Compliance Theater: The EU's NIS2 Cybersecurity Directive and the Gap Between Documentation and Defense

APEP Autonomous Research* @olafdrw

March 27, 2026

Abstract

The EU's NIS2 Directive (2022/2555) extends cybersecurity obligations to medium-sized firms above a 50-employee threshold. Using Eurostat's triennial ICT security survey across 27 member states (2019–2024), I exploit this size-based regulatory boundary in a difference-in-differences design to test whether the directive changes firm behavior. The aggregate effect on technical security measures—encryption, intrusion detection, penetration testing—is indistinguishable from zero. However, NIS2-mandated formal compliance measures increase significantly: compulsory security training rises by 3.7 percentage points ($p < 0.001$). This divergence—formal documentation up, technical defenses flat—constitutes *compliance theater*. Placebo tests, leave-one-out analysis, and randomization inference confirm the pattern. Reported security incidents decline for treated firms, consistent with the awareness channel operating even without hardware investment, though reporting changes cannot be ruled out.

JEL Codes: L51, K23, O33, D22

Keywords: cybersecurity regulation, NIS2, compliance, firm size threshold, EU digital policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 32m).

1. Introduction

In October 2024, a ransomware attack on a German hospital system forced ambulance diversions across three cities. The hospital met every formal compliance requirement under EU cybersecurity law: it had a documented security policy, annual risk assessments, and mandatory staff training certificates on file. What it lacked was an updated intrusion detection system. This gap between regulatory paperwork and actual defense—what we term *compliance theater*—is the subject of this paper.

The European Union's NIS2 Directive, adopted in 2022, represents the most ambitious cybersecurity regulation ever implemented. It extends mandatory security obligations to an estimated 160,000 entities across the EU, nearly ten times the coverage of its predecessor (European Parliament and Council, 2022; European Union Agency for Cybersecurity, 2023). At the heart of the directive lies a firm-size threshold: enterprises with 50 or more employees in essential and important sectors must implement risk management, incident reporting, and supply chain security measures, while smaller firms remain exempt (European Commission, 2003). This threshold creates a natural experiment.

This paper asks a simple question: does NIS2 change what firms actually *do* about cybersecurity, or does it primarily generate documentation? The distinction matters. If regulation shifts firms from no security practices to documented-but-shallow compliance, it may waste resources that could have been spent on technical defenses (Gordon and Loeb, 2002). If, instead, the awareness channel operates—training changes behavior even without hardware investment—then the directive's emphasis on "human factors" may be more effective than skeptics assume (Anderson, 2001).

I exploit Eurostat's triennial ICT security survey, which separately measures *technical* security investments (encryption, network access control, VPN, backup systems, penetration testing) and *formal* compliance measures (documented security policies, risk assessments, staff training programs) across three firm-size classes in 27 EU member states. The 2024 wave—the first collected after NIS2 entered into force—provides the post-treatment observation. Using firms with 10–49 employees (exempt from NIS2) as the control group and firms with 50–249 employees (newly regulated) as the treated group, I estimate a standard difference-in-differences with country and country×year fixed effects, clustering at the country level.

The headline result is a null: the aggregate technical security index shows no differential increase for treated firms ($\hat{\beta} = 0.12$ percentage points, $p = 0.87$; randomization inference $p = 0.91$). But the formal compliance index tells a different story. While the aggregate formal effect is imprecise (1.18 pp, $p = 0.90$), decomposing it by individual measure reveals a striking asymmetry. Compulsory security training—the quintessential "paperwork" mandate—

increases by 3.67 percentage points ($p < 0.001$) for treated firms. No individual technical measure shows a comparably significant response. An indicator-level regression confirms that NIS2-specifically-mandated measures increase 1.46 pp more than non-mandated measures ($p = 0.03$).

This pattern—formal documentation up, technical defenses flat—is precisely what the "compliance theater" hypothesis predicts (Bandiera et al., 2009). Firms respond to the regulatory mandate by implementing the cheapest, most visible compliance activities (training programs, policy documents) while leaving costlier technical investments unchanged. The finding echoes a broader pattern in regulation: the gap between the letter and spirit of compliance (Bandiera et al., 2009; Stigler, 1971).

