajectory would likely have followed a significantly more favorable path.

5. Robustness analysis

A series of robustness exercises were performed to assess the credibility of the estimated treatment effect. Figure 4 presents the placebo-in-space test, in which the treatment is reassigned to each donor state. For almost all donor units, the post-2014 gaps remain small and fluctuate around zero, contrasting with the large and persistent divergence observed for Rio de Janeiro. Only a few placebos display moderate deviations, and even these remain substantially smaller than Rio's. This pattern indicates that the estimated effect is not driven by shocks common to other Brazilian states.

Figure 4. GDP per capita gaps in RJ and placebo gaps in all twenty-six control states.

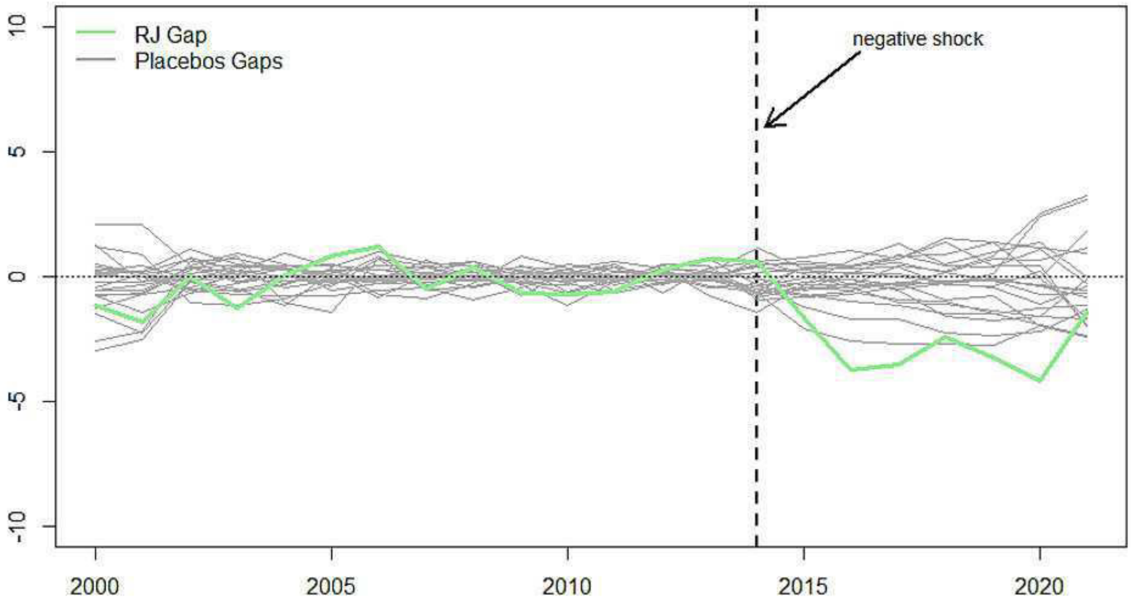

Using post/pre MSPE ratios from placebo reassignments, Rio de Janeiro's permutation-based p-value is 0.231, defined as the share of units with post/pre MSPE ratios at least as large as Rio de Janeiro's. Following the standard permutation test, the p-value equals the fraction of units (placebos plus RJ) whose post/pre RMSPE ratio is at least as large as Rio's. Equivalently, Rio de Janeiro's post/pre MSPE ratio exceeds 20 out of the 25 placebo states (about 80%), indicating that the post-2014 deterioration in fit is relatively large compared with most placebo units.

Figure 5 reports the in-time placebo test, which assigns a false treatment year of 2011 to Rio de Janeiro. The estimated gap remains near zero around the placebo intervention date, and no systematic divergence emerges in the subsequent years. This absence of a pre-treatment break reinforces the interpretation that the observed post-2014 deviation reflects the external oil price shock rather than pre-existing trends.

Figure 5. In time test - GDP per capita in RJ vs. Synthetic RJ.

