Stockholm School of Economics
Department of Economics
BE551 Degree project in economics
Spring 2026

Netflix's Ad-Supported Tier and Online Piracy in the European Union: A Synthetic Control Analysis

Ayham Khalid (25921)

Abstract

Digital piracy persists in the European Union despite the growth of legal subscription video-on-demand alternatives. This thesis investigates whether Netflix's lower-priced advertising-supported subscription tier, launched in November 2022 in France, Germany, Italy, and Spain, was followed by a reduction in online piracy. Using monthly piracy accesses per internet user from the EUIPO (2024) panel covering 2017-2023, the analysis applies the synthetic control method to each treated country separately and supplements the country-by-country estimates with a pooled estimate from the generalized synthetic control method with interactive fixed effects. The estimates are negative in all four treated countries, and the pooled average treatment effect of -1.278 piracy accesses per internet user per month has a 95 percent confidence interval that excludes zero. A leave-one-out donor pool check finds that France's estimate is robust to the exclusion of either dominant donor, while Germany's is not. The magnitude of any reduction cannot be precisely identified, in part because of concurrent enforcement activity and donor-pool contamination. The contribution is therefore narrow: a documented negative association across four countries with identified threats to its causal interpretation.

Keywords: synthetic control method; online piracy; video streaming; price discrimination; European Union

JEL: L82, L86, K42, O34

Supervisor:	Sampreet Goraya
Date submitted:	2026-05-11
Date examined:	2026-05-19
Discussants:	Yicheng Sun and Marthe-Sarah Nkambou Tchunkuo
Examiner:	Johanna Wallenius

Table of Contents

1. Introduction

Digital piracy remains a persistent feature of online media consumption in the European Union. EUIPO (2024) reports that internet users in the EU accessed pirated content approximately ten times per month on average between 2017 and 2023, with television content accounting for the largest share. Aggregate piracy declined gradually over the period but did not approach elimination. The scale of the phenomenon, and its persistence in the face of growing legal streaming alternatives, makes the relationship between legal-content pricing and unauthorised consumption an open empirical question.

Netflix is the dominant subscription video-on-demand platform in the European Union. On November 3, 2022 it launched a lower-priced advertising-supported subscription tier, "Basic with Ads," in France, Germany, and Italy, and on November 10, 2022 in Spain. Twenty-three other EU member states did not receive the tier during the 2017-2023 data window of this study. The contemporaneous launch in four member states and the absence of the tier in the remaining EU panel members creates an opportunity to study the consumer response to a within-platform pricing change, using the synthetic control approach to construct counterfactual piracy trajectories for the treated countries.

The empirical literature on legal streaming and piracy has examined the entry, exit, and experimental introduction of subscription platforms. Lu, Rajavi, and Dinner (2021) study the blocking of Netflix in Indonesia, and Godinho de Matos, Ferreira, and Smith (2018) conduct a household-level randomised experiment in which a telecommunications provider gifts subscription video access to identified piracy users. Both studies find that the availability of legal streaming reduces unauthorised consumption. The contribution of the ad-tier launch is distinct: the platform is already available in all four treated countries, and only the price of legal access changes. Whether a within-platform pricing change generates a substitution response comparable to the introduction of legal access is not addressed by the existing literature.

The design faces several identification challenges. First, the pooled synthetic estimate exhibits a downward drift in the months preceding the November 2022 treatment date, consistent with either consumer anticipation of Netflix's pre-announced launch or features of the interactive fixed effects estimation framework. Second, none of the four treated countries provides a fully clean experimental setting. Spain received a concurrent Netflix password-sharing crackdown in February 2023; France faces both donor-pool exposure to Disney+ and a contemporaneous intensification of site-blocking activity by the French audiovisual regulator ARCOM; Italy and Germany face donor-pool Disney+ exposure. Third, Operation 404, a multi-phase international anti-piracy enforcement operation coordinated from Brazil, conducted two of its six phases within the fourteen-month post-treatment window of this study. These confounds are developed in detail in Chapter 5.

The empirical strategy applies the synthetic control method of Abadie, Diamond, and Hainmueller (2010) separately to each of the four treated countries, with monthly TV and film piracy accesses per internet user as the outcome variable. The donor pool comprises nineteen EU member states. The pre-treatment window spans seventy months from January 2017 to October 2022, and the post-treatment window spans fourteen months from November 2022 to

December 2023. The country-by-country estimates are supplemented by a pooled estimate from the generalized synthetic control method of Xu (2017), which accommodates multiple treated units and time-varying covariates within an interactive fixed effects framework. A leave-one-out donor pool robustness check examines the sensitivity of each country's estimate to the dropping of either of the two dominant donor countries.

The country-by-country estimates yield negative point estimates in all four treated countries. The leave-one-out check finds that France's estimate is robust to the exclusion of either dominant donor, Germany's estimate is not, and Italy and Spain show partial robustness. The evidence is consistent with a reduction in piracy following the ad-tier launch, but the magnitude of any such reduction cannot be precisely identified from this design, and a portion of it is plausibly attributable to concurrent enforcement-side activity rather than to the ad-tier in isolation.

The contribution of this thesis is therefore narrow but specific. It documents a consistent negative association between the introduction of Netflix's advertising-supported subscription tier and online piracy in four European Union member states, identifies the principal threats to a causal interpretation of that association, and shows through a leave-one-out donor check that the association is robust for France and fragile for Germany. To the author's knowledge, no prior published study has applied the synthetic control method to a within-platform pricing intervention in the streaming sector across multiple European Union member states.

The remainder of the thesis is structured as follows. Chapter 2 reviews the relevant literatures on digital piracy, on streaming substitution, and on the synthetic control method. Chapter 3 develops the methodology and describes the data. Chapter 4 reports the country-level and pooled results. Chapter 5 discusses the identification concerns and reports the robustness check. Chapter 6 concludes.

2. Literature Review

This chapter situates the present study within three strands of literature: the empirical and theoretical literature on digital piracy and its determinants, the empirical literature on legal streaming as a substitute for unauthorised consumption, and the methodological literature on the synthetic control approach. The chapter closes by positioning the contribution of this thesis.

2.1 Digital Piracy: Scale and Theoretical Determinants

Digital piracy remains a persistent feature of online media consumption in the European Union. EUIPO (2024) reports that internet users in the EU accessed pirated content approximately ten times per month on average between 2017 and 2023, with television content accounting for roughly half of all accesses and film for a smaller but stable share. Aggregate piracy declined gradually over the sample period but did not approach elimination, and piracy of live sports content increased.

The theoretical literature explains these patterns through a small set of channels. Belleflamme and Peitz (2012), in their chapter on digital piracy theory in the Oxford Handbook of the Digital Economy, decompose the consumer's choice between legal and unauthorised consumption into four components: the relative prices of the two options, the availability of legal content, the perceived risk of enforcement, and a moral or hedonic cost of unauthorised access. A reduction in the price of legal content shifts the marginal consumer from unauthorised toward authorised consumption; an increase in enforcement intensity raises the expected cost of piracy; and the introduction of price discrimination, such as a lower-priced advertising-supported tier, expands the set of consumers for whom legal consumption is preferred to the outside option. The present study focuses on this last channel.

The empirical literature on enforcement provides complementary evidence on the magnitude of the determinants identified by Belleflamme and Peitz (2012). Adermon and Liang (2014) study the Swedish implementation of the European Union's IPRED directive in April 2009, which sharply increased the perceived risk of detection for unauthorised file sharing. Using a difference-in-differences design with Norway and Finland as control units, they report a 16 percent reduction in internet traffic and a 36 percent increase in music sales over the first six months following the reform. The effect attenuated over subsequent months, which the authors interpret as evidence that the perceived risk of enforcement, rather than the legal change itself, is what caused the change. This finding is relevant to the present study because the enforcement-side concerns developed in Chapter 5 discuss a similar mechanism, namely the contemporaneous intensification of national and international anti-piracy activity during the post-treatment window.

2.2 Legal Streaming and Piracy: Empirical Evidence on Substitution

The closest empirical precedent for the present study is Lu, Rajavi, and Dinner (2021), who exploit Netflix's January 2016 expansion to 130 countries and the contemporaneous blocking of the service in Indonesia by the country's principal telecommunications provider. Using the synthetic control method with 40 Asian countries that received Netflix as the donor pool, the authors estimate that the blocking of Netflix in Indonesia produced a 19.7 percent increase in piracy-related search activity relative to the synthetic counterfactual. The study identifies the availability of legal streaming as a substitute for unauthorised consumption and provides the methodological template that the present analysis adapts: a treated unit with a clear treatment date, a donor pool of comparable units, and SCM as the identification strategy. The present study departs from Lu et al. in two respects. The treatment of interest is a pricing change within an existing platform rather than the entry or exit of the platform itself, and the design exploits four treated countries simultaneously rather than one, which permits pooled estimation through the generalized synthetic control framework discussed below.

A second strand of evidence comes from household-level experimental work. Godinho de Matos, Ferreira, and Smith (2018) partner with a major telecommunications provider to randomly allocate forty-five days of free SVOD-style access to households previously identified as BitTorrent users. They find that the average treated household reduced internet downloads and uploads by approximately 4 percent, and that households whose content preferences aligned with the gifted bundle reduced their BitTorrent usage by 18 percent. The mechanism the authors emphasise is availability and content fit. The present study is consistent with their finding that legal access can substitute for unauthorised consumption, but extends the question by examining whether a reduction in the price of legal access, rather than the introduction of access itself, generates a comparable response among consumers who already had access to Netflix but not at the ad-tier's price point.

Music piracy provides a parallel literature with longer time series. Aguiar and Waldfogel (2018) examine whether the growth of paid streaming services displaces unauthorised music consumption between 2013 and 2015, drawing on aggregate streaming and sales data. They report that losses from displaced sales are roughly outweighed by gains in streaming revenue and that streaming also displaces music piracy. Although the music and video industries differ in piracy modalities and rights structures, the substitution mechanism, that low-price authorised access reduces demand for unauthorised consumption, is shared with the video setting examined here.

2.3 The Synthetic Control Method in Applied Economics

The synthetic control method was introduced by Abadie and Gardeazabal (2003) in a study of the economic costs of conflict in the Basque Country and formalised by Abadie, Diamond, and Hainmueller (2010) in their analysis of the effect of California's Proposition 99 tobacco control program. The method addresses a recurrent problem in comparative case studies, namely that a single treated unit may have no individually similar untreated unit available as a counterfactual. The synthetic control instead constructs a weighted combination of donor units that approximates the treated unit's pre-treatment trajectory, and the post-treatment difference between the actual and synthetic series identifies the treatment effect under assumptions that have been clarified in subsequent methodological work.

Abadie (2021) reviews the literature and emphasises three requirements for credible SCM estimates: a long pre-treatment period during which the treated and donor units can be matched, a donor pool whose members were not themselves affected by the treatment, and a treated unit lying within the convex hull of the donor pool in terms of pre-treatment characteristics. Kaul, Klössner, Pfeifer, and Schieler (2022) examine the predictor-choice problem and recommend against the use of all individual pre-treatment outcomes as predictors when auxiliary covariates are also included, as this combination tends to nullify the predictive role of the covariates and overfit the pre-treatment period. They propose the use of averaged pre-treatment outcome windows together with covariates as a balance between fit quality and identification credibility; this study follows this recommendation in its three-window specification (Section 3.3).

The classical SCM accommodates a single treated unit and a discrete treatment date. Xu (2017) generalises the framework through an interactive fixed effects estimator, producing the generalized synthetic control method (GSC). The GSC pools information across multiple treated units, accommodates time-varying covariates and unobserved common factors, and yields a single pooled average treatment effect on the treated with parametric or bootstrap inference. The present study uses GSC alongside the country-by-country classical SCM, with the pooled estimate reported in Section 4.3.

2.4 Position of the Present Study

The three literatures above intersect in this thesis. From the empirical piracy literature, the analysis adopts the affordability channel identified by Belleflamme and Peitz (2012) as the candidate mechanism through which Netflix's ad-supported tier might reduce piracy, and treats the enforcement channel documented by Adermon and Liang (2014) as a competing explanation to be considered in the limitations of Chapter 5. From the empirical streaming-substitution literature, the analysis adopts the substitution hypothesis tested in Lu, Rajavi, and Dinner (2021) and Godinho de Matos, Ferreira, and Smith (2018), but it differs from those studies in examining a price intervention within an existing platform rather than the entry, exit, or experimental introduction of a platform. From the SCM methodological literature, it adopts the classical Abadie-Diamond-Hainmueller (2010) design supplemented by the Xu (2017) generalized synthetic control, with predictor-window choices guided by Kaul et al. (2022) and the donor-pool requirements articulated in Abadie (2021).

3. Methodology

3.1 Research Design

This study employs the Synthetic Control Method (SCM) to estimate the causal effect of Netflix's ad-supported subscription tier on online piracy in the European Union. The SCM, developed by Abadie and Gardeazabal (2003) and formalised in Abadie, Diamond, and Hainmueller (2010), constructs a data-driven counterfactual for each treated country by computing a weighted combination of untreated countries whose pre-treatment piracy trajectories most closely resemble those of the treated unit. This approach is suited to the setting of this study: a small number of treated units (four EU countries), a large panel of potential comparison countries (the remaining EU member states), and a long pre-treatment period (70 monthly observations).

