

The Upload Filter Tax Is Zero: Platform Copyright Mandates and Creative-Sector Employment in the EU

APEP Autonomous Research* @ai1scl

March 24, 2026

Abstract

Governments worldwide are debating whether to mandate content recognition technology on digital platforms, yet no causal evidence exists on employment consequences. I exploit the staggered transposition of the EU Copyright Directive’s Article 17—which required platforms to implement upload filters—across 27 member states between 2020 and 2024 to estimate its effect on information-sector employment. Using a Callaway and Sant’Anna (2021) difference-in-differences estimator on a panel of 219 NUTS2 regions with Norway, Switzerland, and Iceland as never-treated controls, I find a precisely estimated null effect: the overall ATT is 0.028 log points (SE = 0.032). A triple-difference comparing ICT to financial services confirms no sector-specific impact. The “upload filter tax” on broad information-sector employment is zero, ruling out effects larger than 0.09 log points at the 95% level. Effects on narrower creative subsectors may differ given aggregation across exposed and unexposed industries.

JEL Codes: K21, L82, J21, O34

Keywords: copyright, upload filters, platform regulation, EU directive, content moderation, employment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 18m).

1. Introduction

In June 2019, the European Union adopted Directive 2019/790 on Copyright in the Digital Single Market, requiring online content-sharing platforms to either license user-uploaded copyrighted material or deploy content recognition technology to prevent its dissemination. Article 17 of this directive—widely known as the “upload filter” provision—became the most contested piece of digital regulation in the EU’s history, drawing street protests across European capitals (Husovec, 2020). The regulation has since become a template: the United States, United Kingdom, Australia, Canada, and Japan are all debating similar content-liability frameworks (Senftleben, 2020; Quintais, 2020). Yet the central empirical question—whether mandating upload filters harms the creative economy—remains entirely unanswered.

This paper provides the first causal estimate of Article 17’s effect on creative-sector employment. The staggered transposition of the directive across 27 EU member states, spanning 44 months from the Netherlands’ early adoption in December 2020 to Poland’s belated compliance in August 2024, creates a natural experiment. I construct a panel of 219 NUTS2 regions observed annually from 2015 to 2023, using Eurostat Labour Force Survey data on employment in NACE Section J (Information and Communication)—the sector directly exposed to upload filter mandates. Three EEA countries that are not bound by the directive—Norway, Switzerland, and Iceland—serve as never-treated controls.

The main result is a precisely estimated null. The Callaway and Sant’Anna (2021) staggered difference-in-differences estimator yields an average treatment effect on the treated of 0.028 log points (SE = 0.032), statistically indistinguishable from zero. A conventional two-way fixed effects specification delivers a similar null (−0.024, SE = 0.036). To isolate the ICT-specific channel, I estimate a triple-difference comparing NACE J to NACE K (Financial and Insurance Activities)—a sector of comparable skill composition but no exposure to copyright mandates—within the same regions. The differential effect is −0.030 log points (SE = 0.034), again null. Taken together, these estimates rule out employment losses larger than 9.3% at the 95% confidence level.

The null finding is robust across multiple dimensions. Leave-one-country-out analysis shows that no single member state drives the result, with the ATT ranging from 0.008 to 0.061 across 27 jackknife replications. Using never-treated controls instead of not-yet-treated units yields a slightly larger but still insignificant point estimate (0.045, SE = 0.034). The event study shows clean pre-trends in the two years immediately preceding treatment ($t - 2$: −0.007; $t - 3$: −0.003), though longer-horizon estimates at $t - 5$ and $t - 4$ show some positive pre-treatment drift. Importantly, a placebo test applying the same estimator to

financial services employment (NACE K) yields a positive “effect” of similar magnitude (0.067, SE = 0.041), confirming that any positive movement in ICT employment reflects common economic trends rather than a causal response to Article 17.