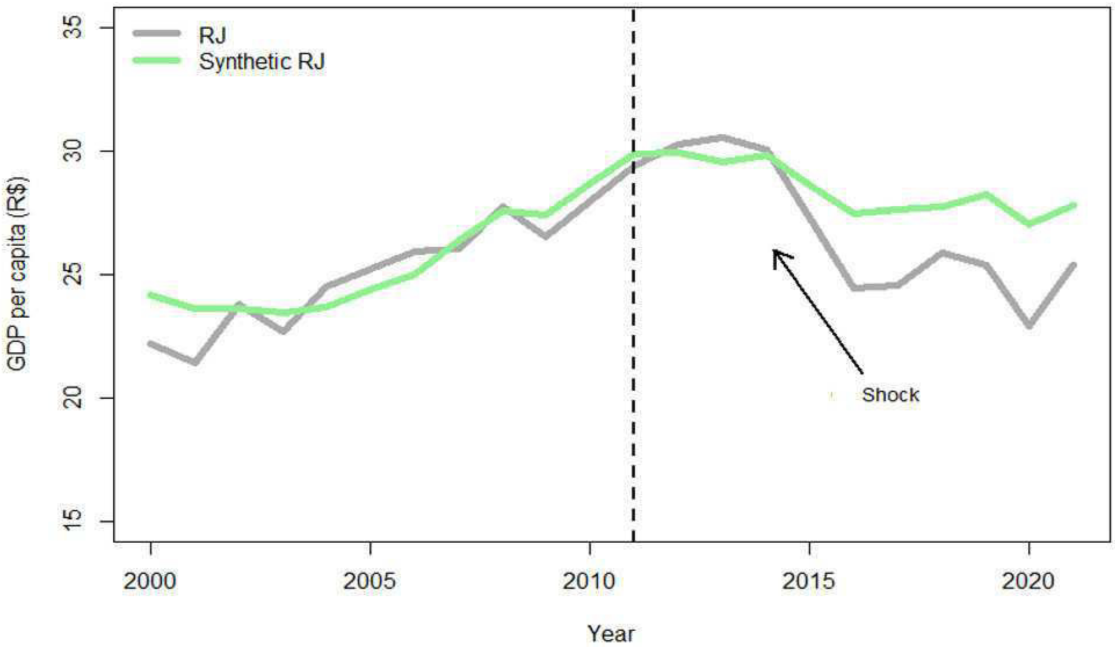

Figure 6 displays the leave-one-out analysis, in which each positively weighted donor state is removed sequentially and the synthetic control is re-estimated. The resulting synthetic paths remain close to the baseline estimate. Although the magnitude of the gap varies slightly depending on which state is excluded, the direction, timing, and persistence of the post-2014 divergence are unchanged. This indicates that the main results are not disproportionately driven by any single donor state.

The donors excluded in Figure 6 are precisely the states that receive positive weights in the baseline synthetic control: Paraná (PR), São Paulo (SP), Santa Catarina (SC), the Federal District (DF), and Espírito Santo (ES). They reflect different sources of pre-2014 comparability: large and diversified economies (SP and PR), Southern states with similar income dynamics (SC and PR), a high-income benchmark with a distinct sectoral structure (DF), and a geographically close state with some oil-related exposure (ES).

To address concerns that some donor states may be indirectly exposed to the oil and gas sector, we re-estimate the SCM using a restricted donor pool that excludes all states with non-negligible oil-royalty exposure in the pre-treatment period. Figure 7 plots Rio de Janeiro and the synthetic control for GDP per capita gap when the donor pool excludes states with pre-2014 oil-royalty shares above 1% (Amazonas, Bahia, Espírito Santo, Rio Grande do Norte, Sergipe, and São Paulo, see Table A1). The pre-treatment fit remains close and the post-2014 divergence persists. The vertical dashed line indicates the 2014 oil-price shock.

Economically, this restriction strengthens the interpretation that the estimated post-2014 gap reflects the oil-price shock hitting Rio de Janeiro through its unusually high fiscal dependence on royalties and its exposure along the oil-related production chain, rather than spillovers affecting other oil-exposed states. By removing donors with direct royalty links, the synthetic control is built from states whose public finances were largely insulated from oil-price movements, making it less likely that common oil-sector dynamics in the donor pool attenuate or confound the estimated effect. The persistence of the divergence under this restriction is therefore consistent with a Rio-specific transmission mechanism operating through subnational fiscal capacity and local activity linked to the energy complex.

Figure 6. Trends in GDP per-capita: Rio de Janeiro vs. leave-one-out synthetic Rio de Janeiro.