The principal alternative, difference-in-differences (DiD), requires that treated and control units follow parallel outcome trends in the absence of treatment. As Abadie (2021) emphasises, this restriction is unlikely to hold when piracy trends across EU countries exhibit heterogeneous trajectories driven by country-specific factors such as income levels, digital infrastructure, and cultural attitudes toward copyright compliance. The SCM relaxes the parallel trends assumption by allowing outcomes to depend on multiple unobserved factors with time-varying coefficients, provided that a weighted combination of control units can approximate the treated unit's pre-treatment trajectory. The country-by-country SCM is supplemented with the Generalized Synthetic Control (GSC) method of Xu (2017), which pools information across all treated and control units to produce a single average treatment effect on the treated (ATT).

3.2 Data

3.2.1 Piracy Data

The dependent variable is monthly online piracy accesses per internet user, measured separately for TV and Film content and summed to produce a combined piracy measure. The data are sourced from the European Union Intellectual Property Office's report "Online Copyright Infringement in the European Union: Films, Music, Publications, Software and TV (2017–2023)" (EUIPO 2024). The underlying data were collected by MUSO, a digital piracy monitoring firm, which tracks traffic to piracy websites and measures accesses associated with the consumption of copyright-infringing content. The report covers all 27 EU member states on a monthly basis from January 2017 to December 2023, yielding 84 time periods.

The EUIPO report presents piracy data exclusively as time-series charts, without accompanying numerical tables. To construct the panel dataset required for econometric analysis, monthly TV and Film piracy values were extracted from these charts using a custom Python pipeline that identifies data points through pixel-level colour analysis, calibrates the axes using gridline and tick-mark detection, and maps pixel positions to numerical values. Verification overlay images were produced for all 27 countries to confirm alignment between extracted data points and the original chart traces. Technical details of the extraction procedure are available upon request. The resulting dataset contains 2,268 observations (27 countries * 84 months) with zero missing values. Further information regarding the Python pipeline is available in the appendices.

3.2.2 Covariates

To improve the pre-treatment fit of the synthetic controls, four covariates are included, sourced from Eurostat. All covariates are measured at annual frequency and merged onto the monthly panel by year, so that each month within a given year receives the same covariate value. GDP per capita in purchasing power standards (tec00114) controls for economic conditions affecting both legal content demand and the opportunity cost of piracy. The youth unemployment rate for ages 15–24 (yth_empl_090) is included based on the EUIPO's (2024) finding that young, unemployed individuals exhibit significantly higher piracy rates. Household internet access (isoc_ci_in_h) proxies for the digital infrastructure that enables both legal streaming and illegal downloading. The Gini coefficient of income inequality (ilc_di12) captures distributional factors that may drive demand for low-cost or free alternatives to paid content.

For the generalized synthetic control estimation, broadband penetration (isoc_ci_it_h), the population share aged 15–29 (demo_pjangroup), tertiary education attainment (edat_lfse_03), and the harmonized index of consumer prices (prc_hicp_aind) are also included. The interactive fixed effects structure of the GSC method can accommodate a richer covariate set without the computational issues that constrain the classical SCM's predictor space.

3.2.3 Panel Structure

Table 1 summarizes the panel structure. Four EU member states received the Netflix ad-tier during the sample window and constitute the treated group.

Table 1: Panel Structure

Role	Countries	N
Treated	FR, DE, IT, ES	4
Excluded (crackdown)	PT	1
Excluded (SCM only)	CY, MT, LU	3
SCM donor pool	AT, BE, BG, CZ, DK, EE, EL, FI, HR, HU, IE, LT, LV, NL, PL, RO, SE, SI, SK	19

Notes: Portugal is excluded because it received the password-sharing crackdown without the ad-tier. Cyprus, Malta, and Luxembourg are excluded from the SCM donor pool because their small size and limited covariate variation prevent the optimization from finding a unique solution. They are included in the GSC panel.

3.3 The Synthetic Control Method

The SCM seeks a vector of non-negative weights $W = (w_1, \ldots, w_J)'$ summing to one, such that the weighted combination of donor countries' pre-treatment characteristics most closely approximates those of the treated country. Let X_1 denote the vector of pre-treatment predictors for the treated unit and X_0 the corresponding matrix for the donor pool. The optimal weights W^* minimise the distance $\|X_1 - X_0W\|$ subject to a diagonal weighting matrix V that reflects the relative importance of each predictor (Abadie, Diamond, and Hainmueller 2010). V is itself selected to minimise the mean squared prediction error (MSPE) of the synthetic control over the pre-treatment outcome series.

The predictor set includes three windows of lagged average piracy accesses and the four Eurostat covariates. The three windows cover time periods 1–24, 25–48, and 49–70, corresponding roughly to the pre-COVID baseline, the COVID era, and the post-lockdown normalisation. Following Kaul et al. (2022), averaged windows are used rather than all individual pre-treatment outcomes to reduce the risk of overfitting while preserving the trajectory's shape.

The treatment effect at post-treatment time t is estimated as the gap between the observed outcome and the synthetic counterfactual: $\hat{\alpha}_{1t} = Y_{1t} - \Sigma w^*_j Y_{jt}$. The model uses a convex combination constraint, which forces the "synthetic" country to stay within the boundaries of the actual donor data. This prevents the model from making unrealistic predictions that fall outside the range of what we observe in the real world. All SCM estimations are conducted using the tidysynth package in R.

3.4 Inference

Since the SCM involves a small number of treated units, conventional large-sample inference is inapplicable. The permutation-based approach of Abadie, Diamond, and Hainmueller (2010) is used, in which the SCM is re-estimated for each donor country as if it had received the treatment. The treated country's estimated effect is then compared to this distribution of placebo effects. To account for differences in pre-treatment fit quality, the ratio of post-treatment MSPE to pre-treatment MSPE for each unit is computed:

$$r_j = MSPE_{post}(j) \ / \ MSPE_{pre}(j)$$

The rank of the treated country's MSPE ratio among all units yields an exact p-value. With 19 donor countries plus one treated country, the pool contains 20 units, and the minimum attainable p-value is 1/20 = 0.050.

3.5 Generalized Synthetic Control

To complement the country-by-country SCM with a pooled estimate, the Generalized Synthetic Control (GSC) method of Xu (2017) is used. The GSC estimates counterfactuals within an interactive fixed effects framework:

$$Y_{it} = \delta_{it} D_{it} + X'_{it}\beta + \lambda'_i f_t + \varepsilon_{it}$$

where Y_{it} is piracy for country i at time t, δ_{it} is the treatment effect itself, D_{it} is a binary treatment indicator, X_{it} is a vector of observed covariates, f_t captures unobserved common factors (such as seasonal streaming demand or EU-wide economic shocks), and λ_i represents country-specific factor loadings. The number of latent factors r is selected via cross-validation; in this estimation, the optimal value is r = 5. The GSC produces a single pooled ATT with bootstrap-based confidence intervals (500 parametric bootstrap replications). Estimation is conducted using the gsynth package in R.

3.6 Identification Strategy

Treatment definition

The treatment event is Netflix's launch of the "Basic with Ads" subscription tier in four EU member states: France (November 3, 2022, €5.99/month), Germany (November 3, 2022, €4.99/month), Italy (November 3, 2022, €5.49/month), and Spain (November 10, 2022, €5.49/month). No other EU-27 country received the Netflix ad-tier during the 2017-2023 data window. A common treatment date of November 2022 (time period 71 of 84) is assigned, as the one-week difference between France/Germany/Italy and Spain is negligible at monthly frequency.

Password-sharing crackdown

A critical identification concern arises from Netflix's concurrent crackdown on password sharing. Spain and Portugal received the crackdown on February 8, 2023, three months after the ad-tier launch. The remaining countries, including France, Germany, Italy, and all donor countries, were subject to the global crackdown from May 2023 onward (time period 77).

This creates two complications. First, Spain's estimated effect after February 2023 reflects a compound treatment: the ad-tier combined with the crackdown. Second, the donor pool is contaminated from May 2023, as control countries are also affected. These concerns are addressed by excluding Portugal entirely, leaving France, Germany, and Italy as the primary cases for isolating the pure ad-tier effect.

Disney+ expansion

Several EU countries received Disney+ on June 14, 2022, approximately five months before the ad-tier launch. This concurrent expansion of legal streaming supply would bias the synthetic control toward lower piracy in donors, thereby understating the estimated ad-tier effect. Seven donor countries already had Disney+ prior to this date (Austria, Belgium, Denmark, Finland, Ireland, the Netherlands, and Sweden), constituting a somewhat "clean" subset for a planned robustness check.

Identifying assumptions

The validity of the design rests on three standard assumptions (Abadie 2021): (i) each treated country lies within the convex hull of the donor pool in terms of pre-treatment characteristics; (ii) the treatment did not affect donor units during the pre-treatment period (SUTVA); and (iii) no concurrent shock confounds the estimated effect.

3.7 Conceptual Framework

Netflix's ad-tier launch and password-sharing crackdown are interpreted as complementary components of a price discrimination strategy. The ad-tier functions as the "carrot": a lower-priced legal alternative for consumers whose willingness to pay falls below the standard subscription price. The crackdown functions as the "stick": raising the effective cost of unauthorised access. Together, these policies expand the legal consumption frontier while narrowing unauthorised alternatives, consistent with Belleflamme and Peitz (2012).

The empirical design exploits the staggered rollout to isolate respective contributions. France, Germany, and Italy received only the ad-tier (not the early crackdown), allowing estimation of the affordability channel in isolation. Spain received both within a three-month window, capturing the combined effect. The closest methodological precedent is Lu, Rajavi, and Dinner (2021), who use a synthetic control design to estimate the effect of Netflix availability on piracy search in Indonesia.

4. Results

This chapter presents the empirical results in three parts. Section 4.1 reports the country-by-country SCM estimates. Section 4.2 summarises cross-country patterns. Section 4.3 presents the pooled GSC estimate.

4.1 Country-by-Country SCM Results

For each treated country, the SCM produces a trends plot (actual vs. synthetic piracy), a gap plot (the period-by-period difference), and a placebo distribution (the treated country's MSPE ratio compared to those of all donor countries). The specification for all four countries includes three lagged-outcome windows, four covariates, and placebos. Table 2 summarises the main results.

Table 2: Summary of SCM Results

Country	MSPE Ratio	Rank	Fisher p	Mean Gap	Pre-RMSPE	Key Donors
Spain	19.01	1/20	0.050	−2.203	0.502	FI, RO, PL, BG, HU
France	8.66	1/20	0.050	−1.266	0.495	HR, SE, IE, LT
Italy	8.15	2/20	0.100	−1.801	0.606	PL, FI, AT, HU
Germany	3.11	4/20	0.200	−0.894	0.515	PL, FI, RO

Notes: Mean Gap is the average of the 14 monthly differences (actual minus synthetic) over the post-treatment period, in piracy accesses per internet user per month. Pre-RMSPE is the root mean squared prediction error over the 70 pre-treatment periods. Fisher p = rank / 20. Key Donors lists countries with SCM weight exceeding 0.05. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

4.1.1 France

France received the ad-tier on November 3, 2022, and was not subject to the early password-sharing crackdown. Among the four treated countries, France exhibits the highest baseline piracy levels, with pre-treatment values ranging from approximately 8 to 14 accesses per user per month.

The synthetic France is constructed primarily from Croatia (w = 0.42), Sweden (w = 0.19), Ireland (w = 0.16), and Lithuania (w = 0.11), with smaller contributions from Slovakia (w = 0.05) and Greece (w = 0.04). The predictor balance is close: the synthetic unit matches France's average piracy in all three pre-treatment windows to within 0.01 accesses per user (13.6 vs. 13.6, 10.4 vs. 10.4, and 9.59 vs. 9.60, respectively). GDP per capita (102 vs. 102 PPS), youth unemployment (34.6 vs. 34.6), and the Gini coefficient (29.1 vs. 29.1) are also closely matched. Internet access shows a small gap (90.1 vs. 86.5). The pre-treatment RMSPE is 0.495, the lowest among the four treated countries.

