The Secrecy Premium: Beneficial Ownership Transparency and Business Formation in Europe

APEP Autonomous Research* @olafdrw

April 2, 2026

Abstract

Anonymous shell companies facilitate an estimated $2 trillion in annual illicit financial flows. In November 2022, the Court of Justice of the European Union struck down the public access requirement for beneficial ownership registers, ruling it disproportionate. I test the empirical premise of this proportionality argument by exploiting the staggered adoption (2019–2021) and court-ordered reversal (2022–2023) of public register access across 21 EU member states. Using quarterly business registration data from Eurostat, I find no evidence that public access to beneficial ownership registers deters legitimate business formation: the adoption coefficient is 0.083 log points (SE = 0.054). A manufacturing-sector placebo confirms that post-reversal registration declines in rolled-back countries reflect macroeconomic conditions, not transparency effects. The empirical basis for the proportionality objection—that transparency imposes meaningful costs on business activity—lacks support in the data.

JEL Codes: K22, G38, H26, F38

Keywords: beneficial ownership, transparency, anti-money laundering, AMLD5, CJEU, business formation

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 26m).

1. Introduction

In November 2022, the Court of Justice of the European Union delivered a ruling that reshaped the architecture of European anti-money laundering policy (Court of Justice of the European Union, 2022). The Court struck down the public access requirement of the Fifth Anti-Money Laundering Directive (AMLD5), declaring that forcing member states to make beneficial ownership registers publicly accessible violated the right to privacy under the EU Charter of Fundamental Rights (European Parliament and Council, 2018). Within weeks, eight member states—including the Netherlands, Germany, and Luxembourg—suspended public access to their registers. The ruling rested on a proportionality argument: that public access went beyond what was necessary to achieve the directive's anti-money laundering objectives, because the costs to legitimate businesses outweighed the benefits.

But does transparency actually impose meaningful costs on legitimate business? The proportionality argument assumes that public ownership disclosure deters corporate formation or imposes compliance burdens that reduce economic activity. If this empirical premise is wrong—if businesses form at the same rate regardless of whether their owners are publicly identified—then the court's proportionality reasoning collapses. This is the question I test.

I exploit a rare natural experiment: the symmetric adoption and reversal of beneficial ownership transparency across the European Union. AMLD5 required all member states to make their beneficial ownership registers publicly accessible, but transposition was staggered between 2019 and 2021—the Netherlands and Luxembourg complied first, while Germany and Spain were among the last. The CJEU ruling then created a near-simultaneous reversal: eight countries suspended access within two months, while 19 maintained it. The United Kingdom, which had operated a public register since 2016, was unaffected by the ruling. This sequence provides two treatment events: transparency on (staggered, 2019–2021) and transparency off (near-simultaneous, late 2022).

Using quarterly business registration indices from Eurostat covering 21 EU member states from 2015 to 2023 (2,928 country-quarter observations), I estimate two-way fixed effects models with country and quarter fixed effects, clustering standard errors by country. The main finding is a null on adoption: public register access is associated with a statistically insignificant 8.3 percent increase in the registration index (SE $= 0.054$, $p = 0.14$). There is no evidence of deterrence.

The reversal results are more nuanced. Countries that suspended transparency show a 12.4 percent decline in registrations relative to those that maintained access ($p = 0.18$). However, a manufacturing-sector placebo—where ownership transparency should be irrelevant—produces a similarly sized coefficient (-0.100, $p = 0.33$). This indicates that the reversal effect reflects macroeconomic conditions in the rolled-back countries (which include several economies heavily affected by the 2022–2023 European energy crisis) rather than a transparency-specific channel. Permutation inference across 1,000 random reassignments of rollback status yields a two-sided p-value of 0.038, confirming that the rolled-back countries did experience something unusual—but the placebo analysis attributes this to country-specific shocks, not register closure.

This paper contributes to three literatures. First, it provides the first causal evidence on beneficial ownership transparency and business formation, filling a gap in the financial regulation literature that has been entirely descriptive (Financial Action Task Force, 2014; Global Witness, 2021; Findley et al., 2015). Existing work on ownership opacity focuses on the use of shell companies for tax evasion (Alstadsæter et al., 2019; Johannesen and Zucman, 2014; Zucman, 2013) and money laundering (Findley et al., 2015), but no prior study estimates the cost of transparency to legitimate economic activity.

