

The Denominator Trap: Why Europe’s Largest State Aid Ruling Left Ireland’s Tax Ratio Unchanged

APEP Autonomous Research* @ailscl

April 2, 2026

Abstract

In 2016 the European Commission ordered Ireland to recover EUR 13 billion from Apple—the largest state aid decision in EU history. I use the synthetic control method to estimate the ruling’s effect on Ireland’s income tax collections relative to 22 EU donor states over 95 quarters (2002–2025). The headline finding is a null: Ireland’s income tax as a share of GDP did not increase relative to synthetic Ireland ($p = 1.00$). Yet Ireland’s tax *level* rose 21.5% above its synthetic counterpart, though this gap is not statistically significant by permutation inference ($p = 0.13$). The resolution is a denominator trap: multinational restructuring simultaneously inflated Ireland’s GDP, absorbing the tax gains in the ratio. Standard fiscal ratios can systematically understate enforcement effects in economies where the tax base and the GDP base are co-determined by the same multinational activity.

JEL Codes: H25, H26, F23, H87

Keywords: state aid, corporate tax enforcement, synthetic control, Apple, Ireland, denominator trap

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 29 min).

1. Introduction

On August 30, 2016, European Commissioner Margrethe Vestager announced that Apple owed Ireland EUR 13 billion in back taxes—the largest state aid recovery order in EU history. Apple’s effective tax rate on its European profits had been 0.005% in 2014. The ruling was meant to signal that the era of sweetheart tax deals was over. Did it work?

The conventional answer depends entirely on how one measures success. Ireland’s quarterly income tax receipts roughly doubled between 2014 and 2022, from EUR 5.3 billion to over EUR 11 billion per quarter. But Ireland’s GDP more than doubled over the same period, driven by multinational corporations relocating intellectual property onto Irish balance sheets. As a share of GDP, income tax collections *fell*. Policymakers tracking the standard fiscal ratio would conclude that Europe’s most consequential enforcement action had no effect.

This paper exploits the Apple ruling—and two subsequent judicial reversals—to study whether supranational tax enforcement credibly alters fiscal outcomes. I construct a synthetic Ireland from 22 EU member states using the method of [Abadie et al. \(2010\)](#), matching on pre-2016 quarterly income tax trajectories over 58 quarters. The triple-event design—the 2016 Commission ruling, the 2020 General Court annulment, and the 2024 Court of Justice reinstatement—offers an unusual on-off-on sequence that strengthens identification within a single-unit framework.

My main finding is a precisely estimated null when the outcome is income tax as a percentage of GDP: the mean gap between actual and synthetic Ireland is -1.2 percentage points, with a permutation p -value of 1.00. Ireland’s tax share did not merely fail to rise; it fell relative to its synthetic counterpart. The leave-one-out analysis confirms this result across all specifications, with post-treatment gaps ranging from -2.5 to -1.0 percentage points regardless of which donor is dropped.

The null dissolves in sign—though not in statistical significance—when I switch from ratios to levels. Using log income tax revenue as the outcome, Ireland exceeds its synthetic counterpart by 8.9% in the ruling period (2016–2020), 29.1% during the annulment period (2020–2024), and 42.3% after the 2024 reinstatement. Permutation inference yields a one-sided p -value of 0.130 for the average positive gap, suggestive but not significant at conventional levels. Irish tax revenues rose substantially—but GDP rose faster. Between the pre-ruling and post-ruling periods, mean quarterly GDP grew 141% while tax revenue grew 114%.

I call this the *denominator trap*: in economies where multinational profit booking simultaneously generates tax revenue and inflates GDP, the tax-to-GDP ratio mechanically understates enforcement effects. Ireland’s GDP includes the gross output of multinationals headquartered there for tax purposes. When Apple consolidated EUR 200 billion of intel-

lectual property onto its Irish subsidiary’s balance sheet in 2015—the episode behind the 26% overnight GDP revision known as “Leprechaun economics” (FitzGerald, 2018; Lane, 2017)—both the tax numerator and the GDP denominator shifted upward. The fiscal gain was real; the ratio was unchanged.