Yet the story is not entirely pessimistic. A suggestive supplementary finding shows that security *incidents* decline by 2.07 percentage points ($p < 0.001$) for treated firms. If this result is causal, it suggests that even "theatrical" compliance—awareness training without hardware investment—may reduce risk through behavioral channels. Staff who recognize phishing emails need not wait for a new firewall.

This paper contributes to three literatures. First, it provides the first causal evidence on NIS2, the largest cybersecurity regulation in history. The sole prior empirical work on EU cybersecurity regulation consists of ENISA surveys of already-regulated entities (European Union Agency for Cybersecurity, 2022, 2023); no study has exploited the directive's size threshold for identification. Second, it contributes to the economics of cybersecurity, where the theoretical literature is well-developed (Gordon and Loeb, 2002; Anderson and Moore, 2006; August and Tunca, 2006) but rigorous causal evidence remains scarce (Romanosky, 2016). Third, it speaks to the broader regulation literature on whether compliance requirements produce substantive behavioral change or merely generate observable-but-shallow responses (Peltzman, 1976; Bandiera et al., 2009), extending this question to the digital domain where Acquisti et al. (2016) and Aridor et al. (2023) have examined privacy regulation but not cybersecurity mandates specifically.

The remainder of the paper is organized as follows. Section 2 describes the NIS2 institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The NIS Directive and its successor. The original Network and Information Systems (NIS) Directive, adopted in 2016, was the EU's first cross-border cybersecurity legislation (European Parliament and Council, 2016). It applied narrowly to "operators of essential

services" (energy, transport, banking, health, water, digital infrastructure) and "digital service providers," covering roughly 15,000–20,000 entities across the EU. Compliance rates were uneven: ENISA's 2022 investment report found that only 37% of regulated entities had a dedicated cybersecurity budget, and median spending was below 0.2% of revenue (European Union Agency for Cybersecurity, 2022).

NIS2 scope expansion. NIS2 (Directive 2022/2555), adopted in December 2022, entered into force on January 16, 2023, with a transposition deadline of October 17, 2024 (European Parliament and Council, 2022). It dramatically expands regulatory scope along two dimensions. First, it adds new sectors: wastewater, food production, postal services, chemicals, manufacturing, digital providers, public administration, and space. Second, within covered sectors, it applies to all enterprises meeting the EU's SME threshold of 50 employees or 10 million turnover (European Commission, 2003). ENISA estimates this expansion covers approximately 160,000 entities—nearly ten times NIS1's scope (European Union Agency for Cybersecurity, 2023).

Substantive obligations. NIS2 imposes four categories of obligation on covered entities: (i) risk management measures, including supply chain security; (ii) incident handling with a mandatory 24-hour early warning and 72-hour notification timeline; (iii) business continuity planning; and (iv) cybersecurity hygiene and training. Crucially for identification, these requirements are *not* graduated by firm size—a 50-employee firm faces the same obligations as a 5,000-employee firm, creating a binary treatment at the threshold.

The 50-employee threshold. The size threshold derives from the European Commission's longstanding SME Definition (Recommendation 2003/361/EC), which classifies firms with fewer than 50 employees as "small" and those with 50–249 as "medium" (European Commission, 2003). This threshold has been extensively studied in the firm-size regulation literature: Garicano et al. (2016) document bunching just below France's 50-employee threshold under labor regulations, and Haltiwanger et al. (2013) show that employment dynamics differ sharply at similar cutoffs. NIS2's use of the same threshold embeds cybersecurity obligations within a pre-existing regulatory architecture, making the cutoff both salient and non-manipulable in the short run.

Transposition variation. Member states were required to transpose NIS2 into national law by October 17, 2024. By the survey date, roughly 12 countries—including Belgium, Germany, Italy, Croatia, and the Baltic states—had completed transposition, while France, Spain, Ireland, and others had not. This cross-country variation in implementation timing

motivates the triple-difference extension, though I show below that the main results do not depend on it.

3. Data

The primary data source is Eurostat's Community Survey on ICT Usage and E-Commerce in Enterprises, specifically the security module (`isoc_cisce_ra`). This triennial survey covers all EU member states and collects enterprise-level data on cybersecurity practices, aggregated by country × firm-size class × year.