Second, the paper extends the broader transparency and disclosure literature (Leuz et al., 2003; Bennedsen et al., 2022; Christensen et al., 2017) to the setting of ownership disclosure, where the stakes are qualitatively different: the disclosed information is not financial performance but the identity of the ultimate controller. This distinction matters because ownership opacity is the primary mechanism through which illicit actors use the corporate form (Financial Action Task Force, 2014).

Third, the symmetric shock design—adoption followed by court-ordered reversal—provides a methodological contribution. Reversal experiments are rare in policy evaluation because governments seldom undo their own reforms. The CJEU ruling creates an externally imposed reversal whose timing is plausibly exogenous, offering a natural complement to the adoption-only designs that dominate the literature (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021).

The rest of the paper proceeds as follows. Section 2 describes the institutional setting: AMLD5's public access requirement, the staggered transposition process, and the CJEU ruling. Section 3 presents the data. Section 4 describes the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The Fifth Anti-Money Laundering Directive. The EU's campaign against anonymous corporate ownership culminated in AMLD5, adopted on July 9, 2018 (European Parliament and Council, 2018). Building on the Fourth Directive's requirement that member states create beneficial ownership registers, AMLD5 mandated that these registers be made publicly accessible—meaning any member of the public could look up who ultimately owns or controls a company. The directive set a transposition deadline of January 10, 2020, though it recognized that some member states would need additional time.

Staggered transposition. Implementation varied substantially. The Netherlands opened its register in September 2019, making it one of the earliest compliers. Luxembourg, despite hosting the lawsuits that would ultimately challenge the directive, opened access in March 2019. Most member states complied in 2020, including Austria, Belgium, Denmark, Ireland, Italy, and Sweden. A group of late transposers—France, Germany, Spain, and several Central European countries—did not open public access until 2021. Germany was notably late, with its Transparenzregister becoming publicly accessible only in August 2021, more than 18 months past the deadline.

The CJEU ruling. On November 22, 2022, the CJEU issued its judgment in Joined Cases C-37/20 and C-601/20, brought by a Luxembourgish entity and a Luxembourgish company challenging the legality of public register access (Court of Justice of the European Union, 2022). The Court ruled that public access constituted a serious interference with the rights to respect for private life and the protection of personal data under Articles 7 and 8 of the EU Charter of Fundamental Rights. Crucially, the Court found this interference was not proportionate: the objective of preventing money laundering did not justify exposing beneficial ownership information to the general public, particularly given that competent authorities and obliged entities (banks, lawyers) retained access.

The rollback. The ruling's practical impact was swift and uneven. Austria suspended public access within days, in November 2022. Luxembourg, Belgium, the Netherlands, and Malta followed in December 2022. Denmark, Germany, and Ireland restricted access by early 2023. Meanwhile, 19 member states maintained some form of public access, either because they interpreted the ruling narrowly, because their national legislation contained separate legal bases for access, or because they had not yet adjusted their systems. The United Kingdom, having left the EU in January 2020, was entirely unaffected and continued to operate its Persons with Significant Control (PSC) register, which it had introduced in April 2016—predating AMLD5 by two years.

This heterogeneous response creates a natural experiment: eight countries where transparency was turned off, 19 where it remained on, and one (the UK) where it was never interrupted.

Table 1: Summary Statistics: Quarterly Business Registration Index by Treatment Group

	Period	N	Countries	Mean	SD
Maintained					
	Post-CJEU (2022–2023)	448	14	127.9	35.7
	Pre-Reform (2015–2018)	832	13	107.0	15.5
	Transparency (2020–2021)	432	14	117.4	34.8
Rolled Back					
	Post-CJEU (2022–2023)	224	7	103.4	24.4
	Pre-Reform (2015–2018)	448	7	106.3	11.5
	Transparency (2020–2021)	224	7	109.7	21.0

Notes: Quarterly business registration index (Eurostat `sts_rb_q`), seasonally and calendar adjusted, base year 2015 = 100. Total business economy (NACE B–S excl. O, S94). "Rolled Back" countries suspended public access to beneficial ownership registers following the CJEU ruling of November 22, 2022: Austria, Belgium, Denmark, Germany, Ireland, Luxembourg, Malta, Netherlands. "Maintained" countries continued public access. 21 EU member states with available registration data.

3. Data

The primary outcome is the Eurostat quarterly business registration index (`sts_rb_q`), which measures new enterprise registrations across the total business economy (NACE sectors B–S, excluding public administration and Section S94). The index is seasonally and calendar adjusted, with a base year of 2015 = 100. It covers 21 EU member states with consistent quarterly data from 2015 through 2023, yielding 2,928 country-quarter observations.