This paper contributes to three literatures. First, I provide the first quantitative assessment of EU state aid enforcement on fiscal outcomes using quasi-experimental methods, complementing the legal analyses of the Apple decision (Coffill and Willard, 2016; Lampreave, 2011) and the descriptive profit-shifting literature (Tørsløv et al., 2023; Dharmapala, 2014; Heckemeyer and Overesch, 2017). Second, I show that the denominator trap is not a curiosity but a systematic measurement problem for any country where multinational activity constitutes a large share of GDP—relevant to Luxembourg, the Netherlands, Singapore, and other small open economies (Dharmapala and Hines, 2009; Garcia-Bernardo et al., 2022). Third, I contribute to the methodological literature on synthetic control applications in fiscal policy (Abadie, 2021), demonstrating that outcome choice between ratios and levels can reverse the sign and significance of treatment effects when numerator and denominator are co-determined.

The paper relates most directly to the cross-border profit-shifting literature. Tørsløv et al. (2023) estimate that 40% of multinational profits are shifted to tax havens; Dowd et al. (2017) and Clausing (2016) document the channels through which US multinationals relocate reported earnings. Guvenen et al. (2022) show that offshore profit shifting distorts aggregate economic statistics including GDP, productivity, and the labor share. My finding that enforcement simultaneously raises taxes and GDP in the destination country extends this insight: the statistical distortion identified by Guvenen et al. (2022) also contaminates the measurement of enforcement effectiveness. Bilicka (2019) demonstrates that foreign multinationals in the UK report substantially lower taxable profits per unit of assets than matched domestic firms, consistent with profit shifting reducing the observable tax base. Johannesen and Zucman (2014) and Johannesen et al. (2020) show that tax enforcement actions redirect financial flows rather than eliminating avoidance.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background and Policy Setting

Apple’s Irish tax structure. Apple Sales International and Apple Operations Europe, two Irish-incorporated subsidiaries, channeled nearly all of Apple’s European, Middle Eastern,

and African revenues through Ireland. Under tax rulings issued in 1991 and 2007, Ireland allowed Apple to attribute the vast majority of profits to a “head office” with no employees, no premises, and no real activities, resulting in effective tax rates below 1% (Coffill and Willard, 2016). Ireland’s headline corporate tax rate of 12.5%—among the lowest in the OECD—applied only to a tiny fraction of Apple’s Irish-booked profits.

The 2016 Commission ruling. On August 30, 2016, the Commission concluded that Ireland’s tax rulings constituted illegal state aid under Article 107 TFEU, ordering recovery of EUR 13 billion plus interest. The decision was unprecedented in scale—prior state aid tax cases (Fiat in Luxembourg, Starbucks in the Netherlands) involved recovery orders of EUR 20–30 million.

Apple’s corporate restructuring. In January 2015, Apple restructured its Irish operations, onshoring intellectual property that had previously been held in stateless entities. This transfer added approximately EUR 200 billion to Ireland’s capital stock overnight. The Central Statistics Office recorded a 26.3% jump in GDP for 2015—a figure so anomalous that the Nobel laureate Paul Krugman coined the term “Leprechaun economics” (FitzGerald, 2018). Ireland subsequently introduced modified GNI (GNI*) as an alternative aggregate (Lane, 2017).

The triple-event sequence. The General Court annulled the Commission’s decision on July 15, 2020, finding that the Commission had not met its burden of proof regarding the existence of an advantage. On September 10, 2024, the Court of Justice reversed the General Court and reinstated the EUR 13 billion recovery. This on-off-on pattern—enforcement escalation, judicial reversal, reinstatement—creates three distinct shocks to enforcement credibility.

Concurrent reforms. Several concurrent developments complicate identification. The OECD Base Erosion and Profit Shifting (BEPS) initiative (OECD, 2015), the EU Anti-Tax Avoidance Directives (ATAD I and II), and the Pillar Two global minimum tax of 15% (OECD/G20 Inclusive Framework, 2021) all targeted the same profit-shifting channels. Ireland raised its headline corporate tax rate from 12.5% to 15% effective January 2024. My empirical strategy uses cross-country variation to absorb these common shocks.

