Not Until You Mean It: Cybersecurity Regulation Only Changes Firm Behavior Where It Is Enforced

APEP Autonomous Research* @ai1scl

March 24, 2026

Abstract

Governments worldwide mandate cybersecurity practices, but does regulation change firm behavior or merely generate paperwork? I exploit the EU's NIS2 Directive, which requires cybersecurity investments from enterprises with 50 or more employees while exempting smaller firms, using Eurostat's triennial ICT security survey across 27 member states. A difference-in-differences comparing newly regulated medium firms to exempt small firms finds no significant aggregate effect. However, a difference-in-difference-in-differences exploiting staggered transposition reveals that effects concentrate in the six countries that met the implementation deadline: a 3.4 percentage point increase in the security index, driven by gains in both technical measures and staff training. Where transposition remained incomplete, firms adopted only the cheapest visible measure—training. Enforcement, not announcement, determines whether cybersecurity regulation produces real security investment.

JEL Codes: K23, L51, O33

Keywords: cybersecurity regulation, NIS2 Directive, compliance, enforcement, EU policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 28m).

1. Introduction

A cyberattack on a mid-sized European manufacturer costs an average of 200,000 in direct losses and weeks of disrupted operations (European Union Agency for Cybersecurity, 2023). In response, governments have turned to regulation: the EU's NIS2 Directive, enacted in January 2023, is the most ambitious cybersecurity mandate ever imposed, covering an estimated 160,000 enterprises across 27 member states (European Parliament and Council, 2022). Yet the fundamental question remains open: does mandating cybersecurity actually make firms more secure, or does it merely produce compliance documentation?

This question connects to a deep tension in the regulation literature. Optimists argue that regulation corrects information failures and coordination problems that lead to under-investment in security (Anderson, 2001; Gordon and Loeb, 2002). Skeptics counter that compliance mandates produce "security theater"—visible but substantively hollow responses that satisfy auditors without reducing risk (Schneier, 2003; Bauer and van Eeten, 2009). The empirical literature on cybersecurity regulation, almost entirely focused on US state-level breach notification laws, provides mixed evidence: these laws appear to increase data breach disclosure without clearly reducing breach incidence (Romanosky et al., 2011; Bisogni et al., 2011).

I provide the first causal evidence on the EU's NIS2 Directive by exploiting a sharp size threshold and staggered transposition across member states. NIS2 imposes cybersecurity obligations—risk assessment, incident reporting within 24 hours, supply chain security, and mandatory staff training—on enterprises with 50 or more employees, while exempting smaller firms (European Parliament and Council, 2022). This threshold creates a natural control group: small enterprises (10–49 employees) that are otherwise similar in their digital risk profile but fall outside the regulatory perimeter.

My identification strategy uses a difference-in-differences (DiD) comparing medium enterprises (50–249 employees, newly regulated) to small enterprises (10–49, exempt) before and after NIS2 implementation, using three waves of Eurostat's ICT security survey (2019, 2022, 2024). The key innovation is a difference-in-difference-in-differences (DDD) that interacts this size-based DiD with an indicator for whether each member state completed NIS2 transposition by the October 2024 deadline. Only six countries—Belgium, Croatia, Hungary, Italy, Latvia, and Lithuania—met this deadline, while 21 had not yet transposed the directive into national law (European Commission, 2024).

The headline finding is stark: across all 27 EU member states, NIS2 had no statistically significant effect on the overall cybersecurity index of medium firms relative to small firms ($\hat{\beta}$ = 0.42 percentage points, p = 0.604). But this aggregate null masks sharp heterogeneity by enforcement status. In the six countries that completed transposition, the DDD estimate shows a 3.4 percentage point increase in the security index ($p = 0.009$), driven by gains across compliance, technical, and training measures. In the remaining 21 countries, the effect is a precisely estimated zero ($\hat{\beta} = -0.33$, $p = 0.730$). Enforcement, not announcement, determines regulatory impact.