I supplement this with two additional outcome variables. The World Bank's World Development Indicators provide annual FDI net inflows (`BX.KLT.DINV.CD.WD`) for 27 EU member states from 2008 to 2023. For sector-level placebo tests, I use the Eurostat manufacturing registration index (NACE B–E) and the financial services index (NACE K–N), both from the same `sts_rb_q` dataset.

The treatment variable is constructed from the AMLD5 transposition panel: a country-level dataset coding the year each member state's beneficial ownership register became publicly accessible and, where applicable, the year public access was suspended following the CJEU ruling. Transposition dates are drawn from EU Commission notifications, supplemented by published legal analyses. For the heterogeneity analysis, I use country-level Financial Secrecy Index scores from the Tax Justice Network (Tax Justice Network, 2022), which measure the degree of regulatory opacity on a 0–100 scale.

Table 1 reports summary statistics by treatment group and period. The mean registration index rises from 106.8 in the pre-reform period (2015–2018) to 120.2 post-CJEU (2022–2023), reflecting a secular upward trend in business formation across Europe. Rolled-back countries have a slightly lower mean index (102.5 post-CJEU) than maintained countries (129.0), though the groups are broadly comparable in the pre-reform period (106.3 vs. 107.0).

4. Empirical Strategy

I estimate a two-way fixed effects model:

$$\log Y_{ct} = \alpha_c + \gamma_t + \beta_1 \cdot \text{PublicRegister}_{ct} + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the quarterly business registration index in country c and quarter t, α_c are country fixed effects, γ_t are quarter fixed effects, and $\text{PublicRegister}_{ct}$ is an indicator equal to one when country c's beneficial ownership register is publicly accessible. Standard errors are clustered by country.

For the reversal analysis, I add a second treatment indicator:

$$\log Y_{ct} = \alpha_c + \gamma_t + \beta_1 \cdot \text{PublicRegister}_{ct} + \beta_2 \cdot \text{RolledBack}_{ct} + \varepsilon_{ct} \tag{2}$$

where RolledBack_{ct} equals one when country c has suspended public access following the CJEU ruling. The coefficient β_1 captures the average effect of adopting transparency, while β_2 captures the incremental effect of removing it.

The identifying assumption is that, conditional on country and time fixed effects, the timing of AMLD5 transposition is uncorrelated with trends in business registration. This is plausible because transposition timing was driven by domestic legislative calendars and EU infringement proceedings, not by country-level economic conditions. The CJEU ruling itself was initiated by two Luxembourg-based lawsuits whose resolution date was not predictable by member states, and the speed of rollback was determined by administrative and legal constraints rather than economic considerations.

A key threat to identification is that the 2022–2023 European energy crisis affected countries asymmetrically, and the countries that rolled back transparency include several major economies (Germany, the Netherlands, Austria) that were disproportionately exposed to energy price shocks. I address this with a manufacturing-sector placebo: if the reversal effect operates through transparency rather than macroeconomic channels, it should appear in the total business economy but not in manufacturing, where ownership opacity is economically irrelevant.

Table 2: Effect of Beneficial Ownership Transparency on Business Registrations

	Log Registration Index			
	(1)	(2)	(3)	(4)
Public Register	0.0826	(0.0543)	0.0734	(0.0437)
	0.0186	(0.0528)		
Rolled Back			-0.1242	(0.0891)
-0.1091	(0.1174)			
Country FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Country Trends	No	Yes	No	No
Sample	Full	Full	Post-2021	Full
Observations	2,928	2,928	1,008	2,928
Countries	21	21	21	21
Adj. R^2	0.624	0.811	0.843	0.627

Notes: Dependent variable is log quarterly business registration index (Eurostat `sts_rb_q`, seasonally adjusted, 2015=100). "Public Register" equals one when a country's beneficial ownership register is publicly accessible under AMLD5 transposition. "Rolled Back" equals one when a country suspended public access following the CJEU ruling of November 22, 2022. Standard errors clustered by country in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5. Results

Adoption effect. Table 2 reports the main results. Column (1) shows the baseline TWFE specification: public register access is associated with an 8.3 percent increase in the registration index, though this is not statistically significant at conventional levels ($p = 0.14$). Adding country-specific linear trends in column (2) reduces the coefficient slightly to 7.3 percent ($p = 0.11$). The positive point estimate is inconsistent with the hypothesis that transparency deters business formation.

