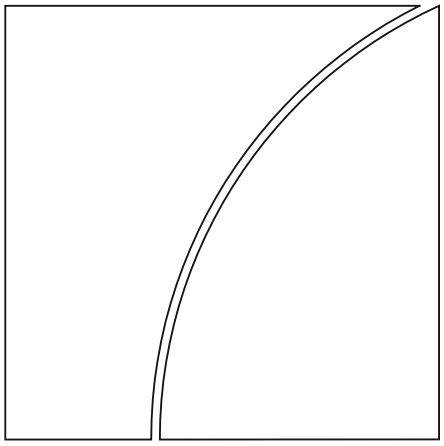

BIS Working Papers

No 1351

Disciplining digital risk: evidence from cyber stress tests

by Nordine Abidi, Leonardo Gambacorta, Christoffer Kok, Leonardo Madio, Ixart Miquel-Flores and Alberto Partida

Monetary and Economic Department

May 2026

JEL classification: G21, G28, G32, L86, K23

Keywords: cyber risk, bank supervision, stress test, IT investment

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Disciplining Digital Risk: Evidence from Cyber Stress Tests

Nordine Abidi* Leonardo Gambacorta† Christoffer Kok‡

Leonardo Madio§ Ixart Miquel-Flores¶ Alberto Partida**

Abstract

Investment in cybersecurity in an interconnected banking system has public-good proper-ties: positive externalities can generate systemic underinvestment. Using confidential supervi-sory data from the European Central Bank, we first identify "laggard" European banks that underinvest relative to their cyber-risk profiles, and then examine how supervisory scrutiny affects their incentives to invest. We exploit the 2024 ECB Cyber Resilience Stress Test (CyRST) as a quasi-natural experiment. In a difference-in-differences design, we find that following the CyRST announcement, laggard banks increased cybersecurity investment by about 80% relative to their peers. The response is stronger among laggards subject to high-intensity supervisory oversight, consistent with scrutiny exerting a disciplining effect. Overall, the results suggest that targeted supervisory scrutiny may help mitigate underinvestment incentives and strengthen banks' operational risk management.

Keywords: Cyber Risk, Bank Supervision, Stress Test, IT Investment.
JEL Codes: G21, G28, G32, L86, K23.

*International Monetary Fund. E-mail: NAbidi@imf.org
†Bank for International Settlements and CEPR. E-mail: Leonardo.Gambacorta@bis.org
‡European Central Bank. E-mail: christoffer.kok@ecb.europa.eu
§University of Padua. E-mail: leonardo.madio@unipd.it
¶European Central Bank. E-mail: Ixart.Miquel_Flores@ecb.europa.eu
‖Frankfurt School of Finance & Management. E-mail: I.MiquelFlores@fs.de
**European Central Bank. E-mail: alberto.partida@ecb.europa.eu

We thank the participants at the 2025 Federal Reserve Stress Testing Research Conference, the MCM Quantm Technical Seminar at the IMF and the ECB Banking Supervision Research Seminar for valuable feedback. We are especially grateful to Sophia Kazinnik (discussant), Stephen Cecchetti, George Pennacchi, and Azamat Abdymomunov for insightful comments and suggestions. We thank Márton Barta for his valuable research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect those of the European Central Bank (ECB), the International Monetary Fund (IMF) or the Bank for International Settlements (BIS). The dataset employed in this paper contains confidential statistical information. Its use for the analyses described in the text has been approved by the relevant ECB decision-making bodies. All necessary measures have been taken to ensure the physical and logical protection of the information.

1 Introduction

Cyber risk has emerged as a primary operational and systemic threat to the global financial system. High-profile incidents, such as the ransomware attack that disrupted ICBC's access to the U.S. Treasury market[1] and the data loss at the service provider CloudNordic,[2] illustrate how localized attacks can propagate rapidly across financial networks. While these episodes resonate with classic models of financial contagion (Allen and Gale, 2000; Acemoglu et al., 2015), cyber risk introduces unique challenges. Specifically, the resilience of the financial system is disproportionately threatened by its most vulnerable institutions, which can become entry points for shocks with cascade effects (e.g., Duffie and Younger, 2019; Gogolin et al., 2021; Eisenbach et al., 2022). Amid these concerns, public attention to cyber-related risks has grown exponentially.

While cybersecurity provides strong private operational benefits, its systemic dimension creates a classic public-good problem[3]. Disruptions at a single institution or critical service provider can propagate across interconnected financial networks, potentially affecting multiple institutions at once. This makes cyber resilience different from many other operational investments: while the benefits of cybersecurity are partly private, they are also partly system-wide. A bank that is better protected is less likely to become an entry point for broader disruption. Because banks internalise only part of the broader benefits of cyber resilience, they may invest less than is optimal from a system-wide perspective (see, e.g., Kashyap and Wetherilt, 2019; Aldasoro et al., 2023; Anand et al., 2024). In principle, this underinvestment problem creates a role for supervision. But most supervisory tools operate through hard incentives, such as capital consequences or market discipline through disclosure.

This paper first identifies underinvestment and then studies how a policy imposing supervisory scrutiny affects banks' investment decisions. We present evidence that *targeted supervisory scrutiny*, implemented through a non-capital-based stress test, can discipline underinvestment in cyber resilience.

We draw on a unique confidential dataset and analyze the European Central Bank (ECB)'s 2024 *Cyber Resilience Stress Test* (CyRST), a novel exercise designed to assess

[1] *Cyber attacks reveal fragility of financial markets*, Financial Times, 2024

[2] *Cyber Tzar Planet: Threat dashboards reveal growing systemic risk*, Financial Times, 2025

[3] Cybersecurity in financial networks has the properties of a quasi-public good. Defensive investments generate positive externalities because stronger protection at one institution lowers the probability that it becomes an entry point for contagion. However, cybersecurity is neither fully non-rival, since defensive resources can become congested during large-scale incidents, nor fully non-excludable, as some protections (e.g., proprietary encryption or internal access controls) remain private. As a result, banks internalise only part of the systemic benefits of their investments, leading to underinvestment relative to the socially optimal level.

a bank's ability to respond to and recover from a sophisticated cyberattack. There are several features that make the CyRST a suitable quasi-natural experiment. First, the CyRST was first announced to the public on March 9, 2023, which serves as our primary treatment announcement.[4] Second, the exercise was purely qualitative, with no direct implications for Pillar 2 capital requirements. This is particularly important, as it isolates the "capital channel" common in traditional stress tests (e.g., Acharya et al., 2018b; Gropp et al., 2019) from other channels. Third, individual bank results were kept confidential, thereby muting the "disclosure channel" through which markets discipline banks (Goldstein and Leitner, 2018; Flannery, 2018). The CyRST design thus provides an ideal setting to study a third, less-studied channel: the "scrutiny channel," where the credible threat of direct examination may induce change in investment strategies by increasing the perceived cost of continued underinvestment (see Kok et al. 2023).

To guide our empirical analysis, we rely on a simple theoretical framework in which targeted supervisory scrutiny, even in the absence of formal capital penalties, mitigates underinvestment in a public-good setting. The model yields two testable predictions: (i) the policy announcement induces an aggregate increase in resources allocated to cybersecurity; and (ii) the response is driven by banks that were largely underinvesting prior to the policy, which we term "laggards."

To test these predictions, we use a two-step empirical design based on an ECB supervisory dataset on banks' cybersecurity investment. First, using only pre-treatment data, we estimate the expected level of cybersecurity investment for each bank as a function of a rich set of bank-specific characteristics. We then classify as "laggards" those banks whose average investment residual over 2020–2021 falls below the sample median. Second, using a difference-in-difference design, we exploit the public announcement of the 2024 CyRST in March 2023 to examine whether investment behavior changed after the policy announcement, between laggards and non-laggards.

We find that, in the pre-CyRST period, cybersecurity investment is more strongly associated with structural balance-sheet characteristics, such as capital and leverage ratios, than with past cyber incidents. In a before-and-after analysis, the CyRST announcement is associated with an average increase in cybersecurity investment of approximately 45% across the sector. In our main difference-in-differences specification, this response is concentrated among laggards, whose investment rises by 81% relative to non-laggards. This pattern is consistent with the CyRST disproportionately affecting banks that had previously invested less in cybersecurity than comparable peers.

[4]See Section 3 for more details. This announcement was widely reported by major financial news outlets, thus serving as a key public signal of supervisory intent.

To identify the mechanism at play, we exploit cross-sectional variation in supervisory scrutiny generated during the active conduct of the CyRST. In particular, during the implementation phase of the exercise, supervisors engaged in intensive, bank-specific interactions, issuing substantive findings and data-quality flags that reflect the depth and intrusiveness of supervisory follow-up. Exploiting variation in supervisory intensity, we find that the investment response is concentrated among laggard banks that faced intensive supervisory follow-up under the CyRST. By contrast, laggards that received limited supervisory attention show no statistically significant change. This heterogeneity is consistent with a scrutiny mechanism: when supervisors credibly raise the likelihood and consequences of detailed review, the expected cost of continued underinvestment increases, prompting banks to strengthen cyber resilience despite the presence of positive externalities.

Our results are robust to alternative proxies for supervisory scrutiny and to more stringent definitions of laggards. We detect no differential pre-trends prior to the announcement, supporting the parallel-trends assumption. Finally, we address a range of potential confounders; taken together, these analyses support the interpretation that the estimated effects are associated with the CyRST rather than broader investment dynamics.

This paper makes three contributions. First, it provides novel causal evidence on the effects of a cyber stress test on bank investment behavior. Second, by exploiting a setting with no direct capital consequences and no public disclosure of bank-level results, it offers evidence consistent with supervisory scrutiny affecting behavior through a distinct channel. Third, it provides empirical support for Anand et al. (2024), who argue that regulatory intervention is necessary to move the financial system from a fragile, low-investment equilibrium to a more resilient one. More broadly, the findings suggest that targeted scrutiny may complement traditional regulatory tools in settings where operational risks are difficult to quantify and fast-moving.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 outlines the characteristics of the ECB's CyRST. Section 4 presents our theoretical model. Section 5 describes our data. Section 6 details the econometric approach. Section 7 presents our findings. Section 8 concludes.

2 Related Literature

This paper contributes to three research streams: cybersecurity as a source of systemic financial risk, the effects of stress testing on bank behavior, and the broader role of supervision in shaping financial intermediation.

Cybersecurity as a Systemic Financial Risk. Recent work emphasizes that cyber incidents can propagate across interconnected institutions and service providers, generating disruptions that extend beyond the directly affected firm (Duffie and Younger, 2019; Gogolin et al., 2021; Eisenbach et al., 2022). A related theoretical literature highlights that cybersecurity investment may exhibit positive externalities, creating incentives for private institutions to invest less than is desirable from a system-wide perspective (Kashyap and Wetherilt, 2019; Aldasoro et al., 2023; Anand et al., 2024). Ahnert et al. (2024) model how underinvestment also arises from a principal-agent problem, where the unobservability of security investment by clients reduces firms' incentive to invest. We build on this literature by providing empirical evidence on how supervisory scrutiny affects cybersecurity investment in a setting with system-wide externalities. We study whether a specific supervisory intervention changes investment behavior, especially among banks that have invested less than comparable peers ex ante.

The Effects of Stress Testing on Bank Behavior. Recent studies examine the impact of stress testing on bank decision-making. Most existing work focuses on large-scale, disclosure-based exercises with direct capital consequences, such as the U.S. Comprehensive Capital Analysis and Review (CCAR) and European banking stress tests [5] [6] (Berger and Bouwman, 2013; Hirtle and Lehnert, 2015; Schäfer et al., 2016; Kohn and Liang, 2019; Cortés et al., 2020). Recent studies on EU stress tests likewise show that public disclosure leads to market discipline and balance sheet adjustments by participating banks (see, e.g., Petrella and Resti, 2013; Schäfer et al., 2016). Other studies found that these programs induce banks to de-risk, adjust lending policies, and increase capitalization (e.g., Acharya et al., 2018b; Goldstein and Leitner, 2018).

Our work departs from this literature by analysing a different type of supervisory exercise—one focused on operational resilience, with no public disclosure of firm-level results and no ex-ante link to capital requirements. This setting allows us to isolate the "scrutiny channel." The effectiveness of such non-capital tools is supported by recent theory; for example, and as mentioned before, Ahnert et al. (2024) also show that imposing minimum investment standards or firm liability rules for breaches can resolve underinvestment problems. Kok et al. (2023) document disciplining effects of supervisory scrutiny within traditional stress tests. However, they cannot fully disentangle this channel from the simultaneous threat of capital add-ons (capital channel) and market reactions (market discipline channel). By exploiting the ECB's CyRST, we can study the scrutiny channel in a setting where the capital and

[5] *Stress Tests, Board of Governors of the Federal Reserve System*
[6] *EU-wide stress testing, Stress Test 2025*

disclosure channels are substantially muted.

The Effects of Banking Supervision. A growing literature shows that supervisory action can discipline banks and induce organizational change, particularly in risk management and governance structures. Recent papers show that supervisors could affect banks by improving internal risk controls and governance (Passalacqua et al., 2021), for instance by acquiring greater weight as stakeholders within the bank's decision-making process (Gopalan et al., 2021), or by generating new information that triggers corrective actions (Ivanov and Wang, 2019). This ultimately alters risk management functions, including staffing, reporting lines, and investment decisions (Schneider et al., 2025). Our paper contributes to the banking supervision literature by examining how qualitative oversight affects cyber-related investment decisions and related organizational adjustments.

3 Institutional Framework

This section outlines the institutional design of the ECB's CyRST. In response to the rapid digitalization of the banking sector, the ECB's 2024–2026 Supervisory Priorities identified strengthening banks' operational resilience as a key strategic objective.[7] The primary objective of the CyRST was to assess a bank's ability to *respond to and recover from* a sophisticated cyberattack, marking a shift in focus from prevention to resilience. The exercise simulated a disruptive scenario in which an attacker compromises a bank's critical IT systems, requiring the institution to activate its emergency protocols and demonstrate its ability to restore core operations from backups within a specified timeframe. The design of the exercise makes it useful for studying supervisory scrutiny while substantially muting two other channels common in stress testing:

1. *No Direct Capital Channel:* The CyRST was a purely qualitative exercise and it entailed neither capital depletion estimates nor pass/fail thresholds, and the ECB provided explicit public assurances that there would be no *direct* impact on bank-specific Pillar 2 capital guidance.[8]

2. *No Market Discipline Channel:* No bank-level results or rankings were publicly disclosed. The ECB published only a high-level, anonymized summary of aggregate findings, shielding individual institutions from market discipline. However, the confidential findings for

[7] See *ECB Banking Supervision: Supervisory Priorities for 2024-2026*.

[8] The Pillar 2 requirement is a bank-specific capital requirement that supplements the minimum capital requirement. The ECB gave explicit public assurances of no direct capital impact. See *ECB to stress test banks' ability to recover from cyberattack*, ECB Press Release, January 3, 2024.

each bank were incorporated into its annual Supervisory Review and Evaluation Process (SREP). This integration serves as the primary enforcement mechanism; significant deficiencies identified in the CyRST could lead to future supervisory measures, including higher Pillar 2 Requirements (P2R), thereby rendering the exercise a credible threat.[9]

The exercise encompassed the universe of 109 Significant Institutions (SIs), the largest banking groups in the euro area under the ECB's direct supervision. All participants completed a comprehensive, 395-item self-assessment questionnaire covering their IT architecture, governance, and contingency planning. A subset of 28 banks was selected for an *enhanced assessment*, with selection based on criteria such as systemic importance and business model diversity.[10]

Banks in the enhanced-assessment cohort were subjected to significantly more intensive and intrusive supervisory oversight. While the standard assessment relied on banks' own attestations, the enhanced assessment entailed direct verification by supervisory teams. This process included On-Site Quality Assurance Reviews (OSQARs), supervisory deep dives to validate internal procedures, and a live IT recovery test, requiring banks to physically demonstrate their ability to restore critical systems from backups.

The CyRST was first publicly signaled on March 9, 2023, in remarks by the Chair of the ECB Supervisory Board. We treat this disclosure as our primary treatment announcement.[11]

The timeline of the CyRST, illustrated in Figure 1, is central to our identification strategy. The process began with the initial public announcement on March 9, 2023, which served as a broad signal of regulatory intent. The active implementation phase, spanning from the official launch on January 3, 2024, to the private dissemination of assessments in July 2024, constitutes the core of the supervisory intervention. During this window, the European Central Bank engaged in an intensive, iterative dialogue with participating institutions. This

[9] The SREP is the core process of European banking supervision, where the ECB assesses a bank's strategies, processes, and risks, and decides on any necessary supervisory measures, such as firm-specific capital requirements (P2R and P2G) and other qualitative actions. While deficiencies identified in the CyRST may inform subsequent SREP assessments, the exercise itself had no direct or automatic capital consequences. This distinction is central to our analysis, as it allows us to isolate the scrutiny channel from the traditional capital channel.

