Disciplining Digital Risk: Evidence from Cyber Stress Tests

Nordine Abidi* Leonardo Gambacorta † Christoffer Kok ‡

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

August 29, 2025

Preliminary Draft — Please do not share without the authors' authorization

Abstract

Investment in cybersecurity within the interconnected banking sector exhibits public good properties, where positive externalities can lead to systemic underinvestment. Using confidential supervisory data from the European Central Bank (ECB), we analyze the 2024 ECB Cyber Resilience Stress Test (CyRST), a novel supervisory initiative with no capital implications and no public disclosure, as a quasi-natural experiment. We follow a two-step procedure. First, we identify "laggard banks" as those banks that systematically underinvest relative to their cyber risk profile. Second, using a difference-in-difference approach, we show that the CyRST induced laggard banks to increase their cybersecurity investments by 34% relative to their non-laggard peers. This response is driven by a "scrutiny channel," as the effect is mainly concentrated among laggards under high-intensity supervision, who exhibit an additional investment increase of 57%. Our findings provide the first causal evidence that targeted supervisory scrutiny is an effective tool for resolving coordination failures and disciplining cyber risk.

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. 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 Márton Barta for his valuable research assistance. *Disclaimer:* This paper should not be reported as representing the views of the European Central Bank (ECB), the International Monetary Fund (IMF), or the Bank for International Settlements (BIS). The views expressed are those of the authors and do not necessarily reflect those of the ECB, the IMF or the BIS. The dataset employed in this Working Paper contains confidential statistical information. Its use for the purpose of the analysis 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 principal operational and, increasingly, systemic threat to the 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. This resonates with classic models of financial contagion (Allen and Gale, 2000; Acemoglu et al., 2015), yet cyber risk introduces unique challenges. The resilience of the financial system may be dictated by its "weakest link," where under-defended institutions serve as 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, as shown by the Google search dynamics in Figure 1.

While cybersecurity has strong private good elements, its systemic dimension, where one bank's failure can create an externality for the entire network, exhibits the features of a classic public good problem. A bank's cybersecurity investment creates positive externalities for all its counterparties. Consequently, banks have incentives to underinvest and free-ride on the efforts of others (see, e.g., Kashyap and Wetherilt, 2019; Anand et al., 2024), a market failure that provides a strong rationale for regulatory intervention.[3]

This paper provides, to the best of our knowledge, the first causal evidence that *targeted supervisory scrutiny*, implemented through a non-capital-based stress test, can effectively discipline under-investment and correct this coordination failure. We analyze the European Central Bank's 2024 *Cyber Resilience Stress Test* (CyRST), a novel exercise designed to assess a bank's ability to respond to and recover from a sophisticated cyberattack.[4] The unique characteristics of the CyRST create a quasi-natural experiment.[5] First, the exercise was purely qualitative, with no direct implications for Pillar 2 capital requirements.[6] This neutralizes the "capital channel" common in traditional stress tests (e.g., Acharya et al.,

[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] See, e.g., World Economic Forum (2024), "Global Cybersecurity Outlook 2024," which highlights the persistent gap between cybersecurity needs and budget allocation. Formally, Bouveret (2018) models this as an agency problem within the firm.

[4] The initiative was a direct outcome of the ECB's 2024-2026 Supervisory Priorities, which established strengthening banks' operational resilience as a key strategic objective. See *ECB Banking Supervision: Supervisory Priorities for 2024-2026*.

[5] The CyRST was first signaled to the public on March 9, 2023, by the Chair of the ECB Supervisory Board, which serves as our primary treatment announcement (see Section 3). This announcement was widely reported by major financial news outlets, serving as a key public signal of supervisory intent.

[6] 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.

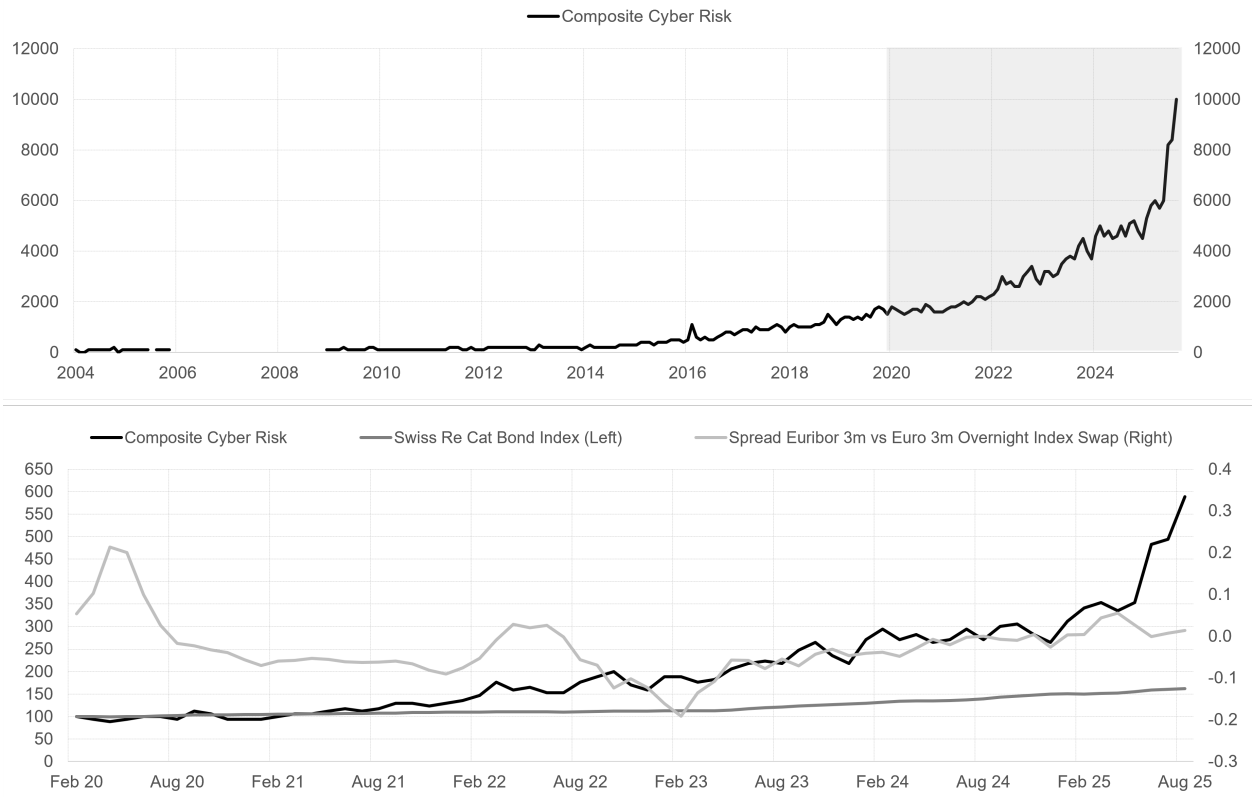

Figure 1: **Public Attention to Cyber Risk and Cybersecurity.**
The upper panel presents normalized Google Trends search activity for cybersecurity-related topics from January 2004 to July 2025. The solid black line displays a composite cyber risk index—averaging three normalized series: "cyber" "cyber incidents", and "cyber risk". A vertical line marks March 2023, corresponding to the public announcement of the cyber stress test initiative. The lower panel plots the indexed performance of the Swiss Re Cat Bond Index (left), alongside the Euribor–OIS 3-month spread (right), from February 2020 to August 2025. The series reflect investor sentiment toward insurance-linked securities and interest rate risk, respectively.
Source: Google Trends and Bloomberg, (authors' calculations).

2018b; Gropp et al., 2019). Second, 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 allows us to isolate a third, less-studied channel: the "scrutiny channel," where the credible threat of direct examination compels behavioral change (Kok et al., 2023).

To guide our empirical work, we develop a stylized theoretical framework that shows how targeted scrutiny can shift the banking system from a vulnerable, low-investment equilibrium to a Pareto-superior, high-investment one. The model delivers two important and testable predictions: the policy announcement leads to an increase in aggregate investment, and this effect is driven by banks that were underinvesting pre-policy, which we term "laggards."

To test these predictions, we implement a two-step methodology using a unique confidential supervisory dataset. First, using only pre-treatment data, we model the expected level of cybersecurity investment for each bank, conditional on a rich vector of its characteristics. We then classify banks as *laggards* if their average investment residual over the 2020-2021 period falls below the median of the sample distribution. Second, we exploit the 2023 public announcement of the CyRST in a difference-in-differences framework.

Our main empirical specification employs a Difference-in-Differences approach in the context of a Poisson Pseudo-Maximum Likelihood (PPML) model, which is robust to the skewed, non-negative nature of investment data. Our results show that the CyRST announcement acted as a coordinating signal, inducing an average investment increase of around 45% across all banks. We find that this effect is driven by laggard banks, which increased their cybersecurity investments by a statistically and economically significant 34% relative to their non-laggard peers. Exploiting variation in oversight intensity within the CyRST design, our triple-difference estimate reveals that this response is concentrated among laggards subjected to high-intensity supervision, who exhibit an additional investment increase of 57%. This provides strong evidence for the scrutiny channel.

The credibility of our findings is confirmed by a battery of robustness tests that rule out alternative explanations. Our main coefficient of interest remains stable when we include a rich set of additional bank-level controls and demanding fixed effects. The results are also robust to using alternative measures of supervisory scrutiny and more stringent definitions of laggard banks. Furthermore, placebo tests show no effect in non-treatment periods, and we provide evidence that the policy induced a persistent shift in investment behavior. Overall, our results remain very similar, which gives us additional confidence that we are picking up the right effect.

This paper makes three main contributions. First, we provide the first causal evidence on the effects of a cyber stress test. This unique setting allows us to isolate a pure "scrutiny channel," showing it is a strong driver that compels under-investing "laggard" banks to reallocate resources toward cyber resilience. Second, our findings offer direct implications for regulators, showing that targeted, qualitative supervision is a powerful tool for enhancing systemic stability in the face of growing cyber-related risks. Third, we offer empirical evidence to the work of Anand et al. (2024), which state that regulatory intervention is necessary to move the financial system from a fragile, low-investment equilibrium to a more resilient one.

The remainder of the paper is organized as follows. Section 2 presents the related literature. Section 3 provides an overview 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 provides concluding remarks.

2 Related Literature

This paper connects to three active streams of research: the economics of cybersecurity as a systemic financial risk, the real effects of bank stress testing, and the broader impacts of banking supervision. We contribute to each by providing the first evidence, to the best of our knowlege, on how targeted, non-capital-based supervisory scrutiny can discipline under-investment in a critical non-financial risk area: cyber.