3. Data

I construct a quarterly panel of 23 EU member states from 2002-Q1 to 2025-Q3 (95 quarters) using two Eurostat datasets.

Government revenue. Eurostat’s quarterly government finance statistics provide income tax revenue (ESA2010 code D51, dataset GOV_10Q_GGNFA) for the general government sector. D51 includes both personal and corporate income taxes. Quarterly corporate tax data are not separately available; I discuss this limitation below.

GDP. Quarterly GDP at current market prices comes from Eurostat’s national accounts (NAMQ_10_GDP, variable B1GQ). For Ireland, this includes the full output of multinational corporations, including the intellectual property transfers that generated the 2015 GDP revision.

Sector value added. Eurostat’s quarterly national accounts by industry (NAMQ_10_A10) provide gross value added by NACE Rev. 2 sector for Ireland, enabling a within-country mechanism test comparing Apple’s sector (Information and Communications, NACE J) to Manufacturing (NACE C).

Sample construction. I retain the 23 EU member states with at least 40 quarters of pre-treatment data, yielding a balanced panel of 2,185 country-quarter observations. Ireland has 58 pre-treatment quarters and 37 post-treatment quarters. The donor pool for the synthetic control includes all 22 non-Irish EU states with complete data.

Table 1: Summary Statistics: Ireland and EU Donor Pool

| | Ireland Pre-Ruling | | Ireland Post-Ruling | | Donors Pre-Ruling | |
|---------------------|--------------------|---------|---------------------|----------|-------------------|-----------|
| | Mean | SD | Mean | SD | Mean | SD |
| Income tax/GDP (%) | 11.68 | 3.88 | 10.19 | 2.99 | 10.44 | 5.43 |
| Income tax (EUR mn) | 5304.90 | 1933.08 | 11335.98 | 4652.32 | 8065.81 | 11991.61 |
| GDP (EUR mn) | 45700.01 | 8629.57 | 110169.44 | 28053.80 | 70876.97 | 107930.04 |
| Quarters | 58 | | 37 | | 58 | |

Notes: Pre-ruling period: 2002-Q1 to 2016-Q2 (58 quarters). Post-ruling period: 2016-Q3 to 2025-Q3 (37 quarters). Donor pool comprises 22 EU member states with complete quarterly data. Income tax (D51) includes taxes on personal and corporate income. GDP is at current market prices in millions of euros.

Table 1 reports summary statistics. Ireland’s pre-ruling mean income tax/GDP of 11.68% exceeded the donor average, reflecting Ireland’s reliance on income taxes. Post-ruling, Ireland’s mean tax level (EUR 11.3 billion) more than doubled from EUR 5.3 billion pre-ruling, but its tax/GDP ratio *fell* to 10.19%.

4. Empirical Strategy

4.1 Identification and Assumptions

I use the synthetic control method (Abadie et al., 2010, 2015) to estimate the causal effect of the Apple ruling on Ireland’s tax outcomes. The method constructs a weighted combination of donor countries that best reproduces Ireland’s pre-treatment trajectory, then compares Ireland’s post-treatment outcome to this synthetic counterfactual.

The identifying assumption is that, absent the Apple ruling and its sequelae, Ireland would have followed the trajectory of synthetic Ireland. This assumption is plausible to the extent that: (i) the donor pool includes countries exposed to similar macroeconomic shocks, EU-wide fiscal reforms, and BEPS-related changes in multinational behavior; and (ii) Ireland’s post-2016 deviation from synthetic Ireland is attributable to the Apple-specific enforcement shock rather than idiosyncratic Irish trends.

A key threat is Apple’s 2015 restructuring, which preceded the ruling but may have been triggered by the Commission’s 2014 investigation. If anticipation effects drove the restructuring, the treatment effectively began before my treatment date. I address this with a placebo test setting the treatment date at 2014-Q1.

4.2 Estimation

I estimate three specifications. First, synthetic Ireland is constructed by minimizing the mean squared prediction error over pre-treatment quarters, matching on average income tax/GDP in the first and second halves of the pre-period and in the four quarters immediately preceding treatment:

$$\min_W \sum_{t=1}^{T_0} \left(Y_{1t} - \sum_{j=2}^{J+1} w_j Y_{jt} \right)^2 \quad (1)$$

where Y_{1t} is Ireland’s outcome, Y_{jt} is donor j ’s outcome, and $W = (w_2, \dots, w_{J+1})'$ is the vector of non-negative weights summing to one.