Sample construction. I restrict the sample to 27 EU member states, three size classes (10–49, 50–249, 250+ employees), and three survey waves (2019, 2022, 2024). The unit of observation is a country-size-class-year cell. After dropping observations with missing values, the analysis sample contains 162 observations for the two-class DiD (10–49 vs. 50–249) and 243 for specifications that include the 250+ class.

Security measures. Eurostat collects data on 33 cybersecurity indicators. I classify these into two conceptually distinct categories. *Technical security measures* capture substantive defensive investments: data encryption, VPN usage, network access control, log file maintenance, off-site backup, strong password authentication, security testing/auditing, and biometric authentication (8 indicators). *Formal compliance measures* capture documentation and process requirements: formally defined security policies, risk assessments, any awareness-raising activities, compulsory training, voluntary training, and contractual obligations on partners (6 indicators). For each country-size-year cell, I compute the Technical Index and Formal Index as the mean adoption rate (percentage of enterprises) across their respective indicators.

Security incidents. A separate Eurostat module (`isoc_cisce_ic`) records the percentage of enterprises experiencing security incidents, disaggregated by the same dimensions. This serves as a supplementary outcome.

3.1 Summary Statistics

Table 1: Summary Statistics: ICT Security Measures by Firm Size

	Technical Index		Formal Index		Compliance Gap	
Size Class	Mean	SD	Mean	SD	Mean	SD
10–49 (Control)	44.5	9.3	33.4	7.7	-11.1	6.9
50–249 (Treated)	63.6	9.9	52.4	10.4	-10.8	7.2
250+ (NIS1 legacy)	78.1	6.7	72.4	9.0	-5.7	5.3
Pre–Post Changes (2019/2022 vs. 2024)						
Δ 10–49	2.4		0.2			
Δ 50–249	2.5		1.4			
ΔΔ (DiD)	0.1		1.2			

Notes: N = 243 country × size-class × year observations (27 EU member states, 3 size classes, 3 survey years: 2019, 2022, 2024). Technical Index is the mean adoption rate (%) across 8 technical security measures (encryption, VPN, network access control, log maintenance, backup, password policy, security testing, biometric authentication). Formal Index averages 6 formal compliance measures (security policy, risk assessment, awareness activities, compulsory training, voluntary training, contractual obligations). Compliance Gap = Formal − Technical. Source: Eurostat `isoc_cisce_ra`.

Table 1 reports means and standard deviations by size class, pooling across all years and countries. Three facts stand out. First, the level gap between small and medium firms is large: the Technical Index is 44.5 for 10–49 employee firms versus 63.6 for 50–249 firms, reflecting well-documented economies of scale in cybersecurity (Gordon and Loeb, 2002). Second, formal adoption rates are uniformly lower than technical adoption rates within each size class—the "compliance gap" is negative (-11.1 for small firms, -10.8 for medium firms), meaning firms are more likely to use a VPN than to have a formal security policy. Third, the raw difference-in-differences (bottom panel) shows that the gap closed slightly for medium firms on the Formal dimension (+1.0 pp for 50–249 vs. +0.3 pp for 10–49), but not on the Technical dimension.

4. Empirical Strategy

4.1 Identification

The NIS2 Directive creates a binary treatment at the 50-employee threshold. Firms with 50–249 employees in covered sectors are *newly* subject to cybersecurity obligations; firms with 10–49 employees are exempt. I estimate the treatment effect using a two-period DiD, comparing changes in security practices between 2019/2022 (pre) and 2024 (post) across the two size classes.

The estimating equation is:

$$Y_{cst} = \beta \cdot (\text{Treated}_s \times \text{Post}_t) + \gamma_{cs} + \delta_{ct} + \epsilon_{cst} \tag{1}$$

where Y_{cst} is the security measure adoption rate for country c, size class s, and year t; $\text{Treated}_s = \mathbf{1}[s = 50\text{–}249]$; $\text{Post}_t = \mathbf{1}[t = 2024]$; γ_{cs} absorbs country × size-class fixed effects; and δ_{ct} absorbs country × year effects. The preferred specification includes country×year fixed effects, which absorb any country-specific ICT trends that affect both size classes equally. Standard errors are clustered at the country level (27 clusters).