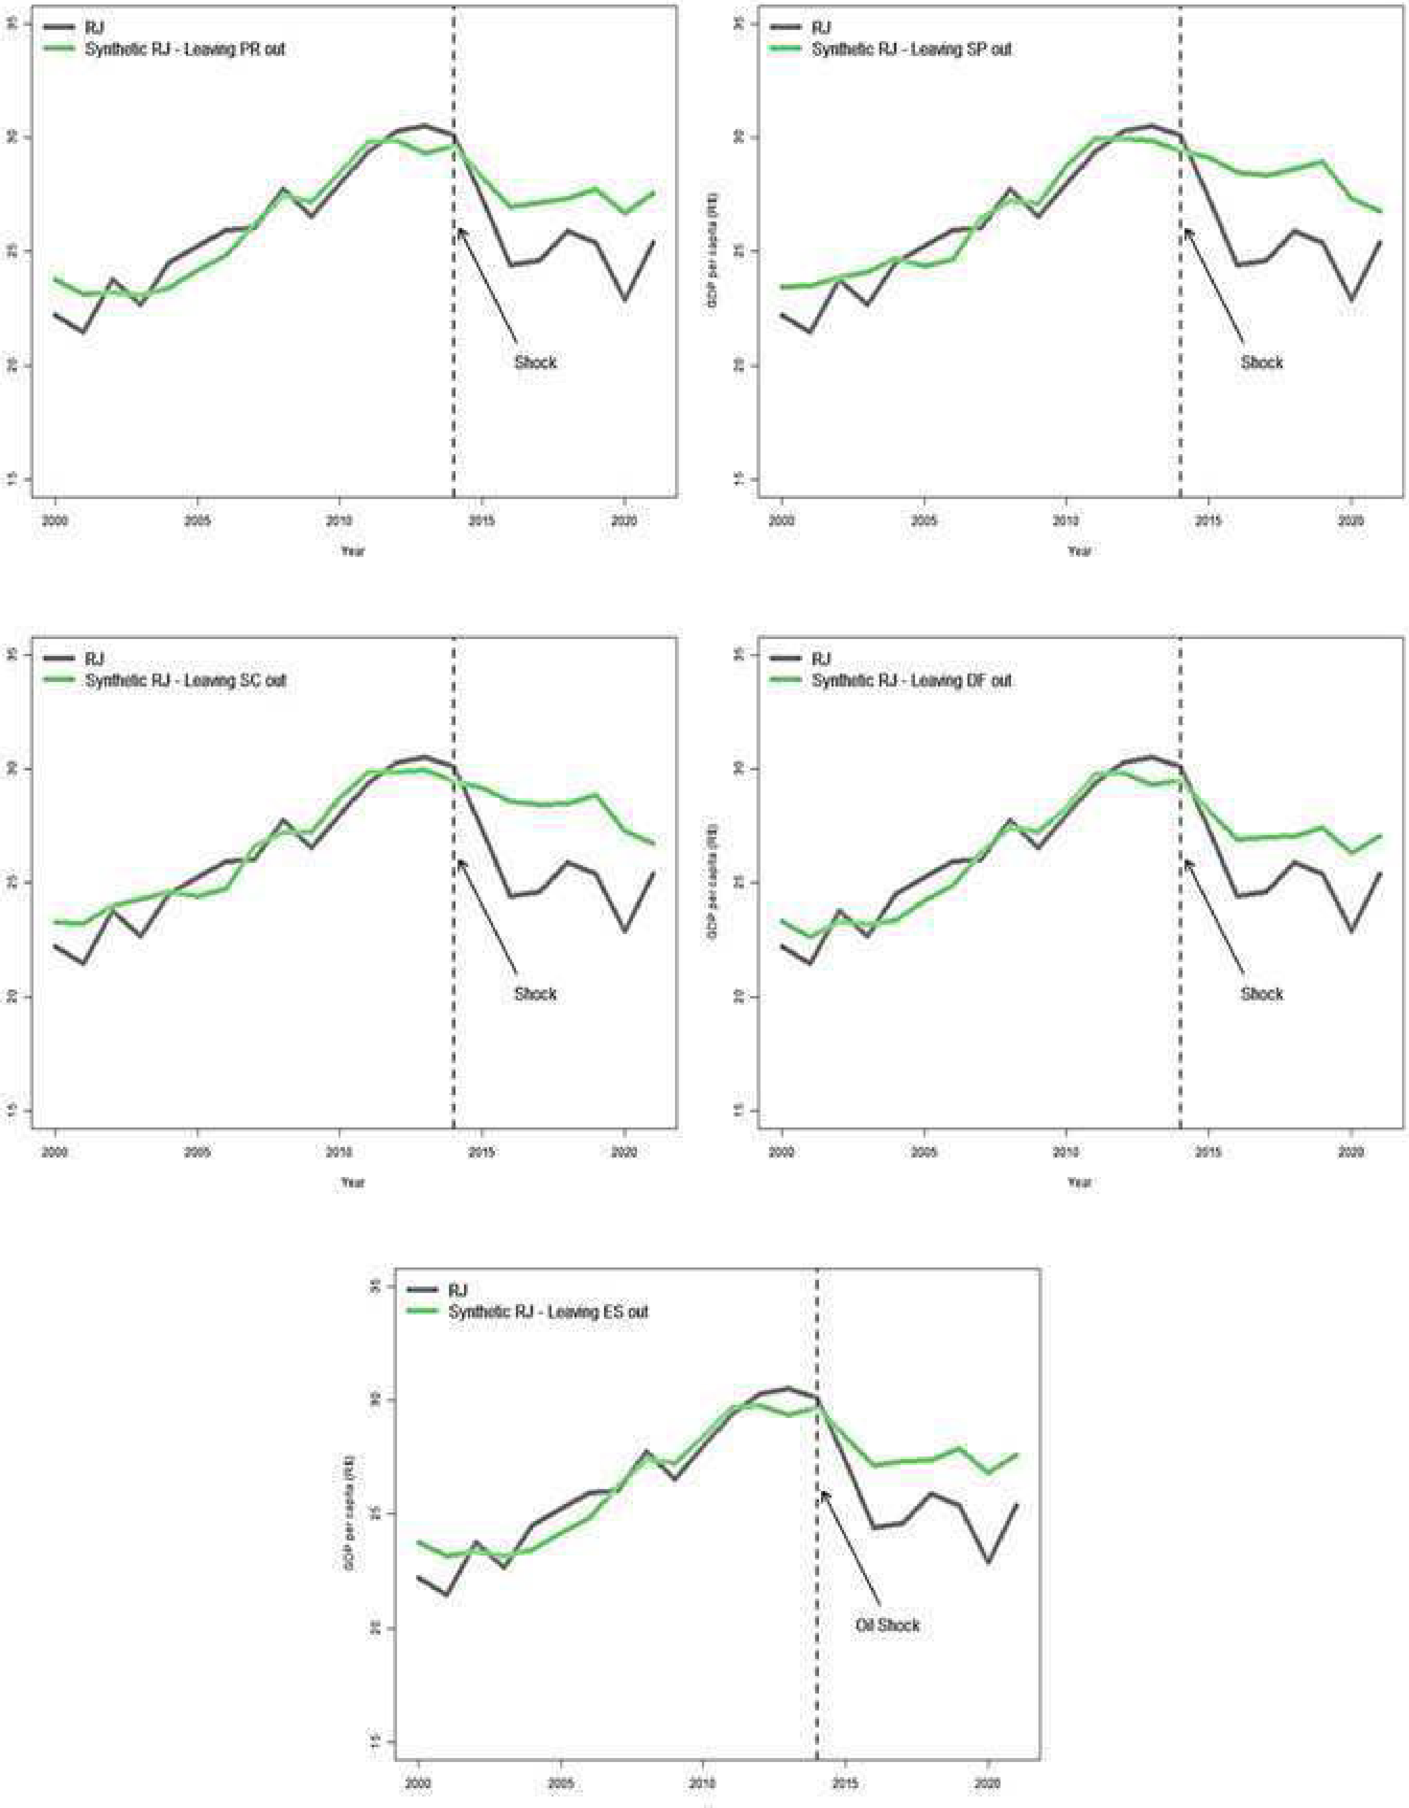

Figure 7. Robustness to donor composition: excluding oil-exposed donors.