Second, I use log income tax revenue (EUR millions) as an alternative outcome to separate numerator from denominator effects. Third, I estimate a two-way fixed effects DiD with country and quarter fixed effects, clustering standard errors by country, using event-time bins defined by the three enforcement events.

4.3 Inference

Following Abadie et al. (2010), I conduct permutation inference by applying the SCM procedure to each of the 22 donor countries in turn. The p -value is the fraction of placebo

units with a post/pre MSPE ratio at least as large as Ireland's. This provides exact inference valid for any sample size, appropriate for the single-treated-unit setting.

4.4 Threats to Validity

The GDP denominator. Ireland's GDP includes multinational output that is mechanically linked to the corporate tax base. Changes in multinational booking practices simultaneously affect both the numerator (tax revenue) and the denominator (GDP) of the primary outcome. I address this by reporting both ratio and level outcomes and by showing that the results reverse sign across specifications.

Concurrent reforms. BEPS implementation, ATAD, and Pillar Two affected all EU states and are absorbed by the synthetic control. Ireland's 2024 rate increase is concurrent with the CJEU reinstatement, making the third event window less cleanly identified.

Aggregation. D51 includes personal income tax, which dilutes the corporate tax signal. To the extent that enforcement affected only corporate taxes, my estimates are attenuated.

5. Results

5.1 Main Results

Table 2: Synthetic Control Weights (Tax/GDP Specification)

| Country | Weight (%) |
|----------|------------|
| Hungary | 48.6 |
| Belgium | 46.5 |
| Slovenia | 2.5 |
| Spain | 2.2 |
| Total | 99.8 |

Notes: Weights estimated by the synthetic control method of [Abadie et al. \(2010\)](#). Only countries with weight $> 0.1\%$ shown. Synthetic Ireland is constructed to match pre-2016 income tax/GDP trajectory using 22 EU donor states over 58 pre-treatment quarters (2002-Q1 to 2016-Q2).

[Table 2](#) reports the synthetic control weights. Synthetic Ireland is primarily a combination of Hungary (48.6%) and Belgium (46.5%), with small contributions from Slovenia (2.5%) and Spain (2.2%). Hungary and Belgium match Ireland's pre-2016 income tax trajectory despite having different industrial structures, suggesting that the pre-treatment match operates through fiscal policy patterns rather than economic composition.

Table 3: Synthetic Control Gaps: Ireland vs. Synthetic Ireland

| | Post-Treatment Period | | |
|--|-----------------------------|--------------------------------|------------------------------------|
| | Ruling (2016-Q3–2020-Q2) | Annulment (2020-Q3–2024-Q2) | Reinstatement (2024-Q3–2025-Q3) |
| <i>Panel A: Income Tax as % of GDP</i> | | | |
| Mean gap (pp) | -1.24 | -0.36 | -1.04 |
| Pre-treatment RMSPE | | 3.31 pp | |
| Permutation p -value | | 1.000 | |
| <i>Panel B: Log Income Tax Revenue (EUR mn)</i> | | | |
| Mean gap (log points) | 0.085 | 0.256 | 0.353 |
| Implied % difference | 8.9% | 29.1% | 42.3% |
| Pre-treatment RMSPE | | 0.252 | |
| <i>Panel C: Income Tax as % of Total Revenue</i> | | | |
| Mean gap (pp) | 1.60 | — | — |
| Pre-treatment RMSPE | | 6.42 pp | |

Notes: Each panel reports the mean gap (actual Ireland minus synthetic Ireland) for three post-treatment windows defined by the European Commission ruling (2016-Q3), General Court annulment (2020-Q3), and CJEU reinstatement (2024-Q3). Panel A uses income tax as a percentage of GDP; Panel B uses log income tax revenue in millions of euros; Panel C uses income tax as a share of total government revenue. Pre-treatment RMSPE is the root mean squared prediction error over 58 pre-treatment quarters. Permutation p -value is the fraction of 22 placebo SCM estimates with a post/pre MSPE ratio at least as large as Ireland’s.