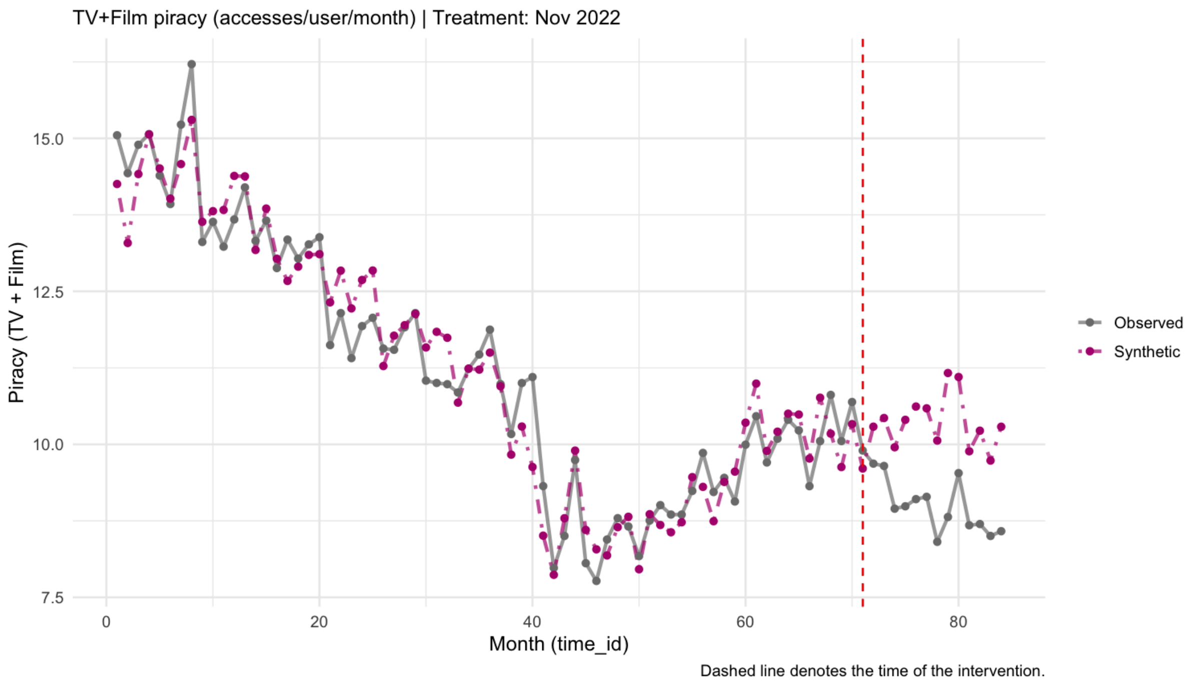

Figure 1: Synthetic France vs France. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Following treatment, actual piracy in France falls below the synthetic counterfactual. The gap is slightly positive in the first post-treatment month (+0.30), turns negative from December 2022 onward (-0.62 in month 72), and reaches its largest magnitude of -2.39 in July 2023 (month 79). The mean post-treatment gap across all 14 months is -1.27 accesses per user per month.

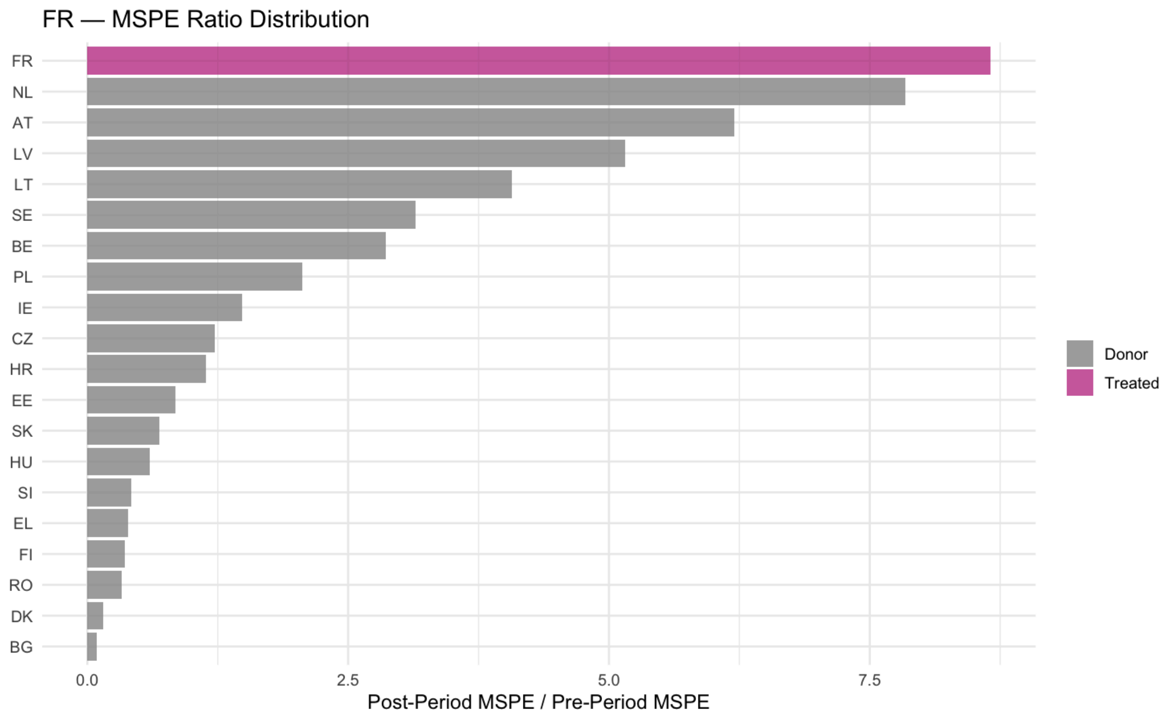

Figure 2: FR MSPE Ratio Distribution. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

In the placebo distribution, France ranks first among all 20 units with an MSPE ratio of 8.66 (pre-MSPE = 0.243, post-MSPE = 2.10), yielding a Fisher exact p-value of 0.050. The Netherlands ranks second among the donors with an MSPE ratio of 7.84.

4.1.2 Spain

Spain received the ad-tier on November 10, 2022, and the password-sharing crackdown on February 8, 2023. The estimated effect from time period 74 onward therefore reflects a compound treatment rather than a pure estimate of the Netflix ad-tier introduction.

The synthetic Spain draws primarily on Finland (w = 0.44), Romania (w = 0.19), and Poland (w = 0.15), with smaller weights on Bulgaria (w = 0.09), Hungary (w = 0.08), and the Netherlands (w = 0.04). The pre-treatment RMSPE is 0.502. The predictor balance is close for the piracy windows (7.12 vs. 7.12, 5.76 vs. 5.73, 5.57 vs. 5.62) but the synthetic unit underestimates Spain's youth unemployment (45.0 vs. 22.7) and Gini coefficient (32.9 vs. 29.2).

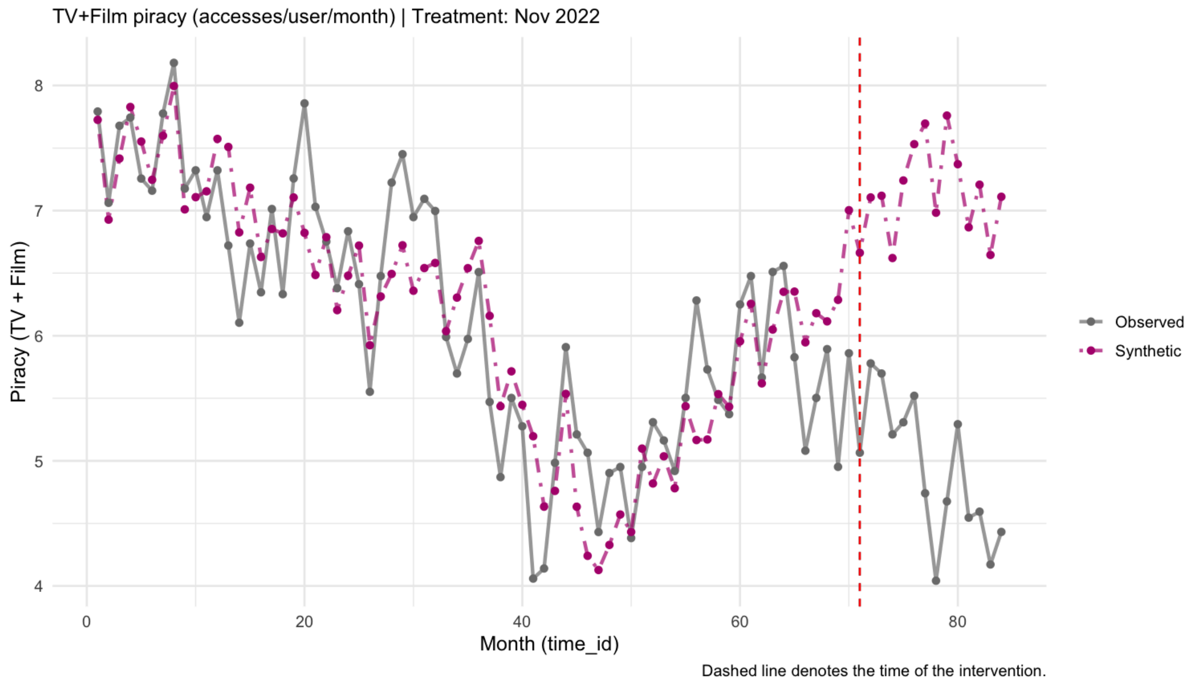

Figure 3: Synthetic Spain versus Spain. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

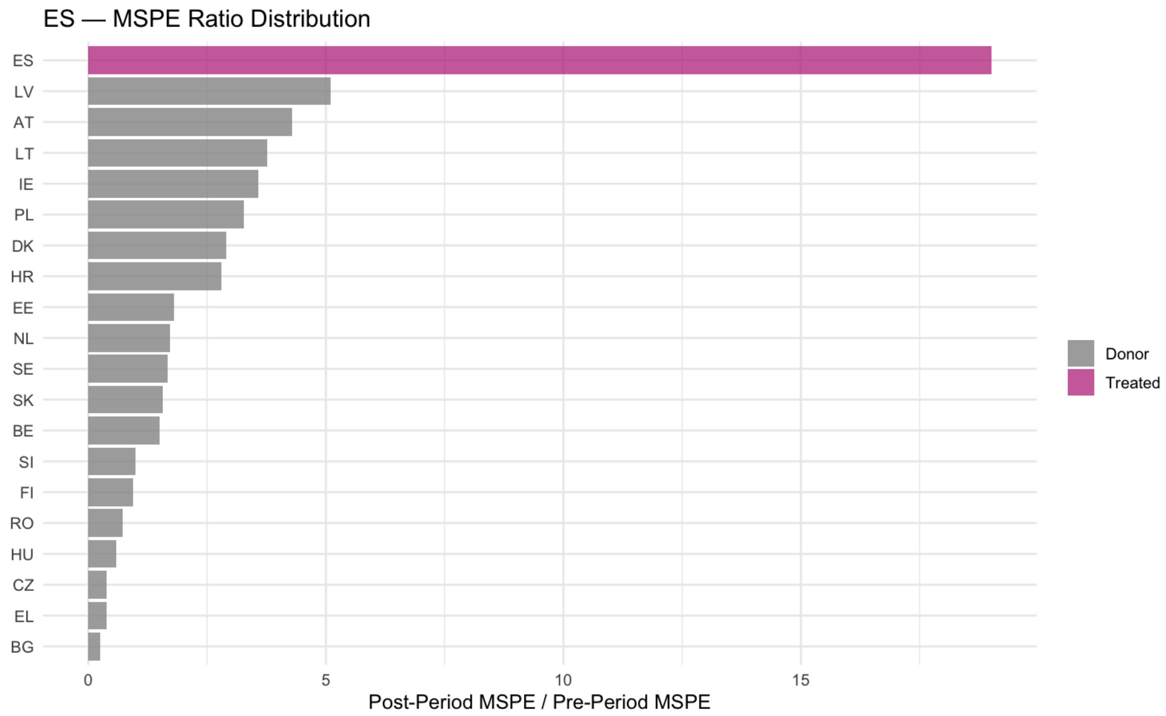

Figure 4: ES MSPE Ratio Distribution. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The mean post-treatment gap is -2.20 accesses per user per month, the largest among the four countries. The gap reaches its largest magnitude of -3.08 in July 2023 (month 79). Spain ranks first in the MSPE ratio distribution with a ratio of 19.01 (pre-MSPE = 0.284, post-MSPE = 5.40). The Fisher exact p-value is 0.050.

4.1.3 Italy

Italy received the ad-tier on November 3, 2022, and was not subject to the early crackdown. The synthetic Italy is composed of Poland (w = 0.37), Finland (w = 0.30), Austria (w = 0.23), and Hungary (w = 0.11). The pre-treatment RMSPE of 0.606 is the highest among the four countries. The predictor balance is mixed: the synthetic unit matches Italy's three piracy windows closely (6.29 vs. 6.46, 5.54 vs. 5.34, 4.67 vs. 4.94) and reproduces GDP per capita almost exactly (96.3 vs. 96.2 PPS), but underestimates youth unemployment (40.5 vs. 19.5 percent) and the Gini coefficient (32.8 vs. 27.1), reflecting the donor pool's concentration in lower-inequality, lower-unemployment Northern and Central European donors.