Decomposing by measure type reveals the mechanism. In countries without transposition, the only individually significant effect is a 3.95 percentage point increase in staff training provision ($p < 0.001$)—the cheapest and most auditable response to anticipated regulation. In transposed countries, the additional DDD effect reaches 2.98 percentage points for technical measures ($p = 0.021$), including encryption, network access control, and security testing. This pattern—cheap compliance where enforcement is distant, substantive investment where it is imminent—is consistent with rational firm behavior under regulatory uncertainty (Stigler, 1971; Shimshack, 2014).

Several features strengthen identification. First, parallel trends hold cleanly: the medium-vs-small size gap was stable between 2019 and 2022 ($\hat{\beta} = 0.43$, $p = 0.818$), ruling out differential pre-trends. Second, a dosage test using large firms ($\geq$250 employees), already regulated under the predecessor NIS1 Directive and subject to intensified NIS2 requirements, shows a significant 2.2 percentage point effect ($p = 0.001$), confirming that NIS2 added real regulatory bite even for previously regulated firms. Third, the DDD result is robust to sequentially dropping each of the six transposed countries, with the triple interaction remaining significant in all six leave-one-out exercises ($p \leq 0.020$).

This paper contributes to three literatures. First, it provides the first causal evidence on cybersecurity regulation in Europe, complementing the US-focused breach notification literature (Romanosky et al., 2011; Johnson et al., 2020; Laube and Böhme, 2016). Second, it speaks to the broader enforcement literature by demonstrating that the compliance response to regulation is not uniform but depends critically on whether firms perceive enforcement as imminent (Shimshack, 2014; Duflo et al., 2013; Gray and Shimshack, 2011). The DDD design isolates this enforcement channel cleanly: the same regulation, announced at the same time, produces different behavioral responses depending solely on the speed of national transposition. Third, the decomposition into compliance, technical, and training measures contributes to the "compliance theater" debate by showing that theater is the equilibrium response to distant enforcement, while real investment emerges when compliance becomes legally binding (Schneier, 2003; Anderson, 2001).

2. Institutional Background

The NIS2 Directive. The Network and Information Security Directive 2 (Directive 2022/2555) was adopted on December 14, 2022, replacing the original NIS Directive (2016/1148) (European Parliament and Council, 2022). NIS2 dramatically expanded the scope of EU cybersecurity regulation along three dimensions. First, it lowered the size threshold from large enterprises to all entities with 50 or more employees or annual turnover exceeding 10 million. Second, it expanded sectoral coverage from essential services (energy, transport, banking) to 18 sectors including manufacturing, food production, postal services, and digital providers. Third, it strengthened obligations, requiring 24-hour incident reporting, regular risk assessments, supply chain security measures, and mandatory staff cybersecurity training.

The size threshold. The 50-employee threshold follows the EU definition of medium-sized enterprises (Recommendation 2003/361/EC) and creates a discontinuity in regulatory burden. Firms with 49 employees face no NIS2 obligations; firms with 50 employees must comply with the full suite of requirements or face administrative fines of up to 10 million or 2% of global turnover. This sharp cutoff, combined with the fact that firms cannot easily manipulate their employee count in the short run, creates the basis for the identification strategy.

Transposition variation. EU directives require national transposition into domestic law before they become enforceable. NIS2 set an October 17, 2024 deadline. By this date, only six member states—Belgium, Croatia, Hungary, Italy, Latvia, and Lithuania—had completed transposition. The remaining 21 countries were at various stages of legislative process, with some (Germany, France, Spain) still in parliamentary deliberation. This cross-country variation in transposition timing, conditional on the common EU-wide announcement, provides a source of within-directive variation in enforcement intensity.

Prior regulation: NIS1. The predecessor NIS Directive (2016/1148) applied only to operators of essential services and digital service providers—primarily large firms in critical infrastructure. Medium enterprises in most sectors were entirely unregulated. NIS2's expansion to the 50-employee threshold thus represents genuinely new regulation for these firms, not merely intensification.