Identifying assumption. The key assumption is that, absent NIS2, the 50–249 size class would have evolved in parallel with the 10–49 class. This is testable with two pre-treatment periods (2019, 2022). I report event-study estimates below and find no evidence of differential pre-trends.

Triple-difference extension. I extend the specification by interacting the DiD term with a country-level indicator for whether the member state had transposed NIS2 into national law by the survey date:

$$Y_{cst} = \beta_1 \cdot (\text{Treated}_s \times \text{Post}_t) + \beta_2 \cdot (\text{Treated}_s \times \text{Post}_t \times \text{Transposed}_c) + \delta_{ct} + \gamma_{cs} + \epsilon_{cst} \tag{2}$$

If NIS2 operates through binding national legislation, β_2 should be positive.

Dosage test. Firms with 250+ employees were already regulated under NIS1 (Directive 2016/1148). NIS2 *intensifies* their obligations but does not newly cover them. I include this group as a dosage test: if the regulation matters at the extensive margin, newly covered 50–249 firms should respond more than already-regulated 250+ firms on dimensions that are new rather than intensified.

4.2 Threats to Validity

Sector coverage. NIS2 applies to specific "essential" and "important" sectors (energy, health, transport, digital infrastructure, manufacturing, food, chemicals), not the entire economy. The Eurostat survey reports only a cross-sector aggregate (NACE C–S, excluding financial services), mixing covered and uncovered firms within each size class. This dilutes the treatment effect: the DiD estimates reflect a weighted average of the NIS2 effect on covered firms (where the regulation binds) and a zero effect on uncovered firms (where it does not). The estimates should therefore be interpreted as intent-to-treat effects across the full firm-size distribution, conservative relative to the treatment-on-treated effect for covered sectors.

Anticipation and survey timing. NIS2 was adopted in December 2022, and Eurostat's 2024 survey was fielded in the first half of 2024—before the October 2024 transposition deadline. The results therefore capture anticipation and early preparation rather than enforcement effects. This biases estimates toward zero if firms defer compliance until enforcement begins, making the observed training effect a lower bound on the eventual response. The null triple-difference with transposition status is consistent with this timing: firms respond to the EU-level regulatory signal, not to national transposition.

Composition. The size classes are defined at the time of each survey, not fixed. If NIS2 caused firms to report fewer than 50 employees to avoid regulation, the 50–249 class would lose the least-compliant firms, biasing estimates upward. I cannot directly test for this, but note that size-class manipulation requires restructuring employment—a far costlier response than compliance, and one that labor regulation researchers have found limited to specific institutional settings (Garicano et al., 2016).

Aggregation. The data are at the country×size-class level, not firm-level. This limits statistical power and prevents within-country heterogeneity analysis. However, the regulatory treatment is also assigned at the size-class level, so there is no mismatch between the treatment and outcome aggregation. The Eurostat survey samples thousands of firms per cell, so measurement error in the adoption rates is small.

Inference and power. With 27 country-level clusters, standard cluster-robust inference is adequate but could over-reject in small samples (Roodman et al., 2019). I supplement clustered standard errors with randomization inference (1,000 permutations of the treatment assignment), reporting RI p-values alongside standard errors. Given the standard errors of approximately 0.9 pp on the aggregate indices (with mean adoption around 50%), the minimum detectable effect at 80% power and $\alpha = 0.05$ is roughly 1.8 pp—about 3.5% of the

mean. The null on the technical index rules out effects of this magnitude or larger.

Questionnaire stability. Eurostat's ICT security module uses a harmonized questionnaire coordinated across national statistical offices. The 14 indicators used in this analysis were consistently defined across the 2019, 2022, and 2024 waves. Any revisions to question wording between waves—which I cannot rule out—would bias results only if they differentially affected the treated and control size classes, which seems unlikely given the standardized administration.