[10] We do not exploit assignment to the enhanced assessment as an exogenous shock to supervisory scrutiny. Selection into the enhanced-assessment cohort was not random, but mechanically driven by ex ante supervisory criteria, most notably systemic importance, as all G-SIB Significant Institutions headquartered in SSM Member States were mandatorily included, alongside considerations of business-model and geographical diversity. As a result, enhanced-assessment banks differ systematically from other institutions along dimensions directly correlated with size, complexity, and baseline supervisory intensity. Using enhanced-assessment status as a treatment would therefore conflate the effect of supervisory scrutiny with pre-existing institutional characteristics, undermining causal interpretation.

[11] *Interview with Andrea Enria, Chair of the Supervisory Board of the ECB*, ECB 2023. The announcement was widely reported by major financial news outlets (e.g., Reuters, *ECB to test banks for cyber resilience*, March 9, 2023), serving as a key public signal of supervisory intent.

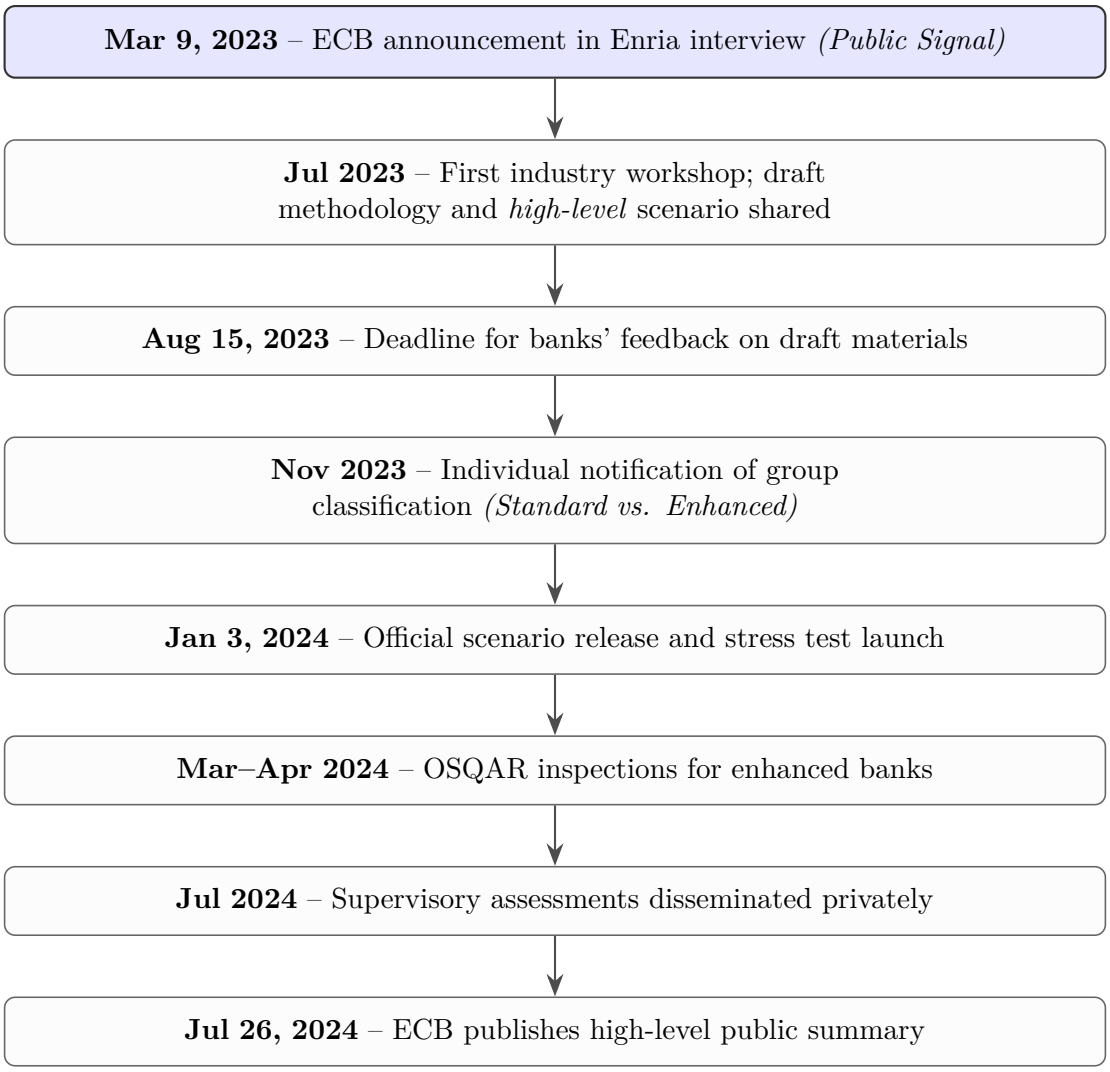

Figure 1: **Timeline of the ECB 2024 Cyber Resilience Stress Test**

phase was characterized by targeted information requests and real-time evaluative assessments. Importantly, these interactions generated the bank-level data that we exploit as proxies for scrutiny intensity.

4 Conceptual Framework

This section presents a stylized theoretical setting that informs our empirical analysis. The model captures the feature of cyber risk as a public-good problem (Kashyap and Wetherilt, 2019; Ahnert et al., 2024).[12] Crucially, we depart from traditional models of system reliability

[12]One might argue that private certification intermediaries (e.g., cyber rating agencies) could mitigate this friction. However, private solutions are insufficient in this context due to two distinct failures. First, *verifiability*: Unlike financial statements, effective cyber resilience relies on deep internal protocols that

that are based on strict weakest-link assumptions or binary actions (e.g., Varian, 2004), which would entail coordination games and multiple equilibria.

The starting point is a simple public-good problem. Banks invest in cybersecurity, and these investments generate not only private benefits for the investing bank, but also positive spillovers for the broader financial system.[13] Because each bank internalizes only part of the total benefit created by its investment, equilibrium investment is below the socially efficient level. This is the familiar underinvestment problem that arises with positive externalities. Banks are heterogeneous: some are "high-type" institutions that can build cyber resilience at relatively low marginal cost, while others are "low-type" institutions for which additional cybersecurity investment is more expensive. Conditional on type, all banks underinvest relative to the socially desirable benchmark, but low-type banks invest less in equilibrium because their marginal costs are higher.

Denote the equilibrium effort for each type $i = \{L, H\}$ as $e_i^\star$, then it is immediate that $e_L^\star < e_H^\star$. Moreover, let θ_i^{SO} indicate the level of effort that would be social desirable conditional on a type i, then $e_L^{SO} > e_L^\star$ and $e_H^{SO} > e_H^\star$. Thus, each bank would underinvest relative to the level that would be socially desirable conditional on a type.

We then introduce a stylized supervisory intervention. The supervisor sets a qualitative benchmark for cybersecurity investment or resilience effort. Falling short of that benchmark triggers a penalty or other consequences of being flagged. We note this benchmark as $\bar{e}$, which is independent on the bank type. Suppose the benchmark is binding for low-type banks but not for high-type banks, i.e., $\bar{e} > e_L^\star$ and $\bar{e} < e_H^\star$. In other words, the supervisor sets the threshold above the pre-policy investment level of laggards, but below that of banks that were already investing relatively more.

We assume that the penalty for non-compliance is sufficiently high that affected banks optimally increase investment to avoid it. Let $e_i^{\star\star}$ denote the post-supervision effort level. The implication is immediate: high-type banks continue to choose their unconstrained optimum, $e_H^{\star\star} = e_H^\star > \bar{e}$, while low-type banks *catch up* by bunching at the benchmark, so that $e_L^{\star\star} = \bar{e} > e_L^\star$.

In the model, low types are firms for which increasing cyber investment is more difficult or more costly at the margin. These firms are also the ones whose pre-policy investment lies furthest below the benchmark and who therefore respond most strongly to supervisory

are difficult for private third parties to verify without the intrusive legal powers of a supervisor (e.g., on-site inspections). Second, *appropriability*: Because the benefits of resilience are largely systemic (public), individual banks have insufficient private incentive to pay for high-quality, expensive auditing, leading to a "lemons problem" in the certification market.

[13] We do not attempt to solve a full network equilibrium under cyber contagion. Appendix II provides formal proofs and derivations.

scrutiny. In the data, we do not observe these structural costs directly. Instead, we identify banks whose investment lies persistently below predicted levels conditional on observables. We interpret these banks as empirical counterparts of firms that face larger adjustment needs under supervisory scrutiny.

Our model identifies two main predictions:

P1 (Aggregate Investment): *The announcement of the CyRST induces an aggregate increase in cybersecurity investment across the banking sector.*

P2 (Heterogeneous Effects): *The investment response is non-uniform; the increase is concentrated among laggard banks, which exhibit a significantly larger response compared to their more compliant peers.*

5 Data, Variables, and Summary Statistics

5.1 Data Sources and Sample Construction

Our primary analysis relies on a proprietary supervisory panel dataset from the ECB, providing a granular view of cybersecurity investment and operational risk for banks in our dataset. The panel tracks the 109 SIs that participated in the 2024 Cyber Stress Test from 2019 to 2024. The data are derived from the annual ECB IT Risk Questionnaire (ITRQ), which is collected in the first quarter of year t but reflects bank outcomes for year $t-1$. Because the CyRST was announced on March 9, 2023, midway through the reporting window for the 2022 reference year, these 2022 observations are potentially contaminated.

As illustrated in Figure A.1 in Appendix III, some institutions reported prior to the policy signal, while others reported after. To ensure a sharp treatment boundary and avoid anticipation bias, we conservatively exclude 2022 from our sample. The pre-treatment period is defined as 2019–2021, and the post-treatment period comprises 2023–2024, capturing the announcement and implementation phases, respectively. Our final balanced panel consists of 96 banks with continuous SI status and complete reporting histories (475 bank-year observations).[14] This dataset integrates three distinct regulatory sources:

1. ECB IT Risk Questionnaire (ITRQ). We draw upon the ITRQ, a mandatory and confidential annual data collection conducted by the European Central Bank as part of the SREP assessment of ICT and operational resilience. The questionnaire provides a structured

[14]Comparison of pre-treatment observables (size, ROE, CET1) indicates no statistically significant selection bias between the final sample and the 13 excluded institutions.

and comprehensive view of banks' ICT risk profile, technology setup, and cyber-resilience capabilities. Each wave is administered in the first quarter of the year but captures banks' IT risk management, governance, and expenditure data for the preceding year. The ITRQ reports, for instance, normalized IT security expenditures (our dependent variable), the frequency and severity of cyberattacks, IT staffing intensity (including vacancy and turnover rates), and operational resilience metrics such as detection and recovery times. It also tracks forward-looking indicators regarding cyber insurance coverage, innovation projects, and legacy IT risks (e.g., end-of-life systems). Crucially, we use distinct proxies to quantify the misalignment between a bank's self-assessment and the supervisor's benchmark, allowing us to track whether banks systematically over- or under-estimate their cybersecurity posture relative to regulatory expectations.

2. 2024 CyRST Archive. We use confidential supervisory records from the ECB's 2024 CyRST. The data are generated during the active conduct of the stress test, between its formal launch on January 3, 2024 and its completion on July 26, 2024, and reflect supervisory interactions, requests, and assessments taking place in real time. Unlike standard financial disclosures, these records capture the iterative dialogue between the supervisor and the bank during the CyRST. These data allow us to construct three, bank-level indicators of scrutiny: (i) $hsti\_sevfind_i$, captures the cumulative intensity of substantive weaknesses or findings identified during the exercise, (ii) $hsti\_flag_i$, proxying for data reporting integrity, this measure aggregates data quality and plausibility flags issued by supervisors during the validation phase; and (iii) $HighScrutiny_i$, a composite measure of elevated supervisory attention. This indicator identifies institutions subjected to heightened supervisory pressure due to either material cyber-security deficiencies or reporting opacity.[15]

3. ECB Supervisory Bank Dataset. We merge the ITRQ data with standard regulatory reports (FINREP and COREP). These filings provide harmonized, audited data on bank financials, enabling us to control for time-varying bank characteristics such as size, profitability, and capitalization.

5.2 Variable Definitions

Our analysis rests on the identification of bank underinvestment, a characteristic that is unobservable to the econometrician. To address this, we proxy for this latent trait by modeling the *expected* level of cybersecurity investment conditional on a bank's risk profile and fundamentals. Banks that systematically invest below this benchmark are classified as "laggards."

[15]Detailed construction steps for each indicator are provided in Appendix I.

This classification forms the basis of our first-stage analysis. Table 1 provides an overview of the variables used in this procedure. We categorize them into four distinct groups.

The first group, *Cyber Risk Exposure and Controls*, captures both the threat environment and the bank's operational capacity. This includes direct measures of past risk realization (*attack* and *attack_losses*) and key indicators of resilience (*detectiontime*, *recoverytime*). We also include detailed metrics on human capital and governance, such as the "three lines of defense" structure, IT staff vacancy and turnover rates, and the share of permanent IT staff. Crucially, we leverage the bank's self-reported control and risk scores alongside supervisory benchmarks to construct proxies for misalignment in risk control (*effort* or *resources* dimension) and risk level (*outcome* or *vulnerability* dimension).[16]

The second group, *Cyber Insurance*, captures reliance on external risk transfer as a complement to or substitute for internal spending. We observe the extensive margin, whether a bank holds a policy (*insurance_d*), and the intensive margin, measured by the direct monetary outlay for coverage (*insurancecontractsdirectcosts*) and the policy's retained loss via the deductible (*insurancecontractsdeductibleamount*).

The third group, *Cyber Innovation*, reflects forward-looking risk management strategies. We record whether a bank reports any innovation initiative in the IT/cyber domain (*innovationprojects_d*) and the scale of that pipeline, distinguishing between planned projects (*innovationprojectstobeimplemented*) and projects under execution (*innovationprojectsongoing*). Economically, these variables proxy for the modernization of detection, response, and recovery capabilities that may not be fully reflected in contemporaneous operating expenditure.

The fourth group, *Legacy Infrastructure and Risk*, captures structural frictions that heighten vulnerability and absorb resources. We include the log number of critical IT change programs (*log_n_criticalprojects*) and their associated spending (*log_criticalprojectsexp*), alongside indicators of technical obsolescence: the stock of end-of-life systems (*log_numbereolsystems*), the planned remediation share (*share_eol_to_be_replaced*), and the planning gap (*share_eol_gap_ratio*). These measures capture both the scale of transformation work and the execution risk that can crowd out discretionary cyber investment.

[16] *y_m_RC_DT: IT risk control level distance (bank–regulator gap).* This variable measures the gap between a bank's reported control maturity in detection/recovery and the regulator's benchmark expectation. Constructed as the difference between the bank's self-reported score and the supervisory benchmark, a positive gap indicates that the bank assesses its control environment more favorably than the regulator, while a negative gap suggests under-reporting. *y_m_RR_DT_reb: IT residual risk distance (bank–regulator gap).* This variable captures misalignment in perceived residual risk (post-control exposure). Analogously defined as the bank's reported level minus the regulator's reference level, a positive distance implies that the bank perceives lower residual risk than the regulator (suggesting underestimation of vulnerability), while a negative distance suggests the bank perceives higher risk. See Figures A.2 and A.3. These indices are constructed as bank-specific averages over 2020–2021 and held fixed, treating them as pre-determined, slow-moving controls.

Finally, we include a set of *Controls* to absorb differences in size, complexity, and financial capacity: the log count of IT systems (*log_numberitsystems*), the log of total assets (*log_TotalAssets*), an IT complexity ratio (*itcomplexityratio*), leverage and profitability (*LeverageRatio*, *ROE*), operating efficiency (*CIR*), and capital adequacy (*C_CET1CapitalRatio*). These controls ensure our laggard classification is not picking up level effects unrelated to cyber risk management per se.

In our analysis, we define our main dependent variable as the logarithm of *Cybersecurity Investment*, measured as a bank's annual IT security operating expenditures in Euros.