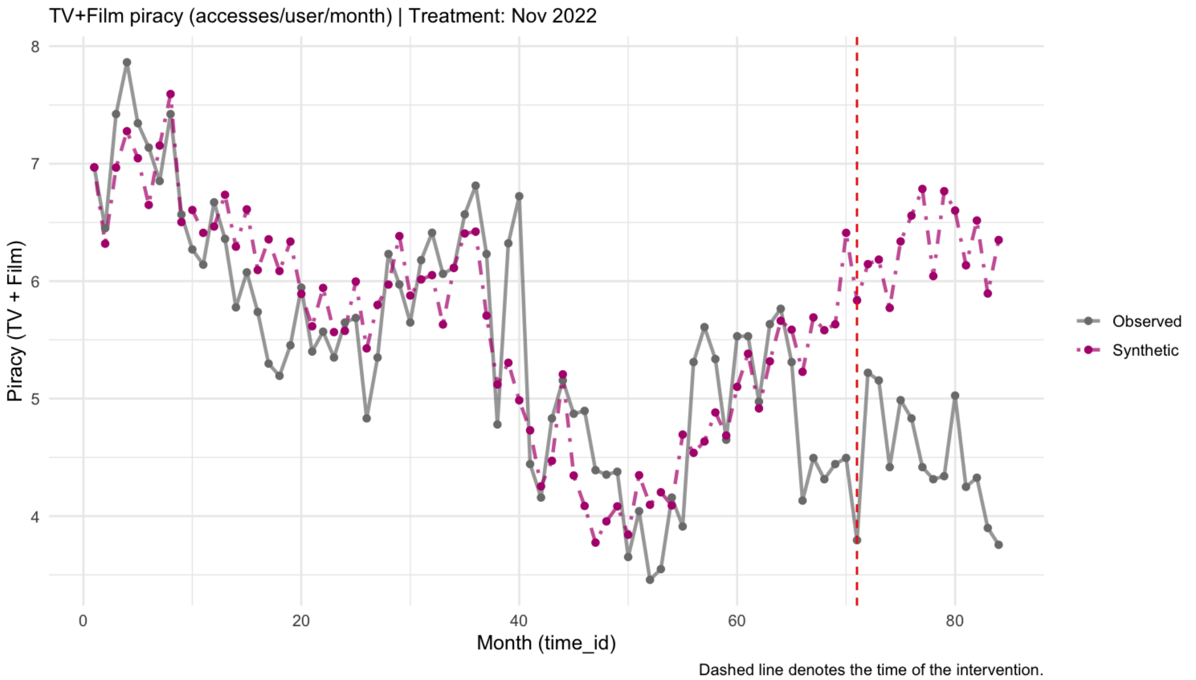

Figure 5: Synthetic Italy vs Italy. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

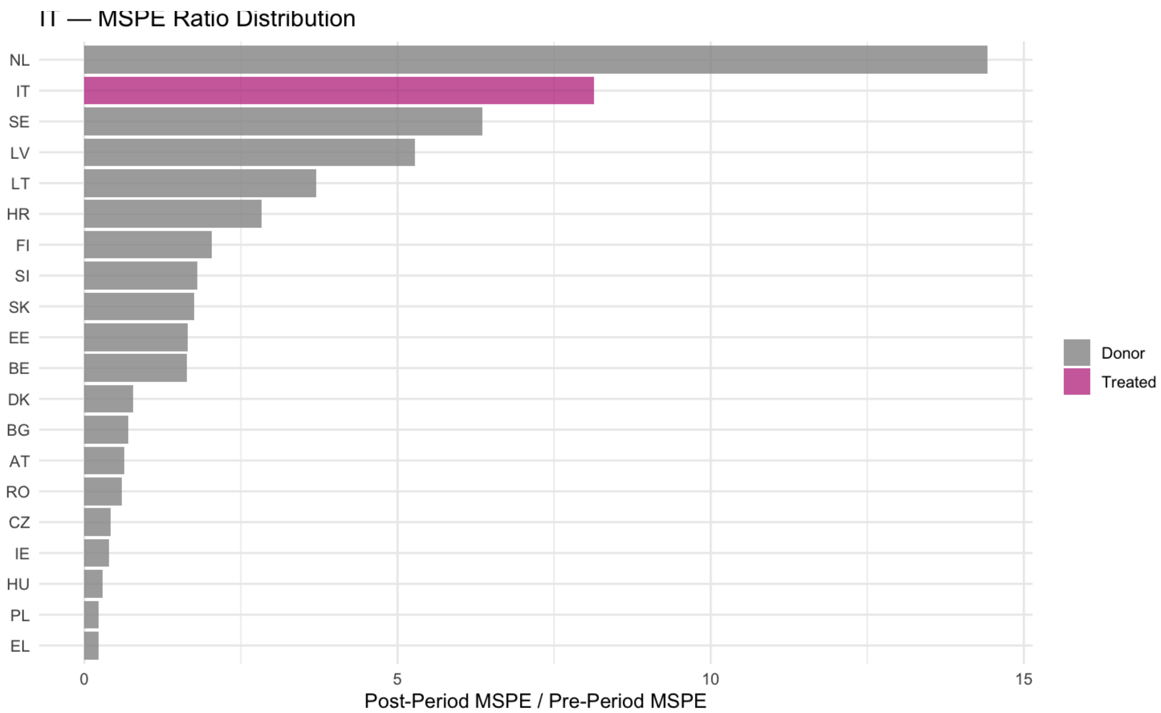

Figure 6: IT MSPE Ratio Distribution. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The mean post-treatment gap is -1.80 accesses per user per month. Italy ranks second in the MSPE ratio distribution with a ratio of 8.15 (pre-MSPE = 0.421, post-MSPE = 3.43). The Netherlands, a donor country, ranks above Italy with an MSPE ratio of 14.4, resulting in a Fisher exact p-value of 0.100 for Italy.

4.1.4 Germany

Germany received the ad-tier on November 3, 2022, at €4.99/month. The synthetic Germany is the most concentrated: Poland (w = 0.58), Finland (w = 0.31), and Romania (w = 0.11). The pre-treatment RMSPE is 0.515. The predictor balance is uneven: the synthetic unit matches Germany's three piracy windows closely (6.59 vs. 6.53, 4.28 vs. 4.64, 4.27 vs. 4.31), but underestimates GDP per capita (123 vs. 85.0 PPS) and overestimates youth unemployment (10.3 vs. 19.6 percent), reflecting the absence of a high-income, low-unemployment donor with a comparable piracy trajectory.

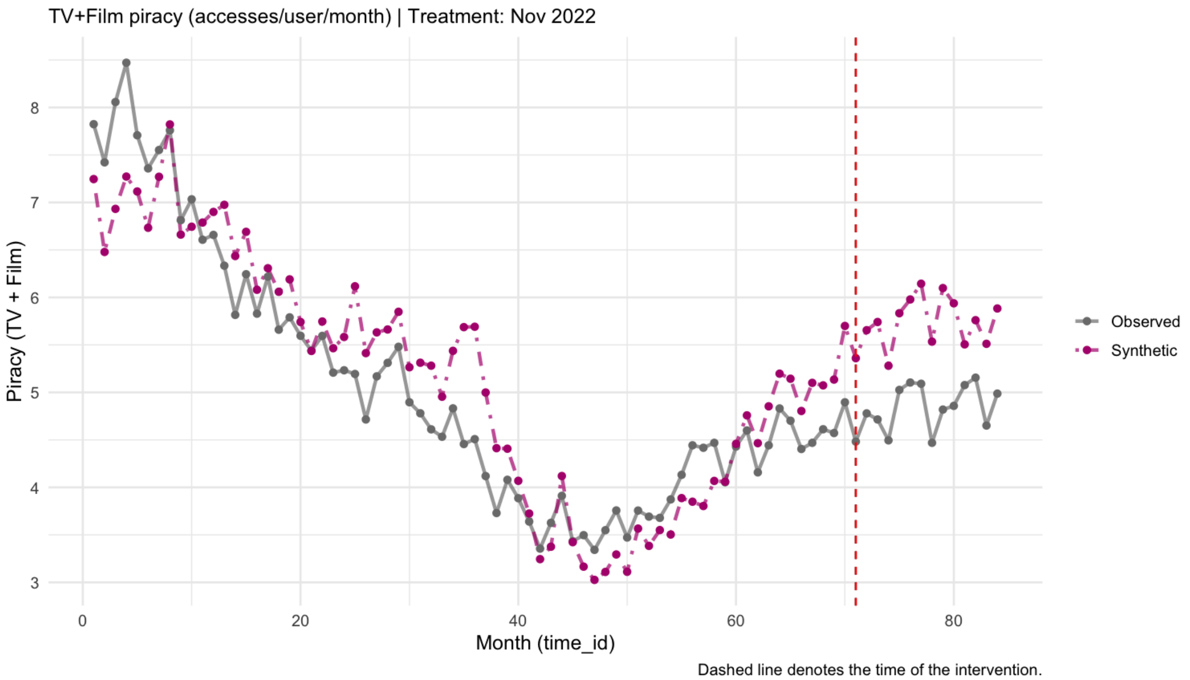

Figure 7: Synthetic Germany vs Germany. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The mean post-treatment gap is -0.89 accesses per user per month. Germany ranks fourth in the MSPE ratio distribution with a ratio of 3.11, behind Latvia (5.20), Lithuania (3.42), and Denmark (3.11). The Fisher exact p-value is 0.200.

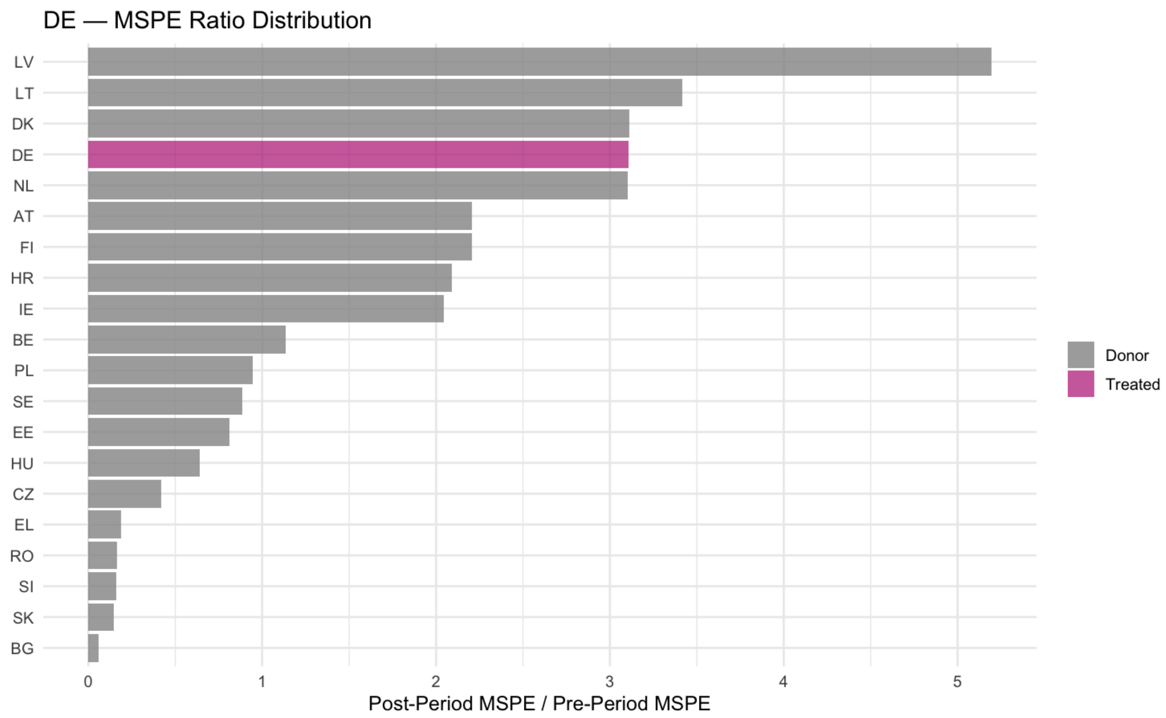

Figure 8: DE MSPE Ratio Distribution. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

4.2 Cross-Country Patterns

All four treated countries exhibit negative post-treatment gaps. The magnitude ranges from -2.20 (Spain) to -0.89 (Germany). The Netherlands consistently appears as a high-MSPE donor: it ranks first in Italy's distribution (ratio 14.4) and second in France's (7.84). The gap plots show a common temporal pattern in which the divergence between actual and synthetic piracy is small in the first one to two months and grows over the subsequent year. After May 2023 (time period 77), when the global password-sharing crackdown reached all donor countries, the donor piracy levels decline, which makes the synthetic counterfactual lower and biases the estimated treatment effect toward zero.

4.3 Generalized Synthetic Control Results

The GSC pools all four treated countries and 22 donor countries (EU-27 excluding Portugal). The cross-validation procedure selects five latent factors (r = 5). The model includes all eight covariates: GDP per capita, youth unemployment, internet access, Gini coefficient, broadband penetration, young population share, tertiary education, and HICP. None of the eight covariate coefficients are individually significant at the 5 percent level.