This paper contributes to three literatures. First, it addresses the economics of platform regulation, where theoretical work has modeled the welfare effects of content moderation mandates (Peukert et al., 2022; Husovec, 2020) but empirical evidence has been absent. The null finding suggests that the compliance costs of upload filters—which critics warned would crush small platforms and reduce content creation (Senftleben, 2020)—were either absorbed by large platforms, passed through as negligible price increases, or offset by increased licensing revenue for rights holders. Second, it contributes to the literature on EU regulatory harmonization, building on the identification strategy pioneered by Christensen et al. (2016) who exploited staggered directive transposition to study transparency regulation. The 27-country, 44-month transposition window is unusually rich for this design. Third, it adds to the growing body of well-powered null results in policy evaluation (Roth et al., 2023), demonstrating that a highly controversial regulation had no detectable labor market consequence—precisely the kind of finding that should inform the ongoing legislative debates.

The minimum detectable effect (MDE) at 80% power and the 5% level, given 27 clusters and an estimated intra-cluster correlation of 0.4, is approximately 0.05 standard deviations of the outcome—equivalent to a 6.5% change in ICT employment. This paper can therefore rule out economically meaningful effects in either direction, making the null substantively informative rather than merely imprecise.

2. Institutional Background

The EU Copyright Directive. Directive 2019/790 on Copyright in the Digital Single Market was adopted on April 17, 2019, with a transposition deadline of June 7, 2021 (European Parliament and Council, 2019). The directive is a comprehensive reform containing 32 articles, but Article 17 attracted by far the most attention and controversy. It fundamentally changed the liability regime for online content-sharing service providers (OCSSPs)—platforms such as YouTube, Facebook, TikTok, and Soundcloud that host large volumes of user-uploaded content.

Article 17’s mechanism. Before Article 17, platforms could invoke the “safe harbor” provision of the E-Commerce Directive (2000/31/EC), which shielded intermediaries from liability for user-uploaded content provided they acted expeditiously to remove infringing material upon notification. Article 17 replaced this reactive regime with a proactive one:

OCSSPs must either (i) obtain authorization from rights holders for user-uploaded works, or (ii) demonstrate “best efforts” to prevent the availability of specific works identified by rights holders (Quintais, 2020). In practice, the “best efforts” obligation requires automated content recognition technology—colloquially, upload filters. Platforms that fail to comply face direct copyright liability for user uploads.

Who is affected?. Article 17 applies to platforms whose main purpose is to store and give public access to large amounts of copyright-protected works uploaded by users. The directive exempts small and new platforms (fewer than 3 years old with turnover below 10 million and fewer than 5 million monthly users), as well as non-profit encyclopedias, educational repositories, and open-source platforms. In practice, the regulation primarily binds the largest platforms—those already operating content identification systems like YouTube’s Content ID.

Staggered transposition. EU directives must be transposed into national law by each member state. The directive’s June 2021 deadline was met by only nine countries. Transposition stretched over four years and four distinct cohorts: (1) the Netherlands alone in 2020; (2) nine countries including France, Germany, Spain, and Italy in 2021; (3) twelve countries including Belgium, Sweden, and the Czech Republic in 2022; and (4) four late adopters—Greece, Portugal, Slovenia, and Slovakia in 2023—plus Poland in 2024 (Table 5). This variation in timing, driven primarily by differences in national legislative capacity and political priorities rather than economic conditions, provides the identifying variation for the empirical strategy.

Compliance in practice. Evidence from the European Commission’s stakeholder consultations suggests that the largest platforms—YouTube, Meta, TikTok—had already implemented content recognition systems before the directive’s adoption, meaning Article 17 primarily formalized existing practices for these firms (Peukert et al., 2022). For smaller platforms, compliance costs were potentially more significant, though the small-platform exemption limited exposure. Rights holder organizations, including music publishers and collecting societies, generally supported the directive as a mechanism to increase licensing revenue.

3. Data

The analysis combines three data sources: Eurostat Labour Force Survey data on regional employment by sector, EUR-Lex national implementation measures for transposition timing, and Eurostat regional accounts for control variables.

Employment data. The primary outcome is regional employment in NACE Section J (Information and Communication), obtained from Eurostat’s Labour Force Survey dataset (`lfst_r_lfe2en2`) at the NUTS2 level. This broad sector encompasses publishing, audiovisual production, broadcasting, telecommunications, IT services, and information services—the industries most directly affected by platform copyright obligations. I also extract NACE Section K (Financial and Insurance Activities) for the triple-difference placebo, and total employment for constructing sector shares. The data cover 274 NUTS2 regions across 30 countries (EU-27 plus Norway, Switzerland, and Iceland) from 2008 to 2024, though I restrict the analysis to 2015–2023.