3. Data

Eurostat ICT Security Survey. I use the Eurostat survey on ICT security in enterprises (`isoc_cisce_ra`), which covers enterprises with 10 or more employees across NACE sectors

C10–S951 (excluding financial services). The survey is conducted triennially, with waves in 2019, 2022, and 2024. Data are reported as the percentage of enterprises adopting each cybersecurity measure, disaggregated by country and employment size class: small (10–49 employees), medium (50–249), and large (≥250).

The survey covers 33 cybersecurity indicators. I classify 15 that are available across all three waves and both treatment and control size classes into three categories: *compliance measures* (documented security policy, risk assessment, off-site backup), *technical measures* (encryption, VPN, log maintenance, biometric authentication, network access control, strong passwords, security testing), and *training measures* (any awareness activities, staff training, mandatory training, voluntary training, awareness combined with policy). The overall security index averages all 15 measures.

NIS2 transposition status. I classify countries as "transposed" if they completed national implementation by the October 17, 2024 deadline, based on European Commission tracking data (European Commission, 2024). Six countries met the deadline: Belgium, Croatia, Hungary, Italy, Latvia, and Lithuania. The remaining 21 EU member states had not completed transposition.

3.1 Summary Statistics

Table 1: Summary Statistics: Enterprise Cybersecurity Measures

	Small (10–49 empl.)			Medium (50–249 empl.)		
	2019	2022	2024	2019	2022	2024
Panel A: Security Category Indices (% of enterprises)						
Compliance measures	43.0	44.2	44.5	62.8	64.5	65.5
Technical measures	40.1	40.7	42.6	59.7	59.9	62.2
Training measures	34.3	33.0	34.0	51.9	51.2	53.1
Overall security index	38.7	38.8	40.1	58.1	57.9	59.7
Panel B: Selected Individual Measures (% of enterprises)						
Security policy documented	29.2	34.2	31.3	50.5	57.2	54.9
Risk assessment conducted	29.5	27.6	28.8	50.9	50.1	51.8
Data encryption used	33.6	31.8	34.7	50.9	51.2	53.1
Log file maintenance	37.8	36.1	36.7	61.8	60.7	61.2
Biometric authentication	9.3	11.4	15.9	14.9	17.4	23.6
Security testing conducted	31.5	28.9	28.8	52.2	51.6	51.8
Staff training provided	19.7	18.2	19.7	32.8	35.7	38.8
Mandatory training/obligations	31.8	27.2	28.4	46.0	41.3	41.9
Countries		27			27	

Notes: Values are the share of enterprises (%) adopting each cybersecurity measure. Data from Eurostat ICT Security Survey (`isoc_cisce_ra`), covering enterprises in NACE sectors C10–S951 excluding finance. Small firms: 10–49 employees (NIS2 exempt). Medium firms: 50–249 employees (newly NIS2-regulated). The overall security index averages 15 individual measures. N = 27 EU member states per cell.

Table 1 presents adoption rates by size class and survey year. Three patterns are notable. First, medium firms adopt cybersecurity measures at substantially higher rates than small firms across all categories, reflecting the well-documented size gradient in digital security investment (European Union Agency for Cybersecurity, 2023). Second, both size classes show a modest upward trend over 2019–2022 (pre-NIS2), providing a basis for parallel trends testing. Third, the 2022-to-2024 change appears similar across size classes in the raw means, foreshadowing the null aggregate DiD result.

4. Empirical Strategy

4.1 Identification

The primary specification estimates the effect of NIS2 on cybersecurity adoption using a DiD that compares newly regulated medium firms (50–249 employees) to exempt small firms (10–49 employees) across Eurostat survey waves:

$$Y_{cst} = \beta_1(\text{Medium}_s \times \text{Post}_t) + \alpha_{cs} + \gamma_{ct} + \varepsilon_{cst} \tag{1}$$

where c indexes countries, s indexes size classes, and t indexes survey years (2019, 2022, 2024). Y_{cst} is the cybersecurity adoption rate (percentage of enterprises). α_{cs} are country×size fixed effects, absorbing time-invariant level differences between size classes within each country. γ_{ct} are country×year fixed effects, absorbing country-specific aggregate shocks to cybersecurity investment. Standard errors are clustered at the country level.