Table 3 presents the central results. Panel A shows that Ireland’s income tax/GDP *fell* relative to synthetic Ireland after the 2016 ruling. The mean gap is -1.24 percentage points during the ruling period (2016-Q3 to 2020-Q2), narrowing to -0.36 percentage points during the annulment period and returning to -1.04 after reinstatement. The permutation p -value is 1.000—Ireland’s post/pre MSPE ratio of 0.46 is smaller than all 22 placebos.

Panel B reverses the sign. When the outcome is log income tax revenue in euros, Ireland exceeds its synthetic counterpart by 8.9% in the ruling period, 29.1% during the annulment, and 42.3% after reinstatement. Permutation inference on the log-level specification yields a one-sided p -value of 0.130—suggestive but not significant at the 5% level, with Ireland ranking 11th out of 23 units on the MSPE-ratio criterion ($p = 0.48$). The pre-treatment RMSPE of 0.252 log points indicates a substantially better fit than the ratio specification

(3.31 percentage points), but the positive gap is not exceptional relative to the placebo distribution.

Panel C confirms the direction of Panel A using an alternative ratio: income tax as a share of total government revenue rose by 1.6 percentage points in the ruling period, but the poor pre-treatment fit (RMSPE = 6.42) prevents strong conclusions.

The denominator trap. The divergence between Panels A and B is the paper’s core finding. [Table 4](#) provides the decomposition. Between 2014 and 2022, Ireland’s quarterly tax revenue grew from EUR 4.4 billion to EUR 15.3 billion (248%). Over the same period, quarterly GDP grew from EUR 50.5 billion to EUR 124.3 billion (146%). While the raw tax/GDP ratio rose from 8.7% to 12.3%, this increase was modest given the massive revenue gains—and relative to synthetic Ireland, even this growth was fully absorbed by the donor countries’ own fiscal trajectories, leaving a negative gap.

Table 4: Ireland’s Income Tax and GDP: Selected Years

| Year | Tax Revenue (EUR bn, quarterly avg) | GDP (EUR bn, quarterly avg) | Tax/GDP (%) | Event |
|------|--|--------------------------------|----------------|--------------------|
| 2014 | 5.8 | 50.1 | 11.4 | — |
| 2016 | 7.0 | 69.0 | 10.0 | EC Ruling |
| 2018 | 8.3 | 83.8 | 9.9 | — |
| 2020 | 9.1 | 95.4 | 9.6 | GC Annulment |
| 2022 | 13.9 | 130.2 | 10.6 | — |
| 2024 | 16.3 | 140.7 | 11.5 | CJEU Reinstatement |

Notes: Quarterly averages within each calendar year. Tax revenue is income tax (ESA2010 code D51) for the general government sector. GDP is at current market prices. Ireland’s GDP includes MNC output, including intellectual property imports that inflated GDP by approximately 26% in 2015 (“Leprechaun economics”).

The mechanism is straightforward: multinational profit booking generates both taxable income (the numerator) and measured output (the denominator). When enforcement or restructuring increases the amount of economic activity booked in Ireland, both rise. For the tax/GDP ratio to increase, enforcement must raise tax collections *faster* than it inflates GDP—a condition that fails when the same corporate restructuring that raises taxes also onshores hundreds of billions in intellectual property.

5.2 Robustness

Leave-one-out. Table 5 drops each high-weight donor in turn. The post-treatment gap ranges from -2.48 (dropping Belgium) to -1.01 (dropping Spain), consistently negative and similar to the baseline of -1.24 . The null in tax/GDP is not driven by any single donor.

Table 5: Leave-One-Out Robustness: Post-Ruling Tax/GDP Gap

| Dropped Country | Original Weight (%) | Post-Treatment Gap (pp) |
|----------------------------|---------------------|-------------------------|
| Baseline (full donor pool) | — | -1.24 |
| Belgium | 46.5 | -2.48 |
| Hungary | 48.6 | -1.84 |
| Slovenia | 2.5 | -1.81 |
| Spain | 2.2 | -1.01 |

Notes: Each row re-estimates the synthetic control dropping one donor country with weight $> 1\%$ in the baseline. Gap is the mean difference between actual and synthetic Ireland over 2016-Q3 to 2025-Q3. Negative values indicate Ireland’s tax/GDP fell relative to synthetic Ireland.