Table 3: Generalized Synthetic Control - Pooled ATT

	ATT	S.E.	CI Lower	CI Upper	p-value
Pooled ATT	−1.278	0.596	−2.445	−0.111	0.032

Notes: Parametric inference with 500 bootstrap replications. Panel: 26 countries (EU-27 excluding Portugal), 84 monthly periods. Five latent factors selected by cross-validation. Treatment: November 2022 (time period 71). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

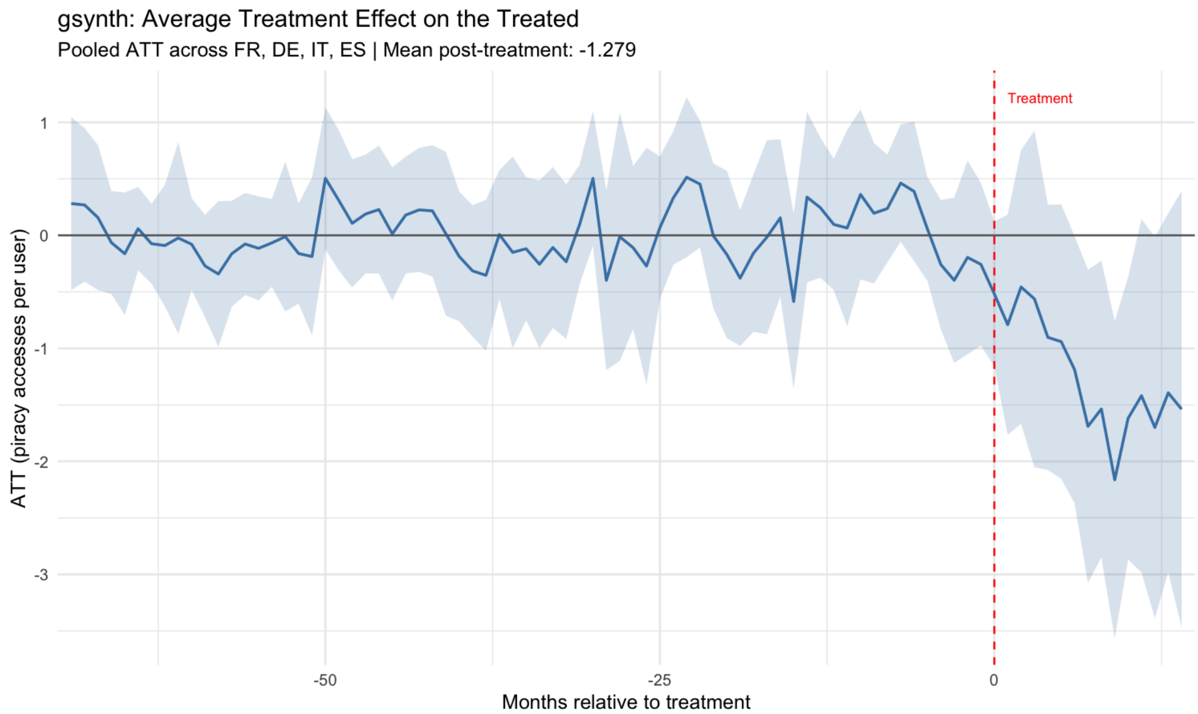

Figure 9: ATT for DE, ES, FR and IT. Figure 9 header shows the unrounded mean from an earlier R run (-1.279); the corrected value to three decimals is -1.278. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The pooled ATT is -1.278 accesses per internet user per month (S.E. = 0.596, 95% CI [-2.445, -0.111], p = 0.032). The confidence interval excludes zero at the 5 percent level. The ATT plot shows that during the pre-treatment period, the ATT fluctuates around zero with no systematic pattern. The ATT starts turning negative around the treatment month and grows increasingly negative over the post-treatment window. Individual post-treatment periods become significant from month 6 onward (ATT = -1.19, p = 0.050). The ATT reaches its largest magnitude at month 9 (-2.16, p = 0.002).

5. Discussion

5.1 Interpretation of Results

This study cannot establish that Netflix's ad-supported tier caused a reduction in piracy. The estimated effects, while consistently negative across the four treated countries, are subject to identification concerns that prevent a causal interpretation. The remainder of this chapter examines those concerns in detail. This section first sets out what the evidence does and does not support on its own terms.

First, the sign of the estimated effect is the same in all four treated countries: France (-1.27), Spain (-2.20), Italy (-1.80), and Germany (-0.89). With four country-level estimates, this cross-country consistency is the most informative feature of the results. Second, the pooled GSC estimate (ATT = -1.278, p = 0.032) yields a 95 percent confidence interval [-2.445, -0.111] that excludes zero, though only narrowly. Third, the individual post-treatment ATTs become statistically distinguishable from zero from the sixth month after launch (ATT = -1.19, p = 0.050) and reach their largest magnitude in month nine (ATT = -2.16, p = 0.002), a temporal pattern that is at least consistent with a behavioural response to a new lower-priced subscription option.

These observations are tempered by several features of the estimates themselves. The pre-treatment root mean squared prediction error is approximately 0.5 across all four synthetic units, comparable in magnitude to Germany's mean post-treatment gap of -0.89 and not negligible relative to France's -1.27. Abadie (2021) emphasises that the credibility of an SCM estimate depends on pre-treatment fit quality; a pre-RMSPE of this magnitude against post-treatment gaps of 1 to 2 accesses per user indicates a low signal-to-noise ratio in the country-by-country estimates. The Fisher exact test reaches p = 0.050 only for France and Spain, and this reflects the floor imposed by a donor pool of nineteen units (1/20 = 0.050) rather than unusually strong evidence against the null; for Italy and Germany, the p-values of 0.100 and 0.200 indicate that the treated countries' MSPE ratios are not extreme within the placebo distribution. The Netherlands appears prominently in both Italy's and France's placebo distributions, a further reminder that the estimates are noisy.

The interpretation that follows from these patterns is therefore narrow. The evidence is consistent with a reduction in piracy following the November 2022 ad-tier launch in the four treated countries, but the magnitude of any such reduction cannot be precisely identified from this design. Several specific threats to a causal interpretation are present: pre-treatment dynamics in the pooled GSC estimate, the absence of a clean treated unit, and concurrent enforcement activity unrelated to the ad-tier. The problems are further addressed in the subsections that follow.

5.2 Pre-Treatment Pattern

Inspection of the pooled GSC estimate in Figure 9 indicates that the ATT begins to drift downward in the months immediately preceding the November 2022 treatment date, rather than at the treatment date itself. While the pre-treatment ATT fluctuates around zero for most of the sample window, a visible downward movement appears in the final months of the pre-treatment period. Two interpretations of this pattern deserve consideration.

The first interpretation is that the pattern reflects genuine anticipation. Netflix first announced plans for an ad-supported tier in April 2022 (Netflix 2022a), confirmed the Microsoft advertising partnership in July 2022 (Netflix 2022b), and published the November launch dates and country-specific prices on October 13, 2022 (Netflix 2022c). These announcements received substantial international press coverage. To the extent that some users in France, Germany, Italy, and Spain adjusted their consumption in anticipation of an imminent lower-priced legal alternative, by deferring acts of piracy they expected to substitute toward the ad-tier, the November 2022 boundary would mischaracterise the timing of the behavioural response. Under this interpretation, the treatment effect captured in the post-treatment window represents the response to the announcement-and-launch sequence as a whole, and the country-by-country SCM estimates would understate the true effect because they treat the pre-period as untreated when behaviour was already shifting.

The second interpretation is that the pre-treatment drift in Figure 9 reflects features of the GSC estimation rather than a genuine behavioural response. The country-by-country SCM gap plots, which are constrained by donor weighting and a fixed predictor set, do not show comparable systematic pre-treatment movement. France's gap is in fact slightly positive (+0.30) in the first post-treatment month and turns negative only in December 2022, a pattern more consistent with a delayed treatment response than with anticipation. The pooled GSC, by contrast, uses an interactive fixed effects framework with five latent factors selected by cross-validation, which introduces additional flexibility that may absorb idiosyncratic pre-treatment variation. Abadie (2021) notes that the credibility of an SCM estimate depends in part on the cleanliness of the pre-treatment fit, and the country-by-country and pooled designs differ in this respect.

These two interpretations cannot be cleanly distinguished within the available data window. The country-level evidence, however, weighs against anticipation as the potential primary driver of the pattern in Figure 9: the post-treatment gaps emerge in the months following the launch rather than the months preceding it. The pre-treatment drift in the pooled GSC is therefore best interpreted as a feature of the method rather than of the underlying data. The November 2022 treatment date is retained as the appropriate boundary for interpreting the country-by-country estimates, while the pooled GSC magnitude should be read with the additional caveat that the interactive fixed effects framework permits some pre-treatment movement that the donor-weighted SCM does not.

5.3 Sample Contamination Across the Four Treated Units

Another concern is that none of the four treated countries provides a fully clean experimental setting. Each is contaminated in a different way, and the strength of any positive interpretation therefore rests not on the cleanliness of any single estimate but on the consistency of the pattern across all four.

Spain is the most contaminated case, as already acknowledged in Section 3.6. The estimated effect from February 2023 onward reflects the compound impact of the ad-tier and the password-sharing crackdown launched on February 8, 2023, three months after the ad-tier introduction. Spain's mean post-treatment gap of -2.20, the largest among the four treated countries, therefore cannot be attributed to the ad-tier alone, and the difference between Spain's gap and France's, equal to -0.93 accesses per user, provides only an upper bound on the contribution of the crackdown component.

France is uncontaminated by the Netflix password-sharing crackdown until May 2023, when the global crackdown reached all donor countries simultaneously, and its pre-treatment fit is the strongest among the four units (RMSPE 0.495, predictor balance to within 0.01 accesses per user across all three pre-treatment outcome windows). Two distinct contamination channels nonetheless apply. The first runs through the donor pool: Sweden (w = 0.19) and Ireland (w = 0.16), which together account for 35 percent of synthetic France's weight, both received Disney+ in June 2022, five months before the ad-tier launch. The Disney+ expansion would have reduced piracy in these donor countries, lowering the synthetic counterfactual and biasing France's estimated treatment effect toward zero. The second channel runs through France itself. The French audiovisual regulator ARCOM, established on January 1, 2022 from the merger of the CSA and Hadopi, ramped up site-blocking activity substantially during the post-treatment window of this study, particularly in sports piracy. Under the loi Bachelot of October 2021, ARCOM blocked approximately 835 illicit sports piracy sites during the first nine months of 2022 (ARCOM 2022). ARCOM's subsequent assessment estimated that the audience of the targeted piracy galaxies, defined as the originally blocked sites and their mirror sites, fell by 23 percent between October 2022 and March 2023, a window overlapping almost exactly with the post-treatment period studied here (ARCOM 2023, cited in ACCES 2023). Because no donor country has a regulator with comparable authority and activity over this period, ARCOM blocking does not difference out through the synthetic control.

Italy is in a similar position to France with respect to treatment timing but with weaker pre-treatment fit (RMSPE 0.606) and a placebo distribution in which the Netherlands ranks above Italy (MSPE ratio 14.4 versus 8.15), yielding a Fisher exact p-value of 0.100. Italy's synthetic counterfactual is constructed primarily from Poland (w = 0.37), Finland (w = 0.30), Austria (w = 0.23), and Hungary (w = 0.11). Austria and Finland both had Disney+ prior to the ad-tier launch, again biasing the synthetic toward lower piracy in donors and the estimated treatment effect toward zero.

Germany is the least contaminated by the Disney+ donor problem, since its synthetic is composed of Poland (w = 0.58), Finland (w = 0.31), and Romania (w = 0.11), and only Finland had Disney+ pre-treatment. Germany's weak result, however, has its own interpretive challenges. The mean gap of -0.89 is comparable to the pre-RMSPE of 0.515, the Fisher exact p-value is 0.200, and the treated unit ranks fourth in the placebo distribution. Whether this reflects a genuinely small effect in Germany, the smallest country-specific price differential between standard and ad-tier subscriptions, or simply estimation noise cannot be determined from this design alone.

The implication is that the design does not deliver a single clean treatment effect. What it does deliver is a consistent negative sign across four countries with different contamination profiles.

Each treated country is confounded by a distinct combination of factors: Spain by the Netflix password-sharing crackdown, France by donor-pool Disney+ exposure and contemporaneous ARCOM site-blocking activity, Italy by donor Disney+ exposure and a relatively weak placebo-distribution rank, and Germany by a small effect size near the noise floor. No single country provides a clean estimate. Reading these four estimates together, rather than relying on any one of them, is the basis on which any positive interpretation of the results must rest.

5.4 The Anti-Piracy Enforcement Environment

A third concern is that the post-treatment window of this study coincides with sustained anti-piracy enforcement activity unrelated to the Netflix ad-tier. The most prominent example is Operation 404, an ongoing multi-phase campaign launched in November 2019 by Brazil's Ministry of Justice and Public Security. Although nominally a Brazilian operation, it is coordinated with international partners including the U.S. Department of Justice, the City of London Police, the EU Intellectual Property Office, and industry associations such as the MPA, ACE, IFPI, and ALIANZA. The sites and applications blocked or seized are typically international rather than purely Brazilian, and many served European audiences.

Within the 2017-2023 sample window, Operation 404 conducted six distinct phases: Phase 1 (November 2019, 210 sites and 100 streaming applications), Phase 2 (November 2020, 252 sites and 65 streaming applications), Phase 3 (July 2021, 334 sites and 94 streaming applications), Phase 4 (June 2022, focused on the metaverse and music streaming with 461 music applications removed), Phase 5 (March 2023, 199 streaming and gaming sites plus 63 music applications and 128 additional domains blocked), and Phase 6 (November 2023, 606 sites in Brazil and approximately 368 in other participating countries) (MJSP 2024). Two of these phases, Phase 5 and Phase 6, fall within the fourteen-month post-treatment window of this study. The cumulative scale of these actions across all six phases exceeds 1,900 sites and 700 applications.

Can Operation 404 be controlled for in the generalized synthetic control? In principle, it could be: a time-varying indicator coded by phase month, or a cumulative-takedown count interacted with country-specific exposure, that would enter the GSC covariate set in the same way as the Eurostat variables already included in Section 3.2.2. In practice, no such indicator is available in the current specification, and constructing one would require imputation of country-specific user exposure to the takedowns, which the Brazilian government does not publicly disclose.