Cybersecurity as a Systemic Financial Risk. Our paper contributes primarily to the literature that conceptualizes cybersecurity not as an idiosyncratic IT issue, but as a first-order threat to financial stability. A burgeoning body of research documents the channels through which cyber shocks propagate, identifying direct operational linkages, such as payment systems (Eisenbach et al., 2022), and indirect contagion through corporate supply chains (Crosignani et al., 2023). These disruptions can be amplified by second-round effects on confidence, potentially triggering liquidity events that transform an operational incident into a financial crisis (Duffie and Younger, 2019; Gogolin et al., 2021).

This literature converges on a core market failure: cybersecurity in the financial sector exhibits strong public good characteristics. Anand et al. (2024), in particular, provide a relevant theoretical framework, modeling the sector as a network where the incentive to free-ride on others' security investments can lead to a fragile, low-investment equilibrium. Complementing this view, Ahnert et al. (2024) model how under-investment also arises from a principal-agent problem, where the unobservability of security investment by clients gives firms an incentive to under-invest. While this work provides a robust theoretical rationale for regulatory intervention, the effectiveness of specific policy tools remains an open empirical question. Our paper provides the first test of this framework's predictions in the context of a major supervisory intervention. We empirically validate the existence of under-investing firms (our "laggards") and show that targeted supervision can, as their model suggests, act as a coordination device to shift the system to a more resilient state.

The Real Effects of Stress Testing. Second, we inform the extensive literature on the real effects of bank stress tests. This research has traditionally focused on large-scale, disclosure-based exercises with direct capital consequences, such as the U.S. Comprehensive Capital Analysis and Review (CCAR). A parallel and extensive body of work has examined the effects of the EU-wide stress tests, similarly finding 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). Seminal studies find that these programs induce banks to de-risk their balance sheets, adjust lending policies, and increase capitalization (e.g., Acharya et al., 2018b; Goldstein and Leitner, 2018). Our work departs from this literature by analyzing a fundamentally 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 unique institutional setting allows us to isolate the "scrutiny channel." The efficacy of such non-capital tools is supported by recent theory; for example, Ahnert et al. (2024) show that imposing minimum investment standards or firm liability for breaches can resolve under-investment problems. While Kok et al. (2023) find that the supervisory scrutiny embedded within traditional stress tests has a disciplining effect, their analysis cannot disentangle this channel from the simultaneous threat of capital add-ons and market reactions. We show that the scrutiny channel is strong enough to shift firm behavior even in the absence of such threats.

The Real Effects of Banking Supervision. Finally, our research contributes to the literature examining how day-to-day banking supervision shapes the internal operations of financial institutions. Prior work has shown that regulatory scrutiny can prompt tangible organizational changes, such as the hiring of more skilled risk-management personnel in response to poor supervisory ratings (Schneider et al., 2025). Our paper provides supporting evidence for this channel and traces the bank-level response to regulatory pressure in granular detail, quantifying a large investment response to a purely qualitative supervisory exercise. In doing so, we provide an empirical counterpart to theoretical models like Ahnert et al. (2024), where such supervisory mandates are shown to directly correct the market failures that lead to under-investment in security.

3 Institutional Framework

This section details the institutional design of the European Central Bank's (ECB) 2024 Cyber Resilience Stress Test (CyRST). This initiative was a direct outcome of the ECB's 2024–2026 Supervisory Priorities, which, in response to the rapid digitalization of the banking sector, established strengthening banks' operational resilience as a key strategic objective.[7]

The key objective of the CyRST was to assess a bank's ability to *respond to and recover from* a sophisticated cyberattack, a deliberate shift in focus from prevention to resilience. The

[7]See ECB Banking Supervision (2023), "ECB supervisory priorities for 2024–2026," December; and ECB (2024), "ECB concludes cyber resilience stress test," Press Release, 26 July 2024.

exercise simulated a disruptive scenario in which an attacker compromises a bank's critical IT systems, inducing the institution to activate its emergency protocols and demonstrate it could restore core operations from backups within a specified timeframe. The test was designed to isolate the effects of pure *supervisory scrutiny* by explicitly neutralizing the two other channels common in stress testing:

1. *No Direct Capital Channel:* The CyRST was a purely qualitative exercise. It did not involve calculations of capital depletion or a pass/fail threshold, and the ECB gave explicit public assurances that there would be no *direct* impact on bank-specific Pillar 2 capital guidance.[8] [9]

2. *No Market Discipline Channel:* No bank-level results or rankings were publicly disclosed. The ECB released only a high-level, anonymized summary of aggregate findings, shielding individual institutions from market reactions. However, the confidential findings for each bank were incorporated into its annual Supervisory Review and Evaluation Process (SREP). This is the crucial channel for enforcement; significant deficiencies identified in the CyRST could lead to future supervisory measures, including higher Pillar 2 Requirements (P2R), thereby making the exercise a credible threat.[10]

The exercise included all 109 Significant Institutions (SIs), the largest banking groups in the euro area under the ECB's direct supervision. All participants completed a granular, 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 factors including systemic importance and business model diversity. This feature represents the key source of variation for our mechanism tests.

Banks in the enhanced-assessment cohort were subjected to significantly more intensive and intrusive supervisory oversight. While the standard assessment relied on the bank's own attestations, the enhanced assessment involved direct verification by supervisory teams. This included On-Site Quality Assurance Reviews (OSQARs), supervisory deep dives to

[8] ECB (2024), "ECB to stress test banks' ability to recover from cyberattack," Press Release, 3 January 2024. The official statement reads: "This predominantly qualitative exercise will not have an impact on capital through the Pillar 2 guidance."

[9] Importantly, while the CyRST was explicitly designed as a qualitative exercise with no direct capital implications or pass/fail outcomes, the confidential supervisory findings were integrated into the annual SREP. This implies that, although there was no mechanical capital depletion channel, the exercise could have had marginal second-order effects via adjustments in the operational risk element of Pillar 2 Requirements (P2R). Our empirical analysis focuses on the direct scrutiny channel, but we acknowledge this indirect pathway.

[10] The SREP is the core process of European banking supervision, where the ECB assesses a bank's strategies, processes, and risks, and decide on any necessary supervisory measures, such as firm-specific capital requirements (P2R and P2G) and other qualitative actions.

validate internal procedures, and a live IT recovery test, which required banks to physically demonstrate their ability to restore critical systems from backups.

The CyRST was first signaled to the public on March 9, 2023, by the Chair of the ECB Supervisory Board, which serves as our primary treatment announcement.[11] The full timeline is detailed in Figure 2.

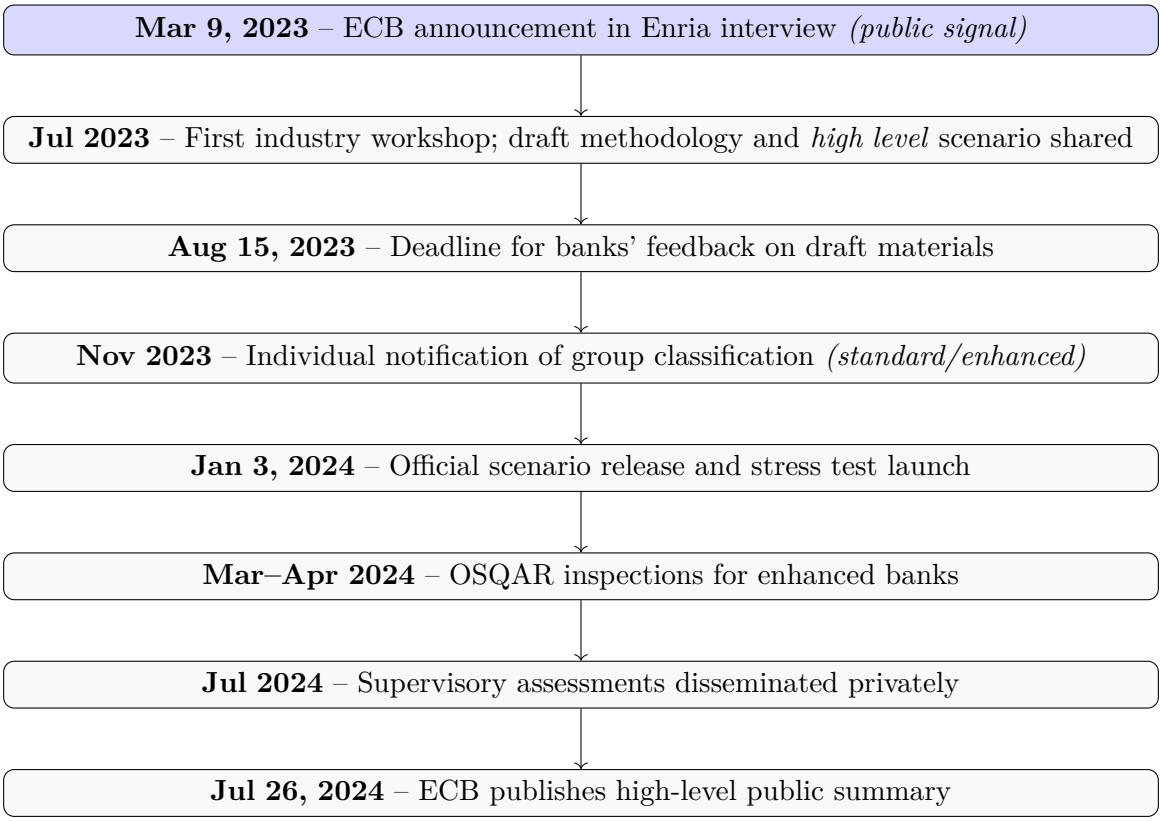

Figure 2: **Timeline of the ECB 2024 Cyber Resilience Stress Test**

This timeline underpins our identification strategy. The sequence of events—namely, the CyRST policy announcement in March 2023, the confidential notification of high-intensity scrutiny assignments in November 2023, and the formal launch of the CyRST in January 2024 with its associated supervisory engagement, creates a well-defined quasi-experimental setting. It enables us to empirically separate the market-wide behavioral adjustments following the general announcement from the targeted responses of banks subjected to the most intrusive supervisory examination.

[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.

4 Conceptual Framework

This section develops a stylized theoretical framework that provides microfoundations for our empirical analysis by formalizing the coordination failure inherent in cybersecurity investment. Our analysis of strategic interaction draws upon several key ideas: the conceptualization of systemic cyber risk as a public good problem with free-rider incentives, as in Anand et al. (2024); the broader literature on financial externalities and regulatory responses, like in Stein (2012); and the role of policy as a coordinating device or focal point, following Schelling (1980). We model how targeted supervision can act as an equilibrium selection device to resolve this system-wide cyber problem. We outline the main economic intuitions here; the Appendix contains all formal proofs and derivations.