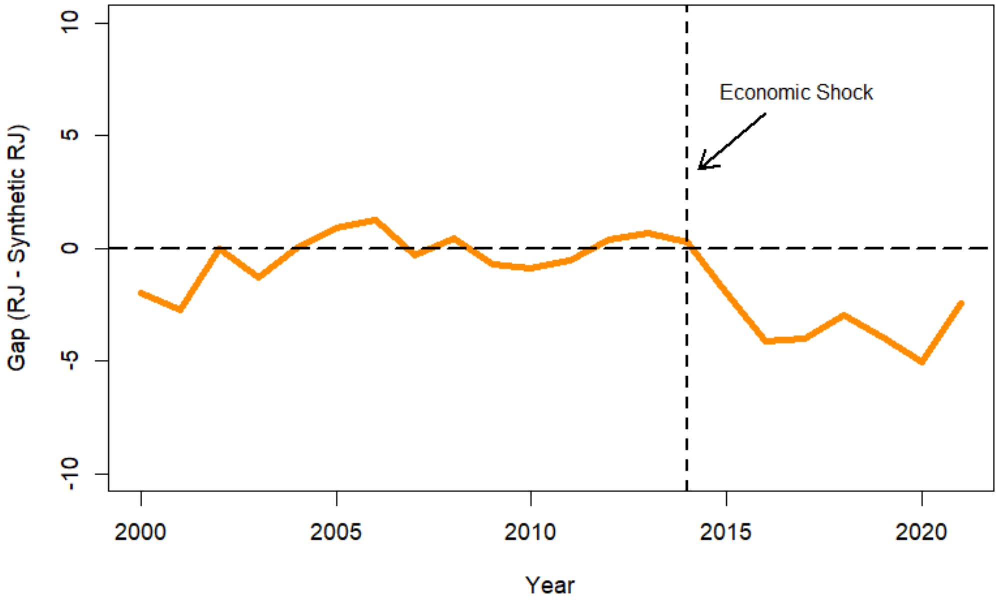

Finally, as an additional robustness check, we re-estimate the model using the Augmented Synthetic Control Method (ASCM) proposed by Ben-Michael et al. (2021). ASCM combines the SCM weighting scheme with an outcome model to reduce bias when pre-treatment fit is imperfect and provides complementary inference in small samples. We use the same treated unit, intervention year (2014), and panel structure as in the baseline SCM. Figure 8 reports the ASCM path of estimated effects; the point estimates remain negative throughout most of the post-2014 period and preserve the qualitative pattern of a persistent post-shock deterioration, although uncertainty bands are wide. Overall, ASCM estimates corroborates that the post-2014 divergence is not an artefact of a particular optimization routine or predictor choice.

Figure 8. Augmented synthetic control estimates for Rio de Janeiro.