Operation 404 is therefore not directly controlled for, and the analysis instead relies on a structural argument. The takedowns affect treated and donor countries through the same channel, since sites and applications shut down by Brazilian and international authorities become inaccessible from all countries simultaneously. To the extent that EU users in the treated and donor countries accessed the same disrupted sites in similar proportions, the resulting reduction in piracy partially differences out through the synthetic control. The estimated treatment effect would be biased toward zero only if the disrupted sites were disproportionately used by donor-country audiences, in which case donor piracy would fall by more than treated-country piracy through the enforcement channel alone, and the synthetic counterfactual would understate the level of piracy that would have prevailed in the treated countries absent the ad-tier. The opposite scenario, in which the disrupted sites were disproportionately used by treated-country audiences, would bias the estimate away from zero.

This Operation 404 confound is distinct from the May 2023 Netflix password-sharing crackdown. The two events are global in scope but operate through different mechanisms, the former through site availability and the latter through subscription compliance, and both fall within the post-treatment window. A more complete treatment of the enforcement environment would also include the country-specific activity discussed in Section 5.3, including French ARCOM blocking and analogous measures in Italy, Germany, and Spain. Within the design and time horizon of this study, the most that can be said is that the negative gaps reported in Section 4 should be interpreted as joint estimates of the ad-tier effect and the contemporaneous enforcement environment, with the relative contribution of each component not separately identified. A recommendation for future work, returned to in the Conclusion, is to construct a country-month enforcement-intensity covariate combining international takedown phases with national blocking activity, and to re-estimate the model with that covariate included.

5.5 Other Limitations

Beyond the identification concerns addressed in Sections 5.1 through 5.4, three additional limitations of the design need acknowledgment.

First, the piracy data are restricted in scope. The MUSO data underlying EUIPO (2024) measure traffic to known piracy websites and the consumption of copyright-infringing content accessible through them. They do not capture piracy delivered via illicit IPTV services, encrypted messaging applications such as Telegram, peer-to-peer protocols, stream-rippers, or social-media-based redistribution. To the extent that consumers in the treated countries shifted across these channels in response to the ad-tier launch or to concurrent enforcement, the estimated effect captures only the website-accessible component of piracy and may understate, overstate, or fail to detect substitution toward unmeasured channels.

Second, the Eurostat covariates are reported at annual frequency, while the piracy outcome is observed monthly. Each month within a given year therefore receives the same covariate value, so within-year variation in GDP per capita, youth unemployment, internet access, the Gini coefficient, and the additional GSC covariates is set to zero by construction. The covariates contribute to the donor match at annual resolution only, and month-on-month variation in the synthetic counterfactual is driven by the three lagged piracy windows used as predictors rather than by the covariates themselves.

Third, the post-treatment window of fourteen months ends in December 2023. This horizon is long enough to capture the within-year evolution of the ATT reported in Section 4.3 but short relative to the timescale of consumer adjustment to a new pricing tier or the longer-run dynamics of platform substitution. Whether the estimated effect persists, attenuates, or reverses over a multi-year window cannot be determined from the current data, and the magnitudes reported in Section 4 should be read as short-run rather than long-run estimates.

5.6 Robustness: Leave-One-Out Donor Pool

The donor concentration concern raised in Section 5.3, that Poland and Finland together account for the majority of the synthetic weights in three of the four treated units, motivates a leave-one-out (LOO) check. The country-by-country SCM was re-estimated with the same Tier 1 specification (three pre-treatment outcome windows and four covariates) under two alternative donor pools: one excluding Poland, and one excluding Finland. Table 4 reports the mean post-treatment gap, the pre-treatment RMSPE, and the MSPE ratio for each combination, together with the resulting top three donor weights.

Table 4: Summary of LOO

Scenario	Country	Mean Gap	Pre-RM SPE	Post-RM SPE	MSPE Ratio	Top Donor 1	Top Donor 2	Top Donor 3
Baseline	FR	-1.266	0.495	1.399	7.99	HR (0.42)	SE (0.19)	IE (0.16)
Baseline	DE	-0.894	0.515	0.917	3.17	PL (0.58)	FI (0.31)	RO (0.11)
Baseline	IT	-1.801	0.606	1.867	9.48	PL (0.37)	FI (0.29)	AT (0.23)
Baseline	ES	-2.203	0.502	2.280	20.65	FI (0.44)	RO (0.19)	PL (0.15)
drop_PL	FR	-1.220	0.494	1.352	7.51	HR (0.44)	IE (0.17)	SE (0.17)
drop_PL	DE	-2.162	1.001	2.174	4.72	FI (0.48)	AT (0.28)	RO (0.24)
drop_PL	IT	-1.979	0.665	2.055	9.55	AT (0.78)	FI (0.22)	—
drop_PL	ES	-1.996	0.482	2.074	18.53	FI (0.36)	AT (0.35)	RO (0.17)
drop_FI	FR	-1.196	0.493	1.328	7.27	HR (0.45)	IE (0.17)	SE (0.16)
drop_FI	DE	-0.300	0.650	0.370	0.32	PL (0.74)	AT (0.16)	DK (0.11)
drop_FI	IT	-1.270	0.717	1.357	3.58	AT (0.51)	PL (0.37)	DK (0.12)
drop_FI	ES	-1.389	0.530	1.495	7.97	AT (0.42)	PL (0.29)	DK (0.25)

Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Three patterns emerge. First, France's estimated gap is essentially unchanged across all three specifications. The baseline mean gap of -1.266 moves to -1.220 when Poland is excluded and to -1.196 when Finland is excluded, and the pre-treatment RMSPE remains at approximately 0.49 in all three scenarios. This is consistent with the donor weights reported in Table 2: Poland and Finland did not contribute meaningfully to synthetic France in the baseline specification, which drew primarily on Croatia, Sweden, and Ireland. The France result is robust to dropping either dominant donor.

Second, Germany's estimate is fragile. Dropping Finland reduces the mean gap from -0.894 to -0.300, and the post-treatment MSPE ratio falls from 3.17 to 0.32, which means the post-treatment fit becomes tighter than the pre-treatment fit and the gap can no longer be distinguished from prediction error. Dropping Poland produces a gap of -2.162, but the pre-treatment RMSPE simultaneously doubles from 0.515 to 1.001, so the larger magnitude reflects a worse pre-treatment fit rather than a stronger treatment signal. In neither LOO specification does Germany produce a defensible negative effect, and the baseline result reported in Section 4.1.4 should be regarded as not robust to this check.

Third, Italy and Spain show partial robustness. Italy's gap declines from -1.801 to -1.270 when Finland is dropped, and the MSPE ratio falls from 9.48 to 3.58. Spain's gap declines from -2.203 to -1.389 under the same scenario, with a smaller change in pre-treatment RMSPE. Both treated units retain a clear negative sign across all three specifications, but the magnitude in each case depends meaningfully on Finland's inclusion in the donor pool.

The LOO results sharpen the conclusion of Section 5.3 rather than overturning it. France remains the most defensible single estimate after all of the contamination concerns discussed in this chapter, including the donor-concentration concern tested here; Germany cannot be defended on robustness grounds; and the magnitudes for Italy and Spain are sensitive to Finland's inclusion in the donor pool, while their signs are not. The cross-country sign consistency that any positive interpretation of the results rests on, as argued in Section 5.3, survives this check at the level of sign but not uniformly at the level of magnitude.

6. Conclusion

This study set out to estimate the effect of Netflix's November 2022 ad-supported subscription tier on online piracy in the four European Union member states where the tier was launched. Using country-by-country synthetic control estimation supplemented by a generalized synthetic control with interactive fixed effects, the analysis produced negative point estimates in all four treated countries and a pooled ATT of -1.278 piracy accesses per internet user per month, with a 95 percent confidence interval that excludes zero.

The interpretation of these estimates is constrained in several ways that Chapter 5 develops at length. None of the four treated countries provides a fully clean experimental setting. France, the country with the strongest pre-treatment fit, is contaminated by both donor-pool exposure to Disney+ and contemporaneous intensification of ARCOM site-blocking activity. Spain is contaminated by the Netflix password-sharing crackdown that began three months after the ad-tier launch. Italy and Germany face donor-pool Disney+ exposure as well. In addition, Operation 404 and other international enforcement activity unrelated to the ad-tier proceeded throughout the post-treatment window without an explicit covariate to control for it. The pre-treatment drift visible in the pooled GSC estimate may reflect either genuine consumer anticipation of the ad-tier launch or features of the interactive fixed effects estimation. A leave-one-out donor-pool check, reported in Section 5.6, finds that France's estimate is robust to dropping either of the two dominant donors in the broader pool, while Germany's estimate collapses when Finland is excluded and the magnitudes for Italy and Spain shrink meaningfully under the same exclusion. The evidence reported in this thesis is consistent with a reduction in piracy following the ad-tier launch, but the magnitude of that reduction cannot be precisely identified, and a portion of it is plausibly attributable to concurrent enforcement-side activity rather than to the ad-tier in isolation.

Two extensions would address the most binding remaining limitations. First, the construction of a country-month enforcement-intensity covariate combining international takedown phases such as Operation 404 with national blocking activity by ARCOM and analogous regulators, included in the GSC's interactive fixed effects framework, would allow the relative contribution of the ad-tier and enforcement channels to be partially separated. Second, re-estimation with a longer post-treatment window once additional EUIPO data are released would clarify whether the estimated effect persists, attenuates, or reverses beyond the fourteen-month horizon available here.

The contribution of this thesis is therefore narrow but specific. It documents a consistent negative association between the introduction of Netflix's ad-supported tier and online piracy in four European Union member states, identifies the principal threats to a causal interpretation of that association, and shows through a leave-one-out donor check that the association is robust for France and fragile for Germany. Whether the ad-tier reduces piracy in the European Union, by how much, and through which channels, remain open questions to which this study contributes a partial and qualified answer.

7. References

Academic sources

Abadie, Alberto. 2021. "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects." *Journal of Economic Literature* 59 (2): 391-425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 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, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review* 93 (1): 113-132.

Adermon, Adrian, and Che-Yuan Liang. 2014. "Piracy and Music Sales: The Effects of an Anti-Piracy Law." *Journal of Economic Behavior and Organization* 105: 90-106.

Aguiar, Luis, and Joel Waldfogel. 2018. "As Streaming Reaches Flood Stage, Does It Stimulate or Depress Music Sales?" *International Journal of Industrial Organization* 57: 278-307.

Belleflamme, Paul, and Martin Peitz. 2012. "Digital Piracy: Theory." Chapter 18 in *The Oxford Handbook of the Digital Economy*, edited by Martin Peitz and Joel Waldfogel. Oxford: Oxford University Press.

Godinho de Matos, Miguel, Pedro Ferreira, and Michael D. Smith. 2018. "The Effect of Subscription Video-on-Demand on Piracy: Evidence from a Household-Level Randomized Experiment." *Management Science* 64 (12): 5610-5630.

Kaul, Ashok, Stefan Klößner, Gregor Pfeifer, and Manuel Schieler. 2022. "Standard Synthetic Control Methods: The Case of Using All Preintervention Outcomes Together With Covariates." *Journal of Business & Economic Statistics* 40 (3): 1362-1376.

Lu, Shijie, Koushyar Rajavi, and Isaac Dinner. 2021. "The Effect of Over-the-Top Media Services on Piracy Search: Evidence from a Natural Experiment." *Marketing Science* 40 (3): 548-568.

Xu, Yiqing. 2017. "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models." *Political Analysis* 25 (1): 57-76.

Data and report sources

ACCES (Association des Chaînes Conventionnées Éditrices de Services). 2023. *De nouveaux chemins de croissance pour la télévision thématique en France*. Report presented at the Rencontres Annuelles de l'ACCES, Paris, November 2023.
https://www.acces.tv/wp-content/uploads/2023/11/231113-Etude-ACCES-V7.pdf

ARCOM (Autorité de régulation de la communication audiovisuelle et numérique). 2022. *L'impact du blocage des services illicites de sport*. Paris, October 2022.
https://www.arcom.fr/en/find-out-more/studies-and-data/studies-reviews-and-reports/impact-blocking-illegal-sports-services-quantitative-study-report

ARCOM. 2023. *Bilan sur le blocage des sites miroirs 2022*. Paris, Spring 2023. https://www.arcom.fr/en/find-out-more/studies-and-data/studies-reviews-and-reports/arcoms-first-fight-agains-piracy-2022-results

EUIPO (European Union Intellectual Property Office). 2024. *Online Copyright Infringement in the European Union: Films, Music, Publications, Software and TV (2017-2023)*. Alicante: EUIPO. https://www.euipo.europa.eu/en/publications/online-copyright-infringement-in-the-european-union-films-music-publications-software-and-tv-2017-2023.

MJSP (Ministério da Justiça e Segurança Pública). 2024. "Operação internacional contra pirataria tira do ar 675 sites e 14 aplicativos de streaming." Federative Republic of Brazil, September 19, 2024. https://www.gov.br/mj/pt-br/assuntos/noticias/operacao-internacional-contra-pirataria-tira-do-ar-675-sites-e-14-aplicativos-de-streaming.

Netflix, Inc. 2022a. *Q1 2022 Letter to Shareholders*. April 19, 2022. https://s22.q4cdn.com/959853165/files/doc_financials/2022/q1/FINAL-Q1-22-Shareholder-Letter.pdf.

Netflix, Inc. 2022b. *Q2 2022 Letter to Shareholders*. July 19, 2022. https://s22.q4cdn.com/959853165/files/doc_financials/2022/q2/Netflix,-Inc.pdf.