4.1 The Main Setting and Pre-Policy Equilibrium

We model an economy with a continuum of risk-neutral banks, indexed by $i \in [0, 1]$. Each bank privately observes its cost, $c_i \in [c_L, c_H]$, of implementing a critical cybersecurity measure, represented by the investment decision $e_i \in \{0, 1\}$. The cost c_i is private information to the bank and is drawn from a common knowledge distribution with support bounded by a minimum cost, c_L, and a maximum cost, c_H. System-wide resilience, Ω, is a "weakest-link" public good: the system is secure ($\Omega = 1$) if and only if all banks invest, an outcome that yields a shared benefit, B, to every institution.

$$\Omega = \min_{i \in [0,1]} \{e_i\} \tag{1}$$

This weakest-link assumption, while stylized, captures the essence of a regulator's concern about contagion, where a single compromised institution can serve as an entry point that threatens the entire network. This premise reflects stated concerns from financial regulators that systemic cyber risk is determined not by the strongest defenses but by the most vulnerable point of entry (Financial Stability Board, 2018; New York State Department of Financial Services, 2020).

A bank's utility depends on its investment choice, the collective outcome, and any regulatory penalty, P, imposed if it is inspected ($\mathbb{I}_i = 1$) and found to be non-compliant ($e_i = 0$). The bank's utility function is thus:

$$U_B = \Omega \cdot B - e_i \cdot c_i - (1 - e_i) \cdot \mathbb{I}_i \cdot P \tag{2}$$

This setup creates strong incentives to free-ride. Under a baseline supervisory regime with

a low, uniform probability of inspection, q_0, no single bank can ensure the public good is provided. Given that the private cost of investing, c_i, exceeds the expected penalty from not investing, $q_0 P$, for all bank types (see Assumption A3 in the Appendix), the market failure is acute, leading to the following pre-policy equilibrium.

Proposition 1 (Pre-Policy Free-Rider Equilibrium)**.** *Under the baseline supervisory regime, the unique Bayesian Nash Equilibrium is for no bank to invest ($e^*(c) = 0$ for all types c), resulting in system-wide vulnerability ($\Omega = 0$).*

4.2 Supervisory Intervention as an Equilibrium Selection Device

To resolve this coordination failure, the supervisor introduces a policy of targeted scrutiny. The supervisor operates under asymmetric information, being unable to perfectly observe a bank's cost c_i or its investment action e_i, but obtains a noisy signal of its type, $s_i^{\text{sup}} = c_i + \varepsilon_i$. This signal allows for a policy where the inspection probability is conditional on the bank's likely propensity to underinvest. Banks with signals suggesting higher costs face an increased probability of inspection, $q_{\text{target}} > q_0$. The resulting type-dependent inspection probability, $q(c)$, is given by:

$$q(c) = q_{\text{target}} - (q_{\text{target}} - q_0)\Phi(\rho(s - c)) \tag{3}$$

where s is the supervisor's signal threshold and $\rho \equiv 1/\sigma$ is the precision of its monitoring technology.

The policy announcement acts as a focal point (Schelling, 1980), altering the strategic environment. Each bank must now assess whether the policy is potent enough to induce *all other banks* to invest. If so, its own action becomes *pivotal* for the provision of the public good. In this pivotal state, the net utility of investing becomes:

$$\Delta U_{\text{post}}(c) = \underbrace{(B - c)}_{\text{Private Net Benefit}} + \underbrace{q(c)P}_{\text{Avoided Expected Penalty}} \tag{4}$$

The supervisor's policy design problem is to set the targeting parameters (s, q_{target}) to ensure that $\Delta U_{\text{post}}(c) \geq 0$ for all bank types, making universal investment an equilibrium. A key insight from our model is a "Precision Paradox": supervisory monitoring that is too precise can be counterproductive. If the inspection probability function $q(c)$ becomes too steep, it can generate strategic uncertainty and support multiple equilibria, as highly precise monitoring may make low-cost banks overly confident that they will not be inspected.[12] A

[12]This result echoes the finding in the global games literature that excessively precise public information can hinder coordination (Morris and Shin, 2002). For the policy to be effective, a degree of "constructive ambiguity" in supervision can be optimal. The formal condition is derived in the Appendix.

well-designed policy must be sufficiently targeted to incentivize laggards without creating strategic instability.

Proposition 2 (Policy-Induced Disciplined Equilibrium)**.** *A supervisory policy that credibly increases the inspection probability for high-cost (laggard) banks can uniquely implement a Pareto-superior equilibrium where all banks invest ($e^*(c) = 1$ for all types c).*

Figure 3 illustrates this mechanism. The policy intervention, through the threat of a higher expected penalty $q(c)P$, shifts the net utility of investing from being universally negative to non-negative, making investment the individually rational choice for all banks.

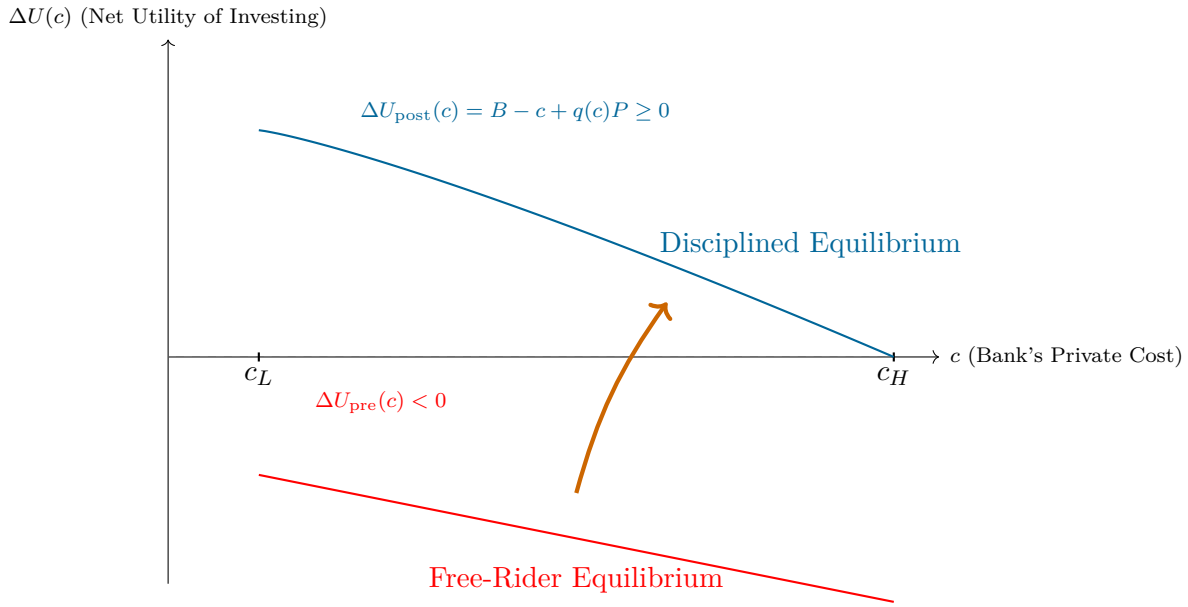

Figure 3: **Mechanism of the Policy-Induced Equilibrium Shift**
This figure illustrates how supervisory scrutiny resolves the underinvestment equilibrium. The horizontal axis captures a bank's private cost of investment, c. The vertical axis is the net utility from investing, $\Delta U(c)$. In the pre-policy state (red line), the net utility is negative for all types, leading to a universal free-rider equilibrium. The policy intervention, by increasing the type-dependent expected penalty for non-compliance, $q(c)P$, shifts the utility curve upward (blue line), making investment individually rational ($\Delta U_{\text{post}}(c) \geq 0$) and shifting the system to the Disciplined Equilibrium.

4.3 Testable Predictions

Our theoretical framework generates two key predictions that we take to the data. The model's unobservable high-cost types ($c_i \to c_H$) map to the banks we empirically classify as "laggards" in Section 6. The policy announcement corresponds to the 2023 CyRST announcement.

P1 (Aggregate Investment): *The announcement of the cyber stress test leads to an increase in cybersecurity investment across the banking sector.*

P2 (Heterogeneous Effects): *The increase in cybersecurity investment is disproportionately larger for laggard banks compared to their more compliant peers.*

We test both predictions. Our theoretical model implies that the treatment effect of the CyRST announcement should be concentrated in the laggard group, for whom the upward shift in expected regulatory penalties provides the decisive incentive to invest.

5 Data, Variables, and Summary Statistics

5.1 Data Sources and Sample Construction

Our empirical analysis leverages a proprietary supervisory panel dataset constructed at the ECB, which provides a uniquely granular view into the operational risk and investment decisions of euro area banks. The sample covers the universe of 109 Significant Institutions (SIs) under the ECB's direct supervision that participated in the 2024 Cyber Resilience Stress Test. The annual data span from 2019 to 2024. We define 2019–2021 as the pre-treatment period and 2023–2024 as the post-treatment period. To ensure our estimates are not biased by policy anticipation effects, we exclude all data from the year 2022, a methodological choice detailed further in Section 6. We require banks to be continuously classified as SIs and have complete data for our variables, which yields a final balanced panel of 96 banks and 475 bank-year observations.[13] This dataset integrates four distinct sources:

1. ECB IT Risk Questionnaire (ITRQ). The foundation of our analysis is the ITRQ, a mandatory and non-public annual data collection. Each wave is administered in the first quarter of a given year, but refers to banks' IT risk management, governance, and expenditure indicators for the preceding year. Its supervisory nature allows us to peer inside the "black box" of bank operational risk management, providing uniquely granular coverage across multiple dimensions. The ITRQ reports, for example, banks' normalized IT security expenditures (our dependent variable), the number and severity of cyberattacks, IT staffing intensity, vacancy and turnover in cyber teams, and preparedness indicators such as detection and recovery times. It also includes forward-looking variables on cyber insurance contracts, innovation project activity, legacy IT risks (e.g., end-of-life systems), as well as IT system complexity and infrastructure exposure measures such as the number of IT systems. Furthermore, we as well utilize two proxies to quantify misalignment between banks'

[13]The reduction from the full universe of 109 SIs is due to banks that were not classified as SIs for the entire sample period or had incomplete data for our main variables. A comparison of key pre-treatment characteristics (e.g., size, profitability) reveals no statistically significant differences between the banks in our final sample and those excluded.

self-assessment and the supervisor's benchmark assessment, and further allow us to track whether banks systematically over- or under-estimate their cybersecurity position relative to supervisory benchmarks.