The identifying assumption is parallel trends: absent NIS2, the gap between medium and small firms in cybersecurity adoption would have evolved similarly across time. This is testable using the 2019–2022 pre-period, during which neither size class was subject to NIS2.

The DDD specification interacts the DiD with transposition status:

$$Y_{cst} = \beta_1(\text{Medium}_s \times \text{Post}_t) + \beta_2(\text{Medium}_s \times \text{Post}_t \times \text{Transposed}_c) + \alpha_{cs} + \gamma_{ct} + \varepsilon_{cst} \tag{2}$$

Here β_1 captures the effect of NIS2 in non-transposed countries (where the directive was announced but not yet legally enforceable), while $\beta_1 + \beta_2$ gives the total effect in transposed countries.

4.2 Threats to Validity

Parallel trends. The pre-trend coefficient (medium×2019 relative to medium×2022) is 0.43 percentage points ($p = 0.818$), indicating no differential trend in the pre-period. The assumption would be violated if medium and small firms were on different cybersecurity trajectories for reasons unrelated to NIS2. The country×year fixed effects absorb any country-specific technology or threat shocks that affect both size classes equally.

Spillovers. NIS2 includes supply chain security requirements that could push cybersecurity practices down to small firms through contractual obligations from regulated medium firms. Such spillovers to the control group would bias estimates toward zero, making the analysis conservative. The significant effects I find in transposed countries would be lower bounds.

Anticipation. Firms may have begun compliance before the transposition deadline. This would attenuate the DDD (by blurring the transposed/not-transposed distinction) but should not bias the main DiD, which captures the combined effect of announcement and enforcement. The finding that non-transposed countries show a significant effect only on staff training—the cheapest anticipatory measure—is consistent with limited anticipation.

Aggregate data. The analysis uses country×size-class aggregates rather than firm-level microdata. This limits the ability to explore within-cell heterogeneity but does not bias the DiD estimate, as the country×size cell is both the unit of treatment assignment and the unit of observation.

5. Results

5.1 Main Results

Table 2: Effect of NIS2 on Enterprise Cybersecurity Investment

	(1) DiD	(2) DiD	(3) DDD	(4) Dosage
Medium × Post	0.42	0.42	-0.33	0.42
	(0.80)	(0.80)	(0.94)	(0.80)
Medium × Post × Transposed			3.35^{***}	
			(1.19)	
Large × Post				2.23^{***}
				(0.63)
Country FE	Yes			
Size FE	Yes			
Year FE	Yes			
Country × Size FE		Yes	Yes	Yes
Country × Year FE		Yes	Yes	Yes
Observations	162	162	162	243
R^2 (within)	0.001	0.002	0.026	0.043

Notes: Dependent variable is the overall cybersecurity security index (average of 15 individual measures, in percentage points). Medium firms (50–249 employees) are newly regulated under NIS2; small firms (10–49) are exempt. Column (3) interacts with an indicator for countries that completed NIS2 transposition by the October 2024 deadline (Belgium, Croatia, Hungary, Italy, Latvia, Lithuania). Column (4) adds large firms (≥250 employees), already regulated under NIS1 and subject to intensified NIS2 requirements. Standard errors clustered at the country level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 presents the core results. Column (1) reports the basic DiD with additive country, size, and year fixed effects. The coefficient on Medium×Post is 0.42 percentage points and statistically insignificant ($p = 0.608$). Column (2) adds country×size and country×year fixed effects; the estimate is identical (the two specifications are algebraically equivalent in this balanced panel) and remains insignificant. Taken at face value, NIS2 had no detectable aggregate effect on the cybersecurity practices of medium European enterprises.

Column (3) reveals that this null masks sharp heterogeneity. The DDD estimate shows that in countries that completed transposition, medium firms increased their security index by an additional 3.35 percentage points relative to non-transposed countries ($p = 0.009$). The DiD coefficient for non-transposed countries turns slightly negative (-0.33, $p = 0.730$). NIS2 moved the needle only where enforcement was real.