Table 1: Variable Definitions

Variable	Definition	Explanation	Group
Dependent Variable			
norm_itsecurityexp	Normalized IT security expenditure *(e.g., Log(inv +1), as % of OPEX, IT Running Expenses or IT Running and IT Change Expenses)*	Dependent variable (investment behavior)	dependent
Cyber Risk Exposure and Controls			
attack	Count of successful cyberattacks	Exposure to realized cyber risk	baseline_vars
attack_losses	Losses due to successful attacks	Severity of past realized risk	baseline_vars
itpermstaffintens	IT/cybersecurity staff intensity	Proactive investment in risk resources	baseline_vars
itvacancyrate	Cyber/IT job vacancy rate	Indicator of resourcing gaps	baseline_vars
itturnindex	Turnover index for IT/cyber staff	Organizational friction, staff churn	baseline_vars
fte1lod	IT First Line of Defense FTEs share	Frontline cyber risk management staff	baseline_vars
fte2lod	IT Second Line of Defense FTEs share	Risk oversight staffing	baseline_vars
fte3lod	IT Third Line of Defense FTEs share	Audit/assurance capacity	baseline_vars
recoverytime	Average time to fully recover from incidents	Key preparedness indicator	baseline_vars
detectiontime	Average time to detect incidents	Key preparedness indicator	baseline_vars
y_m_RC_DT	IT risk control level distance *(bank-regulator gap)*	Risk control misalignment proxy	baseline_vars
y_m_RR_DT_reb	IT residual risk distance *(bank-regulator gap)*	Risk level misalignment proxy	baseline_vars
Cyber Insurance			
insurance_d	Dummy: bank has cyber insurance	Risk transfer strategy	insurance_vars
insurancecontractsdirectcosts	Direct cost of insurance contracts	Monetary investment in risk transfer	insurance_vars
insurancecontractsdeductibleamount	Deductible amount	Depth of coverage (self-insurance)	insurance_vars
Cyber Innovation			
innovationprojects_d	Dummy: any innovation project	Strategic innovation indicator	innovation_vars
innovationprojectstobeimplemented	Planned cyber innovation projects	Forward-looking cyber maturity	innovation_vars
innovationprojectsongoing	Ongoing cyber innovation projects	Execution of strategic change	innovation_vars
Legacy Infrastructure and Risk			
log_numbercriticalprojects	Log of # of critical infra projects	Baseline IT criticality	legacy_vars
log_criticalprojectsexp	Log of critical infra investment	Resource allocation to critical IT	legacy_vars
log_criticalprojectseol	Log of EOL-tagged projects	Legacy risk indicator	legacy_vars
log_criticalprojectseolexp	Log of EOL infra spending	Attempted mitigation of legacy risk	legacy_vars
sd_numbereolsystems	Standardized # of EOL systems	Intensity of legacy risk	legacy_vars
sd_share_eol_to_be_replaced	Share of EOL systems to be replaced (std.)	Remediation planning intensity	legacy_vars
sd_eol_gap_ratio	Share of EOL systems without plan (std.)	Gap in strategic IT planning	legacy_vars
Controls			
log_numberitsystems	Log of total IT systems	System complexity & infrastructure exposure	control_vars
logA_TotalAssets	Log of total assets	Proxy for bank size and scale	control_vars
LeverageRatio	Tier 1 capital / total exposure measure	Capital adequacy and risk buffer	control_vars
ROE	Return on Equity	Bank profitability	control_vars
CIRatio	Cost-to-Income ratio	Operational efficiency	control_vars
C_CET1CapitalRatio	Common Equity Tier 1 capital ratio	Core solvency metric	control_vars

5.3 Descriptive Evidence

Table 2 presents summary statistics for our bank panel. To strictly adhere to confidentiality requirements regarding sensitive supervisory data, we suppress absolute values for critical variables. Instead, we report two disclosure-safe measures: the coefficient of variation (CV), which captures relative cross-sectional dispersion, and an indexed measure of levels, which normalizes the pre-treatment (2019–2021) average to 100 and expresses the post-treatment mean relative to this baseline.[17]

Panel A documents the pre-treatment outcome (2019–2021). Banks spent resources to cybersecurity, yet cross-sectional variation was pronounced, with CV exceeding 2.0 for both IT security expenses and realized cyberattack losses. Operational indicators, such as detection times and recovery times, display significant heterogeneity ($CV > 3.0$), suggesting the presence of asymmetries within the banking sector. In terms of fundamentals (Panel C), the average bank entered the period in solid condition, reporting a return on equity (ROE) of 4.85% and a CET1 capital ratio of 19.23%.

The post-treatment period (2023–2024) shows a marked structural shift in resource allocation. Cybersecurity investment increased by over 40% relative to the pre-treatment baseline. This increase was accompanied by a rise in IT staff intensity and a decline in vacancy rates, suggesting that additional outlays were channeled into building internal capacity rather than merely inflating budget lines.

Importantly, these inputs coincide with tangible improvements in resilience. While the frequency of successful cyberattacks declined moderately, the financial severity of incidents fell sharply, and average detection times improved. However, the data also highlight the growing sophistication of threats: the average recovery time increased post-treatment. Finally, we observe a deepening of the cyber insurance market, evidenced by a doubling of average deductible amounts, consistent with a shift toward higher risk retention and coverage limits.

6 Empirical Strategy

We estimate the effect of the ECB's CyRST on banks' cybersecurity investment. A central identification challenge is that a bank's classification as a "laggard" is not randomly assigned; rather, it reflects endogenous strategic choices shaped by unobservables such as managerial risk preferences and corporate culture. Because these same factors are plausibly correlated

[17] This approach preserves the ability to analyze cross-sectional heterogeneity and temporal shifts while safeguarding the confidentiality of raw supervisory data points.

Table 2: **Summary Statistics**

Variable	Mean	SD	P10	P90	CV	Index
Pre-treatment (2019–2021)						
Panel A: Cybersecurity investment and incidents						
IT security expenses (EUR)	—	—	—	—	2.30	100
Log IT security expenses	—	—	—	—	0.32	100
IT sec./OPEX (%)	1.23	3.82	0.11	1.81	—	—
IT sec./IT running exp. (%)	9.11	22.36	1.49	14.30	—	—
IT sec./IT run.+change exp. (%)	5.80	15.53	0.87	8.13	—	—
Successful cyberattacks	—	—	—	—	1.68	100
Cyberattack losses (EUR)	—	—	—	—	4.72	100
Recovery time (days)	—	—	—	—	3.43	100
Detection time (days)	—	—	—	—	3.15	100
Panel B: Staffing, governance, insurance, innovation, legacy IT						
IT staff intensity (%)	9.44	5.62	3.96	16.92	—	—
IT vacancy rate (%)	7.11	8.71	0.02	16.27	—	—
IT turnover index (%)	30.36	35.68	0.00	71.10	—	—
1st LoD IT FTE share (%)	6.20	6.45	0.18	14.29	—	—
2nd LoD IT FTE share (%)	0.03	0.08	0.00	0.08	—	—
3rd LoD IT FTE share (%)	0.16	0.36	0.01	0.28	—	—
Cyber insurance (dummy)	0.75	0.43	0	1	—	—
Insurance direct costs (EUR)	—	—	—	—	2.16	100
Insurance deductible (EUR)	—	—	—	—	3.92	100
Innovation flag (dummy)	0.92	0.27	1	1	—	—
Std. dev. no. EoL systems	0.06	1.20	-0.31	0.31	—	—
Panel C: IT complexity and financials						
Log no. IT systems	6.66	1.65	4.64	8.54	—	—
Log total assets	25.26	1.32	23.37	27.30	—	—
Leverage ratio (%)	6.88	2.38	4.30	10.83	—	—
ROE (%)	4.85	4.31	0.02	10.29	—	—
Cost-to-income ratio (%)	60.57	15.09	40.05	80.11	—	—
CET1 capital ratio (%)	19.23	6.42	13.79	29.27	—	—
Post-treatment (2023–2024)						
Panel A: Cybersecurity investment and incidents						
IT security expenses (EUR)	—	—	—	—	1.79	141.72
Log IT security expenses	—	—	—	—	0.22	108.81
IT sec./OPEX (%)	1.36	1.02	0.27	2.61	—	—
IT sec./IT running exp. (%)	9.21	6.02	3.05	16.40	—	—
IT sec./IT run.+change exp. (%)	5.58	3.44	1.60	9.89	—	—
Successful cyberattacks	—	—	—	—	2.21	86.51
Cyberattack losses (EUR)	—	—	—	—	6.04	45.25
Recovery time (days)	—	—	—	—	2.85	147.90
Detection time (days)	—	—	—	—	2.85	94.13
Panel B: Staffing, governance, insurance, innovation, legacy IT						
IT staff intensity (%)	11.60	6.31	4.31	19.70	—	—
IT vacancy rate (%)	6.56	7.08	0.11	14.56	—	—
IT turnover index (%)	29.97	31.14	3.57	67.20	—	—
1st LoD IT FTE share (%)	8.61	7.49	0.59	18.66	—	—
2nd LoD IT FTE share (%)	0.03	0.06	0.00	0.09	—	—
3rd LoD IT FTE share (%)	0.16	0.24	0.03	0.31	—	—
Cyber insurance (dummy)	0.83	0.37	0	1	—	—
Insurance direct costs (EUR)	—	—	—	—	1.44	209.80
Insurance deductible (EUR)	—	—	—	—	2.20	226.72
Innovation flag (dummy)	0.94	0.24	1	1	—	—
Std. dev. no. EoL systems	-0.09	0.59	-0.31	0.21	—	—
Panel C: IT complexity and financials						
Log no. IT systems	7.53	2.13	4.95	10.40	—	—
Log total assets	25.33	1.33	23.46	27.36	—	—
Leverage ratio (%)	7.26	2.29	4.74	11.06	—	—
ROE (%)	9.75	4.67	3.55	16.62	—	—
Cost-to-income ratio (%)	49.81	13.10	31.74	66.86	—	—
CET1 capital ratio (%)	19.63	5.97	14.54	26.96	—	—

Note: This table reports descriptive statistics for the main variables used in our analysis. For variables deemed sensitive, we suppress the raw values of the mean, standard deviation, and percentiles. Instead, we report two confidentiality-safe measures: (i) the coefficient of variation (CV), defined as the ratio of the standard deviation to the mean (SD/Mean); and (ii) an indexed measure of levels, constructed by normalizing the pre-treatment (2019–2021) average to 100 and expressing the post-treatment (2023–2024) average relative to this baseline.

with a bank's responsiveness to supervisory scrutiny, simple comparisons would suffer from selection bias.

To address this, we implement a two-stage empirical strategy that exploits the CyRST announcement as a quasi-experimental shock. The first stage constructs a time-invariant, pre-determined proxy for a bank's latent propensity to underinvest. The second stage uses this classification within a difference-in-differences (DiD) framework to estimate the causal effect of the CyRST on investment behavior.

6.1 Identification Strategy

Our identification exploits the public announcement of the CyRST in March 2023 as plausibly exogenous variation. The exercise entailed no direct capital consequences and no public disclosure of bank-level results, limiting the capital and market-discipline channels emphasized in the stress-testing literature (e.g., Acharya et al. 2018b, Goldstein and Leitner 2018). This design allows us to isolate a scrutiny channel, whereby the credible prospect of detailed supervisory examination disciplines banks.

A critical feature of our setting is the data collection calendar: observations for the 2022 fiscal year were collected in early 2023, overlapping with the policy announcement. To eliminate potential contamination from anticipatory adjustments or strategic reporting, we exclude all 2022 data from our main analysis. Thus, we define the pre-treatment period as 2019–2021 and the post-treatment period as commencing in 2023 (see also Figure A.1).

6.2 First Stage: Identification of "Laggards"

We define "laggards" as banks that systematically underinvest in cybersecurity relative to the level predicted by their observable fundamentals and risk profile. To identify this group, we decompose observed investment into an "expected" structural component and an idiosyncratic residual that captures discretionary deviations from the benchmark. Formally, we estimate the following two-way fixed effects model on the pre-treatment panel (2019–2021):

$$\log(\text{Investment}_{it}) = \alpha_i + \lambda_t + \mathbf{X}'_{it}\boldsymbol{\beta} + \varepsilon_{it}, \tag{1}$$

where α_i and λ_t absorb time-invariant bank heterogeneity and common shocks, respectively, and $\mathbf{X}_{it}$ includes detailed controls for cyber risk exposure, operational capacity, technological sophistication, and financial condition. The residuals $\hat{\varepsilon}_{it}$ measure the discretionary component of investment. Averaging over 2020–2021 yields a stable pre-treatment type measure, $\bar{\varepsilon}_i = \frac{1}{2}\sum_{t=2020}^{2021} \hat{\varepsilon}_{it}$. We classify bank i as a laggard if $\bar{\varepsilon}_i$ falls below the median of the sample

distribution:

$$\text{Laggard}_i = \mathbf{1}\left[\bar{\varepsilon}_i \leq \text{P50}(\bar{\boldsymbol{\varepsilon}})\right]. \tag{2}$$

In Section 7.3.1, we relax this assumption and show that our main results hold when considering alternative measures of "laggard" banks.

A potential concern is that a persistently negative residual might reflect unobserved operational efficiency or a simpler business model rather than underinvestment. We argue that this interpretation is unlikely for three reasons. First, the inclusion of bank fixed effects (α_i) and detailed time-varying controls ($\mathbf{X}_{it}$) explicitly accounts for observable drivers of investment efficiency and business model complexity. Second, as we show in Section 7.2, banks classified as laggards are statistically indistinguishable from their peers on general financial dimensions but display materially weaker cyber-resilience metrics (e.g., longer detection times), a pattern consistent with genuine underinvestment rather than superior efficiency. Finally, the finding that this specific group exhibits the largest investment response to the CyRST (as shown in our second-stage results) provides support for our interpretation: efficient firms (i.e., those with high θ) would have little need to increase investment in response to supervisory scrutiny.[18]

6.2.1 Difference-in-Differences Specification

For our main second-stage analysis, we estimate a Poisson Pseudo–Maximum Likelihood (PPML) model within a difference-in-differences framework. The PPML estimator specifies the conditional mean in multiplicative form, is robust to general forms of heteroskedasticity, and naturally incorporates zero-valued observations without requiring arbitrary transformations.[19] Our baseline specification is:

$$\mathbb{E}[\text{Investment}_{it} \mid \alpha_i, \lambda_t, \mathbf{X}_{it-1}] = \exp\left(\alpha_i + \lambda_t + \beta_{ATT} \cdot (\text{Laggard}_i \times \text{Post}_t) + \mathbf{X}'_{it-1}\boldsymbol{\delta}\right), \tag{3}$$

where β_{ATT} measures the average treatment effect on the treated (laggards) in the post-CyRST period. Interpreted as a semi-elasticity, the effect size is approximately $(\exp(\beta_{ATT}) - 1) \times 100\%$. Standard errors are two-way clustered at the bank and year levels to account for serial correlation and common time shocks.

Our identification strategy rests on two core assumptions. First, the Stable Unit Treat-

[18] We also verify that our 'Laggard' classification is distinct from general management quality. We find that a bank's laggard status is not statistically correlated with its overall SREP score (the ECB's comprehensive assessment of management and financial soundness). This confirms that we are capturing a specific domain-mismatch in cyber resilience, rather than a proxy for broadly weak bank management.

[19] This approach avoids the inconsistency and bias issues inherent to log-linear OLS in the presence of heteroskedasticity and zero values (Silva and Tenreyro, 2006, 2011).

ment Value Assumption (SUTVA) requires that the potential outcomes of any bank are unaffected by the treatment status of its peers. Spillovers across banks may attenuate the estimated treatment effects. A potential concern is the presence of spillovers across banks. Because treated and control institutions operate within the same financial system, investment responses by treated banks may affect the behavior of control banks through competitive pressures, benchmarking, or shared supervisory expectations. To the extent that control banks also increase cybersecurity investment in response to the CyRST or to the actions of treated peers, the estimated treatment effect will be attenuated. In this sense, our difference-in-differences estimates should be interpreted as lower bounds on the overall investment effect.[20]

Second, the parallel trends assumption requires that, in the absence of the CyRST, the investment trajectories of laggard and non-laggard banks would have evolved similarly. This implies the absence of confounding shocks, contemporaneous with the CyRST announcement, that differentially affect laggards through channels unrelated to the stress test. We are not aware of an institutional mechanism likely to generate such a confounder.

We provide evidence for the parallel trends assumption by estimating a PPML-based event-study specification, that is:

$$\mathbb{E}[\text{Investment}_{it} \mid \cdot] = \exp\left(\alpha_i + \lambda_t + \sum_{k=-3, k\neq -2}^{2} \beta_k \cdot (\text{Laggard}_i \times \mathbb{1}[t = 2023 + k]) + \mathbf{X}'_{it-1}\boldsymbol{\delta}\right), \tag{4}$$

where we normalize the coefficient for 2021 ($k = -2$, the last pre-treatment year in our sample given the exclusion of 2022) to zero. The coefficients $\beta_{k<0}$ serve as a falsification test; their statistical insignificance would support the parallel trends assumption, while $\beta_{k\geq 0}$ capture the dynamic post-announcement response.[21]

[20] At the same time, several features of the European banking market suggest that such spillovers are likely to be limited in magnitude. Retail and commercial banking activities remain largely segmented along national lines, and cross-border competition among significant institutions is relatively limited. As a result, direct competitive interactions between treated and control banks are often weak, reducing the scope for strong cross-bank investment spillovers. Taken together, while spillovers cannot be fully ruled out, they are more likely to attenuate the estimated effects rather than generate spurious treatment differences.