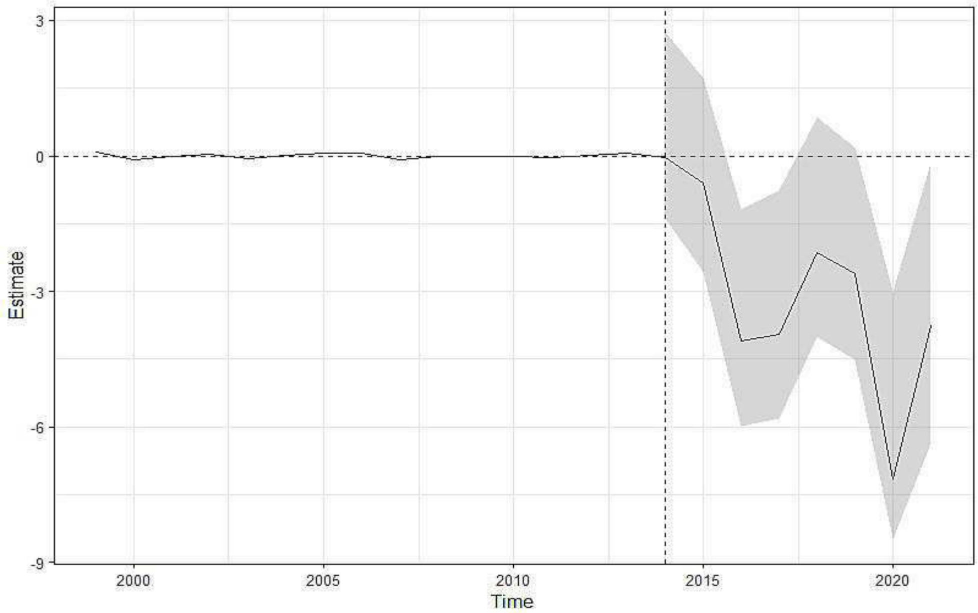

The estimated average treatment effect (ATT) post-2014 is −3.043. Year-by-year estimates indicate sizeable negative effects in 2016 (−4.092; 95% CI [−5.982, −1.189]), 2017 (−3.944; 95% CI [−5.803, −0.762]), 2020 (−7.164; 95% CI [−8.487, −3.016]), and 2021 (−3.752; 95% CI [−6.349, −0.168]). Although uncertainty remains non-negligible in some years (e.g., 2014–2015 and 2018–2019), the augmented-SCM results corroborate the baseline finding of a persistent post-2014 deterioration in Rio de Janeiro's GDP per capita relative to its synthetic counterfactual.

Taken together, the robustness checks consistently support the baseline findings: the 2014 oil price collapse produced a substantial and persistent negative effect on Rio de Janeiro's GDP per capita.

6. Concluding remarks

This paper estimates the causal impact of the 2014 collapse in international oil prices on Rio de Janeiro's GDP per capita using the Synthetic Control Method. We find a sizable and persistent post-shock divergence: Rio's GDP per capita falls about 3% below its synthetic counterpart in 2015, reaches roughly 9% in 2016–2017, remains negative in 2018–2019, and peaks around 14% in 2020, with only partial narrowing in 2021. Economically, the magnitude and persistence of the divergence are consistent with a propagation mechanism operating through Rio's unusually high fiscal reliance on royalty-linked revenues and the contraction of the oil-related production chain, combined with limited subnational capacity to smooth large revenue shocks.

The study contributes to the literature by providing comparative-case causal evidence on the effects of a major externally driven commodity-price shock on a Brazilian subnational economy. More broadly, the evidence highlights how resource-dependent regional economies can experience sizable and persistent losses when exposed to sharp international price declines.

Like most applications of the synthetic control method, the analysis relies on the quality of pre-treatment matching and assumes the absence of unobserved confounders that change discontinuously at the time of the shock. Future work could incorporate more granular sectoral data or alternative identification strategies to further explore the mechanisms through which commodity price shocks propagate across regional economies.

Overall, the results underscore the vulnerability of resource-dependent subnational economies to external shocks and point to the importance of strengthening local fiscal structures to mitigate such risks.

Appendix

Table A1. Pre-2014 distribution of state-level oil royalties (share of national total)

State	Royalties (BRL billions, nominal)	Share (%)	State	Royalties (BRL billions, nominal)	Share (%)
Rio de Janeiro (RJ)	53.236	64%	Paraná (PR)	0.098	0%
Espírito Santo (ES)	8.090	10%	Minas Gerais (MG)	0.068	0%
Rio Grande do Norte (RN)	4.752	6%	Pará (PA)	0.017	0%
Bahia (BA)	4.176	5%	Amapá (AP)	0.003	0%
São Paulo (SP)	3.794	5%	Acre (AC)	0.000	0%
Sergipe (SE)	2.871	3%	Distrito Federal (DF)	0.000	0%
Amazonas (AM)	2.858	3%	Goiás (GO)	0.000	0%
Alagoas (AL)	0.907	1%	Mato Grosso do Sul (MS)	0.000	0%
Rio Grande do Sul (RS)	0.712	1%	Mato Grosso (MT)	0.000	0%
Ceará (CE)	0.612	1%	Piauí (PI)	0.000	0%
Pernambuco (PE)	0.459	1%	Rondônia (RO)	0.000	0%
Santa Catarina (SC)	0.454	1%	Roraima (RR)	0.000	0%
Paraíba (PB)	0.114	0%	Tocantins (TO)	0.000	0%
Maranhão (MA)	0.098	0%	**Total**	83.320	100%

Figure A1. Sensitivity of covariates (Lagged GDP).