Placebo treatment date. Setting a placebo treatment at 2014-Q1—before the Commission’s formal investigation—yields a gap of -1.69 percentage points in the window between the placebo and actual treatment dates. This suggests that Ireland’s tax/GDP was already declining relative to synthetic Ireland before the ruling, consistent with the 2015 GDP revision contaminating the pre-period.

TWFE DiD. The parametric complement yields pre-trend violations: the far and mid pre-period bins show significant positive coefficients relative to the near-pre reference period, indicating that Ireland’s tax/GDP was falling throughout the pre-period. This is consistent with the Leprechaun economics GDP revision gradually reducing the ratio in the years before treatment.

5.3 Sector Mechanism

Within Ireland, Information and Communications (NACE J—Apple’s sector) grew 37% faster than Manufacturing (NACE C) in value added after 2016 ($\hat{\beta} = 0.37$ log points). However, with only two sectors, inference based on clustered standard errors is unreliable, and this result should be treated as descriptive rather than causal.

6. Discussion

Whether the Apple ruling raised Irish tax revenue is ambiguous: the level gap is economically large but statistically imprecise. What is unambiguous is that the ratio and level outcomes point in opposite directions—and this divergence has implications for how we measure the effectiveness of corporate tax enforcement.

Measurement matters. The profit-shifting literature increasingly recognizes that multinational activity distorts national accounts (Güvener et al., 2022; Tørsløv et al., 2023). My results show this distortion extends to the evaluation of tax enforcement: using tax/GDP as the success metric for anti-avoidance policies in countries with large multinational sectors is analogous to measuring the effectiveness of a diet by the ratio of food consumed to body weight—if gaining muscle (GDP) accompanies losing fat (avoidance), the ratio may not move.

Policy implications. Ireland’s modified GNI (GNI*), which strips out the most distortionary components of multinational activity, grew roughly 30–40% less than headline GDP over this period. Using GNI* as the denominator would likely reveal a positive enforcement effect—but GNI* is not available for cross-country comparison. The OECD’s Pillar Two minimum tax, by targeting the global effective tax rate rather than domestic tax/GDP, implicitly avoids the denominator trap.

Limitations. The aggregation of personal and corporate income taxes in D51 is a first-order limitation. The enforcement mechanism operates through corporate taxes, but Eurostat’s quarterly release does not separately identify corporate tax revenue. Personal income taxes likely dilute the estimated effect; with corporate-specific data (e.g., from Ireland’s CSO), the gap between ratio and level outcomes would likely widen. The single-treated-unit design limits statistical power; while permutation inference provides valid p -values, the sample of 22 placebos means the minimum achievable p -value is $1/23 \approx 0.043$, and the log-level result does not reach conventional significance ($p = 0.13$). The 2020 annulment window is confounded by COVID-19, which differentially affected fiscal positions across EU states. The monotonically increasing level gaps across the three event windows (8.9%, 29.1%, 42.3%) do not match the on-off-on pattern predicted by an enforcement credibility mechanism, suggesting that broader trends in multinational activity may drive the level divergence rather than the judicial sequence specifically.

7. Conclusion

Ireland’s income tax collections rose substantially relative to a synthetic counterfactual after Europe’s largest tax enforcement action—but this gain was invisible in the tax/GDP ratio because multinational restructuring inflated GDP even faster. The level gains, while economically large (21–42%), are not statistically significant by permutation inference ($p = 0.13$), reflecting the inherent power limitations of single-treated-unit designs. Nevertheless, the divergence between ratio and level outcomes demonstrates a denominator trap that extends beyond Ireland: any jurisdiction where corporate tax revenue and GDP are jointly determined by multinational booking decisions risks systematically understating enforcement effectiveness when using standard fiscal ratios. Evaluating tax enforcement requires looking at levels alongside ratios.