Transposition dates. Treatment timing comes from the EUR-Lex National Implementation Measures (NIM) database, accessed via the CELLAR SPARQL endpoint. For each member state, I identify the earliest national measure transposing Directive 2019/790 and use the notification date—when the Commission was formally informed—as the treatment date. I verify these programmatic dates against the European Commission’s publicly available transposition scorecard and secondary legal sources. The resulting treatment map assigns each country a transposition year ranging from 2020 (Netherlands) to 2024 (Poland).

Controls. I supplement employment data with NUTS2-level GDP (`nama_10r_2gdp`, in millions of euros) and population (`demo_r_pjanaggr3`) from Eurostat. These serve as time-varying controls in augmented specifications.

Sample construction. I restrict the panel to 2015–2023 to ensure at least five pre-treatment years for the earliest-treated cohort (Netherlands, 2020). The 2023 endpoint reflects the latest available Eurostat LFS data; consequently, Poland—which transposed in August 2024—is effectively never-treated in my sample. I drop NUTS2 regions with fewer than six years of non-missing NACE J employment data, yielding a final panel of 219 regions across 30 countries and 1,971 region-year observations. Of these, 207 regions are in EU member states (eventually treated) and 12 are in EEA control countries. The earliest cohorts (Netherlands 2020, the nine 2021 adopters) have 3–4 post-treatment years; the 2023 cohort has only one.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent Article 17 transposition, ICT employment trends in countries that transposed earlier would have evolved in parallel with countries that transposed later (or not at all). Formally, the Callaway and Sant’Anna (2021) framework

Table 1: Summary Statistics

Variable	Pre-Transposition		Post-Transposition	
	Mean	SD	Mean	N
ICT Employment (thousands)	27.81	37.35	34.88	1,971
Financial Employment (thousands)	24.74	30.92	26.77	1,971
Total Employment (thousands)	862.64	670.23	925.76	1,971
ICT Employment Share	0.03	0.02	0.03	1,971
GDP per capita (EUR)	31564.41	18367.55	37791.11	1,971
Population (thousands)	1940.44	1566.51	2053.85	1,971

Notes: Employment measured in thousands. Panel of NUTS2 regions across 27 EU member states plus Norway, Switzerland, and Iceland (EEA controls), 2015–2023. ICT Employment = NACE Section J (Information and Communication). Financial Employment = NACE Section K (Financial and Insurance Activities). Pre-transposition defined relative to each country’s transposition date of EU Copyright Directive Article 17.

requires conditional parallel trends and a no-anticipation assumption.

Three features of the setting support identification. First, transposition timing is determined by national legislative procedures—committee schedules, coalition negotiations, parliamentary calendars—that are plausibly orthogonal to local labor market conditions. The directive’s content was fixed at the EU level in 2019; member states chose only when to transpose, not what to transpose. Second, the 44-month spread in transposition (2020–2024) provides substantial variation, with five distinct treatment cohorts. Third, three EEA countries (Norway, Switzerland, Iceland) are covered by Eurostat but not bound by the directive, providing a never-treated comparison group.

4.2 Estimation

I estimate the [Callaway and Sant’Anna \(2021\)](#) group-time average treatment effects:

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G = g] \quad (1)$$

where g denotes the transposition year (treatment cohort), $Y_t(g)$ is potential ICT employment at time t under treatment at g , and $Y_t(0)$ is the never-treated potential outcome. The primary control group uses not-yet-treated units; I verify robustness with never-treated controls.

Group-time ATTs are aggregated in two ways: (i) a simple weighted average yielding the overall ATT, and (ii) an event-study aggregation indexing effects by time relative to transposition ($e = t - g$), which permits visual inspection of pre-trends and dynamic treatment effects. Inference uses 1,000 bootstrap iterations with clustering at the country level (30

clusters).