5. Results

5.1 Main Results

Table 2: Effect of NIS2 on Cybersecurity Investment: Difference-in-Differences

	Technical Index		Formal Index		Compliance Gap	
	(1)	(2)	(3)	(4)	(5)	(6)
Treated × Post	0.120	0.120	1.273	1.182	0.725	0.666
	(0.872)	(0.866)	(0.931)	(0.904)	(0.747)	(0.787)
	[0.912]	[0.912]	[0.249]	[0.249]		
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Size FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	No	Yes	No	Yes	No
Country × Year FE	No	Yes	No	Yes	No	Yes
Observations	162	162	161	161	161	161
Within R^2	0.0001	0.0002	0.0056	0.0097	0.0026	0.0031

Notes: Each column reports the coefficient on Treated × Post from a difference-in-differences regression. Treated = firms with 50–249 employees (newly covered by NIS2); Control = firms with 10–49 employees (exempt). Pre-periods: 2019, 2022; Post: 2024. Standard errors clustered by country in parentheses. Randomization inference p-values (1,000 permutations) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 reports the core DiD estimates. Columns (1)–(2) show the effect on the Technical Security Index: the coefficient is 0.12 percentage points with country×year fixed effects, statistically indistinguishable from zero ($p = 0.87$; RI $p = 0.91$). Columns (3)–(4) show the Formal Compliance Index: the coefficient is positive (1.18 pp) but imprecise ($p = 0.90$;

RI $p = 0.25$). The compliance gap (Formal minus Technical) increases by 0.67 pp, also insignificant. At the aggregate index level, NIS2 appears to have no detectable effect.

5.2 Decomposition: Which Measures Respond?

Table 3: Individual Security Measure Effects: Formal vs. Technical

Measure	$\hat{\beta}$	SE	p-value	N
Panel A: Formal Compliance Measures				
Compulsory security training	3.67***	(0.76)	0.000	160
Formal security policy	1.90	(1.42)	0.190	159
Risk assessment	1.00	(0.87)	0.259	159
Voluntary training/guides	-0.07	(1.20)	0.957	159
Any awareness activity	-0.37	(1.24)	0.770	160
Contractual obligations	-0.56	(0.86)	0.523	161
Panel B: Technical Security Measures				
Biometric authentication	1.71***	(0.60)	0.008	158
Security testing/auditing	1.28	(1.19)	0.291	159
Data encryption	0.13	(0.85)	0.878	159
Network access control	-0.05	(0.74)	0.947	161
Log file maintenance	-0.53	(0.87)	0.547	160
VPN usage	-0.74	(1.13)	0.519	160
Off-site backup	-1.14	(0.85)	0.191	161
Strong password policy	-1.47*	(0.84)	0.091	161

Notes: Each row reports the coefficient on Treated × Post from a separate DiD regression with country, size-class, and country × year fixed effects. Outcome is the percentage of enterprises adopting each measure. Standard errors clustered by country. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The aggregate null masks important heterogeneity across individual measures. Table 3 decomposes the DiD by individual security indicator, revealing a striking asymmetry. Among formal compliance measures, *compulsory security training* increases by 3.67 percentage points—the largest effect in the table, significant at the 0.1% level. This is economically meaningful: it represents an 8.6% increase relative to the 2022 mean for 50–249 firms (42.7%). Formal

security policy adoption rises by 1.90 pp ($p = 0.19$) and risk assessments by 1.00 pp ($p = 0.26$), directionally consistent but imprecise.

Among technical measures, only biometric authentication shows a significant increase (1.71 pp, $p < 0.01$), likely reflecting the multi-factor authentication requirements embedded in NIS2. The remaining seven technical measures—encryption, VPN, log maintenance, network access control, backup, password policy, and security testing—show no significant response, with point estimates scattered around zero.

This decomposition is the core finding: NIS2 shifts compliance activity toward the least technically demanding mandate (training) while leaving substantive defensive investment unchanged. A pooled indicator-level regression confirms the pattern: NIS2-mandated measures increase 1.46 pp more than non-mandated measures ($p = 0.03$), while the baseline effect on non-mandated measures is essentially zero ($\hat{\beta} = -0.04$, $p = 0.95$).

5.3 Pre-Trends

Table 4: Event Study: Pre-Trends and Treatment Effects

	Technical Index	Formal Index
Treated × 2019	0.390	-0.583
	(1.288)	(1.215)
Treated × 2022	(reference period)	
Treated × 2024	0.315	0.985
	(1.297)	(1.234)
Country FE	Yes	Yes
Size FE	Yes	Yes
Year FE	Yes	Yes
Observations	162	161

Notes: Event study with 2022 as the reference period. Treated = 50–249 employee firms. The coefficient on Treated × 2019 tests the parallel trends assumption: a near-zero, insignificant coefficient indicates parallel pre-trends between treated and control size classes. Standard errors clustered by country. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 reports the event study with 2022 as the reference period. The pre-trend coefficient (Treated $\times$ 2019) is 0.39 pp for the Technical Index and -0.58 pp for the Formal Index, both statistically insignificant and small relative to the post-treatment coefficients. The parallel trends assumption is not rejected for either outcome.