[21] Our identification strategy estimates the direct, within-bank response to supervisory scrutiny. While cybersecurity investment generates positive externalities, these spillovers are likely to operate with lags through counterparties' risk assessments, funding conditions, and network repricing. By abstracting from such general equilibrium effects, our estimates should be interpreted as partial-equilibrium treatment effects and therefore as lower bounds on the full systemic impact of the CyRST.

7 Main Results

This section presents the main empirical findings on the effects of the ECB's CyRST. Our analysis proceeds in four steps. First, we provide descriptive evidence of the aggregate impact of the policy announcement on bank investment using a before-after analysis. Second, we validate our classification of "Laggard" banks as an empirical proxy for firms with propensity to underinvest. Third, we present our main difference-in-differences estimates and an event-study analysis. Fourth, we provide evidence for the supervisory "scrutiny channel" as the primary mechanism and conduct a series of robustness tests.

7.1 Aggregate Investment

We begin by testing Prediction 1 focusing on the change in the investment level in IT security after the announcement of the CyRST. The specification is a panel fixed-effects model estimated via PPML: $\mathbb{E}[\text{Investment}_{it}|\cdot] = \exp(\alpha_i + \beta \cdot \text{Post}_t + \mathbf{X}'_{it-1}\boldsymbol{\delta})$, where the coefficient of interest, β, captures the average change in investment in the post-announcement period (2023–2024), controlling for bank fixed effects (α_i) and a vector of lagged, time-varying bank controls ($\mathbf{X}_{it-1}$).

Table 3 reports the results. The coefficient on the *Post* indicator is positive and statistically significant across all specifications. As we progressively saturate the model with controls and fixed effects, the estimate remains stable in both magnitude and significance. Our most preferred specification (Column 6), which includes bank fixed effects and the full vector of time-varying controls, shows that, on average, banks increased their cybersecurity investment by around 45% following the ECB's announcement. This response is consistent with Prediction 1.

Table 3: **Aggregate Effect of the Policy on Investment**

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: IT Security Investment					
Post	**0.349***	**0.366*****	**0.435*****	**0.433*****	**0.415*****	**0.370*****
	(0.188)	(0.119)	(0.141)	(0.103)	(0.148)	(0.050)
Bank Controls	No	Yes	Yes	Yes	Yes	Yes
Country FE	No	No	Yes	Yes	No	Yes
Business Model FE	No	No	No	Yes	No	Yes
Bank FE	No	No	No	No	Yes	Yes
Observations	475	475	475	475	465	465

Note: This table reports the estimated effect of the ECB Cyber Resilience Stress Test announcement (*Post*) on cybersecurity investment using a Poisson Pseudo-Maximum Likelihood (PPML) model. Column (1) presents the baseline specification. Columns (2)–(6) progressively introduce controls and fixed effects. *Post* is a dummy variable equal to one for the years 2023–2024 and zero for the pre-treatment period (2019–2021). Bank controls include the natural logarithm of Total Assets, Leverage Ratio, ROE, Cost-to-Income Ratio, and CET1 Capital Ratio. Robust standard errors, clustered at the bank level, are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7.2 First Stage Analysis

In our baseline setting, we define a bank as a "Laggard" if its cybersecurity spending, averaged over the 2020–2021 pre-treatment period, falls below the median of the residuals derived from a benchmark investment model. We report these results in Table A.1 in the Appendix. Among them, reactive measures of past risk realization, such as the number of previous cyberattacks or associated losses, are not statistically significant predictors. These patterns suggest that pre-treatment investment variations are less likely to be driven by incident response rates and more by structural factors such CET1 capital ratio, leverage ratio.[22].

Table 4 reports that, in the pre-treatment period, "Laggard" and non-laggard banks are statistically indistinguishable across financial dimensions, including size, profitability, and capitalization. We find no statistically significant difference in the Leverage Ratio, ROE, or CET1 Capital Ratio. On operational metrics, the only significant difference is on the IT security expenses, with no significant differences on average detection rates, cyberattack losses, index of successfull cyberattacks, or recovery time.

Panel B indicates that there are no major differences in staffing structures or governance capacity. A notable divergence, however, appears in risk transfer: laggards select insurance policies with relatively higher deductibles, consistent with a potential strategy of minimizing

[22]Figure A.4 to **??** in the Appendix plot actual versus model-predicted normalized IT security investment in 2020–2021, illustrating how residual variation is used to identify under- and over-investing banks ex ante. Figure A.6 shows as we saturate the model with controls, the dispersion of the residuals tightens considerably, indicating that our model effectively isolates the idiosyncratic component of investment behavior.

Table 4: **Pre-treatment differences between laggard and non-laggard banks (confidentiality-compliant)**

Variable	Non-Laggard Mean / Index	Laggard Mean / Index	Std. diff.	p-value
Panel A: Cybersecurity investment and incidents				
IT security expenses (Index)	100	52.74	-0.27	0.18
Log IT security expenses (Index)	100	86.15	-0.55	0.01
IT sec./OPEX (%)	1.58	0.85	-0.24	0.24
IT sec./IT running exp. (%)	12.56	5.47	-0.42	0.04
IT sec./IT run.+change exp. (%)	8.03	3.47	-0.39	0.06
Successful cyberattacks (Index)	100	87.50	-0.10	0.62
Cyberattack losses (Index)	100	66.44	-0.13	0.52
Recovery time (Index)	100	66.75	-0.18	0.38
Detection time (Index)	100	115.46	0.05	0.80
Panel B: Staffing, governance, insurance, innovation, legacy IT				
IT staff intensity (%)	8.99	9.84	0.16	0.45
IT vacancy rate (%)	6.46	7.82	0.19	0.36
IT turnover index (%)	30.81	30.28	-0.02	0.92
1st LoD IT FTE share (%)	6.28	6.09	-0.03	0.88
2nd LoD IT FTE share (%)	0.03	0.03	0.08	0.70
3rd LoD IT FTE share (%)	0.18	0.14	-0.14	0.49
Cyber insurance (dummy)	0.66	0.83	0.41	0.05
Insurance direct costs (Index)	100	84.02	-0.11	0.59
Insurance deductible (Index)	100	373.89	0.46	0.03
Innovation flag (dummy)	0.94	0.90	-0.22	0.29
Planned innovation projects (count)	16.11	20.78	0.11	0.61
Ongoing innovation projects (count)	26.46	27.74	0.02	0.92
Std. dev. no. EoL systems (count)	0.15	-0.04	-0.18	0.38
Log no. critical projects	2.89	2.97	0.07	0.74
Log critical project expenditure	16.12	16.73	0.15	0.46
Log EoL project count	1.27	1.29	0.02	0.91
Log EoL project expenditure	10.71	12.64	0.30	0.14
Panel C: IT complexity and financials				
Log no. IT systems	6.47	6.79	0.22	0.29
Log total assets	25.20	25.22	0.02	0.92
Leverage ratio (%)	7.14	6.74	-0.18	0.39
ROE (%)	4.67	5.13	0.13	0.52
Cost-to-income ratio (%)	60.18	60.74	0.04	0.85
CET1 capital ratio (%)	19.39	19.14	-0.04	0.85

Notes: Confidentiality-sensitive variables (IT security expenditure in levels and logs, cyber incidents, operational disruption metrics, and cyber insurance amounts) are expressed as indices normalised to 100 for non-laggard banks in the pre-period. Percentage variables are multiplied by 100. Standardised differences are computed from underlying raw values. p-values correspond to two-sample t-tests of mean differences.

explicit outlays. Finally, Panel C also shows that the two groups are similar across standard balance sheet characteristics.

Table 5: **Post-treatment levels and changes: laggard vs. non-laggard banks**

	Pre (2019–2021)		Post (2023–2024)		Change (Post–Pre)			
Variable	NL	L	NL	L	ΔNL	ΔL	ΔL−ΔNL	p-value
Panel A: Cybersecurity investment and incidents								
IT security expenses (Index)	100	100	114.74	199.91	14.74	99.91	85.18	0.065
Log IT security expenses (Index)	100	100	102.72	116.23	2.72	16.23	13.52	0.010
IT sec./OPEX (%)	1.58	0.85	1.38	1.34	-0.20	0.49	0.69	0.287
IT sec./IT running exp. (%)	12.56	5.47	9.19	9.23	-3.37	3.76	7.13	0.056
IT sec./IT run.+change exp. (%)	8.03	3.47	5.44	5.73	-2.59	2.27	4.86	0.056
Successful cyberattacks (Index)	100	100	96.77	76.04	-3.23	-23.96	-20.74	0.523
Cyberattack losses (Index)	100	100	32.53	66.42	-67.47	-33.58	33.89	0.457
Recovery time (Index)	100	100	155.82	141.55	55.82	41.55	-14.27	0.619
Detection time (Index)	100	100	74.14	110.94	-25.86	10.94	36.80	0.557
Panel B: Staffing, governance, insurance, innovation, legacy IT								
IT staff intensity (%)	8.99	9.84	11.20	12.00	2.22	2.17	-0.05	0.928
IT vacancy rate (%)	6.46	7.82	6.51	6.60	0.05	-1.22	-1.26	0.359
IT turnover index (%)	30.81	30.28	29.38	30.56	-1.43	0.28	1.70	0.788
1st LoD IT FTE share (%)	6.28	6.09	8.76	8.46	2.48	2.38	-0.10	0.904
2nd LoD IT FTE share (%)	0.03	0.03	0.04	0.02	0.01	-0.01	-0.02	0.104
3rd LoD IT FTE share (%)	0.18	0.14	0.19	0.13	0.01	-0.01	-0.01	0.543
Cyber insurance (dummy)	0.66	0.83	0.79	0.88	0.13	0.05	-0.08	0.170
Insurance deductible (Index)	100	100	508.04	156.44	408.04	56.44	-351.60	0.369
Insurance direct costs (Index)	100	100	203.02	224.95	103.02	124.95	21.93	0.959
Innovation flag (dummy)	0.94	0.90	1.00	0.88	0.06	-0.02	-0.08	0.112
Planned innovation projects (count)	16.11	20.78	20.53	23.31	4.42	2.53	-1.89	0.589
Ongoing innovation projects (count)	26.46	27.74	29.47	29.80	3.01	2.06	-0.95	0.889
Std. dev. no. EoL systems	0.15	-0.04	-0.03	-0.14	-0.18	-0.10	0.09	0.659
Log no. critical projects	2.89	2.97	3.40	3.27	0.51	0.30	-0.21	0.140
Log critical project expenditure	16.12	16.73	17.54	17.44	1.42	0.70	-0.71	0.325
Log EoL project count	1.27	1.29	1.50	1.65	0.22	0.36	0.14	0.360
Log EoL project expenditure	10.71	12.64	12.63	14.45	1.92	1.81	-0.11	0.920
Panel C: IT complexity and financials								
Log no. IT systems	6.47	6.79	7.84	7.22	1.37	0.43	-0.94	0.009
Log total assets	25.20	25.22	25.33	25.34	0.14	0.11	-0.02	0.692
Leverage ratio (%)	7.14	6.74	7.49	7.02	0.35	0.28	-0.07	0.811
ROE (%)	4.67	5.13	10.17	9.34	5.50	4.21	-1.29	0.158
Cost-to-income ratio (%)	60.18	60.74	48.26	51.36	-11.92	-9.38	2.54	0.212
CET1 capital ratio (%)	19.39	19.14	19.37	19.89	-0.02	0.75	0.77	0.145

Notes: Pre-period corresponds to 2019–2021 and post-period to 2023–2024. Laggards (L) and non-laggards (NL) are defined based on pre-CyRST underinvestment residuals. For confidentiality-sensitive variables (IT security expenditure in levels and logs, cyber incidents, cyberattack losses, detection and recovery times, and cyber insurance amounts), pre- and post-treatment means are expressed as indices normalised to 100 in the pre-period. Percentage variables are multiplied by 100; reported changes for these variables are in percentage points. Changes and associated p-values are computed from the underlying raw (non-indexed) bank-level data and are invariant to the normalisation.

Table 5 reports pre- and post-CyRST changes in cybersecurity investment and related outcomes for laggard and non-laggard banks. The largest difference is in IT security expenditure. Following the CyRST, laggard banks exhibit a larger increase in cybersecurity spending than non-laggard banks. By contrast, differences in realized cyber incidents, losses, and several operational metrics are smaller and often statistically imprecise. The table therefore provides its strongest evidence on differential investment adjustment, with more limited

evidence on other dimensions of cyber resilience.[23]

7.3 The Heterogeneous Effect of the CyRST on Investment

Table 6 presents the estimates from our main difference-in-differences specification. Across all specifications, the coefficient on the interaction term, Post × Laggard, is positive and statistically significant, providing support for Prediction 2.

Our preferred specification (Column 5) saturates the model with bank and year fixed effects alongside time-varying controls. The estimated coefficient of 0.596 implies that the CyRST announcement induced "Laggard" banks to increase their cybersecurity investment by about 81% relative to their non-laggard peers.[24]

Table 6: **The Effect of Cyber Stress Tests on Laggard Bank Investment (DiD)**

	(1)	(2)	(3)	(4)	(5)
	Dependent variable: IT Security Investment				
Post × Laggard	**0.576****	**0.576****	**0.574****	**0.595****	**0.596****
	(0.253)	(0.252)	(0.254)	(0.251)	(0.259)
Post	0.110	0.122	–	0.166*	–
	(0.117)	(0.124)	–	(0.095)	–
Laggard	-0.660	-0.632**	-0.631**	–	–
	(0.461)	(0.288)	(0.290)	–	–
Bank Controls	No	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	Yes
Bank FE	No	No	No	Yes	Yes
Observations	475	475	475	465	465

Note: This table reports difference-in-differences estimates using a Poisson Pseudo-Maximum Likelihood (PPML) model. The dependent variable is IT security investment. Column (1) reports the baseline specification. Columns (2)–(5) progressively introduce controls and fixed effects, with Column (5) representing the fully saturated model. The main "Laggard" effect is absorbed by Bank FE in Columns (4) and (5); the main *Post* effect is absorbed by Year FE in Columns (3) and (5). Bank controls include the natural logarithm of Total Assets, Leverage Ratio, ROE, Cost-to-Income Ratio, and CET1 Capital Ratio. Robust standard errors, clustered at the bank level, are reported in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 2 presents an event-study specification, providing support for the parallel trend assumption. In other words, prior to the announcement, laggards and non-laggards were on

[23] Figure A.7 in the Appendix displays the same pre–post normalized indices reported in Table 5 for confidentiality-sensitive outcomes.

[24] The magnitude of this effect is consistent with the "catch-up" effect documented in Table 5, where laggards started from a significantly lower base, bringing them closer with the resources already deployed by non-laggard peers.

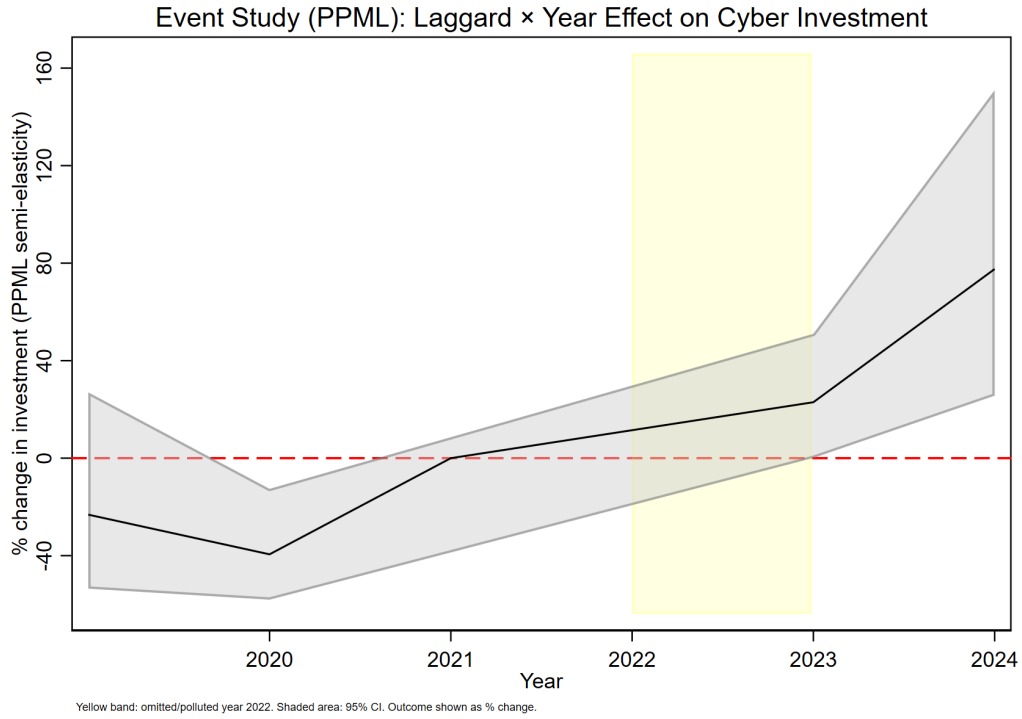

Figure 2: **Event Study Specification**

Note. This figure plots the estimated semi-elasticities of being a cybersecurity "Laggard" on IT security investment relative to the baseline year 2021 ($t = -2$), derived from a PPML event-study model including bank fixed effects and time-varying controls. Coefficients are converted to percentage changes ($100 \times (\exp(\beta) - 1)$). The statistically insignificant coefficients prior to 2023 support the parallel trends assumption, while the sharp, persistent increase from 2023 onwards illustrates the treatment effect. The shaded area represents 95% confidence intervals.

parallel investment trajectories, with the change in investment decisions starting precisely in 2023.