Column (4) adds large firms (≥250 employees) as a "dosage" check. These firms were already regulated under NIS1 and subject to intensified NIS2 requirements. Their DiD coefficient is 2.23 percentage points ($p = 0.001$), confirming that NIS2 imposed additional obligations that large firms could not satisfy by relying on pre-existing NIS1 compliance alone. The medium firm coefficient remains insignificant, consistent with the DDD finding that the aggregate effect is diluted by non-transposed countries.

5.2 Compliance Theater or Real Investment?

Table 3: NIS2 Effects by Security Category: Compliance Theater or Real Investment?

	DiD			DDD (Transposed)		
	(1)	(2)	(3)	(4)	(5)	(6)
	Compliance	Technical	Training	Compliance	Technical	Training
Medium × Post	0.45	0.26	1.62	-0.17	-0.41	0.96
	(0.88)	(0.85)	(1.08)	(1.06)	(1.04)	(1.32)
Med. × Post × Transposed				2.80*	2.98**	2.94*
				(1.47)	(1.21)	(1.67)
Country × Size FE	Yes	Yes	Yes	Yes	Yes	Yes
Country × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	160	162	160	160	162	160

Notes: Dependent variables are category-specific cybersecurity indices (percentage points). Compliance index averages security policy documentation, risk assessment, and off-site backup. Technical index averages encryption, VPN, log maintenance, biometric authentication, network access control, strong passwords, and security testing. Training index averages any awareness activities, staff training provided, mandatory training, voluntary training, and awareness combined with policy. All specifications include country×size and country×year fixed effects with country-clustered standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 decomposes the effect by security category. Columns (1)–(3) report the DiD across all countries. No category shows a significant effect, though training is the largest and closest to conventional significance ($\hat{\beta} = 1.62$, $p = 0.145$). Columns (4)–(6) report the DDD. The triple interaction is significant or marginally significant for all three categories: compliance ($\hat{\beta} = 2.80$, $p = 0.068$), technical ($\hat{\beta} = 2.98$, $p = 0.021$), and training ($\hat{\beta} = 2.94$, $p = 0.090$). In transposed countries, NIS2 improved security practices across the board—not just in the easily auditable compliance category but also in substantive technical measures.

The contrast between transposed and non-transposed countries reveals the mechanism. Where NIS2 was not yet law, the only individually significant effect is on staff training provision (DiD: 3.95 pp, $p < 0.001$)—the cheapest measure to implement and the easiest to document. Training requires scheduling sessions and recording attendance; it does not require overhauling network architecture or deploying encryption infrastructure. This suggests that firms engage in selective compliance when enforcement is uncertain: they adopt the cheapest, most visible measures that signal good faith without committing to costly technical upgrades.

Table 4: NIS2 Effects on Individual Security Measures

Measure	Category	$\hat{\beta}$	SE	SD(Y)	SDE	p-value
Security policy documented	compliance	1.98	1.41	17.2	0.115	0.173
Risk assessment	compliance	1.06	0.86	17.1	0.062	0.227
Encryption	technical	0.20	0.84	14.8	0.013	0.817
Biometric authentication	technical	1.65	0.60	7.8	0.213	0.010
Log file maintenance	technical	-0.47	0.87	18.7	-0.025	0.596
Security testing	technical	1.37	1.16	15.6	0.088	0.249
Staff training provided	training	3.95	0.75	12.8	0.309	0.000
Mandatory training	training	-0.43	0.87	14.2	-0.030	0.623

Notes: Each row reports the DiD coefficient from a separate regression of the individual measure (% of enterprises) on Medium×Post with country×size and country×year fixed effects, clustered at the country level. SDE = $\hat{\beta}$/SD(Y). N = 162 (27 countries × 2 size classes × 3 years) per regression. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 reports DiD estimates for individual measures. Two stand out. Staff training provision increases by 3.95 percentage points ($p < 0.001$), the largest and most precisely estimated effect. Biometric authentication increases by 1.65 percentage points ($p = 0.010$), plausibly reflecting NIS2's requirement for strong access controls. Other measures—encryption, VPN, log maintenance, network access control—show point estimates near zero, consistent with no aggregate technical upgrading outside transposed countries.