2. 2024 CyRST Archive. We use confidential supervisory records from the ECB's 2024 CyRST. These archives provide two key sets of information. First, they contain the ex-ante assignment of each bank to either an *enhanced* or a *standard* assessment group. This classification, which was not revealed to banks until November 2023, generates the quasi-experimental variation in supervisory intensity that is central to our identification strategy. Second, the archive offers granular data on the deficiencies uncovered during the exercise. For each bank, we observe (i) the full set of findings, each rated on a four-point severity scale (from 1=low to 4=critical), and (ii) the set of data-quality flags raised by supervisors, coded as either *Amber* or *Red*. From these raw data, we construct weighted severity scores for both findings and flags, which we then use to create the binary indicators *hsti_sevfind* and *hsti_sevflag* via a median split. Appendix I provides a detailed description of the construction of these variables. This approach allows us to create a clean proxy of supervisory pressure that reflects both the formal intensity of the review and the qualitative depth of supervisory concerns.

3. ECB Supervisory Bank Dataset. We merge the ITRQ data with regulatory reports (FINREP and COREP). These filings provide harmonized, audited data on bank financials, allowing us to control for a rich set of time-varying bank characteristics.

5.2 Variable Definitions

A key empirical challenge is that a bank's propensity to underinvest is unobservable. Our strategy is to proxy for this latent characteristic by first 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.

Table 1 provides detailed definitions for the comprehensive set of variables we use to construct this first-stage model. Our main dependent variable is investment, defined as a bank's annual IT security expenditures in Euros. For our first-stage classification and robustness checks, we use its natural logarithm. The granular ITRQ data allows us to precisely measure this outcome.

The rich breadth of the ITRQ data allows us to use a set of explanatory variables, which we use to comprehensively model the economic determinants of optimal cybersecurity spending. The first group, *Cyber Risk Exposure and Controls*, captures 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 IT staff structure, the IT staff vacancy and turnover rates, as well as the share of IT permanent staff share. Importantly, we as well have access to data in terms of the bank's self-reported control and risk levels score and the benchmark score derived from supervisory reference, which serves us as a proxy for risk control *(effort or resources dimension)* and risk level *(outcome or vulnerability dimension)* misalignment.[14]

The second group, *Cyber Insurance*, captures reliance on external risk transfer as a complement or substitute for internal spending. We observe the extensive margin—whether a bank holds a policy (*insurance_d*), and the intensive margin via the direct monetary outlay for coverage (*insurancecontractsdirectcosts*) and the policy's retained loss through the deductible (*insurancecontractsdeductibleamount*). Together, these variables map the decision to insure and the depth of coverage, allowing us to test whether insurance crowds out, or layers on top of, cybersecurity investment.

The third group, *Cyber Innovation*, reflects forward-looking risk-management strategy. We record whether the bank reports any innovation initiative in the IT/cyber domain (*innovationprojects_d*) and the scale/intensity of that pipeline—planned projects (*innovationprojectstobeimplemented*) and projects under execution (*innovationprojectsongoing*). Economically, these variables proxy for modernization of detection, response, and recovery capabilities that may not show up one-for-one 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 spending (*log_criticalprojectsexp*), plus 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 without a remediation plan, *share_eol_gap_ratio*). These measures capture both the scale of transformation work and the execution risk that can crowd out discretionary cyber invest-

[14] *y_m_RC_DT: IT risk control level distance (bank–regulator gap).* This variable measures the gap between the bank's reported control maturity in detection/recovery (RC/DT dimensions) and the regulator's benchmark/control expectation. In practice, it is constructed as the difference between the bank's self-reported control score and the benchmark score derived from supervisory reference levels. A positive gap indicates that the bank assesses its control environment more favorably than the benchmark, while a negative gap suggests under-reporting of control strength. It serves as a proxy for risk control misalignment. *y_m_RR_DT_reb: IT residual risk distance (bank–regulator gap).* This variable captures the misalignment in perceived residual risk (post-control risk exposure). It is built analogously, taking the bank's reported residual risk levels and subtracting the regulator's reference risk levels. A higher positive distance implies that the bank perceives lower residual risk than the regulator (possible underestimation of vulnerability), while a negative distance suggests the bank perceives higher residual risk. It functions as a proxy for risk level misalignment. See Figures 9 and 10.

ment. 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 ensure our laggard classification is not picking up level effects unrelated to cyber risk management per se. Taken together, the ITRQ's breadth—spanning insurance, innovation, legacy constraints, and core controls—enables a first-stage model of expected cybersecurity spending that is tightly conditioned on risk exposure, organizational capacity, and balance-sheet fundamentals. This richness, unavailable in public sources, is central to credibly identifying residual underinvestment and, hence, cybersecurity laggards.

Our specification is saturated with a standard vector of *bank Controls* from regulatory filings to account for size, profitability, leverage, and efficiency. The final column in Table 1 ("Group") serves as an internal classification for organizing the variables in our estimation procedures.

Table 1: Variable Definitions

Variable	Definition	Intuition	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 Statistics

Table 2 reports the summary statistics for our sample, providing a first look at the stylized facts that motivate our analysis. The table also reflects some confidentiality requirements: for sensitive variables we suppress absolute values and instead present two disclosure-safe measures, i.e. the coefficient of variation (CV), which captures relative dispersion, and an indexed measure of levels, which normalizes the pre-treatment average to 100 and expresses the post-treatment mean relative to this baseline. This approach allows us to highlight cross-sectional heterogeneity and any potential temporal shifts while safeguarding the raw data.

In the pre-treatment period (2019–2021), banks devoted substantial resources to cybersecurity, with IT security spending normalized to an index value of 100. Variation across institutions was pronounced, with CVs exceeding 2.0 in several key metrics. Cyber incidents were frequent, as reflected in the baseline level of successful cyberattacks and lengthy detection times. Organizational indicators such as IT staff intensity, vacancy rates, and turnover further illustrate uneven preparedness, pointing to capacity constraints in parts of the sector. Financially, the average bank exhibited a modest return on equity (ROE) of 4.85% and a CET1 capital ratio close to 19%, suggesting that the sector entered the period in reasonably solid condition.

The post-treatment period (2023–2024) displays a marked shift in resource allocation. Cybersecurity spending rises by more than 40 percent, along with higher IT staff intensity and a reduction in vacancy rates. These developments indicate that the additional outlays were channeled into building internal capacity rather than simply inflating budgets. Importantly, the improvement in inputs coincides with a modest decline in realized threats: the index of successful cyberattacks falls to 86.5, and recovery times lengthen less than proportionally, consistent with greater resilience. Insurance markets also reflect the shift, with higher uptake of cyber insurance and a more than doubling of average deductible amounts.

Finally, these unconditional means must be interpreted in context. The post-treatment period was characterized by a stronger macroeconomic and financial environment, with bank profitability nearly doubling (e.g. ROE rising to 9.75%) and efficiency ratios improving materially. Distinguishing the effect of supervisory pressure from these favorable external conditions is thus the key challenge of our empirical strategy. The descriptive evidence is nonetheless consistent with a sector that reallocated resources toward cybersecurity in response to supervisory scrutiny, while simultaneously benefiting from a more supportive macro-financial backdrop.

Table 2: Summary Statistics

Period	Variable	Mean	SD	P10	P90	CV	Index
Pre (2019–2021)	IT Security Expenses (EUR)	—	—	—	—	2.30	100
	Log IT Security Expenses	—	—	—	—	0.32	100
	IT Security Exp. to OPEX	0.0123	0.0382	0.0011	0.0181	—	—
	IT Security Exp. to IT Running Exp.	0.0911	0.2236	0.0149	0.1430	—	—
	IT Security Exp. to IT Running & Change Exp.	0.0580	0.1553	0.0087	0.0813	—	—
	Successful Cyberattacks	—	—	—	—	1.68	100
	Cyberattack Losses (EUR)	—	—	—	—	4.72	100
	IT Staff Intensity	0.0944	0.0562	0.0396	0.1692	—	—
	IT Vacancy Rate	0.0711	0.0871	0.0002	0.1627	—	—
	IT Turnover Index	0.3031	0.3581	0	0.711	—	—
	1st Line of Defense FTE Share	0.0620	0.0645	0.0018	0.1429	—	—
	2nd Line of Defense FTE Share	0.0003	0.0008	0	0.0008	—	—
	3rd Line of Defense FTE Share	0.0016	0.0036	0.0001	0.0028	—	—
	Recovery Time (days)	—	—	—	—	3.49	100
	Detection Time (days)	—	—	—	—	3.23	100
	Perceived Control Misalignment	0.2600	0.2445	0.0571	0.5143	—	—
	Residual Risk Misalignment	-0.4766	0.4729	-1.100	0.0286	—	—
	Cyber Insurance (Dummy)	0.753	0.432	0	1	—	—
	Cyber Insurance Cost (EUR)	—	—	—	—	2.16	100
	Insur. Deductible (EUR)	—	—	—	—	3.92	100
	Innovation Flag (Dummy)	0.919	0.274	1	1	—	—
	Planned Innovation Projects	18.739	46.009	0	35	—	—
	Ongoing Innovation Projects	27.537	63.775	0	66	—	—
	Std. Dev. of No. EoL Systems	0.0597	1.199	-0.3081	0.3121	—	—
	Log No. Critical Projects	2.933	1.220	1.386	4.344	—	—
	Log Total Critical Exp.	16.450	4.607	14.403	19.696	—	—
	Log EoL Project Count	1.293	1.107	0	2.773	—	—
	Log EoL Project Exp.	11.769	7.398	0	18.421	—	—
	Log No. IT Systems	6.665	1.647	4.644	8.536	—	—
	Log Total Assets	25.256	1.322	23.372	27.297	—	—
	Leverage Ratio	0.0688	0.0238	0.0430	0.1083	—	—
	ROE	0.0485	0.0431	0.0002	0.1029	—	—
	Cost-to-Income Ratio	0.6057	0.1509	0.4005	0.8011	—	—
	CET1 Capital Ratio	0.1923	0.0642	0.1379	0.2927	—	—
Post (2023–2024)	IT Security Expenses (EUR)	—	—	—	—	1.79	141.72
	Log IT Security Expenses	—	—	—	—	0.22	108.81
	IT Security Exp. to OPEX	0.0136	0.0102	0.0027	0.0261	—	—
	IT Security Exp. to IT Running Exp.	0.0921	0.0602	0.0305	0.1640	—	—
	IT Security Exp. to IT Running & Change Exp.	0.0558	0.0344	0.0160	0.0989	—	—
	Successful Cyberattacks	—	—	—	—	2.21	86.51
	Cyberattack Losses (EUR)	—	—	—	—	6.04	45.25
	IT Staff Intensity	0.1160	0.0631	0.0431	0.1970	—	—
	IT Vacancy Rate	0.0656	0.0708	0.0011	0.1456	—	—
	IT Turnover Index	0.2997	0.3114	0.0357	0.6720	—	—
	1st Line of Defense FTE Share	0.0861	0.0749	0.0059	0.1866	—	—
	2nd Line of Defense FTE Share	0.0003	0.0006	0	0.0009	—	—
	3rd Line of Defense FTE Share	0.0016	0.0024	0.0003	0.0031	—	—
	Recovery Time (days)	—	—	—	—	2.85	150.18
	Detection Time (days)	—	—	—	—	2.85	96.23
	Perceived Control Misalignment	0.2630	0.2477	0.0571	0.5143	—	—
	Residual Risk Misalignment	-0.4900	0.4872	-1.114	0.0286	—	—
	Cyber Insurance (Dummy)	0.833	0.374	0	1	—	—
	Cyber Insurance Cost (EUR)	—	—	—	—	1.44	209.80
	Insur. Deductible (EUR)	—	—	—	—	2.20	226.72
	Innovation Flag (Dummy)	0.938	0.243	1	1	—	—
	Planned Innovation Projects	21.922	50.287	0	40	—	—
	Ongoing Innovation Projects	29.635	67.698	1	66	—	—
	Std. Dev. of No. EoL Systems	-0.0877	0.592	-0.3081	0.2130	—	—
	Log No. Critical Projects	3.333	1.201	1.946	5.011	—	—
	Log Total Critical Exp.	17.487	2.864	15.607	20.026	—	—
	Log EoL Project Count	1.575	1.189	0	3.045	—	—
	Log EoL Project Exp.	13.541	6.473	0	18.636	—	—
	Log No. IT Systems	7.532	2.126	4.949	10.402	—	—
	Log Total Assets	25.335	1.329	23.464	27.356	—	—
	Leverage Ratio	0.0726	0.0229	0.0474	0.1106	—	—
	ROE	0.0975	0.0467	0.0355	0.1662	—	—
	Cost-to-Income Ratio	0.4981	0.1310	0.3174	0.6686	—	—
	CET1 Capital Ratio	0.1963	0.0597	0.1454	0.2696	—	—