7.3.1 Alternative Definition of Laggards

To assess whether the baseline results depend on the median-split definition of laggards, we consider a more stringent classification based on the tails of the pre-treatment residual distribution. Specifically, we re-estimate the DiD specification by comparing banks in the bottom quartile of investment residuals (Q1) with banks in the top quartile (Q4).

Table 7 reports the results. In Column 1, the interaction term for the Q1–Q4 comparison is 1.743 and statistically significant. By contrast, Column 2 compares banks in the second quartile (Q2) with those in the third quartile (Q3). The corresponding interaction coefficient is small (0.029) and statistically insignificant. Taken together, these results suggest that the differential investment response is concentrated among the most severe laggards.

Table 7: **DiD Estimates with Alternative Laggard Definitions**

	(1) **Q1 vs. Q4** *(Extreme Laggards)*	(2) **Q2 vs. Q3** *(Moderate Laggards)*
Post × Laggard (Quartile)	**1.743*** **	**0.029**
	(0.426)	(0.153)
Bank & Year FE	Yes	Yes
All Controls	Yes	Yes
Observations	230	235

Note: This table reports PPML DiD estimates using the fully saturated specification. Column (1) defines the treatment group as banks in the bottom quartile (Q1) of residual investment, with the control group composed of banks in the top quartile (Q4). Column (2) compares banks in the second quartile (Q2) to those in the third (Q3). Standard errors, clustered at the bank level, are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Taken together with the event-study evidence in Figure 3, these results suggest that the policy did not affect all banks uniformly. Rather, the response was stronger among banks that appeared more exposed to pre-existing underinvestment.

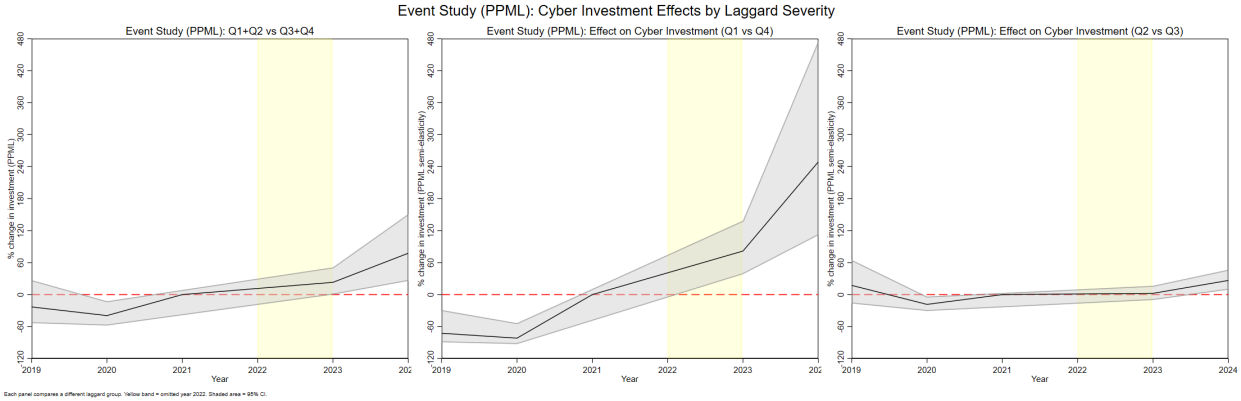

Figure 3: **Event Studies – Cyber Investment Effects by Laggard Severity**

Note. This figure presents three event study specifications estimating the dynamic effect of the CyRST on cybersecurity investment. The left panel shows the baseline median-split result. The center panel contrasts the most severe laggards (Q1) with top performers (Q4). The right panel compares moderate underperformers (Q2) with their peers (Q3). The effect is concentrated among the most extreme laggards. All specifications include bank fixed effects and a full set of controls. The shaded area represents 95% confidence intervals.

7.4 Mechanism: The Supervisory Scrutiny Channel

We next examine the mechanism underlying the observed effects. In particular, we consider whether the response reflects not only the general signalling effect of the public announcement, but also a *"scrutiny channel"*: a response to the prospect of more intensive and direct supervisory examination (Kok et al., 2023).

If this channel is at work, the investment response of "Laggard" banks should be stronger among those facing more intensive supervisory oversight. Conversely, laggards subject to less intensive supervisory engagement should display a more limited response.

7.4.1 Defining Supervisory Scrutiny

To isolate this channel, we leverage cross-sectional variation in bank-level oversight derived from the 2024 ECB Cyber Resilience Stress Test (CyRST). We operationalise supervisory intensity by constructing two indicators: the first aggregates the severity of substantive findings, while the second captures reporting skepticism through data quality and plausibility flags. These dimensions are ultimately synthesized into a composite binary treatment variable identifying banks subject to heightened regulatory pressure.[25] [26]

First, we construct a severity-weighted measure of stress test findings for each bank. We aggregate the ordinal severity scores, $s \in \{1, \ldots, 4\}$, for all findings j issued to bank i. The cumulative score, $Score_i^{Findings} = \sum_{j=1}^{N_i} s_{ij}$, represents the degree to which the supervisor identified material gaps in the bank's cyber defenses during the CyRST. Formally, we rank banks based on this severity-weighted measure and define the high-severity indicator as equal to one for those in the upper half of the distribution. We define $hsti\_sevfind_i$ as a binary indicator equal to one if the bank's aggregate score exceeds the sample median.

Second, we construct an indicator based on data quality and plausibility flags. These flags capture the supervisor's skepticism regarding the bank's internal data governance and reporting transparency. We aggregate the total number of flags issued during the CyRST validation phase. This metric identifies institutions that faced heightened supervisory friction due to reporting opacity or inconsistent risk modeling. We define $hsti\_flag_i$ as a binary indicator for banks in the upper half of the cross-sectional distribution. This measure allows us to isolate institutions that faced intensified scrutiny specifically due to concerns over reporting discipline and information asymmetry between the firm and the regulator.

Finally, we synthesize these dimensions into an overarching indicator of heightened reg-

[25] Specifically, the underlying data are generated during the active conduct of the stress test, from its formal launch on January 3, 2024 through its completion on July 26, 2024, and capture real-time supervisory interactions, targeted information requests, and evaluative assessments.

[26] Detailed construction steps for each indicator are provided in Appendix I

ulatory pressure. This composite treatment variable is defined as the logical union of the two previously constructed sub-indicators. We define $HighScrutiny_i$ as a binary variable equal to one if a bank falls into the high-intensity category for either substantive findings or reporting flags. This composite measure identifies those institutions that, due to either their inherent risk profile or lack of transparency, were subjected to the most intensive regulatory engagement. This variation allows us to separate the general effect of being stress-tested from the specific effect of being subjected to a more intense supervisory pressure.

7.4.2 Treatment Effects Conditional on Scrutiny

Table 8 examines whether the investment response of cybersecurity laggards varies with the intensity of supervisory scrutiny.

Columns (1)–(3) restrict the sample to banks classified as high scrutiny according to the composite indicator $HighScrutiny_i$. Within this group, laggard banks exhibit a large and significant increase in cybersecurity investment following the CyRST. The coefficient on the *Post* × "Laggard" interaction ranges from 0.59 to 0.64 across specifications and remains robust to the inclusion of bank-level controls and bank fixed effects. These estimates imply stronger increases in IT security spending relative to non-laggards facing comparable supervisory attention.

In contrast, Columns (4)–(6), which focus on banks outside the high-scrutiny group, show markedly smaller and generally statistically insignificant interaction effects once fixed effects are included. This pattern indicates that the CyRST does not induce a uniform adjustment among underinvesting banks, but rather exerts its influence where supervisory engagement is more intense.

Columns (7)–(9) further isolate the role of substantive supervisory assessments by restricting the sample to banks with high-severity findings. In this subsample, the estimated *Post* × "Laggard" coefficients remain large and significant, ranging between 0.67 and 0.69. By contrast, Columns (10)–(12), which restrict the sample to banks without high-severity findings, reveal no statistically significant differential investment response among laggards, despite positive point estimates.

Two additional features reinforce the interpretation of these results. First, the laggard indicator enters negatively in the pre-period across most specifications, confirming that the classification captures genuine underinvestment rather than transitory noise. Second, the standalone *Post* coefficient is generally small and insignificant once fixed effects are included, suggesting that the observed adjustments are not driven by system-wide shifts in IT spending but are concentrated among laggards facing elevated supervisory pressure.

Table 8: **The Effect of CyRST on Investment: Separating High-Scrutiny Subsamples**

	Dependent variable: IT Security Investment											
	High Scrutiny			**Low Scrutiny**			**High Sev. Findings**			**Low Sev. Findings**		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Post × Laggard	**0.590****	**0.644****	**0.641****	**0.341***	**0.346**	**0.194**	**0.670***	**0.688***	**0.671****	**0.576**	**0.623**	**0.585**
	(0.285)	(0.289)	(0.296)	(0.202)	(0.273)	(0.151)	(0.347)	(0.354)	(0.318)	(0.420)	(0.393)	(0.370)
Post	0.067	0.125	–	0.533***	0.317	–	0.221***	0.238*	–	–0.140	–0.133	–
	(0.123)	(0.132)	–	(0.152)	(0.243)	–	(0.082)	(0.132)	–	(0.350)	(0.302)	–
Laggard	–0.797	–0.647**	–	–0.124	–0.408*	–	–0.871	–0.742*	–	–0.314	–0.691	–
	(0.509)	(0.317)	–	(0.395)	(0.242)	–	(0.606)	(0.398)	–	(0.692)	(0.441)	–
Bank Controls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Year FE	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Bank FE	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
Observations	304	304	304	171	171	161	229	229	229	246	246	236

Note: This table reports PPML Difference-in-Differences estimates of the effect of the CyRST on investment, separating the sample by pre-determined scrutiny intensity. Columns (1)–(3) restrict the sample to banks with $HighScrutiny_i = 1$; Columns (4)–(6) use $HighScrutiny_i = 0$. Columns (7)–(9) restrict to banks with high-severity findings ($hsti\_sevfind_i = 1$); Columns (10)–(12) use the complement. All specifications include bank controls (log Total Assets, Leverage, ROE, Cost-to-Income, CET1). Columns (3), (6), (9), and (12) include the full set of Bank and Year fixed effects (absorbing the main levels of *Post* and *Laggard*). Robust standard errors clustered at the bank level in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

These findings are consistent with a targeted transmission mechanism. Rather than operating as a uniform regulatory shock, the CyRST had its strongest effects where supervisory interactions revealed more material deficiencies. In particular, laggard banks facing high-severity findings and more intensive follow-up exhibit the largest post-CyRST increases in cybersecurity investment. By contrast, where supervisory engagement was less intensive or focused primarily on data quality rather than substantive weaknesses, investment responses are smaller and statistically less precise. Taken together, these patterns suggest that supervisory scrutiny mattered most where prior underinvestment was greater and supervisory concerns were stronger.

7.5 The Effect of CyRST on Cyber Risk Management

This section examines whether the CyRST was associated with changes in cyber-risk management beyond monetary investment. We first report average post-announcement changes using a before–after design (Table 9). We then estimate difference-in-differences specifications for selected outcomes (Table 10), focusing on whether laggard banks exhibit differential adjustments relative to their peers.

Table 9 presents the average, unconditional adjustments in banks' operational resilience following the announcement of the ECB Cyber Resilience Stress Test (CyRST). The es-

timates suggest a possible reconfiguration of banks' ICT risk management along multiple dimensions. Most notably, banks rebalanced their outsourcing structure: payments to external (non-group) ICT service providers declined by approximately 50.1 percent, while expenditures on intra-group service providers increased by 23.9 percent. This estimate is consistent with a move toward greater internal control and governance over critical ICT services rather than a mechanical expansion of total outsourcing. Moreover, consistent with a forward-looking adjustment, banks also significantly increased planned ICT outsourcing expenditures for the subsequent budget year. This may suggest that the response extended beyond short-term compliance toward medium-term strategic planning.

Banks simultaneously accelerated the modernization of their ICT infrastructure. The number of critical systems classified as end-of-life fell by 41.2 percent, reflecting an active replacement of legacy technologies that are widely recognized as a key source of operational fragility. In parallel, banks strengthened risk mitigation and governance mechanisms. The likelihood of holding cyber insurance rose by 9.4 percent, consistent with greater use of formal risk transfer instruments, while the frequency of board-level reviews of outsourcing-related key performance indicators increased by 6 percent, consistent with heightened senior-level oversight of ICT risks. Finally, we observe a decline in staff turnover within first-line ICT functions on the order of 20.5 percent, suggesting increased retention of specialized human capital precisely in the operational areas most exposed to cyber risk.

Table 9: **Post-CyRST Adjustments in Realised Risk and Cyber Risk Management (Before–After.)**

	(1) # Signif. Attacks	(3) ICT Outs. Ext (Ref)	(4) ICT Outs. Intra (Ref)	(5) ICT Outs. Intra (Budg)	(6) ICT Outs. Ext (Budg)	(7) # Outs. Contracts	(9) # EOL Systems	(12) IT Turn. 3rd LoD	(13) Deductible (EUR)	(14) Cyber Insurance	(15) KPI Review Freq.
Post	-0.213	-0.501***	0.239***	0.334***	0.207***	0.053	-0.412***	-0.198	-0.019	0.094**	0.060*
	(0.142)	(0.017)	(0.061)	(0.050)	(0.026)	(0.067)	(0.157)	(0.184)	(0.365)	(0.046)	(0.032)
Obs.	401	475	281	289	373	475	459	359	402	412	475

Note: This table reports PPML estimates of the aggregate effect of the ECB cyber stress test announcement (*Post*) on outsourcing intensity, legacy systems, IT turnover, and governance outcomes. The specification mirrors Column 5 of Table 3. Dependent variables include: significant cyber-attacks; total and ICT outsourcing expenses (broken down by intra/extra-group and reference/budget year); number of outsourcing contracts; critical end-of-life (EOL) systems; ICT turnover indices for the three lines of defense; cyber-insurance deductible and coverage indicator; and the frequency of ICT KPI reviews. All specifications include bank-level controls (log total assets, leverage ratio, ROE, cost-to-income ratio, CET1 ratio) and fixed effects for country, business model, and bank. The table presents a reduced set of outcomes for readability; the full set of results is reported in the Appendix Table A.2. Robust standard errors clustered at the bank level are in parentheses. See Table A.4 in the Appendix for further details on the variable definitions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 10: **Post-CyRST Adjustments in Realised Risk and Cyber Risk Management (DiD)**

	(1) # Signif. Attacks	(3) ICT Outs. Ext (Ref)	(4) ICT Outs. Intra (Ref)	(5) ICT Outs. Intra (Budg)	(6) ICT Outs. Ext (Budg)	(7) # Outs. Contracts	(9) # EOL Systems	(12) IT Turn. 3rd LoD	(13) Deductible (EUR)	(14) Cyber Insurance	(15) KPI Review Freq.
Post × Laggard	-0.681**	-0.198*	0.191	0.456	-0.343**	-0.763**	-0.453	-0.692*	-1.672***	-0.169**	0.002
	(0.323)	(0.116)	(0.278)	(0.283)	(0.139)	(0.326)	(0.593)	(0.384)	(0.456)	(0.083)	(0.062)
Obs.	401	475	281	289	373	475	459	359	402	412	475

Note: This table reports PPML Difference-in-Differences estimates of the effect of the CyRST on outsourcing, legacy systems, turnover, and governance. The coefficient of interest is the interaction term *Post* × *Laggard*, where "Laggard" denotes banks with pre-treatment underinvestment. The model reproduces the specification in Column 6 of Table 6. Dependent variables are defined as in Table 9. All specifications include the full set of bank-level controls and fixed effects. The table presents a reduced set of outcomes for readability; the full set of results is reported in the Appendix Table A.3. Robust standard errors clustered at the bank level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 10 presents the difference-in-differences estimates.