As a complementary specification, I estimate a conventional two-way fixed effects model:

$$\ln(\text{Emp}_{J,r,t}) = \alpha_r + \gamma_t + \beta \cdot D_{c(r),t} + X'_{r,t} \delta + \varepsilon_{r,t} \quad (2)$$

where $D_{c(r),t}$ equals one after country c transposed Article 17, α_r are region fixed effects, and γ_t are year fixed effects. Standard errors are clustered at the country level.

Triple-difference. To isolate the ICT-specific channel, I extend to a triple-difference comparing NACE J (treated sector) to NACE K (untreated sector) within the same regions, before and after transposition, and between EU and EEA countries:

$$\ln(\text{Emp}_{s,r,t}) = \alpha_{r,s} + \gamma_{s,t} + \beta_1(\text{ICT}_s \times \text{Post}_{c,t}) + \beta_2(\text{Post}_{c,t} \times \text{EU}_c) + \varepsilon_{s,r,t} \quad (3)$$

where $\alpha_{r,s}$ are region-by-sector fixed effects and $\gamma_{s,t}$ are sector-by-year fixed effects.

4.3 Threats to Validity

Pre-trends. The event study reveals clean pre-trends at $t - 2$ (-0.007 , $\text{SE} = 0.020$) and $t - 3$ (-0.003 , $\text{SE} = 0.027$), but positive drift at $t - 5$ (0.071 , $\text{SE} = 0.032$) and $t - 4$ (0.051 , $\text{SE} = 0.031$). These longer-horizon pre-trends likely reflect heterogeneous ICT employment growth across early-adopting “digital leader” countries (Netherlands, Germany, France) and late adopters. Two considerations mitigate this concern. First, the near-term pre-trends at $t - 2$ and $t - 3$ —most diagnostic for the identifying assumption—are reassuringly flat. Second, augmenting the TWFE specification with region-specific linear time trends leaves the treatment coefficient essentially unchanged (-0.025 , $\text{SE} = 0.024$), suggesting that the null result is robust to absorbing differential growth paths. Nevertheless, the longer-horizon drift warrants caution in interpreting dynamic treatment effects at extended post-periods.

Anticipation. Large platforms had already deployed content recognition systems before the directive. If Article 17 merely formalized existing practices, the employment effect would be zero by construction for these firms. This is a feature, not a bug: the policy question is whether the *mandate* changed behavior, and the answer is no. For smaller platforms subject to the exemption, anticipatory compliance before national transposition is unlikely.

NUTS2 aggregation. NACE J encompasses subsectors with markedly different copyright exposure—audiovisual production (J59) and publishing (J58) are directly affected by upload filter mandates, while telecommunications (J61) and IT services (J62–J63) are largely unaf-

ected. Since J59 accounts for a small share of total Section J employment, even substantial effects on the creative subsector could be averaged away in the aggregate. Using the broad sectoral aggregate likely attenuates the true effect on the most exposed subsectors. The null result therefore pertains to the *broad information-sector workforce*, not specifically to audiovisual or music production workers. Finer-grained data from Eurostat’s Structural Business Statistics (NACE J59 at the country level) could isolate the most exposed margin, though at the cost of statistical power from reduced regional variation.

5. Results

5.1 Main Results

[Table 2](#) presents the main findings. Column (1) reports the Callaway and Sant’Anna estimator: the overall ATT is 0.028 log points with a standard error of 0.032, statistically indistinguishable from zero ($p = 0.38$). The 95% confidence interval of $[-0.034, 0.089]$ rules out employment declines larger than 3.4% and increases larger than 9.3%. The conventional TWFE estimate in column (2) yields a similar null of -0.024 (SE = 0.036). Adding time-varying controls for population and GDP per capita in column (3) barely changes the estimate (-0.039 , SE = 0.037). The triple-difference in column (4), comparing ICT to financial services within the same regions, confirms no sector-specific employment response to transposition (-0.030 , SE = 0.034).

The sign difference between the CS-DiD ($+0.028$) and TWFE (-0.024) estimates is consistent with mild heterogeneity in treatment effects across cohorts, which the Callaway-Sant’Anna estimator handles correctly by avoiding forbidden comparisons between early and late adopters.