Note: This table reports descriptive statistics for the key 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) coefficient of variation (CV), defined as the ratio of the standard deviation to the mean (SD/Mean), which captures relative dispersion independent of levels; 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.

6 Empirical Strategy

Our analysis aims to estimate the effect of the ECB's Cyber Resilience Stress Test (CyRST) on banks' cybersecurity investment. A key identification challenge we confront comes from the fact that a bank's classification as a "laggard" is not randomly assigned but is instead the result of endogenous strategic choices shaped by (un)observables such as managerial risk preferences and corporate culture. These same factors are plausibly correlated with a bank's responsiveness to supervisory scrutiny, creating a classic selection bias problem.

To tackle this challenge, we implement a two-stage empirical strategy that exploits the CyRST announcement as a quasi-natural experiment. The first stage constructs a time-invariant, pre-determined proxy of a bank's latent propensity to underinvest in cybersecurity; the second stage embeds this classification in a difference-in-differences (DiD) framework to estimate the effect of the policy shock.

6.1 Identification Strategy

Our source of exogenous variation is the March 2023 public announcement of the CyRST, a novel supervisory exercise focused on operational resilience, with neither capital implications nor public disclosure of bank-level results. This design explicitly neutralizes the "capital" and "market discipline" channels that dominate the stress-testing literature,[15] allowing us to isolate a "scrutiny channel" in which the credible threat of intrusive examination alters bank behavior.[16]

A key institutional feature is the data-collection calendar: 2022 observations were collected in early 2023, overlapping with the policy announcement. To eliminate contamination from anticipatory adjustments or strategic reporting, we exclude all 2022 data from our main analysis. The pre-treatment period is therefore defined as 2019–2021, and the post-treatment period begins in 2023. This ensures strict temporal separation between treatment-group assignment and the policy intervention, bolstering internal validity. Robustness checks in Section 7.5 confirm that our results are robust alternative sample definitions and placebo tests[17]

6.2 First Stage: Identifying Laggards

We define "laggards" as banks that systematically underinvest in cybersecurity relative to the level predicted by their observable fundamentals and risk profile. To operationalize

[15] See, e.g., Acharya et al. (2018b) and Goldstein and Leitner (2018).
[16] See Section 3 for a detailed overview of the CyRST's design and implementation.
[17] Figure 7 illustrates the temporal structure of our identification strategy.

this, we decompose observed investment into an "expected" component and an orthogonal residual capturing discretionary deviations from that benchmark. Formally, we estimate the following two-way fixed effects model on the pre-treatment panel (2019–2021):

$$\log(\text{Investment}_{it}) = \alpha_i + \lambda_t + \mathbf{X}'_{it}\boldsymbol{\beta} + \varepsilon_{it}, \tag{5}$$

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:

$$\text{Laggard}_i = \mathbf{1}\left[\bar{\varepsilon}_i \leq \text{P50}(\bar{\boldsymbol{\varepsilon}})\right]. \tag{6}$$

While a persistently negative residual could, in principle, reflect unobserved efficiency or a simpler business model, several factors make these alternative interpretations unlikely. First, our model for $\mathbf{X}_{it}$ is saturated with operational, technological, and governance controls, including business model fixed effects, that directly account for these observable drivers. Second, as we demonstrate in Section 7.2, laggards in the pre-treatment period are statistically indistinguishable from non-laggards on key financial dimensions but display materially weaker cyber-resilience metrics (e.g., longer detection times), consistent with genuine underinvestment. Finally, the fact that this specific group exhibits the largest post-CyRST investment surge provides strong ex-post validation of our underinvestment interpretation.

6.3 Second Stage: Estimation of the Treatment Effect

6.3.1 A Note on the Two-Stage Estimation Procedure

Our two-stage procedure intentionally employs different estimators for each task, a choice dictated by the distinct objective of each stage. The first stage is a *classification exercise* designed to rank banks and generate a proxy for their latent type. For this purpose, the log-linear OLS model in Equation (5) is a standard and effective tool, as its residuals provide an intuitive, percentage-based measure of relative underinvestment. We do not interpret the coefficients of this first stage causally.

The second stage, in contrast, is a *causal estimation exercise*. Here, the primary objective is to obtain a consistent and unbiased estimate of the treatment effect parameter, β_{ATT}. For this task, the choice of estimator is paramount, leading us to use the Poisson Pseudo-Maximum Likelihood (PPML) model in a DiD framework.

6.3.2 Model: Poisson Pseudo-Maximum Likelihood

We estimate the second-stage DiD model using the PPML estimator. PPML models the conditional mean in multiplicative form, accommodates heteroskedasticity, and retains zero-expenditure observations without ad-hoc transformations, thereby avoiding the well-documented biases of log-linear OLS models in this context.[18]

6.3.3 Baseline Difference-in-Differences Specification

Our baseline specification is:

$$\mathbb{E}[\text{Investment}_{it}|\alpha_i, \lambda_t, \mathbf{X}_{it-1}] = \exp\left(\alpha_i + \lambda_t + \beta_{ATT} \cdot (\text{Laggard}_i \times \text{Post}_t) + \mathbf{X}'_{it-1}\boldsymbol{\delta}\right), \quad (7)$$

where β_{ATT} measures the average treatment effect on laggards post-CyRST, interpreted as a semi-elasticity, $(\exp(\beta_{ATT}) - 1) \times 100\%$. Standard errors are clustered two-ways at the bank and year level.

6.3.4 Triple-Difference Specification for the Scrutiny Channel

To provide a direct test of the scrutiny channel, we exploit the cross-sectional variation in supervisory intensity inherent in the CyRST design. We construct a time-invariant indicator, HighScrutiny_i, to capture substantive supervisory pressure that extends beyond the formal assessment track. This composite indicator equals one if a bank meets at least two of the following three baseline conditions: (i) assignment to the CyRST's high-intensity *enhanced assessment* track; (ii) an above-median weighted severity of supervisory findings *hsti_sevfind*=1; and (iii) an above-median weighted severity of data-quality flags *hsti_sevflag*=1. The rationale here is to create a robust measure that identifies banks facing pervasive supervisory concern, thereby mitigating the risk that our classification is driven by a single, potentially noisy, dimension of the assessment. This indicator combines the regulator's *intended* scrutiny (ex-ante track assignment) with the *realized* scrutiny stemming from the discovery of significant substantive deficiencies and data governance issues. As detailed in Appendix 8, all three components are determined using information from the stress test exercise itself and are therefore pre-determined with respect to our post-announcement outcome variable. By interacting this HighScrutiny_i indicator with our Laggard_i classification and the Post_t dummy, we implement a triple-difference (DDD) specification to test our central hypothesis: whether the investment response of laggard banks was dispropor-

[18] See Silva and Tenreyro (2006) for the seminal critique of log-linear models and Silva and Tenreyro (2011) for evidence on PPML performance.

tionately stronger when they faced the most intense supervisory examination. Our baseline specification is therefore the following:

$$\begin{aligned}\mathbb{E}[\text{Investment}_{it}|\cdot] = \exp(&\alpha_i + \lambda_t + \beta_1(\text{Laggard}_i \times \text{Post}_t) + \beta_2(\text{HighScrutiny}_i \times \text{Post}_t) \\ &+ \beta_3(\text{Laggard}_i \times \text{HighScrutiny}_i) + \beta_{DDD}(\text{Laggard}_i \times \text{Post}_t \times \text{HighScrutiny}_i) \\ &+ \mathbf{X}'_{it-1}\boldsymbol{\delta}). \end{aligned} \tag{8}$$

Our main coefficient of interest is β_{DDD}, which captures the additional effect for laggard banks subjected to high-intensity scrutiny, isolating the causal impact of the scrutiny channel.

6.3.5 Event Study and the Parallel Trends Assumption

The key identifying assumption of our DiD model is that of parallel trends. We provide strong corroborating evidence by estimating a PPML event-study specification. For this assumption to be violated, a confounding shock contemporaneous with the CyRST announcement must have differentially affected laggards for reasons unrelated to the stress test, an eventuality for which we find no institutional evidence. The specification is:

$$\mathbb{E}[\text{Investment}_{it}|\cdot] = \exp\left(\alpha_i + \lambda_t + \sum_{k=-3, k\neq -2}^{2} \beta_k \cdot (\text{Laggard}_i \times \mathbf{1}[t = E + k]) + \mathbf{X}'_{it-1}\boldsymbol{\delta}\right), \tag{9}$$

normalizing β_{-2} (year 2021) to zero. The coefficients $\beta_{k<0}$ serve as a falsification test; as shown in Section 7.3, their statistical insignificance supports the parallel trends assumption. The $\beta_{k\geq 0}$ coefficients trace the dynamic post-announcement investment response.