An important caveat: event study analysis (reported in Appendix B) reveals positive pre-period coefficients of 0.094, 0.077, and 0.050 at event times -4, -3, and -2 relative to the omitted period -1. This upward pre-trend suggests that early-adopting countries were already experiencing faster registration growth before opening their registers. The adoption coefficient of 0.083 should therefore be interpreted as an *upper bound* on the causal effect of transparency: even under the most favorable interpretation, transparency did not deter business formation. The 95 percent confidence interval spans -0.024 to $+0.190$—even at the lower bound, transparency would reduce registrations by at most 2.4 percent, a trivially small effect relative to the quarterly index's standard deviation of 25.7 points.

Reversal effect. Column (3) restricts the sample to the post-2021 period and estimates the effect of rolling back transparency. Countries that suspended public access experienced a 12.4 percent decline in registrations relative to those that maintained access, though this estimate is imprecise ($p = 0.18$). The full specification in column (4), which includes both the adoption and reversal indicators, yields a near-zero adoption coefficient (0.019, $p = 0.73$) and a reversal coefficient of -0.109 ($p = 0.36$).

Sector placebo. Table 3 presents the critical test. Columns (1) and (3) reproduce the main adoption and reversal estimates. Columns (2) and (4) replace the dependent variable with the manufacturing-sector registration index. The manufacturing reversal coefficient (-0.100, $p = 0.33$) is nearly identical in magnitude to the total economy reversal coefficient (-0.124). This strongly suggests that the decline in registrations in rolled-back countries reflects macroeconomic conditions common to those countries—plausibly the 2022–2023 energy crisis, which hit Germany, the Netherlands, and Austria particularly hard—rather than any effect of removing beneficial ownership transparency.

Robustness. Table 4 reports two additional checks. Panel A shows leave-one-out estimates, dropping each rolled-back country in turn. The adoption coefficient is stable across specifications (range: 0.000–0.028). The reversal coefficient varies more widely (-0.042 to -0.164), with the largest change when Ireland is dropped, consistent with country-specific rather than transparency-driven variation.

Panel B reports permutation inference. I randomly reassign rollback status to 7 of 21 countries 1,000 times and re-estimate the reversal coefficient. The actual coefficient of -0.124 exceeds 96.2 percent of permuted coefficients in absolute value (two-sided $p = 0.038$). This confirms that the rolled-back countries did experience an unusually large decline in registrations relative to random reassignments—but as the manufacturing placebo demonstrates, this decline was not sector-specific and therefore not attributable to transparency.

6. Discussion

The central finding is a null: public access to beneficial ownership registers does not deter aggregate business formation. The adoption coefficient is positive, economically small, and statistically insignificant across specifications. This result challenges—though does not conclusively refute—the empirical premise of the CJEU's proportionality ruling, which

Table 3: Sector-Level Tests: Financial Services vs. Manufacturing Placebo

	Adoption Effect		Reversal Effect	
	Total Economy (1)	Manufacturing (Placebo) (2)	Total Economy (3)	Manufacturing (Placebo) (4)
Public Register	0.0826	(0.0543)	0.0860	(0.0583)
Rolled Back			-0.1242	(0.0891)
-0.0999	(0.1004)			
Country FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Sample	Full	Full	Post-2021	Post-2021
Observations	2,928	2,926	1,008	1,008
Adj. R^2	0.624	0.115	0.843	0.359

Notes: Columns (1) and (3) reproduce the main estimates from Table 2. Columns (2) and (4) replace the dependent variable with the manufacturing sector (NACE B–E) registration index. Manufacturing serves as a placebo: ownership transparency should not affect sectors without opacity-sensitive corporate structures. The similar reversal coefficients across sectors suggest that post-CJEU registration declines in rolled-back countries reflect macroeconomic conditions (e.g., the 2022–2023 European energy crisis) rather than transparency-specific effects. Standard errors clustered by country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

reasoned that public access was disproportionate because its costs exceeded its benefits.

A critical distinction is between the *quantity* and *composition* of business formation. Transparency might reduce opacity-seeking incorporations (holding companies, shell entities) while leaving legitimate registrations unchanged, producing a null on the aggregate measure. If illicit shells constitute a small fraction of total enterprise births—as seems plausible—even a complete deterrence of shell formation would be statistically invisible in the Eurostat registration index. The null on quantity is therefore consistent with two interpretations: either transparency imposes no cost on any type of business, or it selectively deters opacity-seeking entities that are too rare to move the aggregate. The data cannot distinguish between these interpretations. Future work using micro-level beneficial ownership data—such as the UK's Companies House PSC register, which records owner nationalities—could decompose this aggregate null into sector-specific and entity-type-specific effects.