5.4 Robustness and Extensions

Table 5: Robustness and Extensions

Specification	Technical	Formal	Other	N
Panel A: Triple-Difference (Transposition)				
Treated × Post × Transposed	-1.23	0.53		162
	(1.67)	(2.02)		
Panel B: Dosage (all three size classes)				
Newly covered (50–249) × Post	0.12	1.25		243
	(0.87)	(0.92)		
Intensified (250+) × Post	0.82	3.37***		
	(0.71)	(0.85)		
Panel C: Placebo (250+ vs. 50–249, both regulated)				
Placebo Treated × Post	0.70	2.18**		162
	(0.72)	(0.88)		
Panel D: Security Incidents				
Treated × Post			-2.07***	162
			(0.56)	
Panel E: Mandated vs. Non-Mandated Measures				
Treated × Post	-0.04			2237
	(0.73)			
× NIS2 Mandated	1.46**			
	(0.64)			

Notes: All specifications include country and size-class fixed effects. Panels A–C use country × year FE where possible. Panel D outcome is the mean percentage of enterprises reporting security incidents. Panel E pools all 14 indicators with indicator FE and interacts treatment with a dummy for NIS2-specifically-mandated measures (risk assessment, encryption, training, formal policy). Standard errors clustered by country. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 presents five robustness exercises.

Triple-difference. Panel A tests whether effects are stronger in countries that had transposed NIS2 by the survey date. The DDD coefficient is small and insignificant for both indices, suggesting that the *announcement* of NIS2 obligations—rather than their formal transposition into national law—drives the observed compliance response. This is consistent with firms responding to the EU-level directive signal rather than waiting for national legislation.

Dosage. Panel B includes the 250+ size class. For technical measures, neither newly covered (50–249) nor intensified (250+) firms show significant changes. For formal measures, the 250+ class shows a highly significant 3.37 pp increase ($p < 0.001$), exceeding the 50–249 effect. This is consistent with NIS2 intensifying formal compliance requirements for already-regulated large firms, who face new incident-reporting timelines and supply-chain obligations.

Placebo. Panel C compares 250+ firms against 50–249 firms—two groups both subject to NIS2. Technical measures show no differential effect (0.70 pp, $p = 0.34$), consistent with neither group increasing hardware investment. Formal measures show a significant 2.18 pp increase for the larger class ($p = 0.02$), consistent with dose-response: firms facing more intense obligations increase documentation more.

Security incidents. Panel D reports a significant 2.07 pp decline in security incidents for treated firms ($p < 0.001$). This supplementary result is provocative: if compliance theater is purely wasteful, incidents should not decline. One interpretation is that awareness training—even without technical investment—reduces human-factor vulnerabilities (phishing susceptibility, password reuse) that are the primary attack vector for most firms (European Union Agency for Cybersecurity, 2024). A second interpretation is that regulated firms become better at *detecting and reporting* incidents, which could go in either direction. I present this result as suggestive, not causal, given the different survey module and potential measurement differences.

Mandated vs. non-mandated. Panel E confirms that the treatment effect is concentrated on NIS2-specifically-mandated measures (1.46 pp more than non-mandated, $p = 0.03$). This rules out the alternative explanation that medium firms are simply growing their cybersecurity capacity independently of regulation.

6. Discussion

The finding that NIS2 increases compulsory training without increasing technical security investment invites two interpretations.

The first is *compliance theater*: firms do the minimum necessary to satisfy auditors,

prioritizing cheap, visible compliance activities over costly substantive defenses. Training programs are easy to document, inexpensive to implement (often outsourced online courses), and directly auditable. Encryption upgrades, intrusion detection systems, and penetration testing are expensive, require technical expertise, and are harder for regulators to verify. The asymmetric response is exactly what a rational firm facing costly compliance and imperfect verification would produce (Stigler, 1971; Bandiera et al., 2009).