7 Main Results

This section presents the main empirical findings on the real effects of the ECB's CyRST. We proceed in four steps. First, we establish the aggregate impact of the policy announcement on bank investment using a before-after analysis. Second, we validate our classification of "laggard" banks, our empirical proxy for firms with a high propensity to underinvest. Third, we present our main difference-in-differences estimates, supported by an event-study analysis, to test the paper's central hypothesis on heterogeneous effects (P2). Fourth, we provide robust evidence for the supervisory "scrutiny channel" as the primary mechanism and conduct a series of tests to corroborate our findings.

7.1 Before-After Analysis: Aggregate Investment Response

We begin by testing our first theoretical prediction (P1), which posits that the announcement of the cyber stress test can act as a powerful coordinating signal, shifting industry norms and prompting a broad-based increase in cybersecurity investment. This test provides a baseline measure of the policy's overall impact before we explore the heterogeneous effects at the heart of our study.[19]

Table 3 reports the results. The coefficient on the *Post* indicator is positive and statistically significant across all specifications, providing strong initial evidence of a sector-wide increase in investment following the CyRST announcement. As we progressively saturate the model with controls and fixed effects, the estimate remains stable in magnitude and significance. Our most demanding specification (Column 6), which includes a full suite of fixed effects and controls, yields a coefficient of 0.370. This estimate implies that, on average, banks increased their cybersecurity investment by approximately 45% ($e^{0.370} - 1$) following the announcement. This response provides clear support for Prediction 1. The economic magnitude of this effect suggests that the supervisory initiative acted as a strong focal point, sufficient to overcome institutional inertia and trigger a significant reallocation of resources toward managing cyber risk across the European banking sector.

Table 3: Aggregate Effect of the Policy on Investment (Before-After)

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: 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 shows the effect of the ECB Cyber Stress Test announcement (*Post*) on cybersecurity investment using a PPML model. Column (1) is the baseline. Columns (2)–(6) progressively add controls and fixed effects. Bank controls include log(Total Assets), Leverage Ratio, ROE, Cost-to-Income Ratio, and CET1 Capital Ratio. Robust standard errors clustered by bank in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

[19]The specification is a panel fixed-effects model estimated via PPML: $\mathbb{E}[\text{Investment}_{it}|\cdot] = \exp(\alpha_i + \lambda_t + \beta \cdot \text{Post}_t + \mathbf{X}'_{it-1}\boldsymbol{\delta})$, 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), year fixed effects (λ_t), and a vector of lagged, time-varying bank controls ($\mathbf{X}_{it-1}$).

7.2 First Stage: Validation of Laggards

Our empirical strategy hinges on a credible, ex-ante classification of banks based on their underlying cybersecurity investment posture. To this end, we construct a proxy for banks that systematically underinvest relative to their peers and risk profiles. We classify a bank as a *Laggard* if its cybersecurity spending, averaged over the 2020–2021 pre-treatment period, falls below the median of residuals from a benchmark investment model.[20] This first-stage model is designed to partial out the component of investment attributable to a set of observable IT risk, operational, and financial characteristics, thereby isolating a stable, pre-treatment measure of a bank's discretionary investment policy.

[20] We test more stringent, quartile-based definitions in Section 7.5.1.

Table 4: Determinants of Normalized Cybersecurity Investment (Residual Benchmark Model)

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)
FTE 1st line of defense	–3.885	(4.302)
FTE 2nd line of defense	27.179	(111.643)
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 4 reports the results for this benchmark model, estimated on the pre-treatment panel. With an R^2 of 0.706, the model explains a substantial portion of the variation in cybersecurity spending. Notably, direct measures of past risk realization, such as the number of prior cyberattacks, are not statistically significant predictors. This finding suggests that investment policies are driven less by reactive responses and more by deep-seated, structural factors related to governance and strategic priorities. The model's strong explanatory power provides confidence that the resulting residuals are not mere statistical noise but a meaningful measure of discretionary spending choices. Figure 4 visually supports this: as controls are added, the dispersion of the residuals tightens considerably, indicating our model effectively isolates the idiosyncratic component of investment policy.

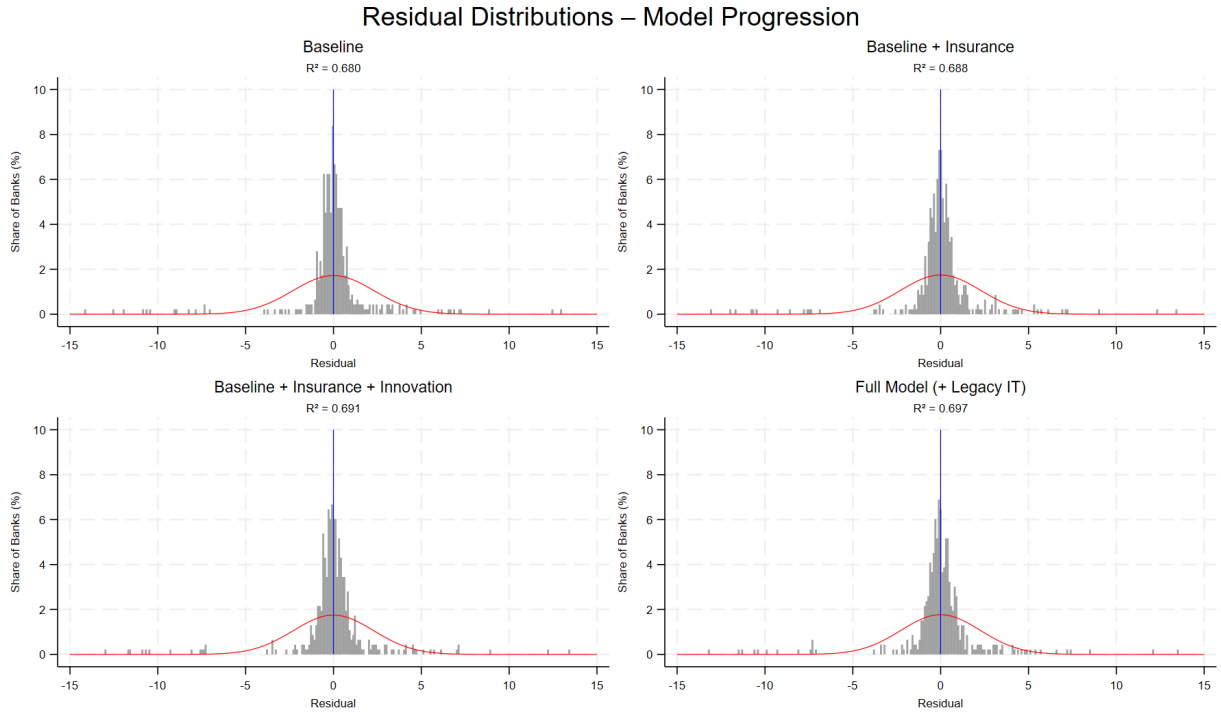

Figure 4: **Distribution of Residualized Cybersecurity Investment.** Each panel shows the histogram of residuals from the benchmark model predicting IT security investment. The models incrementally add controls. The narrowing dispersion of the residuals illustrates how our model explains variation in banks' investment, supporting the robust identification of cybersecurity laggards for our causal analysis.

An important aspect of our research design is that the *Laggard* indicator isolates a strategic choice regarding operational risk, not merely differences in financial health. Table **??** confirms this. In the pre-treatment period, *Laggard* and non-laggard banks are statistically indistinguishable across key financial dimensions, including size, profitability, and capitalization.Formal tests for differences in means confirm this visual inspection. We find no statistically significant pre-treatment differences for key financial ratios such as the Leverage

Ratio, ROE, or CET1 Capital Ratio, whereas operational metrics like 'Detection Time' show a statistically significant divergence. This financial equivalence mitigates concerns that our results are driven by confounding factors such as pre-existing financial constraints. While financially similar, these groups exhibit stark operational differences. Prior to the CyRST, *Laggards* invested only half as much in IT security (€19.8 million vs. €39.2 million) and displayed materially weaker resilience, with significantly longer average times to detect cyber incidents (45.1 vs. 35.5 days). This combination of financial similarity and operational divergence provides strong supportive evidence that our *Laggard* indicator is a meaningful proxy for pre-policy underinvestment.

Table 5: Summary Statistics by Laggard Status: Pre- and Post-Treatment Periods

Note: This table reports descriptive statistics for the key 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) coefficient of variation (CV), defined as the ratio of the standard deviation to the mean (SD/Mean), which captures relative dispersion independent of levels; and (ii) 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.