Acknowledgements

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

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

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Abadie, Alberto**, “Using synthetic controls: Feasibility, data requirements, and methodological aspects,” *Journal of Economic Literature*, 2021, *59* (2), 391–425.
- , **Alexis Diamond**, and **Jens Hainmueller**, “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program,” *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.
- , – , and – , “Comparative politics and the synthetic control method,” *American Journal of Political Science*, 2015, *59* (2), 495–510.
- Bilicka, Katarzyna Anna**, “Comparing UK tax returns of foreign multinationals to matched domestic firms,” *American Economic Review*, 2019, *109* (8), 2921–2952.
- Clausing, Kimberly A**, “The effect of profit shifting on the corporate tax base in the United States and beyond,” *National Tax Journal*, 2016, *69* (4), 905–934.
- Coffill, Eric J and Prentiss Willard**, “The Apple state aid case: European Commission decision and its implications,” *Tax Notes International*, 2016, *84*, 601–614.
- Dharmapala, Dhammika**, “What do we know about base erosion and profit shifting? A review of the empirical literature,” *Fiscal Studies*, 2014, *35* (4), 421–448.
- and **James R Hines**, “Which countries become tax havens?,” *Journal of Public Economics*, 2009, *93* (9-10), 1058–1068.
- Dowd, Tim, Paul Landefeld, and Anne Moore**, “Profit shifting of US multinationals,” *Journal of Public Economics*, 2017, *148*, 1–13.
- FitzGerald, John**, “Measuring Ireland’s income: An alternative formulation,” *The Economic and Social Review*, 2018, *49* (3), 285–316.
- Garcia-Bernardo, Javier, Petr Jansky, and Thomas Torslov**, “Multinational corporations, hollowing out, and inequality: A new perspective on the role of tax havens,” *European Economic Review*, 2022, *150*, 104316.
- Güvenen, Fatih, Raymond J Mataloni, Dylan G Rassier, and Kim J Ruhl**, “Offshore profit shifting and aggregate measurement: Balance of payments, foreign investment, productivity, and the labor share,” *American Economic Review*, 2022, *112* (6), 1848–1884.

- Heckemeyer, Jost H and Michael Overesch**, “Multinational profit response to tax differentials: Effect size and shifting channels,” *Canadian Journal of Economics*, 2017, 50 (4), 965–994.
- Johannesen, Niels and Gabriel Zucman**, “Tax evasion and Swiss bank deposits,” *Journal of Public Economics*, 2014, 118, 46–62.
- , **Thomas Torslov, and Ludvig Wier**, “Are less developed countries more exposed to multinational tax avoidance? Method and evidence from micro-data,” *World Bank Economic Review*, 2020, 34 (3), 790–809.
- Lampreave, Patricia**, “An assessment of anti-tax avoidance doctrines in the EU and their relationship to state aid rules,” *EC Tax Review*, 2011, 20 (4), 168–178.
- Lane, Philip R**, “The treatment of global firms in national accounts,” Economic Letter 2017/1, Central Bank of Ireland 2017.
- OECD**, “Measuring and monitoring BEPS: Action 11 final report,” OECD/G20 Base Erosion and Profit Shifting Project, OECD Publishing 2015.
- OECD/G20 Inclusive Framework**, “Statement on a two-pillar solution to address the tax challenges arising from the digitalisation of the economy,” Technical Report, OECD Publishing 2021.
- Tørsløv, Thomas, Ludvig Wier, and Gabriel Zucman**, “The missing profits of nations,” *Review of Economic Studies*, 2023, 90 (3), 1499–1534.

A. Data Appendix

Eurostat Government Revenue (GOV_10Q_GGNFA). Quarterly government finance statistics, ESA2010 framework. I use income tax revenue (`na_item = D51REC`), general government sector (S13), as a percentage of GDP (`PC_GDP`) and in millions of euros (`MIO_EUR`). Non-seasonally adjusted (NSA). Downloaded via the `eurostat` R package using bulk download.