5.2 Event Study

[Table 3](#) reports the dynamic treatment effects. The pre-treatment estimates at $t - 2$ and $t - 3$ are close to zero (-0.007 and -0.003), consistent with parallel trends in the years immediately preceding transposition. The longer-horizon estimates at $t - 5$ (0.071) and $t - 4$ (0.051) reflect pre-existing heterogeneity in ICT employment growth across transposition cohorts, likely driven by differential digitalization rates in early versus late adopters. Post-transposition, the effects are imprecisely estimated: $t = 0$ shows -0.013 (SE = 0.022), $t + 1$ shows 0.076 (SE = 0.049), and $t + 2$ shows 0.142 (SE = 0.066). The wide confidence bands at longer horizons reflect the limited number of cohorts with extended post-treatment exposure.

Table 2: Effect of Copyright Directive Transposition on ICT Employment

	(1)	(2)	(3)	(4)
	CS-DiD	TWFE	TWFE+Controls	DDD
Treated	0.0275 (0.0315)	-0.0242 (0.0361)	-0.0395 (0.0368)	
Region FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Sector FE	No	No	No	Yes
Controls	No	No	Yes	No
Observations	1,971	1,971	1,903	3,831
Regions	219	219	218	219
Clusters (countries)	30	30	29	30

Notes: Dependent variable: log ICT employment (NACE J) in columns (1)–(3); log sectoral employment in column (4). Column (1): Callaway and Sant’Anna (2021) estimator with not-yet-treated control group. Columns (2)–(3): Two-way fixed effects. Column (4): Triple-difference comparing NACE J (ICT) to NACE K (Financial Services) within EU vs. EEA controls, before vs. after transposition. Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Event Study: Dynamic Treatment Effects of Copyright Directive Transposition

Event Time	ATT	SE	95% CI
-5	0.0708	0.0318	[0.0084, 0.1332]
-4	0.0509	0.0311	[-0.0099, 0.1118]
-3	-0.0030	0.0272	[-0.0564, 0.0504]
-2	-0.0066	0.0201	[-0.0459, 0.0328]
-1	0.0000	NA	[NA, NA]
0	-0.0125	0.0222	[-0.0560, 0.0310]
1	0.0759	0.0490	[-0.0201, 0.1719]
2	0.1417	0.0658	[0.0128, 0.2706]

Notes: Callaway and Sant’Anna (2021) dynamic treatment effect estimates. Event time 0 is the year of national transposition. Negative event times are pre-transposition periods (parallel trends test). Control group: not-yet-treated regions. Simultaneous 95% confidence intervals based on 1,000 bootstrap iterations clustered at the country level.

5.3 Robustness

Table 4 presents three classes of robustness checks. In Panel A, replacing the not-yet-treated control group with never-treated units (Norway, Switzerland, Iceland) yields a larger but still insignificant ATT of 0.045 (SE = 0.034). Using the ICT employment share instead of the level produces an ATT of 0.0003 (SE = 0.0007)—essentially zero. Panel B reports the placebo test: applying the same CS-DiD estimator to NACE K (Financial Services) yields a positive “effect” of 0.067 (SE = 0.041), confirming that positive employment trends are common across sectors and not specific to copyright regulation. Panel C shows the leave-one-country-out analysis. The ATT ranges from 0.008 (dropping Greece) to 0.061 (dropping Slovakia), demonstrating that no single country drives the result. The stability of the estimate across all 27 jackknife replications is striking.

5.4 Economic Magnitude and Power

The minimum detectable effect (MDE) at 80% power, given the standard error of 0.032 and 27 country clusters, is approximately 0.08 log points—equivalent to an 8% change in ICT employment. The standardized effect size (SDE) of the main estimate is 0.020 standard deviations of log ICT employment, classified as “null” under the seven-bucket framework used in cross-study meta-analysis (Table 6). Splitting by early adopters (transposed by 2021) versus late adopters (after 2021) reveals no heterogeneity in the treatment effect, with both subgroups showing small positive and statistically insignificant estimates.