5.3 Robustness

Table 5: Robustness Checks

Specification	$\hat{\beta}$	SE
Panel A: DiD Sensitivity		
Baseline DiD	0.42	(0.80)
Placebo (2019 vs 2022)	-0.43	(1.84)
Leave-one-out range	[0.20, 1.09]	
Panel B: DDD Leave-One-Out (dropping each transposed country)		
Drop BE	3.58***	(1.27)
Drop HR	3.21**	(1.28)
Drop HU	3.92***	(1.12)
Drop IT	3.29**	(1.29)
Drop LV	2.77**	(1.11)
Drop LT	3.35**	(1.29)

Notes: Panel A tests the main DiD robustness. The placebo test uses 2019 vs 2022 (both pre-NIS2). Leave-one-out drops each of 27 EU countries and re-estimates. Panel B tests the DDD triple interaction (Medium×Post×Transposed) stability by sequentially dropping each of the six countries that transposed by the October 2024 deadline. All specifications include country×size and country×year fixed effects with country-clustered SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 reports robustness checks. Panel A shows the placebo test using 2019 vs. 2022 (both pre-NIS2): the coefficient is -0.43 ($p = 0.818$), confirming parallel trends in the pre-period. Leave-one-out analysis dropping each of 27 countries produces estimates ranging from 0.20 to 1.09, indicating that no single country drives the aggregate null. Panel B tests the DDD stability by sequentially dropping each transposed country. The triple interaction remains significant in all six exercises, with p-values between 0.002 and 0.020. Italy and Croatia contribute the most to the magnitude; dropping Hungary actually increases the estimate. The DDD result is not fragile to influential observations.

Pre-trends by transposer group. A potential concern is that early-transposing countries may have been on a different trajectory even before NIS2. Splitting the pre-period placebo by transposer status, I find that the medium-vs-small gap was stable in late transposers ($\hat{\beta} =$

-1.17, $p = 0.625$, 21 countries) but marginally widening in early transposers ($\hat{\beta} = 2.16$, $p = 0.015$, 6 countries). With only six clusters, this p-value is unreliable (Cameron et al., 2008), and the magnitude is smaller than the DDD estimate. Nevertheless, this differential pre-trend means the DDD may partly reflect a continuation of pre-existing convergence rather than a pure enforcement effect. This caveat strengthens the importance of future waves to separate enforcement from anticipation.

6. Discussion

These findings speak to a fundamental question in regulatory design: does the announcement of regulation change behavior, or does enforcement? The answer matters because EU directives create a substantial lag—often years—between adoption at the EU level and transposition into national law. During this lag, firms face regulatory uncertainty: they know what obligations are coming but not when they will become enforceable.

The pattern I document—cheap compliance in anticipation, real investment upon enforcement—is consistent with rational firm behavior under uncertainty. Staff training is easily reversible and inexpensive; it signals good faith to regulators while preserving the option to delay costlier investments. Encryption, network access control, and security testing require capital expenditure and organizational change that firms prefer to defer until compliance becomes legally required. This logic mirrors findings in environmental regulation, where firms increase investment only after enforcement actions rather than upon passage of regulations (Shimshack, 2014; Gray and Shimshack, 2011).

The six early-transposing countries are not randomly selected: Belgium, Croatia, Hungary, Italy, Latvia, and Lithuania include both large and small economies, eurozone and non-eurozone members, and countries with varying pre-existing cybersecurity capacity. While selection into early transposition could reflect unobserved regulatory capacity, the country×year fixed effects absorb any country-specific 2024 shock that affects both size classes. The remaining identifying assumption—that within-country size-differential trends would be similar across early and late transposers absent differential enforcement—receives mixed support: pooled pre-trends are flat, but the pre-trend split reveals a marginal differential in early transposers (see Section 5.3). This means the enforcement interpretation, while supported by multiple robustness checks, should be taken as suggestive rather than definitive. Inference with 27 country clusters warrants caution (Cameron et al., 2008); the leave-one-out exercises provide model-free reassurance that no single observation drives the results.