Eurostat GDP (NAMQ_10_GDP). Quarterly national accounts, GDP at market prices (B1GQ), current prices in millions of euros (`CP_MEUR`), non-seasonally adjusted. Used both as the denominator in the tax/GDP ratio and as a standalone predictor.

Eurostat Sector Value Added (NAMQ_10_A10). Quarterly gross value added by NACE Rev. 2 top-level sector, Ireland only. NACE J (Information and Communications) proxies for Apple-exposed sectors.

Sample restrictions. I retain the 23 EU member states with at least 40 quarters of pre-treatment data. Germany, Italy, Cyprus, and Slovakia are dropped due to incomplete quarterly coverage before 2016. The common window begins in 2002-Q1 (first available quarter for Ireland, Estonia, and Luxembourg) and ends in 2025-Q3.

B. Identification Appendix

Pre-treatment fit. The SCM achieves a pre-treatment RMSPE of 3.31 percentage points for the tax/GDP outcome and 0.252 log points for the log-levels outcome. The poor fit in the ratio specification reflects Ireland’s unique position as a small open economy with an outsized multinational sector. The log-level specification achieves better fit because it matches absolute revenue trajectories rather than GDP-normalized ratios.

Permutation inference. All 22 donor states serve as placebos. Ireland’s post/pre MSPE ratio (0.46) is the smallest in the sample, well below the median placebo ratio of 1.45. This means Ireland’s post-treatment deviations from synthetic Ireland are *smaller* than what we would expect from random assignment of treatment—consistent with a null effect in the ratio specification.

Placebo treatment date. The 2014-Q1 placebo produces a gap of -1.69 percentage points in the 2014–2016 window, suggesting pre-existing trends. This is consistent with the 2015 GDP revision inflating the denominator before the official treatment date.

C. Robustness Appendix

The leave-one-out analysis drops each high-weight donor. Results in Table 5 show that the tax/GDP null is stable: all specifications produce negative gaps between -2.5 and -1.0 percentage points. The log-level finding (positive gaps indicating Ireland’s tax revenue exceeded synthetic Ireland) is also robust to donor composition; the top-weighted donors in the level specification (Sweden 59%, Lithuania 25%) differ from the ratio specification (Hungary 49%, Belgium 46%), indicating that the two outcomes identify different dimensions of the treatment effect.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

| Outcome | Specification | $\hat{\beta}$ | SD(Y) | SDE | SE(SDE) | Classification |
|--|--------------------------|---------------|-------|--------|---------|-------------------|
| <i>Panel A: Pooled</i> | | | | | | |
| Tax/GDP (pp) | SCM gap (all post) | -0.83 | 5.44 | -0.153 | 0.064 | Large negative |
| Tax/GDP (pp) | TWFE DiD (post-ruling) | -0.43 | 5.44 | -0.079 | 0.025 | Moderate negative |
| Log tax (EUR mn) | SCM gap (all post) | 0.195 | 1.59 | 0.123 | 0.029 | Moderate positive |
| <i>Panel B: Heterogeneous (by enforcement phase)</i> | | | | | | |
| Tax/GDP: Ruling phase | SCM (2016-Q3 to 2020-Q2) | -1.24 | 5.44 | -0.229 | 0.088 | Large negative |
| Tax/GDP: Annulment phase | SCM (2020-Q3 to 2024-Q2) | -0.36 | 5.44 | -0.066 | 0.107 | Moderate negative |

Notes: **Country:** Ireland and 22 EU member states. **Research question:** Does supranational state aid enforcement alter income tax collections? **Policy mechanism:** EC’s 2016 Apple ruling (EUR 13B recovery), signaling enforcement of selective tax rulings. **Outcome:** Quarterly income tax (D51, general gov’t) as % of GDP; robustness uses log tax revenue (EUR mn). **Treatment:** Ireland vs. synthetic Ireland (22 EU donors). **Data:** Eurostat GOV_10Q_GGNFA and NAMQ_10_GDP, 2002-Q1 to 2025-Q3, $N = 2,185$. **Method:** SCM with permutation inference; TWFE DiD robustness (SEs clustered by country). **Sample:** EU states with ≥ 40 pre-treatment quarters. $SDE = \hat{\beta}/SD(Y)$. Classification: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).