	Not Laggard (Pre)				**Laggard (Pre)**				**Not Laggard (Post)**				**Laggard (Post)**			
Variable	Mean	SD	CV	Index	Mean	SD	CV	Index	Mean	SD	CV	Index	Mean	SD	CV	Index
IT Security Expenses (EUR)	—	—	2.05	100	—	—	2.49	100	—	—	1.80	114.80	—	—	1.74	165.15
Log IT Security Expenses	—	—	0.19	100	—	—	0.39	100	—	—	0.18	102.45	—	—	0.26	112.73
IT Security Exp. / OPEX	0.0161	0.0514	—	—	0.0089	0.0183	—	—	0.0139	0.0111	—	—	0.0132	0.0093	—	—
IT Security Exp. / IT Run. Exp.	0.1286	0.3125	—	—	0.0565	0.0450	—	—	0.0929	0.0614	—	—	0.0914	0.0599	—	—
IT Security Exp. / IT Run.+Change	0.0820	0.2178	—	—	0.0358	0.0292	—	—	0.0550	0.0348	—	—	0.0567	0.0345	—	—
Successful Cyberattacks	—	—	1.70	100	—	—	1.64	100	—	—	2.10	96.75	—	—	2.36	73.12
Cyberattack Losses (EUR)	—	—	4.30	100	—	—	5.29	100	—	—	4.08	40.82	—	—	8.05	52.60
IT Staff Intensity	0.0880	0.0554	—	—	0.1012	0.0569	—	—	0.1093	0.0585	—	—	0.1233	0.0673	—	—
IT Vacancy Rate	0.0667	0.0812	—	—	0.0759	0.0932	—	—	0.0670	0.0802	—	—	0.0649	0.0613	—	—
IT Turnover Index	0.303	0.423	—	—	0.306	0.283	—	—	0.289	0.301	—	—	0.309	0.325	—	—
1st Line of Defense FTE Share	0.0593	0.0562	—	—	0.0634	0.0719	—	—	0.0840	0.0672	—	—	0.0874	0.0826	—	—
2nd Line of Defense FTE Share	0.00023	0.00045	—	—	0.00040	0.00107	—	—	0.00032	0.00058	—	—	0.00027	0.00064	—	—
3rd Line of Defense FTE Share	0.00183	0.00456	—	—	0.00145	0.00251	—	—	0.00189	0.00327	—	—	0.00139	0.00115	—	—
Recovery Time (days)	—	—	3.31	100	—	—	3.67	100	—	—	2.62	176.34	—	—	3.20	121.32
Detection Time (days)	—	—	3.29	100	—	—	3.15	100	—	—	3.16	77.38	—	—	2.65	108.30
Perceived Control Misalignment	0.272	0.296	—	—	0.254	0.182	—	—	0.278	0.301	—	—	0.253	0.182	—	—
Residual Risk Misalignment	-0.476	0.552	—	—	-0.488	0.384	—	—	-0.500	0.575	—	—	-0.491	0.385	—	—
Cyber Insurance (Dummy)	0.701	0.460	—	—	0.797	0.403	—	—	0.830	0.378	—	—	0.833	0.375	—	—
Cyber Insurance Cost (EUR)	—	—	2.06	100	—	—	2.25	100	—	—	1.38	196.89	—	—	1.54	223.10
Cyber Insurance Deductible (EUR)	—	—	2.98	100	—	—	3.33	100	—	—	2.21	682.43	—	—	2.26	120.81
Innovation Flag (Dummy)	0.942	0.235	—	—	0.895	0.307	—	—	1.000	0.000	—	—	0.875	0.332	—	—
Planned Innovation Projects	17.2	42.5	—	—	20.7	49.6	—	—	21.4	46.9	—	—	22.8	54.1	—	—
Ongoing Innovation Projects	23.1	49.6	—	—	32.4	75.2	—	—	30.6	66.9	—	—	29.2	69.4	—	—
Std. Dev. of No. EoL Systems	0.195	1.523	—	—	-0.065	0.767	—	—	-0.024	0.661	—	—	-0.212	0.194	—	—
Log No. Critical Projects	2.93	1.23	—	—	2.93	1.22	—	—	3.44	1.17	—	—	3.18	1.19	—	—
Log Total Critical Exp.	16.08	5.31	—	—	16.72	3.83	—	—	17.52	3.33	—	—	17.40	2.34	—	—
Log EoL Project Count	1.24	1.22	—	—	1.33	0.98	—	—	1.45	1.27	—	—	1.65	1.06	—	—
Log EoL Project Exp.	10.69	7.94	—	—	12.92	6.60	—	—	12.47	7.19	—	—	14.47	5.57	—	—
Log No. IT Systems	6.49	1.85	—	—	6.81	1.41	—	—	7.79	2.24	—	—	7.20	1.92	—	—
Log Total Assets	25.26	1.39	—	—	25.20	1.23	—	—	25.33	1.39	—	—	25.30	1.24	—	—
Leverage Ratio	0.0711	0.0245	—	—	0.0670	0.0229	—	—	0.0756	0.0238	—	—	0.0703	0.0215	—	—
ROE	0.0478	0.0435	—	—	0.0487	0.0431	—	—	0.104	0.047	—	—	0.091	0.046	—	—
Cost-to-Income Ratio	0.598	0.149	—	—	0.612	0.154	—	—	0.480	0.128	—	—	0.513	0.132	—	—
CET1 Capital Ratio	0.194	0.067	—	—	0.191	0.062	—	—	0.194	0.061	—	—	0.200	0.058	—	—

7.3 The Heterogeneous Effect of the CyRST on Investment

Having established the validity of our *Laggard* classification, we now test our main hypothesis (Prediction 2): the investment response to the CyRST is concentrated among laggard banks. Table 6 reports the results from our main DiD specification.

The estimates provide strong evidence in support of our hypothesis. Across all specifications, the coefficient on the interaction term, Post × Laggard, is positive and statistically significant. Our most comprehensive specification is presented in Column (6), which includes bank, year, business model, and country fixed effects, alongside a full battery of controls. The coefficient of 0.290 indicates that the CyRST announcement induced *Laggard* banks to increase their cybersecurity investment by about 34% relative to their non-laggard peers. This represents a substantial and targeted reallocation of resources, providing powerful support for our second prediction. The finding aligns perfectly with the equilibrium-switching mechanism in our theoretical model, wherein the supervisory intervention provides the decisive incentive for high-cost types to abandon their free-riding strategy.

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

	(1)	(2)	(3)	(4)	(5)	(6)
			Dependent variable: Investment			
Post × Laggard	**0.364***	**0.350***	**0.347***	**0.361***	**0.360***	**0.290***
	(0.192)	**(0.192)**	**(0.195)**	**(0.195)**	**(0.207)**	**(0.168)**
Post	0.137	0.158	–	0.163*	–	–
	(0.123)	(0.127)	–	(0.086)	–	–
Laggard	-0.684	-0.450*	-0.449*	–	–	–
	(0.460)	(0.262)	(0.262)	–	–	–
Bank Controls	No	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	Yes	Yes
Bank FE	No	No	No	Yes	Yes	Yes
Business Model FE	No	No	No	No	No	Yes
Observations	470	470	470	470	470	470

Note: This table presents our main DiD estimation using the PPML model. The dependent variable is the level of IT security investment. Column (1) presents the raw DiD. Columns (2) through (6) progressively add controls and fixed effects. Our preferred specification is Column (6). The main *Laggard* effect is absorbed by Bank FE in Columns (4)-(6); the main *Post* effect is absorbed by Year FE in Columns (3), (5), and (6). Robust standard errors clustered by bank in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

To provide further support for a causal interpretation and to trace out the policy's dynamic effects, we estimate an event-study model. Figure 5 plots the dynamic coefficients. The point estimates for the pre-treatment periods are economically small and statistically indistinguishable from zero, providing strong visual support for the parallel trends assumption. In stark contrast, we observe a sharp structural break beginning precisely in 2023, the year of the announcement. The coefficient becomes positive and grows in magnitude into 2024. The absence of pre-treatment anticipation, combined with the timing and persistence of this break, powerfully corroborates a causal interpretation of our DiD estimates.

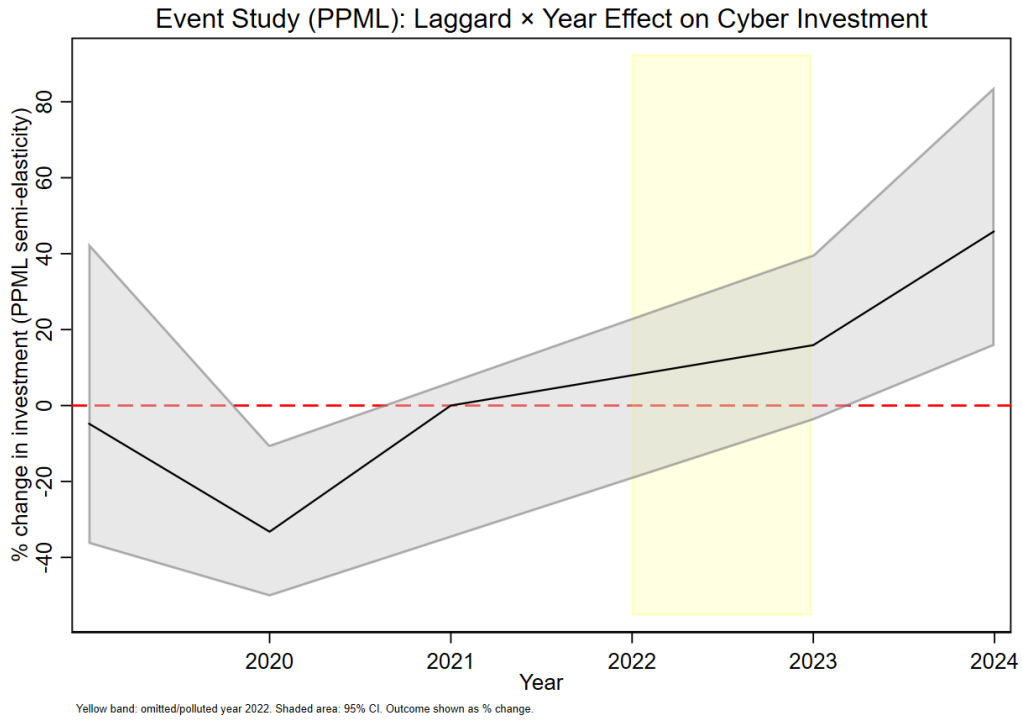

Figure 5: **Investment Response to ECB Cyber Resilience Stress Test (PPML).** This figure presents year-specific semi-elasticity estimates of being a cybersecurity laggard on investment, relative to the baseline year 2021, from a PPML model with bank fixed effects and controls. Coefficients are reported as percentage changes ($100 \times (\exp(\beta) - 1)$). The flat pre-trend followed by a sharp increase in 2023 supports a causal interpretation. The shaded area represents 95% confidence intervals.

7.4 Mechanism: The Supervisory Scrutiny Channel

Having established a large and heterogeneous investment response, we now test the underlying mechanism. We argue that the effect was not driven by the public announcement alone, but by the *"scrutiny channel"*, a behavioral change induced by the credible threat of intensive, direct supervisory examination (Kok et al. (2023)). To test this, we exploit cross-sectional variation in supervisory intensity by using a composite indicator, HighScrutiny_i. As explained in Section 6, this variable is designed to identify banks facing substantively higher pressure and equals one if a bank meets at least two of the following three pre-determined conditions: (i) assignment to the *enhanced assessment* track; (ii) an above-median weighted severity of findings (`hsti_sevfind` $= 1$); and (iii) an above-median weighted severity of data-quality flags (`hsti_sevflag` $= 1$). In our sample, 32 of the 48 laggard banks fall into this high-scrutiny group. This composite design creates a strong measure of supervisory pressure, and as detailed in Appendix I, all its components are determined at baseline, ensuring exogeneity.

Table 7 presents the results of the DDD. The analysis reveals that the entire treatment effect is concentrated among laggard banks subjected to high-intensity scrutiny. The key coefficient of interest is on the triple-interaction term, Post × Laggard × High Scrutiny. Across all specifications, this coefficient is positive and statistically significant. In our most demanding model (Column 6), the coefficient of 0.452 implies that laggard banks under intense oversight had an additional investment increase of 57%.