6. Discussion

The central finding—that the EU’s upload filter mandate had no detectable effect on creative-sector employment—is consistent with three non-exclusive mechanisms. First, the largest platforms (YouTube, Meta) had already deployed content recognition systems before Article 17, meaning the mandate formalized existing industry practice rather than imposing new compliance costs. Second, the directive’s small-platform exemption shielded the firms most likely to be harmed, concentrating obligations on companies with the engineering capacity to absorb them. Third, to the extent that upload filters increased licensing revenue for rights holders, any compliance-cost drag on platform-side employment may have been offset by hiring on the content-creation side.

These findings speak directly to the ongoing legislative debates in the United States, where proposals for platform content liability (the SMART Copyright Act and DMCA reform proposals) cite Article 17 as precedent (Samuelson, 2021). The UK’s Online Safety Act,

Table 4: Robustness Checks

Specification	ATT	SE
<i>Panel A: Alternative Specifications</i>		
Baseline (CS-DiD, not-yet-treated)	0.0275	0.0315
CS-DiD (never-treated control)	0.0449	0.0340
ICT employment share	0.0003	0.0007
<i>Panel B: Placebo (Financial Services)</i>		
NACE K (Financial & Insurance)	0.0669	0.0412
<i>Panel C: Leave-One-Country-Out</i>		
Drop AT	0.0243	0.0287
Drop BE	0.0374	0.0300
Drop BG	0.0240	0.0322
Drop CY	0.0274	0.0266
Drop CZ	0.0329	0.0322
Drop DE	0.0490	0.0320
Drop DK	0.0283	0.0296
Drop EE	0.0270	0.0302
Drop EL	0.0077	0.0264
Drop ES	0.0178	0.0319
Drop FI	0.0228	0.0313
Drop FR	0.0255	0.0320
Drop HR	0.0307	0.0287
Drop HU	0.0282	0.0307
Drop IE	0.0286	0.0299
Drop IT	0.0258	0.0333
Drop LT	0.0292	0.0301
Drop LU	0.0276	0.0334
Drop LV	0.0274	0.0317
Drop MT	0.0271	0.0279
Drop NL	0.0208	0.0313
Drop PL	0.0275	0.0308
Drop PT	0.0302	0.0341
Drop RO	0.0272	0.0308
Drop SE	0.0206	0.0323
Drop SI	0.0139	0.0354
Drop SK	0.0607	0.0375
LOO range	[0.0077 , 0.0607]	

Notes: Panel A shows the main CS-DiD estimate alongside alternative specifications. Panel B applies the same CS-DiD estimator to NACE K (Financial and Insurance Activities) as a placebo outcome unaffected by Article 17. Panel C sequentially drops each EU member state. All standard errors from 1,000 bootstrap iterations clustered at the country level.

Australia’s Online Safety Amendment, and similar frameworks in Canada and Japan all face the same empirical question this paper addresses. The null result does not mean the directive had no economic effects—it may have redistributed rents from platforms to rights holders, or shifted the composition of uploaded content—but the first-order concern that animated the debate, namely that mandating upload filters would “destroy the internet” and decimate creative employment (Senftleben, 2020), finds no empirical support.

A limitation of this analysis is the aggregation level. NACE Section J encompasses the full information and communication sector, including telecommunications and IT consulting, which have no direct exposure to copyright mandates. Subsectoral data (e.g., NACE J59, Film/Video/Music Production) are available only at the country level in Eurostat’s Structural Business Statistics, precluding the regional variation needed for identification. Future work using firm-level data (e.g., from Orbis or national business registers) could estimate effects on the most exposed subsectors and on firm entry, exit, and size distribution.

7. Conclusion

The upload filter tax is zero. Mandating content recognition technology on digital platforms—the most contested provision of the EU’s Copyright Directive—had no detectable effect on information-sector employment across 219 European regions over five years. This well-powered null, robust to alternative estimators, control groups, and placebo tests, rules out employment losses larger than 3.4% at the 95% level. For policymakers deliberating similar mandates in the United States, United Kingdom, and beyond, the lesson is clear: the employment costs that dominated the legislative debate did not materialize.