First, laggard banks exhibited a significant decline in reported significant cyber-attacks relative to non-laggards. While changes in reporting incentives cannot be fully ruled out, the magnitude of the effect is consistent with improvements in cyber posture, such as enhanced detection, containment, or response protocols, that reduced the incidence or severity of reportable events.

Second, laggards undertook a restructuring of their ICT outsourcing. Relative to non-laggards, they reduced total outsourcing expenditures—covering both ICT and non-ICT services. This contraction was driven by a decline in payments to external (non-group) ICT service providers in the reference year, alongside a further reduction in planned extra-group ICT outsourcing for the subsequent budget year.

Third, the increase in investment was accompanied by lower ICT turnover across the third line of defence, indicating an effort to stabilise the institutional knowledge vital for incident response for internal auditors who independently assess IT and cyber controls.

Fourth, relative to non-laggards, laggards reduced their cyber-insurance deductibles sharply, indicating a shift toward contracts offering broader coverage per incident. At the same time, laggards were 15.5 percent less likely to hold cyber-insurance coverage on average. These estimates suggest a reconfiguration of insurance arrangements under supervisory pressure, with laggards concentrating coverage among insured institutions.

8 Conclusion

Cybersecurity in an interconnected banking system presents a classic public good problem: network externalities create strong incentives to free-ride, leading to systemic underinvestment and the persistence of critical network vulnerabilities. This paper provides, to the best of our knowledge, the first evidence that supervisory scrutiny, distinct from capital regulation

or market discipline, can influence cyber-related investment behavior in this setting.

We show that the stress test induced a significant response among "laggard" banks. These institutions increased cybersecurity investment by about 80% relative to their peers, consistent with a catch-up effect among banks that had previously invested less than comparable institutions. The response is heterogeneous and concentrated among laggards subject to more intensive supervisory scrutiny. Taken together, these patterns are consistent with supervisory scrutiny playing an important disciplining role.

Beyond monetary investment, the results are also consistent with broader operational adjustments among laggard banks. Following the CyRST, these institutions reduced some forms of external ICT dependence, showed lower turnover in specialized ICT control functions, and reconfigured aspects of their cyber-insurance arrangements. We also find a decline in reported significant cyber incidents among laggards relative to non-laggards. Taken together, these patterns suggest that the supervisory exercise was effecting in delivering changes in cyber-risk management beyond spending alone.

These findings have broader implications for banking theory and policy. In settings where risks are difficult to quantify and where investments generate cross-firm externalities, qualitative supervisory scrutiny may complement more traditional tools such as capital-based regulation and disclosure. More broadly, the results suggest that non-public supervisory assessments can affect firm behavior even when they do not rely on direct capital add-ons or immediate market discipline. Whether similar effects arise in other institutional settings remains an open question.

Authorities in other critical infrastructure sectors facing similar public-good frictions, such as energy grids, telecommunications, and supply chains, could deploy this qualitative stress-test model as a transferable regulatory technology to strengthen systemic resilience. Our findings also imply that regulation must evolve in tandem, moving beyond a historical reliance on capital requirements toward targeted, scrutiny-based interventions that directly reshape incentives.

For policymakers tasked with safeguarding macroeconomic stability, particularly in interconnected regional economies where cyber contagion can rapidly cross borders, ensuring that financial institutions internalize these operational externalities is critical. More generally, the results highlight that supervisory scrutiny can shape firm behavior even in the absence of direct financial penalties, particularly for institutions that appear to have been further from prevailing investment benchmarks prior to the intervention.

References

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.

Acharya, V. V., Berger, A. N., and Roman, R. A. (2018a). Lending implications of USs bank stress tests: Costs or benefits? *Journal of Financial Intermediation*, 34:58–90.

Acharya, V. V., Berger, A. N., and Roman, R. A. (2018b). The real effects of the bank stress tests. *The Review of Financial Studies*, 31(10):3930–3972.

Ahnert, L., Vogt, P., Vonhoff, V., and Weigert, F. (2020). Regulatory stress testing and bank performance. *European Financial Management*, 26(5):1449–1488.

Ahnert, T., Brolley, M., Cimon, D., and Riordan, R. (2024). Cyber risk and security investment. *CEPR Discussion Paper No. 17403*.

Aldasoro, I., Gambacorta, L., Giudici, P., and Thomas, L. (2023). Operational and cyber risks in the financial sector. *International Journal of Central Banking*, 19(5):341–402.

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.

Anand, K., Duley, C., and Gai, P. (2024). Cybersecurity and financial stability. *Deutsche Bundesbank Discussion Paper*.

Berger, A. N. and Bouwman, C. H. (2013). How does capital affect bank performance during financial crises? *Journal of Financial Economics*, 109(1):146–176.

Berrospide, J. M. and Edge, R. M. (2019). The effects of bank capital buffers on bank lending and firm activity: What can we learn from five years of stress-test results? *Finance and Economics Discussion Series 2019-050*.

Bonfim, D., Cerqueiro, G., Degryse, H., and Ongena, S. (2023). On-site inspecting zombie lending. *Management Science*, 69(5):2547–2567.

Buch, C. M. and DeLong, G. (2008). Do weak supervisory systems encourage bank risk-taking? *Journal of Financial Stability*, 4(1):23–39.

Calem, P., Correa, R., and Lee, S. J. (2020). Prudential policies and their impact on credit in the united states. *Journal of Financial Intermediation*, 42:100826.

Connolly, M. (2021). The real effects of stress testing. *Available at SSRN 3069376*.

Cortés, K. R., Demyanyk, Y., Li, L., Loutskina, E., and Strahan, P. E. (2020). Stress tests and small business lending. *Journal of Financial Economics*, 136(1):260–279.

Crosignani, M., Macchiavelli, M., and Silva, A. F. (2023). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *Journal of Financial Economics*, 147(2):432–448.

Cuzzola, A., Barbieri, C., and Halaj, G. (2025). Gaming the test? window-dressing and portfolio similarity around the EU-wide stress tests. *ECB Working Paper*.

Delis, M. D., Kim, S.-J., Politsidis, P. N., and Wu, E. (2021). Regulators vs. markets: Are lending terms influenced by different perceptions of bank risk? *Journal of Banking & Finance*, 122:105990.

Duffie, D. and Younger, J. (2019). *Cyber runs*. Brookings.

Eisenbach, T. M., Kovner, A., and Lee, M. J. (2022). Cyber risk and the us financial system: A pre-mortem analysis.

Flannery, M. J. (2018). Informing investors about the risks of large financial institutions. In Brandi, M. K. and John, K., editors, *Risk Topography: Systemic Risk and Macro Modeling*, pages 47–64. University of Chicago Press.

Florackis, C., Louca, C., Michaely, R., and Weber, M. (2023). Cybersecurity risk. *The Review of Financial Studies*, 36(1):351–407.

Gogolin, F., Lim, I., and Vallascas, F. (2021). Cyberattacks on small banks and the impact on local banking markets. *Available at SSRN 3823296*.

Goldstein, I. and Leitner, Y. (2018). Stress tests and information disclosure. *Journal of Economic Theory*, 177:34–69.

Gopalan, Y., Kalda, A., and Manela, A. (2021). Hub-and-spoke regulation and bank leverage. *Review of Finance*, 25(5):1499–1545.

Gropp, R., Mosk, T., Ongena, S., and Wix, C. (2019). The real effects of bank distress: Evidence from a banking crisis in Germany. *Journal of Financial Economics*, 132(1):234–254.

Hirtle, B., Kovner, A., Vickery, J., and Bhanot, M. (2016). Assessing financial stability: The capital and loss assessment under stress scenarios (class) model. *Journal of Banking & Finance*, 69:S35–S55.

Hirtle, B. and Lehnert, A. (2015). Supervisory stress tests. *Annual Review of Financial Economics*, 7(1):339–355.

Huang, J., Lin, X., Shi, X., and Zhang, S. S. (2025). Market pressure or regulatory pressure? U.S. small bank pre-emptive IT investment to data privacy regulation. *Journal of Corporate Finance*, 95:102863.

Ivanov, I. T. and Wang, J. Z. (2019). The impact of bank supervision on corporate credit: Evidence from syndicated loan reviews. *Mimeo, Federal Reserve Board, Washington, USA*.

Jamilov, R., Rey, H., and Tahoun, A. (2021). The anatomy of cyber risk. Technical report, National Bureau of Economic Research.

Kandrac, J. and Schlusche, B. (2021). The effect of bank supervision and examination on risk taking: Evidence from a natural experiment. *The Review of Financial Studies*, 34(6):3181–3212.

Kashyap, A. K. and Wetherilt, A. (2019). Some principles for regulating cyber risk. *AEA Papers and Proceedings*, 109:482–487.

Kohn, D. and Liang, N. (2019). Understanding the effects of the us stress tests. In *Federal reserve system conference: Stress testing: A discussion and review*.

Kok, C., Müller, C., Ongena, S., and Pancaro, C. (2023). The disciplining effect of supervisory scrutiny in the eu-wide stress test. *Journal of Financial Intermediation*, 53:101015.

Passalacqua, A., Angelini, P., Lotti, F., and Soggia, G. (2021). The real effects of bank supervision: evidence from on-site bank inspections. *Bank of Italy Temi di Discussione (Working Paper) No*, 1349.

Petrella, G. and Resti, A. (2013). Supervisory and market discipline in a crisis: The case of european banks. *Journal of Financial Stability*, 9(4):525–540.

Rezende, M. and Wu, J. (2014). The effects of supervision on bank performance: Evidence from discontinuous examination frequencies. In *Midwest Finance Association 2013 Annual Meeting Paper*.

Schäfer, A., Stegemann, U., and Weder di Mauro, B. (2016). The effects of the 2010 and 2011 EU-wide stress tests on bank lending. *Journal of Banking & Finance*, 69:153–167.

Schneider, T., Strahan, P. E., and Yang, J. (2023). Bank stress testing: Public interest or regulatory capture? *Review of Finance*, 27(2):423–467.

Schneider, T., Strahan, P. E., and Yang, J. (2025). Bank stress testing, human capital investment and risk management. *Journal of Financial Economics*, 171:104104.

Silva, J. M. C. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Silva, J. M. C. S. and Tenreyro, S. (2011). Further simulation evidence on the performance of the poisson pseudo-maximum likelihood estimator. *Economics Letters*, 112(2):220–222.

Varian, H. (2004). System reliability and network pricing. In Camp, L. J. and Lewis, S., editors, *Economics of information security*, pages 1–14. Springer US.

Appendix I Variable Definitions

Construction of Supervisory Scrutiny Indicators

This section details the three-step aggregation process used to construct our primary measures of supervisory scrutiny: *hsti_sevfind*, *hsti_flag*, and the composite *high_scrutiny* indicator.

1. Substantive Findings Intensity (*hsti_sevfind*)

To identify banks subject to heightened concern regarding cyber-resilience deficiencies, we construct an index based on the 2024 CyRST findings.

(a) For each bank i, we compute a weighted sum of all supervisory findings j. Each finding is assigned an ordinal severity score $s_{ij} \in \{1, 2, 3, 4\}$. The aggregate score is defined as:

$$Score_i^{Findings} = \sum_{j \in \text{Findings}_i} s_{ij} \tag{5}$$

This linear specification ensures the metric captures both the *extensiveness* (frequency) and the *severity* (criticality) of identified vulnerabilities.

(b) To ensure robustness against outliers and to facilitate a non-parametric cross-sectional comparison, we define the binary indicator using a median split:

$$hsti\_sevfind_i = \mathbb{1}\{Score_i^{Findings} > \text{Median}(Score^{Findings})\} \tag{6}$$

2. Data Quality and Reporting Integrity (*hsti_flag*)

A parallel methodology is employed to proxy for supervisory skepticism regarding a bank's internal data governance and reporting accuracy.

(a) For each bank i, we sum the total number of data quality and plausibility flags issued by supervisors during the validation phase. The bank-level reporting score is defined as:

$$Score_i^{Flags} = \sum_{q=1}^{Q_i} \text{Flag}_{iq} \tag{7}$$

where Q_i denotes the total count of flags assigned to the institution. This metric quantifies the extensiveness of supervisory concerns regarding the accuracy and consistency of the submitted data.

(b) **Classification:** We dichotomize this continuous measure based on the cross-sectional sample distribution. The binary indicator is formally defined as:

$$hsti\_flag_i = \mathbb{1}\left\{Score_i^{Flags} > \text{Median}(Score^{Flags})\right\} \tag{8}$$

where $hsti\_flag_i = 1$ identifies banks in the upper half of the distribution for reporting deficiencies, indicating a relatively high level of supervisory friction regarding data integrity.

3. Composite Supervisory Scrutiny (*HighScrutiny*)

Finally, we synthesize the substantive and reporting dimensions into a comprehensive measure of supervisory pressure. This composite indicator identifies institutions that drew significant supervisory attention through either the findings channel or the data-reporting channel. Formally, the indicator is defined as:

$$high\_scrutiny_i = \begin{cases} 1 & \text{if } hsti\_sevfind_i = 1 \text{ or } hsti\_flag_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

By construction, $HighScrutiny_i$ equals one if a bank is classified as having high-severity substantive findings OR a high volume of data-quality flags, and zero otherwise. These indicators are constructed from supervisory interactions during the active conduct of the CyRST and then held fixed in the panel for empirical convenience. cThis ensures that the treatment status is predetermined and remains exogenous to post-stress-test investment outcomes.

Core Model Variables and Controls

- **LogInv**$_{it}$: The natural logarithm of IT security investment for bank i in year t. This is the primary outcome variable in our DiD and event-study models. The underlying measure, Investment, is a bank's total annual expenditure on IT security in Euros.
- **Laggard**$_i$: A binary indicator equal to 1 if bank i was classified as a cybersecurity laggard. The classification is based on its average investment residual from the prediction model (Equation 1) over the 2020–2021 pre-treatment period, as formally defined in Equation 2.
- **Post**$_t$: An indicator equal to 1 for years 2023 and 2024, capturing the post-

announcement period of the Cyber Resilience Stress Test (CyRST). The pre-treatment period is 2019–2021.

- **Bank Controls ($\boldsymbol{X}_{it-1}$)**: A vector of lagged, time-varying bank-level controls, including:
  - `logA_TotalAssets`: Logarithm of total assets.
  - `LeverageRatio`: Tier 1 leverage ratio.
  - `ROE`: Return on equity.
  - `CIRatio`: Cost-to-income ratio.
  - `C_CET1CapitalRatio`: Common Equity Tier 1 capital ratio.
  - `log_numberitsystems`: Logarithm of the total number of IT systems.
- **Cyber Risk, Governance, and Operations**: A comprehensive set of controls from the ITRQ data source, grouped as follows:
  - *Cyber Risk Exposure and Controls:* `attack`, `attack_losses`, `itpermstaffintens`, `itvacancyrate`, `itturnindex`, `fte1lod`, `fte2lod`, `fte3lod`, `recoverytime`, `detectiontime`, `y_m_RC_DT`, `y_m_RR_DT_reb`.
  - *Cyber Insurance:* `insurance_d`, `insurancecontractsdirectcosts`, `insurancecontractsdeductibleamount`.
  - *Cyber Innovation:* `innovationprojects_d`, `innovationprojectstobeimplemented`, `innovationprojectsongoing`.
  - *Legacy Infrastructure and Risk:* `log_numbercriticalprojects`, `log_criticalprojectsexp`, `log_criticalprojectseol`, `log_criticalprojectseolexp`, `sd_numbereolsystems`, `sd_share_eol_to_be_replaced`, `sd_eol_gap_ratio`.

Appendix II Conceptual Framework

This appendix provides the formal derivations and proofs for the theoretical setup presented in Section 4.

II.1 Setup and Analysis

We model an economy populated by a finite number of risk-neutral banks, indexed by $i \in \{1, \ldots, N\}$. Each bank simultaneously chooses a continuous level of cybersecurity investment, $e_i \geq 0$. The private cost of investment is convex and given by $c(e_i, \theta_i) = \frac{e_i^2}{2\theta_i}$, where $\theta_i \in \{\theta_L, \theta_H\}$ denotes the bank's cyber type (with $0 < \theta_L < \theta_H$). The parameter θ_i captures the bank's structural efficiency in building cyber defenses; a high-type type (θ_H) can implement controls at a lower marginal cost than a low-type type (θ_L).