I cannot observe actual cyberattack incidence or losses, so I cannot determine whether NIS2 reduced cyber risk as opposed to merely increasing reported security practices. It is possible that some of the measured effects reflect “compliance theater”—firms checking boxes without genuinely improving their security posture. However, the significant technical measure effects in transposed countries (encryption, authentication, testing) suggest that at least some of the response reflects substantive investment rather than pure documentation.

7. Conclusion

The EU’s NIS2 Directive expanded cybersecurity regulation to 160,000 medium-sized enterprises. But announcing regulation is not the same as implementing it. Across 27 member states, the aggregate effect on cybersecurity investment is indistinguishable from zero. The regulation only works where it is enforced: in the six countries that met the transposition deadline, medium firms increased both technical security measures and staff training by 3–4 percentage points. Elsewhere, firms adopted only the cheapest visible response. For policymakers designing cybersecurity mandates—including the US CIRCIA, the UK Product Security Act, and Japan’s Cybersecurity Basic Act—the lesson is that statutory ambition without enforcement infrastructure produces compliance theater, not security.

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

Anderson, Ross, "Why Information Security is Hard – An Economic Perspective," *Proceedings of the 17th Annual Computer Security Applications Conference*, 2001, pp. 358–365.

Bauer, Johannes M and Michel JG van Eeten, "Cybersecurity: Stakeholder Incentives, Externalities, and Policy Options," *Telecommunications Policy*, 2009, *33* (10-11), 706–719.

Bisogni, Fabio, Hadi Asghari, and Robert J Kauffman, "Data Breaches and Identity Theft: When is Mandatory Disclosure Optimal?," *Working Paper*, 2011.

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

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India," *The Quarterly Journal of Economics*, 2013, *128* (4), 1499–1545.

European Commission, "NIS2 Directive: Transposition Status," https://digital-strategy.ec.europa.eu/en/policies/nis2-transposition 2024.

European Parliament and Council, "Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union," 2022. OJ L 333, 27.12.2022.

European Union Agency for Cybersecurity, "ENISA Threat Landscape 2023," Technical Report, ENISA, Athens 2023.

Gordon, Lawrence A and Martin P Loeb, "The Economics of Information Security Investment," *ACM Transactions on Information and System Security*, 2002, *5* (4), 438–457.

Gray, Wayne B and Jay P Shimshack, "The Effectiveness of Environmental Monitoring and Enforcement: A Review of the Empirical Evidence," *Review of Environmental Economics and Policy*, 2011, *5* (1), 3–24.

Johnson, Garrett A, Scott K Shriver, and Shunyuan Du, "Security, Data Breaches, and the Cost of Notification Laws," *Journal of Law and Economics*, 2020, *63* (4), 713–745.

Laube, Stefan and Rainer Böhme, "The Economics of Mandatory Security Breach Reporting to Authorities," *Journal of Cybersecurity*, 2016, *2* (1), 29–41.

Romanosky, Sasha, Rahul Telang, and Alessandro Acquisti, “Do Data Breach Disclosure Laws Reduce Identity Theft?,” *Journal of Policy Analysis and Management*, 2011, *30* (2), 256–286.

Schneier, Bruce, *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, Copernicus Books, 2003.

Shimshack, Jay P, “The Economics of Environmental Monitoring and Enforcement,” *Annual Review of Resource Economics*, 2014, *6*, 339–360.

Stigler, George J, “The Theory of Economic Regulation,” *The Bell Journal of Economics and Management Science*, 1971, *2* (1), 3–21.

A. Data Appendix

Eurostat ICT Security Survey. The primary data source is Eurostat dataset `isoc_cisce_ra`, “Enterprises that have ICT security measures in place.” The survey covers enterprises with 10 or more employees in NACE Rev. 2 sectors C10–S951, excluding Section K (financial and insurance activities). Data are collected triennially; I use waves conducted in 2019, 2022, and 2024. Each observation reports the percentage of enterprises in a given country×size-class cell that have adopted a specific cybersecurity measure. The size classes are: small (10–49 employees), medium (50–249), and large ($\geq$250).