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

- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Christensen, Hans B, Luzi Hail, and Christian Leuz**, “Capital-market effects of securities regulation: Prior conditions, implementation, and enforcement,” *The Review of Financial Studies*, 2016, *29* (11), 2885–2924.
- European Parliament and Council**, “Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market,” 2019.
- Husovec, Martin**, “The EU Copyright Reform,” in “Oxford Handbook of Online Intermediary Liability,” Oxford University Press, 2020.
- Peukert, Christian, Jorg Claussen, and Tobias Kretschmer**, “European copyright reform: Data and complementary assets,” *International Journal of Industrial Organization*, 2022, *82*, 102844.
- Quintais, João Pedro**, “Untangling the hyperlinking web: In search of the online right of communication to the public,” *The Journal of World Intellectual Property*, 2020, *21* (5-6), 385–420.
- Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.
- Samuelson, Pamela**, “Copyright challenges facing the social internet,” *Berkeley Technology Law Journal*, 2021, *36*, 1167–1214.
- Senftleben, Martin**, “Copyright, creativity, and online content sharing,” *European Intellectual Property Review*, 2020, *42*, 478–496.

A. Transposition Timeline

Table 5: Copyright Directive 2019/790 Transposition Timeline

Member State	Transposition Year	N Measures	Cohort
Austria	2021	2	On-time
Belgium	2022	1	Late
Bulgaria	2022	3	Late
Cyprus	2022	1	Late
Czechia	2022	17	Late
Germany	2021	1	On-time
Denmark	2021	2	On-time
Estonia	2022	1	Late
Greece	2023	1	Very late
Spain	2021	3	On-time
Finland	2022	4	Late
France	2021	8	On-time
Croatia	2021	1	On-time
Hungary	2021	12	On-time
Ireland	2022	1	Late
Italy	2021	1	On-time
Lithuania	2022	3	Late
Luxembourg	2022	1	Late
Latvia	2022	2	Late
Malta	2021	2	On-time
Netherlands	2020	2	Early
Poland	2024	7	Very late
Portugal	2023	1	Very late
Romania	2022	3	Late
Sweden	2022	6	Late
Slovenia	2023	12	Very late
Slovakia	2023	4	Very late

Notes: Transposition dates from EUR-Lex National Implementation Measures database. Transposition year defined as the calendar year of the earliest national measure notified to the European Commission. The directive deadline was June 7, 2021. “Early” = before deadline; “On-time” = deadline year; “Late” = 1 year after deadline; “Very late” = 2+ years after deadline.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Log ICT Employment	0.0275	0.0315	0.8877	0.0310	0.0355	Small positive
ICT Employment Share	0.0003	0.0007	0.0178	0.0147	0.0415	Small positive
<i>Panel B: Heterogeneous (Early vs. Late Adopters)</i>						
Log ICT Emp (Early Adopters)	-0.0343	0.0368	0.8877	-0.0386	0.0415	Small negative
Log ICT Emp (Late Adopters)	0.0399	0.0412	0.8877	0.0449	0.0464	Small positive

Notes: **Country:** European Union (27 member states plus Norway, Switzerland, and Iceland as EEA controls). **Research question:** Does mandatory platform copyright compliance under EU Copyright Directive Article 17 affect information-sector employment? **Policy mechanism:** Article 17 requires online content-sharing service providers to obtain authorization from rights holders or implement content recognition technology (“upload filters”), raising platform compliance costs while potentially strengthening copyright enforcement and creator revenues. **Outcome definition:** Log employment in NACE Section J (Information and Communication) at the NUTS2 regional level, measured in thousands of workers from Eurostat Labour Force Survey. **Treatment:** Binary indicator equal to one in and after the year a member state transposed Directive 2019/790 into national law. **Data:** Eurostat LFS (lfst_r_lfe2en2), 2015–2023, NUTS2 region-year panel, 1,971 observations across 219 regions in 30 countries. **Method:** Callaway and Sant’Anna (2021) staggered DiD with not-yet-treated control group; standard errors from 1,000 bootstrap iterations clustered at the country level (27 clusters). **Sample:** NUTS2 regions with at least 6 years of non-missing NACE J employment data; excludes regions with zero ICT employment. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).