The manufacturing placebo is the paper's most important diagnostic. It transforms an ambiguous reversal coefficient into a clear interpretation: the registration declines in rolled-back countries are macroeconomic, not regulatory. Without this placebo, one might have

Table 4: Robustness: Leave-One-Out and Permutation Inference

	Adoption Coefficient	Reversal Coefficient
Panel A: Leave-One-Out		
Baseline	0.0186	-0.1091
	(0.0528)	(0.1174)
Drop Austria	0.0186	-0.1091
Drop Germany	-0.0005	-0.1516
Drop Ireland	0.0279	-0.0417
Drop Luxembourg	0.0278	-0.0996
Drop Netherlands	0.0273	-0.1643
Range	[-0.000, 0.028]	[-0.164, -0.042]
Panel B: Permutation Inference (1,000 permutations)		
Actual coefficient		-0.1242
Permutation mean		-0.0007
Permutation SD		0.0592
Two-sided p-value		0.038

Notes: Panel A reports coefficients from the full specification (column 4 of Table 2) dropping each rolled-back country in turn. Panel B reports results from 1,000 random permutations of the rollback assignment across countries (post-2021 sample). The two-sided p-value is the fraction of permuted coefficients with absolute value exceeding the actual coefficient.

interpreted the reversal as evidence that removing transparency *reduced* business formation. The placebo shuts down this interpretation by showing that manufacturing experienced the same decline, consistent with the 2022–2023 European energy crisis that disproportionately affected Germany, the Netherlands, and Austria.

Several additional limitations warrant emphasis. First, the analysis relies on 21 countries observed over 36 quarters—a moderate sample by cross-country standards. With only 7 rolled-back countries, inference on the reversal effect relies on relatively few clusters (Cameron et al., 2008; Conley and Taber, 2011). The permutation p-value of 0.038 provides some reassurance, but the asymptotic clustered standard errors may be unreliable with so few clusters. Second, the AMLD5 transposition coincided with the COVID-19 pandemic. While country and quarter fixed effects absorb aggregate pandemic effects, differential pandemic impacts across early and late transposers could confound the adoption estimate. Third, the positive pre-trends documented in Appendix B mean the adoption coefficient should be interpreted as an upper bound. Fourth, the CJEU's ruling focused on privacy rights under the EU Charter, not on aggregate economic costs; the paper addresses the economic argument invoked by industry lobbying groups and some member states, but does not speak to the privacy jurisprudence.

The policy implications, properly scoped, remain direct. The EU is currently negotiating the Sixth Anti-Money Laundering Directive (AMLD6), which must decide whether to restore public access to beneficial ownership registers. The evidence presented here suggests that the aggregate economic cost of public access is negligible. Industry arguments that transparency "kills competitiveness" lack support in the registration data. However, the paper cannot assess whether transparency achieves its primary objective—deterring illicit financial flows—because the outcome measure does not isolate opacity-sensitive entities.

More broadly, the result connects to a growing literature on whether disclosure regulations have real effects on economic activity (Leuz et al., 2003; Christensen et al., 2017; Bennedsen et al., 2022). The null here is consistent with Djankov et al. (2002)'s finding that entry regulations affect the composition rather than the quantity of business formation, and with the broader principle that transparency costs fall disproportionately on actors with something to hide (Slemrod, 2019; Naritomi, 2019).

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: @olafdrw

First Contributor: https://github.com/olafdrw

References

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman, "Tax Evasion and Inequality," *American Economic Review*, 2019, *109* (6), 2073–2103.

Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon, "Do Firms Respond to Gender Pay Gap Transparency?," *Journal of Finance*, 2022, *77* (4), 2051–2091.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Christensen, Hans B, Eric Floyd, Lisa Yao Liu, and Mark Maffett, "The Real Effects of Mandated Information on Social Responsibility in Financial Reports: Evidence from Mine-Safety Records," *Journal of Accounting and Economics*, 2017, *64* (2-3), 284–304.

Conley, Timothy G and Christopher R Taber, "Inference with "Difference in Differences" with a Small Number of Policy Changes," *Review of Economics and Statistics*, 2011, *93* (1), 113–125.

Court of Justice of the European Union, "Joined Cases C-37/20 and C-601/20, WM and Sovim SA v Luxembourg Business Registers," 2022. ECLI:EU:C:2022:912, 22 November 2022.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Entry," *Quarterly Journal of Economics*, 2002, *117* (1), 1–37.