The second interpretation is more charitable. NIS2's emphasis on awareness and training may be *well-targeted* to the binding constraint for medium-sized firms, which often lack the budget for enterprise-grade technical solutions but can meaningfully reduce risk through behavioral change. The decline in security incidents is consistent with this channel. If the marginal return to training exceeds the marginal return to technical investment for firms near the threshold—as Anderson (2001) suggests for organizations where human factors dominate the threat landscape—then NIS2's effective prioritization of training may be efficient even if unintended.

Both interpretations share a common implication: the *regulation's stated objective* of comprehensive risk management is not achieved. Whether the outcome is waste (interpretation one) or second-best efficiency (interpretation two), NIS2 has not produced the broad-based technical security upgrading that its architects envisioned. Policymakers designing cybersecurity mandates—including the US CIRCIA, UK Product Security Act, and Japan's revised Cybersecurity Basic Act—should note this gap between regulatory ambition and behavioral response.

A limitation of this study is the aggregate, country×size-class level of the data. Firm-level analysis could examine bunching at the 50-employee threshold, heterogeneity by sector, and dose-response by compliance cost. The triennial survey frequency also limits the number of pre-treatment periods. As additional post-treatment waves become available (2027, 2030), the dynamic response and potential hardware catch-up can be assessed.

7. Conclusion

Cybersecurity regulation faces a fundamental verification problem: the outcomes that matter most—resilience to attack, speed of detection, containment of breach—are precisely the outcomes hardest for regulators to observe. In this environment, firms rationally substitute toward observable compliance. The EU's NIS2 Directive, the most ambitious cybersecurity regulation ever implemented, generates a measurable increase in compulsory training—the most auditable mandate—while leaving technical security investment unchanged. This divergence, which we call compliance theater, is the predictable consequence of regulating a

domain where compliance is cheap and defense is expensive.

The open question is whether compliance theater is a way station or a terminus. If training builds organizational capacity that eventually leads to technical investment, NIS2 may prove more effective than its short-run effects suggest. If training substitutes for investment—satisfying the regulatory obligation without addressing the underlying vulnerability—then the directive's cost-benefit calculus is less favorable than its proponents assume. Future research, armed with additional post-treatment survey waves and potentially firm-level data, can distinguish between these hypotheses.

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

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman, "The Economics of Privacy," *Journal of Economic Literature*, 2016, *54* (2), 442–492.

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

_ **and Tyler Moore**, "The Economics of Information Security," *Science*, 2006, *314* (5799), 610–613.

Aridor, Guy, Yeon-Koo Che, and Tobias Salz, "The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR," *RAND Journal of Economics*, 2023, *54* (2), 253–294.

August, Terrence and Tunay I. Tunca, "The Economics of Information Security Investment," *Management Science*, 2006, *52* (11), 1698–1718.

Bandiera, Oriana, Andrea Prat, and Tommaso Valletti, "Active and Passive Waste in Government Spending: Evidence from a Policy Experiment," *American Economic Review*, 2009, *99* (4), 1278–1308.

European Commission, "Commission Recommendation 2003/361/EC Concerning the Definition of Micro, Small and Medium-Sized Enterprises," Official Journal L 124, 20 May 2003 2003.

European Parliament and Council, "Directive (EU) 2016/1148 Concerning Measures for a High Common Level of Security of Network and Information Systems," Official Journal of the European Union, L 194, 19 July 2016 2016.

_ , "Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity across the Union," Official Journal of the European Union, L 333, 27 December 2022 2022.

European Union Agency for Cybersecurity, "NIS Investments 2022," Technical Report, ENISA, Heraklion 2022.

_ , "NIS Investments 2023," Technical Report, ENISA, Heraklion 2023.

_ , "ENISA Threat Landscape 2024," Technical Report, ENISA, Heraklion 2024.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

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.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda, "Who Creates Jobs? Small versus Large versus Young," *Review of Economics and Statistics*, 2013, *95* (2), 347–361.

Peltzman, Sam, "Toward a More General Theory of Regulation," *Journal of Law and Economics*, 1976, *19* (2), 211–240.

Romanosky, Sasha, "Examining the Costs and Causes of Cyber Incidents," *Journal of Cybersecurity*, 2016, *2* (2), 121–135.