Let us denote Ω the system-wide resilience. We model this via a linear summation technology, $\Omega = \sum_{j=1}^{N} e_j$, where each bank derives a constant marginal benefit $b > 0$ from aggregate system resilience.[27] Note that Ω exhibits the classic features of a public good: While a bank's investment provides a private operational benefit, it also generates positive externalities for the interconnected network.

The pre-policy payoff for bank i is given by:

$$U_i(e_i, e_{-i} \mid \theta_i) = b \sum_{j=1}^{N} e_j - \frac{e_i^2}{2\theta_i}, \tag{II.1}$$

which is concave in e_i as $\left(\frac{\partial^2 U_i}{\partial e_i^2} = -\frac{1}{\theta_i} < 0\right)$. Differentiating it with respect to e_i yields:[28]

$$\frac{\partial U_i}{\partial e_i} = b - \frac{e_i}{\theta_i} = 0. \tag{II.2}$$

The equilibrium level of investment is given by $e_i^* = b\theta_i$, which constitutes the unique Nash equilibrium in pure strategies.

We compare this with the socially optimal investment level. The social planner maximises

[27] The linear public good structure is a tractable approximation that abstracts from strategic complementarities. Although payoffs depend on $\sum_j e_j$, linearity implies the marginal incentive $\partial U_i/\partial e_i = b - e_i/\theta_i$ is independent of peers' choices; hence the scrutiny shock maps cleanly into a within-bank FOC. While payoff spillovers remain, this absence of strategic interaction justifies our partial equilibrium treatment. In our empirical mapping, we interpret the estimated DiD effects as the direct (within-bank) response to the scrutiny shock, abstracting from slower-moving general equilibrium spillovers that may operate through counterparties' risk assessments or network repricing.

[28] Because the payoff is additively separable in $(e_i, \sum_{j \neq i} e_j)$, we have $\frac{\partial^2 U_i}{\partial e_i \partial e_j} = 0$; thus, there are no strategic complementarities in marginal investment choices.

aggregate welfare, $W(\mathbf{e}) = \sum_{i=1}^{N} U_i$:

$$W(\mathbf{e}) = Nb \sum_{j=1}^{N} e_j - \sum_{i=1}^{N} \frac{e_i^2}{2\theta_i}. \tag{II.3}$$

The planner internalizes that each unit of e_i yields a marginal payoff b to all N banks. The planner's FOC for bank i is:

$$Nb - \frac{e_i}{\theta_i} = 0 \implies e_i^{SO} = Nb\theta_i. \tag{II.4}$$

Comparing the private equilibrium with the social optimum yields $e_i^* = \frac{1}{N} e_i^{SO}$. Furthermore, since $\theta_L < \theta_H$, we have $e_L^* = b\theta_L < b\theta_H = e_H^*$.

Because each bank fails to internalise the positive externality its investment confers on the remaining $N - 1$ institutions, there is underinvestment relative to the social planner's optimum. Moreover, since investments are increasing monotonically in θ_i, low-type banks (θ_L) emerge as "laggards," investing strictly less than their high-type peers ($e_L^* < e_H^*$).

II.2 Supervisory Intervention

We now consider a stylised supervisory intervention where the supervisor implements a targeted cyber stress-testing exercise (i.e., the CyRST). The supervisor acts as a constrained planner: lacking the statutory authority to impose a Pigouvian tax or direct Pillar 2 capital add-ons, the supervisor sets a qualitative benchmark $\bar{e}$, based on observable resilience controls and remediation effort, to induce a higher investment effort by such banks.[29] We assume the supervisor targets such a benchmark.

We microfound the cost of being flagged for falling short of this benchmark through a dynamic, forward-looking scrutiny channel. Suppose the interaction extends to a second period. Banks flagged for underinvestment ($e_i < \bar{e}$) face intense supervisory follow-up and elevated funding spreads from wholesale markets that reflect heightened perceived operational fragility. We abstract from formally modeling the second-period belief updating process and

[29] For instance, a natural benchmark is $\bar{e} = \kappa Nb\theta_L \in (b\theta_L, b\theta_H)$ for some $\kappa \in (1/N, \min\{1, \theta_H/(N\theta_L)\})$, which corresponds to requiring laggards to move part-way toward the planner provision while leaving high types unaffected.

treat R as the non-compliance cost.[30] The bank's post-policy objective function becomes:

$$U_i^{post}(e_i, e_{-i}, \theta_i) = b\sum_{j=1}^{N} e_j - \frac{e_i^2}{2\theta_i} - R \cdot \max\{0, \bar{e} - e_i\}.$$

We assume that R is sufficiently large that it is always optimal for laggards to comply.

For $e_i \geq \bar{e}$, the unconstrained effort level is already above the required level. Therefore, $e_H^{\star\star} = b\theta_H$, which is equal to the pre-policy level. Since R is large enough for low-types, the optimal effort level is constraiend to be equal to $\bar{e}$. Therefore $e_L^{\star\star} = \bar{e} > b\theta_L$.

[30]For example, if the incremental risk premium Δs on wholesale funding F scales with the severity of the identified gap, this implies an effective marginal reputational cost $R = \beta F \Delta s$ per unit of shortfall. We normalize units so that R is measured in payoff per unit of shortfall in e.

Appendix III Tables and Figures

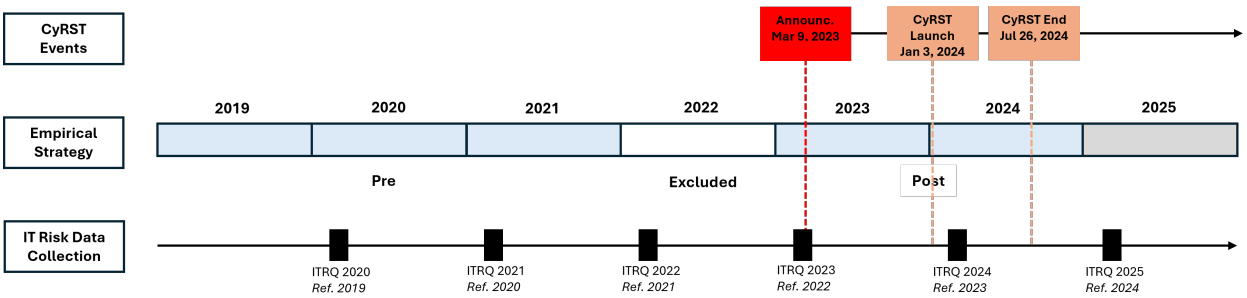

Figure A.1: **Timeline.**
This figure illustrates the temporal structure of our identification strategy. The pre-treatment period (2019–2021) is followed by the exclusion of 2022. This is done because the 2022 observation is collected in Q1 2023, overlapping with the public announcement of the Cyber Resilience Stress Test (CyRST) in March 2023, during which banks had heterogeneous submission timing of the ECB IT Risk Questionnaire (ITRQ). As some banks submitted their 2022 responses prior to the announcement while others responded afterward, the 2022 data point is contaminated by a mixture of pre- and post-announcement information. The post-treatment period begins with the CyRST policy announcement in March 2023, and covers the launch (January 2024) and completion (July 2024) of the CyRST exercise. The lower panel reports the annual delivery of the ECB's IT Risk Questionnaire (ITRQ), which provides the reference-year inputs for our outcome variables.

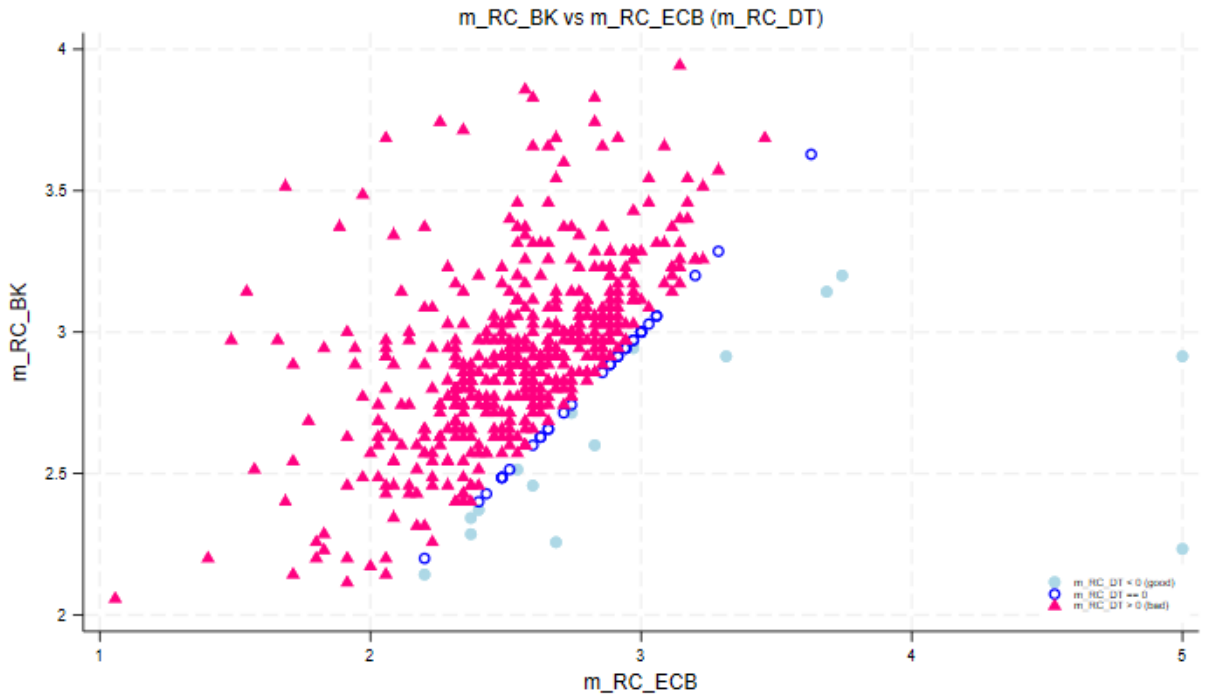

Figure A.2: **Supervisory Misalignment in Risk Control Perception: Risk Control Misalignment.**
This scatterplot compares banks' average self-assessed cyber risk control scores (m_RC_BK) with those assessed by the ECB (m_RC_ECB). Both scores are constructed by reversing original risk control indicators (from 1 = strong to 4 = weak) to ensure higher values denote stronger control setups, and then averaged across all relevant items (the ITRQ collects 35 IT Risk Control Sub-Scores). Points above the 45-degree line reflect banks that perceive their controls to be stronger than the ECB does ($m\_RC\_DT > 0$), signaling potential overconfidence or governance opacity. Marker shapes and colors denote overconfident (pink triangles), underconfident (light blue dots), and aligned (blue circles) bank-year observations.

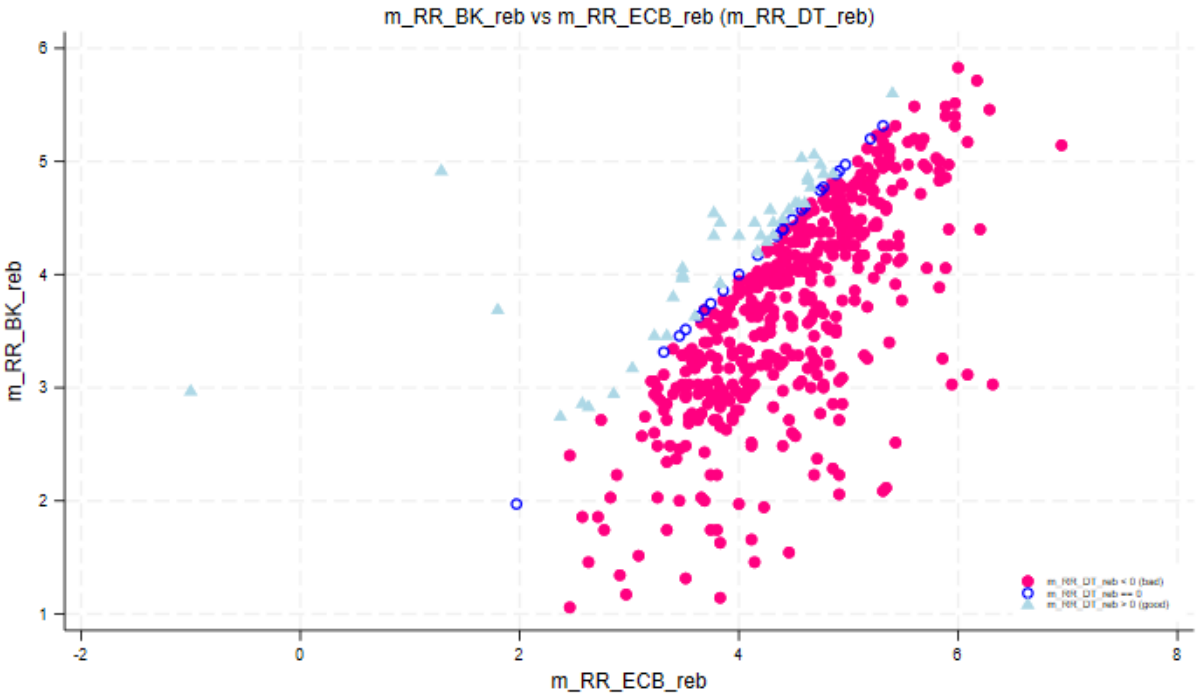

Figure A.3: **Supervisory Misalignment in Risk Level Perception: Risk Level Misalignment.**
This scatterplot depicts banks' perceived residual IT risk (*m_RR_BK_reb*) against the ECB's assessment (*m_RR_ECB_reb*). Residual risk is defined as the difference between perceived risk level and risk control, rescaled to range from 0 (low residual risk) to 8 (high residual risk). The variable *m_RR_DT_reb* captures the distance between these two risk perceptions. Observations where $m\_RR\_DT\_reb > 0$ indicate banks that underestimate their residual cyber risk relative to the ECB's view. These gaps are central to our empirical strategy as proxies for governance misalignment.

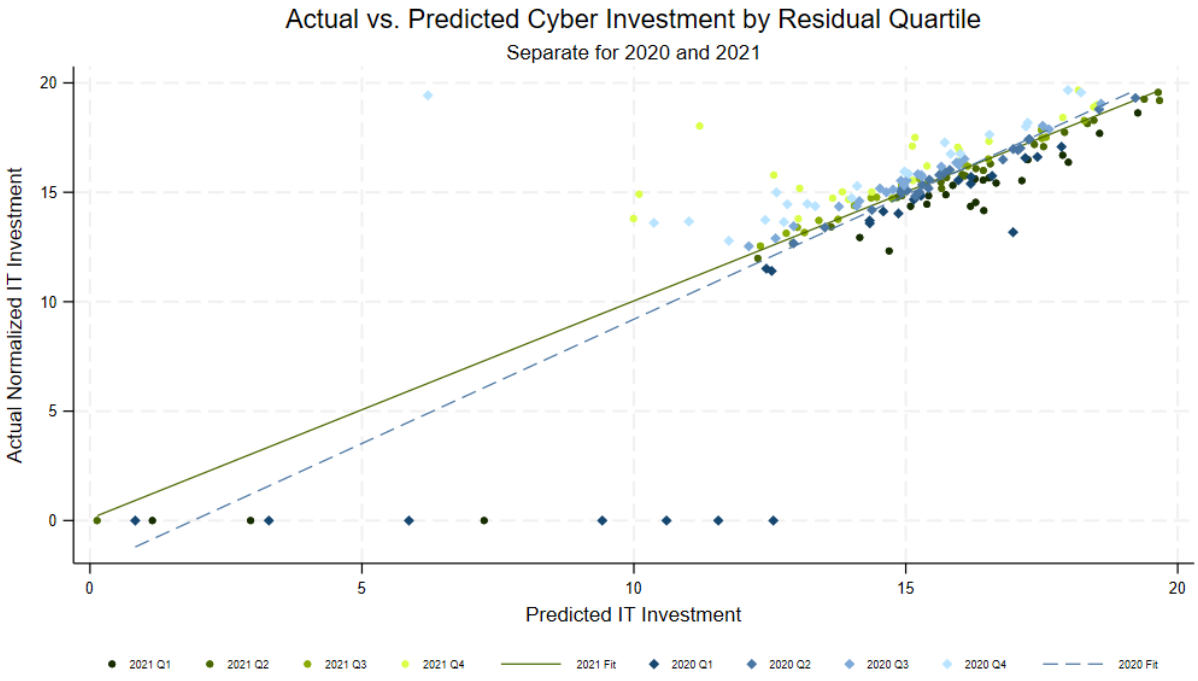

Figure A.4: **Residuals normalized IT security investment and predicted IT security investment.**
Actual vs. predicted normalized IT investment for European banks, split by quarter in 2020 (blue, dashed fit) and 2021 (green, solid fit). Each point shows a bank-residual quartile; the fit lines indicate the model's predicted cyber investment accuracy by year, highlighting residual variation that defines under- or over-investing banks.