Netflix, Inc. 2022c. *Q3 2022 Letter to Shareholders*. October 18, 2022. https://s22.q4cdn.com/959853165/files/doc_financials/2022/q3/FINAL-Q3-22-Shareholder-Letter.pdf

Eurostat Datasets “tec00114, yth_empl_090, isoc_ci_in_h, ilc_di12, isoc_ci_it_h, demo_pjangroup, edat_lfse_03 and prc_hicp_aind”, accessed March 18, 2026 through Eurostat API via R.
https://ec.europa.eu/eurostat/databrowser/view/tec00114/default/table
https://ec.europa.eu/eurostat/databrowser/view/yth_empl_090/default/table
https://ec.europa.eu/eurostat/databrowser/view/isoc_ci_in_h/default/table
https://ec.europa.eu/eurostat/databrowser/view/ilc_di12/default/table
https://ec.europa.eu/eurostat/databrowser/view/isoc_ci_it_h/default/table
https://ec.europa.eu/eurostat/databrowser/view/demo_pjangroup/default/table
https://ec.europa.eu/eurostat/databrowser/view/edat_lfse_03/default/table
https://ec.europa.eu/eurostat/databrowser/view/prc_hicp_aind/default/table

8. Appendices

The figures below correspond to the leave-one-out robustness check reported in Section 5.6 and summarised in Table 4. For each treated country, the synthetic control is re-estimated under two alternative donor pools: one excluding Poland (drop_PL) and one excluding Finland (drop_FI). The solid line shows observed piracy; the dashed line shows the re-estimated synthetic counterfactual under the restricted donor pool. The vertical dashed line marks November 2022 (time period 71). Where the post-treatment gap closes or the pre-treatment fit deteriorates relative to the baseline, the country's baseline estimate in Table 2 should be regarded as donor-dependent rather than robust.

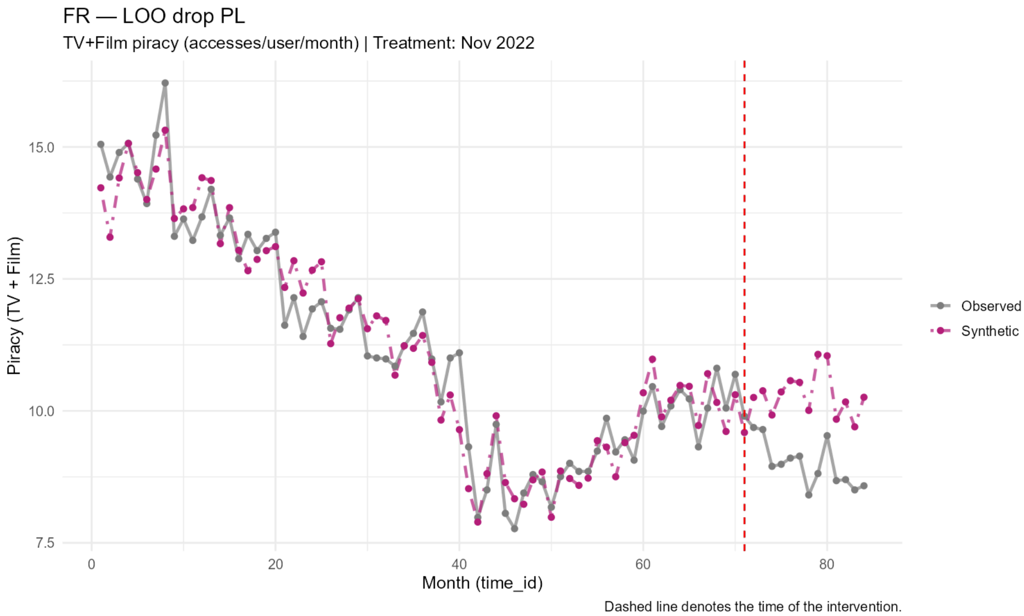

Figure A.1: France, donor pool excludes Poland. Mean post-treatment gap: -1.220 (baseline: -1.266). Top donors: HR (0.44), IE (0.17), SE (0.17). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

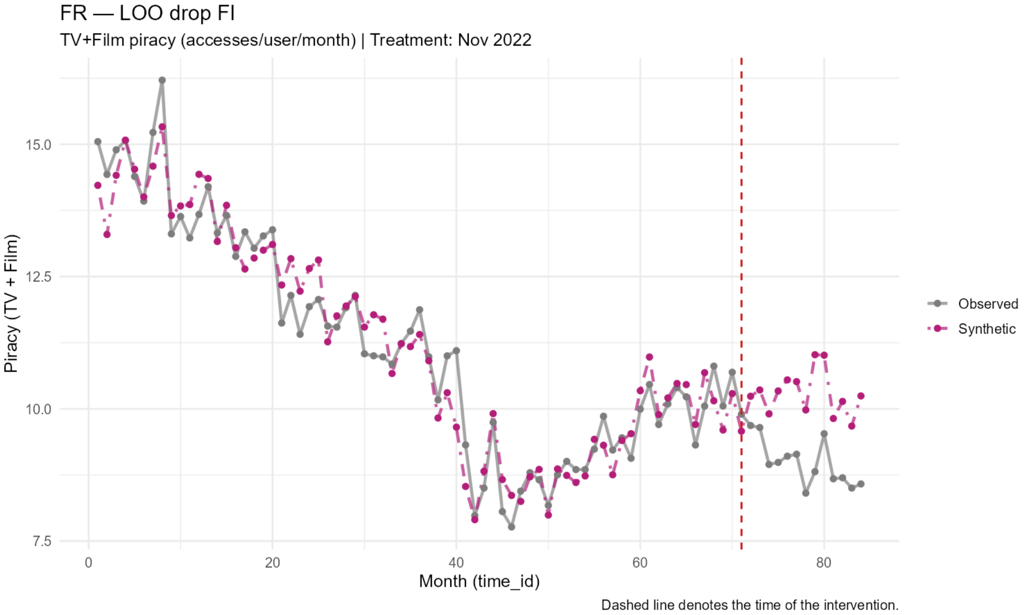

Figure A.2: France, donor pool excludes Finland. Mean post-treatment gap: -1.196 (baseline: -1.266). Top donors: HR (0.45), IE (0.17), SE (0.16). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

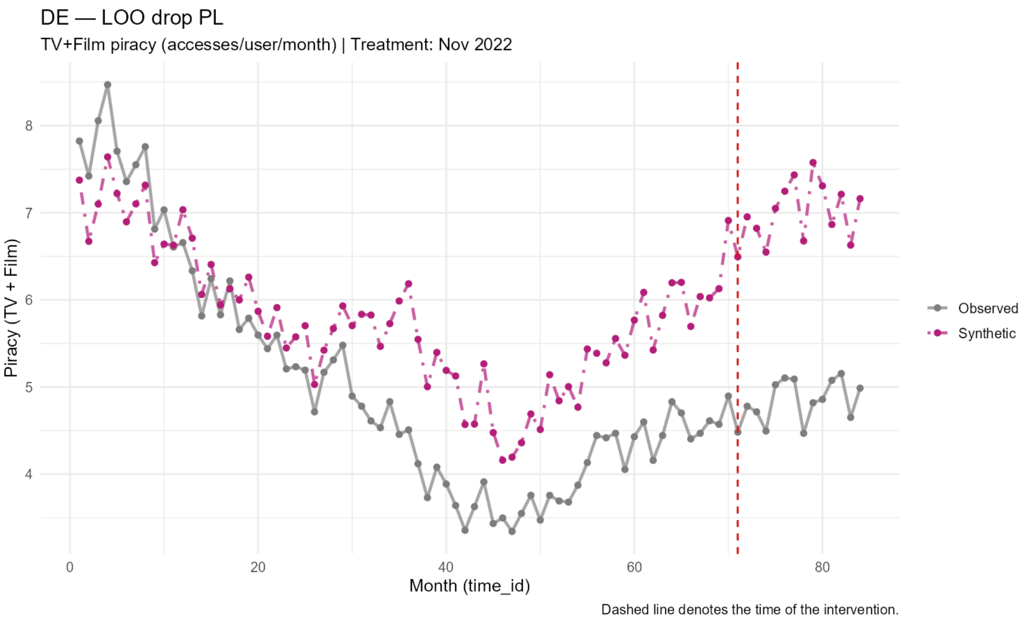

Figure A.3: Germany, donor pool excludes Poland. Mean post-treatment gap: -2.162 (baseline: -0.894). Pre-RMSPE rises from 0.515 to 1.001. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

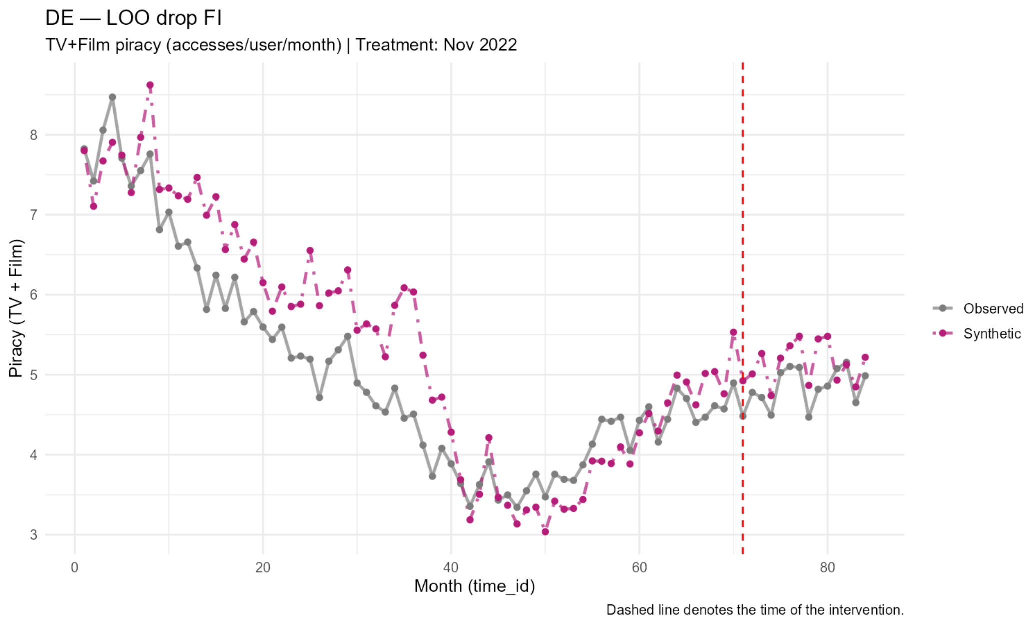

Figure A.4: Germany, donor pool excludes Finland. Mean post-treatment gap: -0.300 (baseline: -0.894). Post-MSPE / pre-MSPE ratio falls to 0.32. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

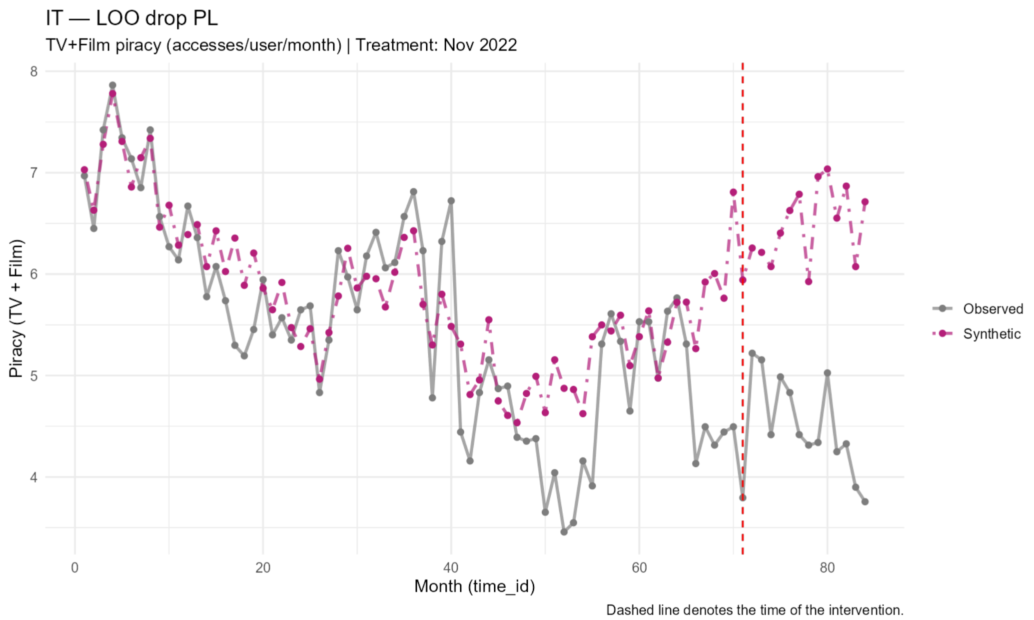

Figure A.5: Italy, donor pool excludes Poland. Mean post-treatment gap: -1.979 (baseline: -1.801). Top donors: AT (0.78), FI (0.22). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