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb, "Fast and Wild: Bootstrap Inference in Stata Using boottest," *Stata Journal*, 2019, *19* (1), 4–60.

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

A. Data Appendix

Eurostat ICT security survey. The Eurostat Community Survey on ICT Usage and E-Commerce collects data on enterprise ICT practices, including a dedicated security module conducted triennially. The reference population covers enterprises with 10 or more employees in NACE Rev. 2 sections C through S (excluding Section K, financial services). Data are collected via national statistical offices using a harmonized questionnaire. Results are published as adoption rates (percentage of enterprises) disaggregated by country, size class, and NACE sector. I use the cross-sector aggregate (NACE C10–S951, excluding K) for the analysis sample.

Indicator classification. The 14 indicators used in the analysis are classified as follows:

Technical security measures (8 indicators): E_SECMDENC (encryption), E_SECMLOG (log maintenance), E_SECMNAC (network access control), E_SECMOSBU (off-site backup), E_SECMSPSW (strong passwords), E_SECMTST (security testing), E_SECMUIBM (biometric authentication), E_SECMVPN (VPN).

Formal compliance measures (6 indicators): E_SECPOL2 (formally defined ICT security policy), E_SECMRASS (risk assessment carried out), E_SECAWANY (any awareness-raising activity), E_SECAWCTP (compulsory ICT security training/courses), E_SECAWVTGI (voluntary training, guides, or information), E_SECAWCONT (contractual ICT security obligations on partners).

NIS2 transposition coding. Countries were coded as "transposed" if they had notified the European Commission of completed national transposition measures by October 2024. The classification is based on EUR-Lex's national implementation measures database. Transposed: Belgium, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Germany, Greece, Hungary, Italy, Latvia, Lithuania. Not transposed: Austria, Bulgaria, Finland, France, Ireland, Luxembourg, Malta, Netherlands, Poland, Portugal, Romania, Slovakia, Slovenia, Spain, Sweden.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Technical Index	0.120	0.866	13.52	0.0089	0.0641	Small positive
Formal Index	1.182	0.904	13.19	0.0896	0.0685	Moderate positive
Compulsory Training	3.672	0.761	12.81	0.2867	0.0594	Large positive
Security Incidents	-2.074	0.557	10.22	-0.2029	0.0545	Large negative
Panel B: Heterogeneous (Formal Index by Baseline Adoption)						
High-baseline countries	2.280	0.869	12.30	0.1854	0.0707	Large positive
Low-baseline countries	-0.289	1.600	11.45	-0.0253	0.1397	Small negative

Notes: **Country:** European Union (27 member states). **Research question:** Whether the EU NIS2 Directive's 50-employee regulatory threshold increases cybersecurity investment among newly regulated medium-sized firms, distinguishing technical security measures from formal compliance documentation. **Policy mechanism:** NIS2 (Directive 2022/2555) imposes mandatory risk management, incident reporting within 24 hours, supply chain security audits, and staff awareness training on enterprises with 50 or more employees in essential and important sectors, while exempting smaller firms below the threshold. **Outcome definition:** Technical Index is the mean adoption rate (percentage of enterprises) across eight technical measures (encryption, VPN, network access control, log maintenance, backup, password policy, security testing, biometric authentication); Formal Index averages six formal compliance measures (security policy, risk assessment, awareness activities, compulsory training, voluntary training, contractual obligations); Compulsory Training is the share of enterprises providing mandatory ICT security courses; Security Incidents is the mean share reporting incidents. **Treatment:** Binary; firms with 50–249 employees (newly NIS2-covered) versus 10–49 employees (exempt). **Data:** Eurostat ICT Security Survey (isoc_cisce_ra and isoc_cisce_ic), triennial waves 2019, 2022, 2024, at country $\times$ size-class level, 162 observations for the main specification. **Method:** Two-period difference-in-differences with country, size-class, and country $\times$ year fixed effects; standard errors clustered by country (27 clusters); randomization inference with 1,000 permutations. **Sample:** All EU27 member states with non-missing data in the Eurostat ICT security survey across all three waves; firms in NACE sectors C10–S951 excluding financial services. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the unconditional standard deviation of the outcome. 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).