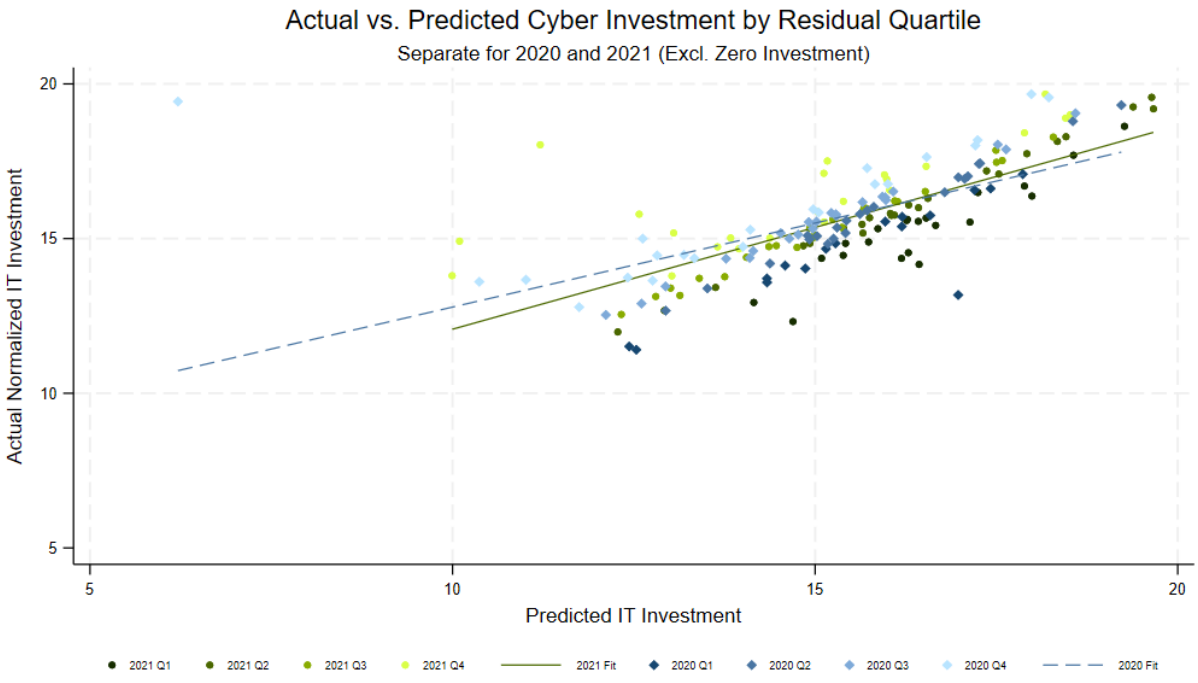

Figure A.5: **Actual and Predicted Cyber Security Investment by Residual Quartile.**

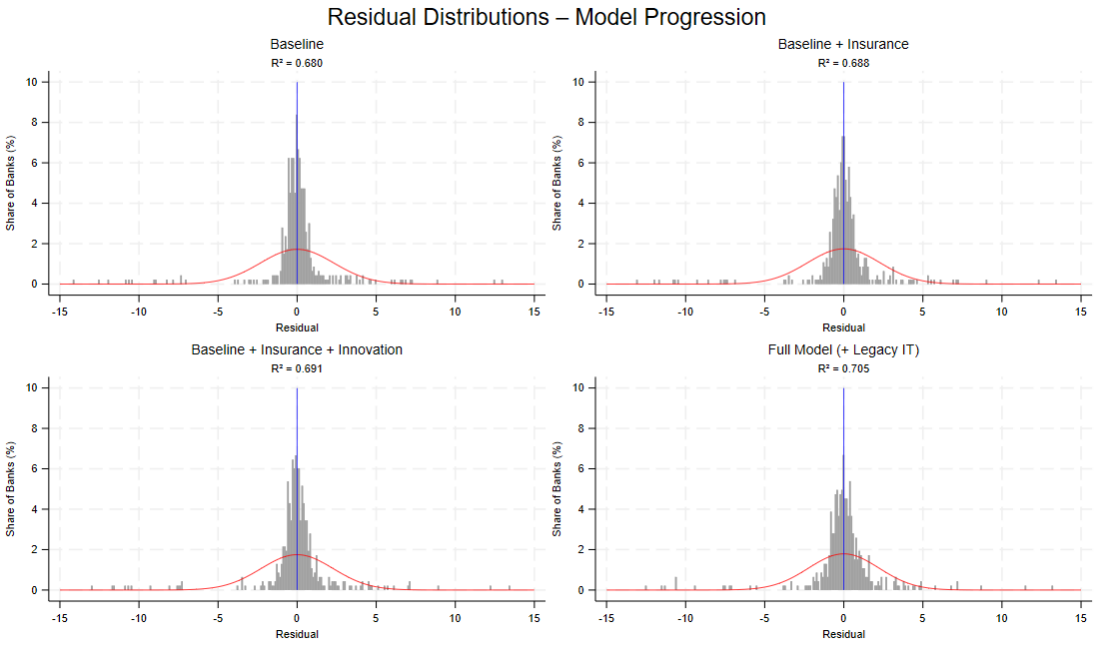

Figure A.6: **Distribution of Residualized Cybersecurity Investment**

Note. Each panel displays the histogram of residuals from the benchmark model predicting IT security investment. The panels sequentially add controls from left to right. The narrowing dispersion illustrates the model's ability to partial out observable variation, supporting the identification of cybersecurity laggards for causal analysis.

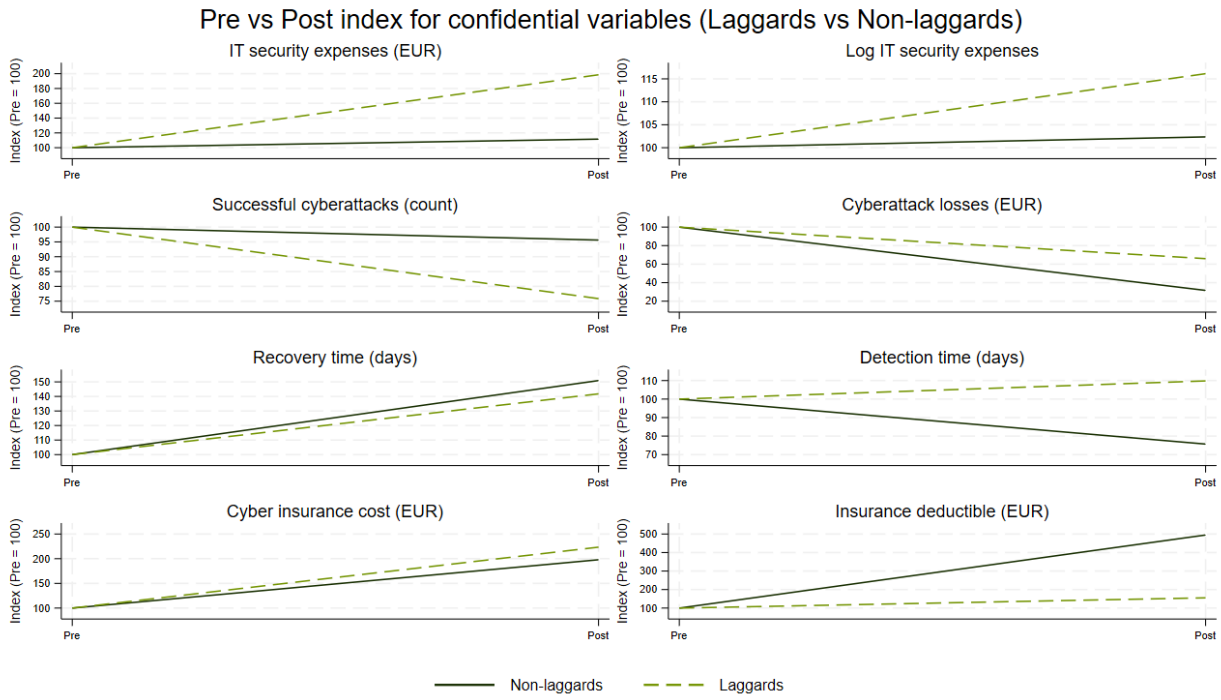

Figure A.7: **Pre- vs. post-CyRST evolution of confidential variables (indexed to pre-period = 100).** This figure presents confidentiality-compliant summaries of key cyber risk and cyber insurance variables for which raw values cannot be disclosed. For each confidential variable, we compute the bank-level mean in the pre-CyRST window (2019–2021) separately for laggard and non-laggard institutions, normalize this value to 100, and express the corresponding post-CyRST mean (2023–2024) as an index relative to that pre-period baseline. The resulting multi-panel figure therefore displays relative changes, rather than levels, between laggards and non-laggards. This approach ensures that supervisory-sensitive information remains hidden while preserving economically meaningful cross-group and over-time comparisons. Laggards are defined ex ante based on systematic underinvestment in cybersecurity relative to model-predicted benchmarks.

Table A.1: **Determinants of Normalised Cybersecurity Investment**

Variable	**Coefficient**	**(Std. Error)**
Cyber Risk and Exposure		
Number of cyberattacks	0.098	(0.065)
Log attack losses	–0.019	(0.018)
Staffing and Governance		
IT staffing intensity	0.452	(7.879)
IT vacancy rate	–1.657	(3.423)
IT turnover index	0.060	(0.397)
IT FTE 1st line of defense	–3.885	(4.302)
IT FTE 2nd line of defense	27.179	(111.643)
IT FTE 3rd line of defense	–109.904*	(61.708)
Preparedness and Misalignment		
Recovery time	–0.000	(0.001)
Detection time	0.002	(0.001)
Risk control misalignment	1.532	(20.059)
Residual risk misalignment	0.329	(18.958)
Insurance		
Insurance coverage (dummy)	1.095	(0.794)
Insurance direct cost	–0.000	(0.000)
Insurance deductible amount	0.000	(0.000)
Innovation		
Innovation project (dummy)	–1.376*	(0.796)
Projects to be implemented	0.002	(0.006)
Ongoing innovation	–0.001	(0.005)
Legacy Infrastructure Risk		
EOL systems (std. dev.)	–0.587	(0.598)
Log number of critical projects	0.721	(0.488)
Log critical project expenditure	–0.016	(0.092)
Log EOL systems	–0.190	(0.301)
Log EOL project spend	0.053	(0.045)
Financial and Scale Controls		
Log total assets	1.153	(1.790)
Leverage ratio	30.471**	(13.574)
Return on equity (ROE)	3.239	(3.584)
Cost-income ratio	–1.204	(1.899)
CET1 capital ratio	–16.075*	(9.288)
Log number of IT systems	0.136	(0.147)
Constant	–15.091	(42.988)
Observations	464	
R-squared	0.706	

Notes: This table reports estimates from the first-stage OLS model used to generate investment residuals, estimated on the 2019–2021 panel. The dependent variable is the natural logarithm of IT security expenditure. All specifications include bank and business model fixed effects (not reported). Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.10.

Table A.2: **Post-CyRST Adjustments in Realised Risk and Cyber Risk Management (Before-After)**

	(1) # Signif. Attacks	(2) Tot. Outs. Expenses	(3) ICT Outs. Ext (Ref)	(4) ICT Outs. Intra (Ref)	(5) ICT Outs. Intra (Budg)	(6) ICT Outs. Ext (Budg)	(7) # Outs. Contracts	(8) # ICT Contracts	(9) # EOL Systems	(10) IT Turn. 1st LoD	(11) IT Turn. 2nd LoD	(12) IT Turn. 3rd LoD	(13) Deductible (EUR)	(14) Cyber Insurance	(15) KPI Review Freq.
Post	**-0.213**	**0.309****	**-0.501*****	**0.239*****	**0.334*****	**0.207*****	**0.053**	**-0.134**	**-0.412*****	**-0.205***	**0.045**	**-0.198**	**-0.019**	**0.094****	**0.060***
	(0.142)	(0.154)	(0.017)	(0.061)	(0.050)	(0.026)	(0.067)	(0.089)	(0.157)	(0.113)	(0.250)	(0.184)	(0.365)	(0.046)	(0.032)
Obs.	401	475	475	281	289	373	475	475	459	441	333	359	402	412	475

Note: This table reports PPML estimates of the aggregate effect of the ECB cyber stress test announcement (*Post*) on outsourcing intensity, legacy systems, IT turnover, and governance outcomes. The specification mirrors Column 5 of Table 3. Dependent variables include: significant cyber-attacks; total and ICT outsourcing expenses (broken down by intra/extra-group and reference/budget year); number of outsourcing contracts; critical end-of-life (EOL) systems; ICT turnover indices for the three lines of defense; cyber-insurance deductible and coverage indicator; and the frequency of ICT KPI reviews. All specifications include bank-level controls (log total assets, leverage ratio, ROE, cost-to-income ratio, CET1 ratio) and fixed effects for country, business model, and bank. Robust standard errors clustered at the bank level are in parentheses. See Table A.4 in the Appendix for further details on the variable definitions. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.3: **Post-CyRST Adjustments in Realised Risk and Cyber Risk Management (DiD)**

	(1) # Signif. Attacks	(2) Tot. Outs. Expenses	(3) ICT Outs. Ext (Ref)	(4) ICT Outs. Intra (Ref)	(5) ICT Outs. Intra (Budg)	(6) ICT Outs. Ext (Budg)	(7) # Outs. Contracts	(8) # ICT Contracts	(9) # EOL Systems	(10) IT Turn. 1st LoD	(11) IT Turn. 2nd LoD	(12) IT Turn. 3rd LoD	(13) Deductible (EUR)	(14) Cyber Insurance	(15) KPI Review Freq.
Post × Laggard	**-0.681****	**-0.417***	**-0.198***	**0.191**	**0.456**	**-0.343****	**-0.763****	**-0.323**	**-0.453**	**-0.178**	**0.141**	**-0.692***	**-1.672*****	**-0.169****	**0.002**
	(0.323)	(0.239)	(0.116)	(0.278)	(0.283)	(0.139)	(0.326)	(0.264)	(0.593)	(0.246)	(0.459)	(0.384)	(0.456)	(0.083)	(0.062)
Obs.	401	475	475	281	289	373	475	475	459	441	333	359	402	412	475

Note: This table reports PPML Difference-in-Differences estimates of the effect of the CyRST on outsourcing, legacy systems, turnover, and governance. The coefficient of interest is the interaction term *Post* × *Laggard*, where "Laggard" denotes banks with pre-treatment underinvestment. The model reproduces the specification in Column 6 of Table 6. Dependent variables are defined as in Table 10. All specifications include the full set of bank-level controls and fixed effects. Robust standard errors clustered at the bank level are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.4: **Description of Dependent Variables in Section 7.5**

No.	Variable (as in tables)	Definition
(1)	# significant attacks	Number of significant cyber incidents formally reported during the year.
(2)	Total outsourcing expenses	Total ICT and non-ICT outsourcing spending in the reference year.
(3)	ICT outsourcing expenses, extra-group (ref.)	Payments to external (non-group) ICT service providers during the reference year.
(4)	ICT outsourcing expenses, intra-group (ref.)	Payments to intra-group ICT service providers during the reference year.
(5)	ICT outsourcing expenses, intra-group (budget)	Planned intra-group ICT outsourcing spending for the next budget year.
(6)	ICT outsourcing expenses, extra-group (budget)	Planned extra-group ICT outsourcing spending for the next budget year.
(7)	# outsourcing contracts	Total number of ICT and non-ICT outsourcing arrangements.
(8)	# ICT outsourcing contracts	Number of ICT-specific outsourcing arrangements.
(9)	# EOL systems	Number of critical ICT systems assessed as end-of-life (EOL).
(10)	IT turnover, 1st LoD	Staff turnover ratio in ICT functions within the first line of defence.
(11)	IT turnover, 2nd LoD	Staff turnover ratio in ICT functions within the second line of defence.
(12)	IT turnover, 3rd LoD	Staff turnover ratio in ICT functions within the third line of defence.
(13)	Deductible (EUR)	Monetary deductible applied per cyber insurance claim.
(14)	Cyber insurance	Indicator variable equal to 1 if the bank holds cyber insurance coverage.
(15)	KPI review frequency	Frequency with which outsourcing-related KPIs and risks are reviewed.