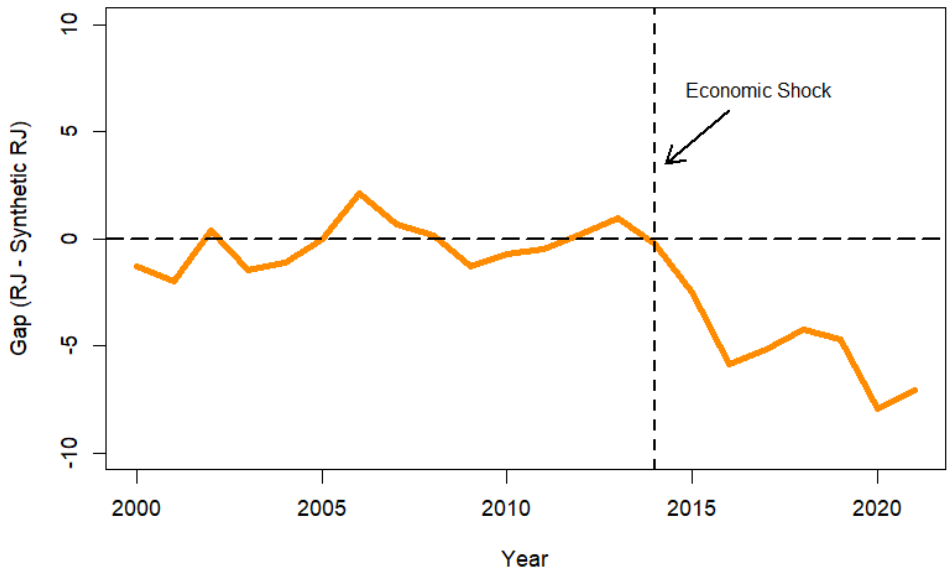

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2), 391–425.

Abadie, A., Diamond, A., & 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.

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2), 495–510.

Abadie, A., & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review*, 93(1), 113–132.

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

Ben-Michael, E., Feller, A., & Rothstein, J. (2021). The Augmented Synthetic Control Method. *Journal of the American Statistical Association*, 116(536), 1789-1803.

Billmeier, A., & Nannicini, T. (2013). Assessing economic liberalization episodes: A synthetic control approach. *Review of Economics and Statistics*, 95(3), 983–1001.

Born, B., Müller, G. J., Schularick, M., & Sedláček, P. (2019). The costs of economic nationalism: evidence from the Brexit experiment. *The Economic Journal*, *129*(623), 2722-2744.

Brollo, F., Nannicini, T., Perotti, R., & Tabellini, G. (2013). The political resource curse. *American Economic Review*, *103*(5), 1759-1796.

Campos, N. F., Coricelli, F., & Moretti, L. (2019). Institutional integration and economic growth in Europe. *Journal of Monetary Economics*, *103*, 88-104.

Cárdenas, M., Mejía, C., & Olivera, M. (2011). The dynamics of commodity prices. *Resources Policy*, 36(3), 188–203.

Cavallo, E., Galiani, S., Noy, I., & Pantano, J. (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics*, *95*(5), 1549-1561.

Caselli, F., & Michaels, G. (2013). Do oil windfalls improve living standards? Evidence from Brazil. *American Economic Journal: Applied Economics*, *5*(1), 208-238.

Combes, J.-L., & Guillaumont, P. (2002). Commodity price volatility, vulnerability and development. *Development Policy Review*, 20(1), 25–39.

Frankel, J. (2012). The natural resource curse: A survey. *NBER Working Paper* No. 15836.

Guillaumont, P. (1999). On the economic vulnerability of low-income countries. *UNU/WIDER Working Paper.*

Kuruc, K. (2022). Are IMF rescue packages effective? A synthetic control analysis of macroeconomic crises. *Journal of Monetary Economics*, *127*, 38-53.

Le Billon, P., & Good, E. (2016). Responding to the commodity bust: Governance challenges in resource-dependent economies. *Extractive Industries and Society*, 3(4), 961–969.

Pellegrini, L., Tasciotti, L., & Spartaco, A. (2021). A regional resource curse? A synthetic-control approach to oil extraction in Basilicata, Italy. *Ecological Economics*, *185*, 107041.

Perez-Sebastian, F., & Raveh, O. (2016). The natural resource curse and fiscal decentralization. *American Journal of Agricultural Economics*, *98*(1), 212-230.

Ploeg, F. V. D. (2011). Natural resources: curse or blessing?. *Journal of Economic literature*, *49*(2), 366-420.

UNCTAD. (2019). State of Commodity Dependence 2019. *United Nations Conference on Trade and Development*.

UNCTAD. (2021). Commodity Dependence and Development Report. *United Nations Conference on Trade and Development*.