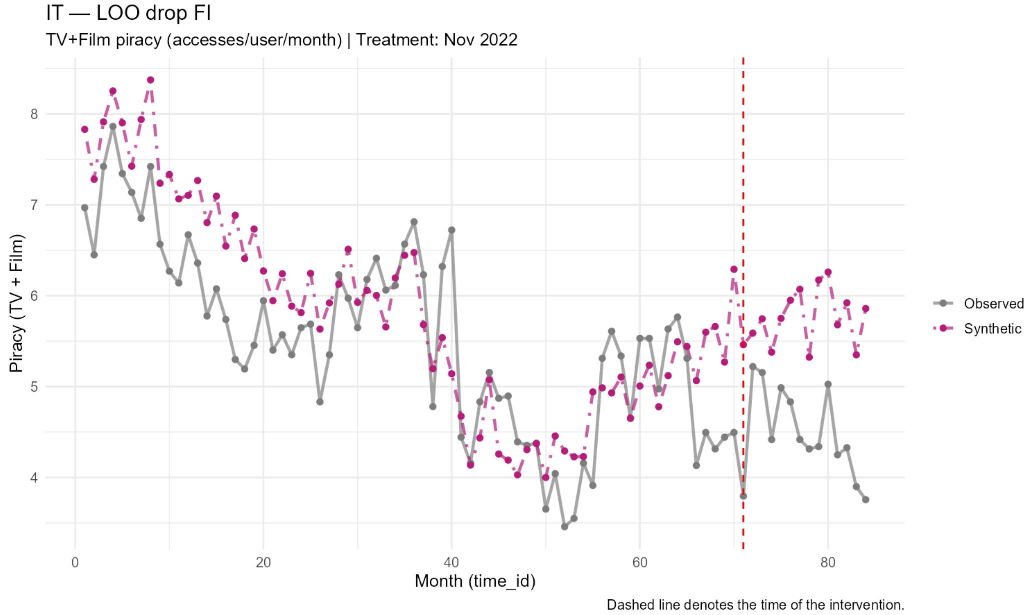

Figure A.6: Italy, donor pool excludes Finland. Mean post-treatment gap: -1.270 (baseline: -1.801). MSPE ratio falls from 9.48 to 3.58. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

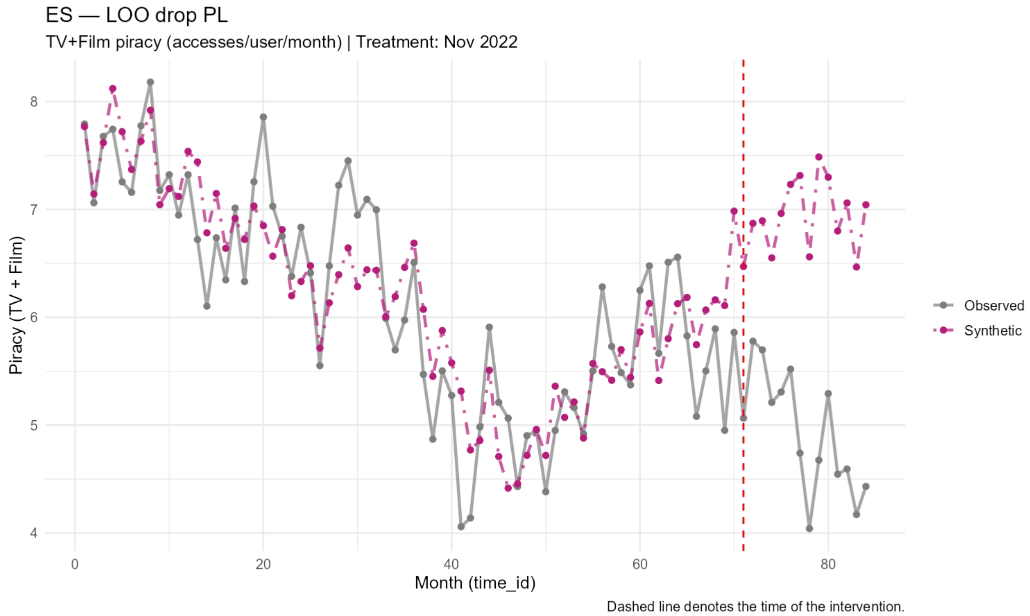

Figure A.7: Spain, donor pool excludes Poland. Mean post-treatment gap: -1.996 (baseline: -2.203). Top donors: FI (0.36), AT (0.35), RO (0.17). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

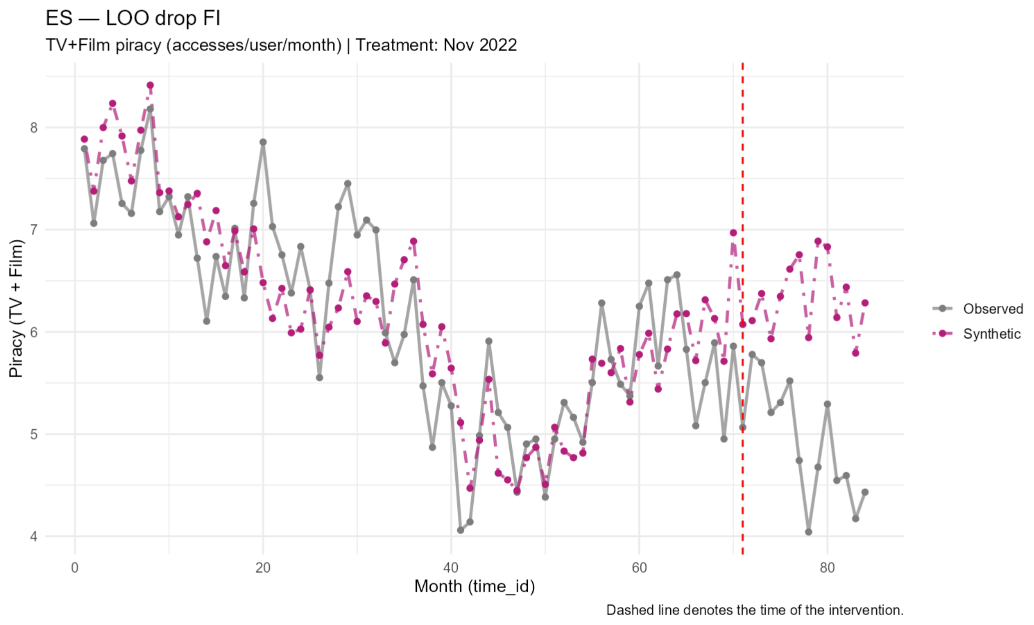

Figure A.8: Spain, donor pool excludes Finland. Mean post-treatment gap: -1.389 (baseline: -2.203). Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Donor	FR	ES	IT	DE
AT	0.002	0.000	0.232	0.000
BE	0.003	0.000	0.000	0.000
BG	0.003	0.089	0.000	0.000
CZ	0.002	0.000	0.000	0.000
DK	0.002	0.000	0.000	0.000
EE	0.004	0.000	0.000	0.000
EL	0.037	0.000	0.000	0.000
FI	0.003	0.444	0.295	0.311
HR	0.422	0.000	0.000	0.000
HU	0.003	0.082	0.106	0.000
IE	0.164	0.000	0.000	0.000
LT	0.105	0.000	0.000	0.000
LV	0.000	0.000	0.000	0.000
NL	0.002	0.040	0.000	0.000
PL	0.002	0.152	0.367	0.582
RO	0.002	0.191	0.000	0.107
SE	0.192	0.000	0.000	0.000
SI	0.002	0.000	0.000	0.000
SK	0.049	0.000	0.000	0.000
Sum	**1.000**	**1.000**	**1.000**	**1.000**

The table reports the full donor weight vector for the baseline SCM estimation in each treated country. Section 4.1 names only donors with weight above 0.05; the table here includes all 19 donor countries. As is typical of synthetic control estimation, three to six donors carry most of the predictive weight in each case, while the remainder receive weights at or near zero. The pattern of donor concentration motivates the leave-one-out check in Section 5.6 and Appendix figures A.1 - A.8

Table A.1: Donor weights, baseline SCM by treated country. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The tables below report the pre-treatment predictor balance for each treated country's baseline SCM. Each row gives the treated unit's value, the synthetic counterfactual's value, and the unweighted mean across the 19 donor pool countries. The three piracy windows are averages of monthly TV plus film piracy accesses per internet user over months 1 to 24 (pre-COVID), 25 to 48 (COVID), and 49 to 70 (post-lockdown). The four covariates are described in Section 3.2.2. Each treated country's balance is also reported in Section 4.1 and is reproduced here for reference.

Predictor balance tables:

Predictor	France	Synthetic	Donor mean
Piracy, months 1–24	13.6	13.6	11.2
Piracy, months 25–48	10.4	10.4	8.70
Piracy, months 49–70	9.59	9.60	7.98
GDP per capita (PPS, EU27 = 100)	102	102	96.0
Youth unemployment, ages 15–24 (%)	34.6	34.6	25.1
Household internet access (%)	90.1	86.5	88.1
Gini coefficient	29.1	29.1	29.0

Table A.2: France predictor balance. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Predictor	Germany	Synthetic	Donor mean
Piracy, months 1–24	6.59	6.53	11.2
Piracy, months 25–48	4.28	4.64	8.70
Piracy, months 49–70	4.27	4.31	7.98
GDP per capita (PPS, EU27 = 100)	123	85.0	96.0
Youth unemployment, ages 15–24 (%)	10.3	19.6	25.1
Household internet access (%)	93.6	89.9	88.1
Gini coefficient	30.1	27.8	29.0

Table A.3: Germany predictor balance. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Predictor	Italy	Synthetic	Donor mean
Piracy, months 1–24	6.29	6.46	11.2
Piracy, months 25–48	5.54	5.34	8.70
Piracy, months 49–70	4.67	4.94	7.98
GDP per capita (PPS, EU27 = 100)	96.3	96.2	96.0
Youth unemployment, ages 15–24 (%)	40.5	19.5	25.1
Household internet access (%)	86.6	90.8	88.1
Gini coefficient	32.8	27.1	29.0

Table A.4: Italy predictor balance. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

Predictor	Spain	Synthetic	Donor mean
Piracy, months 1–24	7.12	7.12	11.2
Piracy, months 25–48	5.76	5.73	8.70
Piracy, months 49–70	5.57	5.62	7.98
GDP per capita (PPS, EU27 = 100)	88.5	89.3	96.0
Youth unemployment, ages 15–24 (%)	45.0	22.7	25.1
Household internet access (%)	91.3	89.9	88.1
Gini coefficient	32.9	29.2	29.0

Table A.5: Spain predictor balance. Source: Author's rendering of EUIPO (2024) and Eurostat data (2026).

The piracy windows are matched closely for all four treated countries, consistent with the use of three averaged outcome windows as the dominant predictors in the Tier 1 specification. The covariate balance is uneven. France's covariates are all matched to within one unit. For Germany, the synthetic unit substantially underestimates GDP per capita (123 versus 85.0 PPS) and overestimates youth unemployment (10.3 versus 19.6 percent), reflecting the absence of a high-income low-unemployment donor with a comparable piracy trajectory. For Italy and Spain, youth unemployment is matched poorly (40.5 versus 19.5 for Italy, 45.0 versus 22.7 for Spain), as a result of the donor weights placing high mass on Northern and Central European countries with much lower youth unemployment.

Eurostat covariate codes

The table below lists the Eurostat datasets used as covariates in this study. The first four are the Tier 1 covariates included in the country-by-country SCM (Section 3.2.2 and 3.3). The remaining four are additional covariates included in the generalized synthetic control estimation (Section 3.5). All series are reported at annual frequency and merged onto the monthly panel by year.

Thesis Label	Eurostat Code	Description	Units	Source URL
gdp_pps	tec00114	GDP per capita	PPS, EU27 = 100	https://ec.europa.eu/eurostat/databrowser/view/tec00114/default/table
youth_unemp	yth_empl_090	Youth unemployment rate, ages 15–24	percent	https://ec.europa.eu/eurostat/databrowser/view/yth_empl_090/default/table
internet	isoc_ci_in_h	Households with internet access	percent of households	https://ec.europa.eu/eurostat/databrowser/view/isoc_ci_in_h/default/table
gini	ilc_di12	Gini coefficient of equivalised disposable income	index, 0–100	https://ec.europa.eu/eurostat/databrowser/view/ilc_di12/default/table
Additional (GSC only)				
broadband	isoc_ci_it_h	Households with broadband internet	percent of households	https://ec.europa.eu/eurostat/databrowser/view/isoc_ci_it_h/default/table
young_share	demo_pjangroup	Population aged 15–29 as share of total population	percent, computed	https://ec.europa.eu/eurostat/databrowser/view/demo_pjangroup/default/table
tertiary	edat_lfse_03	Tertiary educational attainment, ages 25–64	percent	https://ec.europa.eu/eurostat/databrowser/view/edat_lfse_03/default/table
hicp	prc_hicp_aind	Harmonised index of consumer prices, annual	index, 2015 = 100	https://ec.europa.eu/eurostat/databrowser/view/prc_hicp_aind/default/table

AI Disclosure

AI Disclosure. The data extraction pipeline was developed with assistance from a large language model (Claude, Anthropic, specifically Sonnet 4.6, Opus 4.6 and Opus 4.7). The AI generated the Python and R code based on the author's specifications. The author verified the extraction output against the original EUIPO charts for all 27 countries to ensure accuracy. The generated R code was verified against the specifications. While the usage of AI tools have been used to assist the author with code, data extraction, the reading of the literature and the analysis in this thesis was done by the author. The analytical and critical thinking in the writing process has also been the author's own effort.