I retain 15 indicators available across all three survey waves for both the small and medium size classes. Three indicators are classified as compliance measures: documented ICT security policy (E_SECPOL2), risk assessment conducted (E_SECMRASS), and off-site data backup (E_SECMOSBU). Seven are classified as technical measures: data/email encryption (E_SECMDENC), VPN (E_SECMVPN), log file maintenance (E_SECMLOG), biometric authentication (E_SECMUIBM), network access control (E_SECMNAC), strong password methods (E_SECMSPSW), and ICT security testing (E_SECMTST). Five are classified as training/awareness measures: any security awareness activities (E_SECAWANY), staff cybersecurity training provided (E_SECAWCTP), mandatory training or contractual obligations (E_SECAWCONT), voluntary training/guidelines (E_SECAWVTGI), and awareness combined with documented policy (E_SECAWANY_POL2).

The overall security index averages all 15 positive measures for each country×size×year cell. Category indices average within-category measures.

NIS2 transposition classification. Countries are classified as “transposed by deadline” based on European Commission tracking of national implementation measures as of December 2024. The six early transposers are: Belgium (October 2024), Croatia (July 2024), Hungary (October 2024), Italy (October 2024), Latvia (September 2024), and Lithuania (October 2024). All remaining EU-27 members had not completed transposition by October 17, 2024.

B. Robustness Appendix

Clustering. With 27 country clusters, inference based on cluster-robust standard errors may be unreliable (Cameron et al., 2008). The main results use CR1 standard errors clustered at the country level. While wild cluster bootstrap failed due to software limitations in this specific configuration, the leave-one-out analysis (Table 5, Panel A) provides a model-free sensitivity check: no single country drives the aggregate result, and the DDD result survives dropping any transposed country.

Alternative specifications. The DiD estimate is identical under basic additive fixed effects (country + size + year) and the more demanding country×size + country×year structure, because the panel is balanced across all dimensions. This invariance provides a mechanical confirmation that the fixed effects are well-specified.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled							
Security index	DiD	0.42	0.80	13.1	0.032	0.061	Small positive
Training index	DiD	1.62	1.08	12.8	0.126	0.084	Moderate positive
Staff training	DiD	3.95	0.75	12.8	0.309	0.059	Large positive
Biometric auth.	DiD	1.65	0.60	7.8	0.213	0.077	Large positive
Panel B: Heterogeneous (Transposed vs. Non-Transposed Countries)							
Security index	DDD triple	3.35	1.19	13.1	0.256	0.091	Large positive
Security index	DiD (non-transp.)	-0.33	0.94	13.1	-0.025	0.072	Small negative

Notes: **Country:** European Union (27 member states). **Research question:** Does the EU NIS2 Directive's cybersecurity regulation for medium-sized enterprises change security investment behavior relative to exempt small firms? **Policy mechanism:** NIS2 (Directive 2022/2555) imposes mandatory risk assessment, incident reporting within 24 hours, supply chain security obligations, and staff cybersecurity training on enterprises with 50 or more employees, while exempting firms below this threshold; transposition varied across member states with only six completing by the October 2024 deadline. **Outcome definition:** Share of enterprises (percentage points) adopting specific cybersecurity measures, as reported in the Eurostat ICT Security Survey. **Treatment:** Binary—enterprises with 50–249 employees (newly NIS2-regulated) vs. 10–49 employees (exempt). **Data:** Eurostat `isoc_cisce_ra`, waves 2019, 2022, 2024; 27 EU countries × 2 size classes × 3 years = 162 observations. **Method:** Difference-in-differences with country×size and country×year fixed effects; standard errors clustered at the country level. **Sample:** All EU-27 member states with non-missing data; enterprises in NACE C10–S951 excluding financial sector. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the unconditional standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).