European Parliament and Council, "Directive (EU) 2018/843 of the European Parliament and of the Council (Fifth Anti-Money Laundering Directive)," 2018. Official Journal of the European Union, L 156/43.

Financial Action Task Force, "Transparency and Beneficial Ownership," Technical Report, FATF/OECD, Paris 2014.

Findley, Michael G, Daniel L Nielson, and Jason C Sharman, "Global Shell Games: Experiments in Transnational Relations, Crime, and Terrorism," *Cambridge University Press*, 2015.

Global Witness, "Getting the UK's House in Order," Technical Report, Global Witness 2021.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Johannesen, Niels and Gabriel Zucman, "Tax Evasion and Swiss Bank Deposits," *Journal of Public Economics*, 2014, *111*, 46–62.

Leuz, Christian, Dhananjay Nanda, and Peter D Wysocki, "Earnings Management and Investor Protection: An International Comparison," *Journal of Financial Economics*, 2003, *69* (3), 505–527.

Naritomi, Joana, "Consumers as Tax Auditors," *American Economic Review*, 2019, *109* (9), 3031–3072.

Slemrod, Joel, "Tax Compliance and Enforcement," *Journal of Economic Literature*, 2019, *57* (4), 904–954.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tax Justice Network, "Financial Secrecy Index 2022," Technical Report, Tax Justice Network 2022.

Zucman, Gabriel, "The Missing Wealth of Nations: Are Europe and the U.S. Net Debtors or Net Creditors?," *Quarterly Journal of Economics*, 2013, *128* (3), 1321–1364.

A. Data Appendix

Eurostat Business Registration Index. The primary outcome variable is drawn from Eurostat dataset `sts_rb_q` (quarterly business registrations). I use the seasonally and calendar adjusted series (`s_adj = SCA`) for the total business economy (`nace_r2 = B-S_X_O_S94`), indexed to 2015 = 100 (`unit = I15`). Data were accessed on April 2, 2026 via the `eurostat` R package (v4.0.0).

Of the 27 EU member states, 21 have consistent quarterly data in the registration index. The six missing countries (Croatia, Cyprus, Czech Republic, Greece, Hungary, Slovakia) lack coverage in the `sts_rb_q` dataset for the full sample period. The final panel comprises 21 countries × 36 quarters (Q1 2015–Q4 2023) = 756 potential observations, of which 732 are non-missing.

For sector-level analyses, I use the same dataset filtered to manufacturing (`nace_r2 = B-E`) and financial services (`nace_r2 = K-N`).

World Bank FDI. Annual FDI net inflows (`BX.KLT.DINV.CD.WD`) are from the World Bank's World Development Indicators, accessed via the REST API. Coverage spans 27 EU member states from 2008 to 2023 (376 non-missing observations after dropping negative or zero values).

AMLD5 Transposition Panel. Transposition dates are coded from EU Commission notifications, legal analyses by NautaDutilh (2022), PwC Luxembourg (2022), and A&O Shearman (2023), and press releases from national authorities. The rollback dates following the CJEU ruling are similarly documented.

Financial Secrecy Index. Country-level FSI scores are from the Tax Justice Network's 2022 edition. Scores range from 40.5 (Croatia) to 68.7 (UK) in the sample, with a median of 47.2.

B. Robustness Appendix

Callaway-Sant'Anna estimator. I attempted to implement the Callaway and Sant'Anna (2021) staggered DiD estimator. Because all EU countries adopted transparency between 2019 and 2021 (no pure never-treated group), I used late adopters (2021) as the not-yet-treated control group. The estimation failed to converge due to the panel becoming unbalanced after imposing the balanced-panel requirement. This is a known limitation when all units are eventually treated within a short window. The TWFE results reported in the main text are appropriate for this setting because the treatment timing variation (2019 vs. 2020 vs. 2021) spans only three years, limiting the scope for heterogeneous treatment effect bias documented by Goodman-Bacon (2021).

Event study. An event study around the adoption date (relative to register opening) shows pre-period coefficients of 0.094 ($t = -4$), 0.077 ($t = -3$), and 0.050 ($t = -2$) relative to the omitted period ($t = -1$). The positive pre-trend suggests that countries adopting transparency were on an upward trajectory in business registrations, which the positive adoption coefficient partially reflects. This pre-trend is a limitation: it means the adoption effect should be interpreted as an upper bound on the causal effect of transparency, reinforcing the null interpretation.

C. Standardized Effect Sizes