The power of this research design lies in interpreting the lower-order interaction terms as important placebo tests. The coefficient on Post × Laggard (0.043 and statistically insignificant in Column 6) captures the effect for laggards in the *low-scrutiny* group. This null result is a key finding: laggards who did not face the imminent prospect of direct examination did not react. Similarly, the coefficient on Post × High Scrutiny (0.076 and insignificant) measures the effect for *non-laggards* in the high-scrutiny group. This shows that merely being under intense oversight was not sufficient to trigger a differential investment response. Taken together, these results confirm that the CyRST's main impact was to force the weakest links within the most closely-watched group to increase their investment, providing strong evidence for the scrutiny channel.[21]

Table 7: Mechanism: Triple-Difference Estimates of High-Intensity Scrutiny (PPML)

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: Investment					
Post×Laggard×High_Scrutiny	**0.733*** (0.235)**	**0.524** (0.225)**	**0.518** (0.228)**	**0.642** (0.251)**	**0.643** (0.271)**	**0.452** (0.192)**
Post×Laggard	-0.067 (0.129)	0.068 (0.109)	0.071 (0.110)	-0.009 (0.112)	-0.008 (0.112)	0.043 (0.101)
Post×High_Scrutiny	-0.009 (0.198)	0.073 (0.174)	0.073 (0.182)	0.044 (0.163)	0.053 (0.195)	0.076 (0.173)
Laggard×High_Scrutiny	-2.668*** (0.749)	-0.941** (0.413)	-0.939** (0.412)	– –	– –	– –
Bank Controls	No	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	No	Yes	Yes
Bank FE	No	No	No	Yes	Yes	Yes
Business Model FE	No	No	No	No	No	Yes
Observations	470	470	470	460	460	460

Note: This table presents PPML triple-difference estimates. All regressions cluster standard errors at the bank level. Column specifications follow the same progressive inclusion of controls and fixed effects as in Table 6. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

[21]To address the concern that assignment to the high-scrutiny group was not random, we conduct a balancing test comparing the pre-treatment characteristics of banks across both groups. The test, reported in Appendix Table A.1, reveals no evidence of systematic, pre-existing differences on key financial or operational dimensions, lending strong support to the validity of our triple-difference design.

7.5 Robustness and Additional Evidence

To verify the robustness of our main findings, we conduct a series of additional tests. We confirm that our results are not an artifact of our specific classification of laggards by testing more stringent definitions. We also provide evidence on the persistence of the behavioral change we document.

7.5.1 Alternative Definition of Laggards

To ensure our results are not driven by the specific median-split threshold, we test more stringent definitions of laggards that isolate the extremes of the investment distribution. Table 8 presents DiD estimates comparing banks in the bottom investment residual quartile (Q1, the most severe laggards) against those in the top quartile (Q4, the most proactive investors). The result is a stark contrast. The interaction term for the Q1 vs. Q4 comparison is 0.683 and highly significant, implying a substantial 98% relative increase in investment for the most extreme laggards. In contrast, when comparing moderately underperforming banks (Q2) against their moderately overperforming peers (Q3), the coefficient is economically small and statistically insignificant.

Table 8: DiD Estimates with Alternative Laggard Definitions

	(1) **Q1 vs. Q4** *(Extreme Laggards)*	(2) **Q2 vs. Q3** *(Moderate Laggards)*
Post × Laggard (Quartile)	**0.683*** **	**-0.208**
	(0.274)	**(0.184)**
Bank & Year FE	Yes	Yes
All Controls	Yes	Yes
Observations	225	235

Note: This table reports PPML DiD estimates using the most saturated specification (equivalent to Column 6 in Table 6). Column (1) defines the treatment group as banks in the bottom quartile (Q1) of residual investment and the control group as banks in the top quartile (Q4). Column (2) compares banks in the second quartile (Q2) to those in the third (Q3). Robust standard errors clustered by bank in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

The event study in Figure 6 visually corroborates this non-linearity. The center panel (Q1 vs. Q4) shows a flat pre-trend followed by a dramatic post-announcement surge in investment. The right panel (Q2 vs. Q3) remains flat throughout the sample period. Together, these results reveal that the policy did not induce a uniform response among all under-investors; rather, its disciplinary force was precisely targeted at the most extreme laggards. This confirms that our findings are not driven by an arbitrary cutoff but by a gen-

uine behavioral shift among the banks with the most significant pre-existing vulnerabilities, consistent with our theoretical predictions.

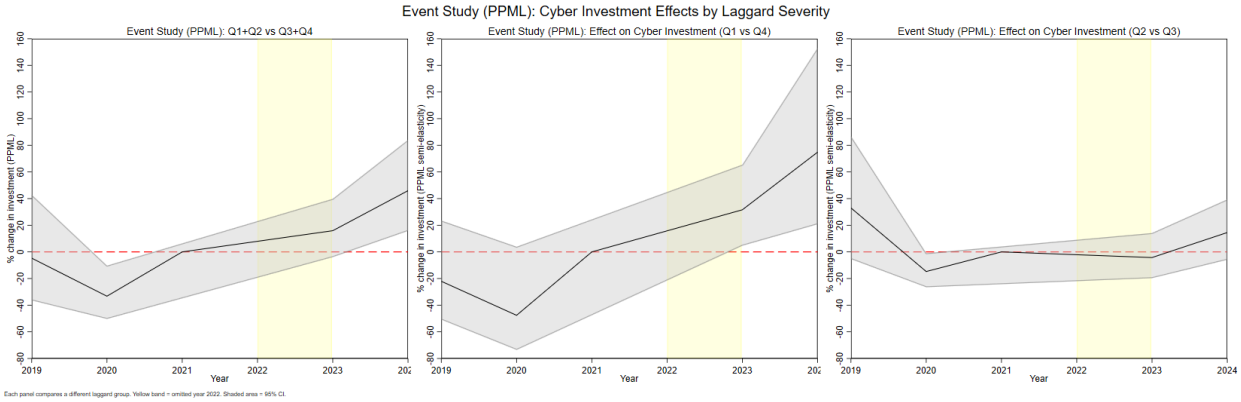

Figure 6: **Event Study Panels – Cyber Investment Effects by Laggard Severity (PPML).** This figure presents three event study specifications estimating the dynamic effect of the CyRST on cybersecurity investment. The left panel shows the baseline median-split result. The center panel contrasts the most severe laggards (Q1) against top performers (Q4). The right panel isolates moderate underperformers (Q2) from their peers (Q3). The effect is clearly concentrated among the most extreme laggards. All specifications include bank fixed effects and a full set of controls.

7.5.2 Persistence of Post-policy Investment

A key question is whether the supervisory intervention induced a persistent change in investment strategy or merely a temporary response. A Probit model confirms that the stress test catalyzed meaningful and lasting adjustments for the majority of laggards, showing a high and statistically significant unconditional exit probability of about 71%.

Interestingly, Table 9 reveals heterogeneity in this dynamic. While the most extreme laggards (Q1) exhibit very high exit rates regardless of scrutiny intensity (over 80%), a different pattern emerges for moderately underperforming laggards (Q2). For this group, those receiving high scrutiny were substantially *less* likely to exit laggard status than their low-scrutiny peers (44.4% vs. 66.7%). This suggests a complex interaction between supervision and firm capacity. For the most severe laggards, the path to compliance is clear (i.e., spend more). For moderately underperforming banks, however, intensive scrutiny may uncover deeper, more persistent structural deficiencies that require multi-year remediation plans. Thus, while these banks increase investment, the discovery of more complex issues may mechanically lower their short-term "exit" rate from the laggard classification.

Table 9: Descriptive Statistics on Exit from Laggard Status

Pre-Treatment Quartile	Supervisory Scrutiny	Stayed Laggard	Exited Laggard	Total	Exit Rate (%)
Q1 (Lowest)	No	2	9	11	81.8
	Yes	2	11	13	84.6
Q2	No	5	10	15	66.7
	Yes	5	4	9	44.4
Total		14	34	48	70.8

Notes: This table reports the share of banks initially classified as cybersecurity laggards (based on 2020–2021 data) that exited laggard status by 2024, conditional on pre-treatment severity (Q1 vs. Q2) and the intensity of supervisory scrutiny.

8 Conclusion

Investment in cybersecurity within the interconnected banking system presents a classic public good problem, where network externalities create incentives for firms to free-ride, leading to systemic underinvestment and "weakest link" vulnerabilities. This paper provides, to the best of our knowledge, the first causal evidence that supervisory action, distinct from traditional capital regulation or market discipline, can resolve this coordination failure. We exploit the 2024 ECB CyRST as a quasi-natural experiment. The exercise's design with no direct capital implications and no public disclosure of individual results provides a unique opportunity to empirically test the effect of a pure "scrutiny channel" on banks' investments.

We find that targeted supervisory scrutiny is a remarkably effective tool for correcting this market failure. Using confidential supervisory data, we first identify a set of "laggard" banks that systematically underinvested in cybersecurity relative to their risk profiles prior to the intervention. Our DiD estimates show that in response to the CyRST, these laggard institutions increased their cybersecurity investment by 34% relative to their non-laggard peers. This strong response is not driven by the policy announcement alone. Exploiting variation in supervisory intensity within a triple-difference framework, we show the effect is mainly concentrated among laggards subjected to high-intensity scrutiny, who exhibit an additional investment increase of 57%. For laggard banks facing only low-intensity oversight, we find no statistically significant effect, confirming the scrutiny channel as the key mechanism. Thus, our findings suggest that the increase in laggards' investments is driven by the credible threat of direct examination.

From a policy perspective, our results show that "soft" supervisory tools can have "hard" real effects, quantifying a large corporate investment response to a non-public, qualitative assessment. Moreover, our results provide a blueprint for a new class of regulatory interventions that extend far beyond banking. The qualitative stress test model we analyze is a transferable regulatory technology. Authorities in other jurisdictions and sectors facing similar coordination failures (e.g., energy grids, telecommunications, and critical supply chains)

can adapt this approach to strengthen resilience against systemic threats. Our findings show that focusing supervisory resources on the identified "weakest links" could be a highly effective and efficient strategy.

As finance becomes intertwined with technology, the sources of systemic risk will continue to evolve. Our findings show that regulation must also evolve, shifting from its historical reliance on capital requirements towards more targeted, scrutiny-based interventions that can directly shape banks' incentives. While our results document a strong immediate response, assessing the long-term persistence of these behavioral changes and their ultimate effect on financial stability remains an important avenue for future research.

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 us 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, T., Brolley, M., Cimon, D., and Riordan, R. (2024). Cyber risk and security investment. Technical Report 3781295, Social Science Research Network (SSRN). Available at SSRN: https://ssrn.com/abstract=3781295.

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.

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

Bouveret, A. (2018). Cyber risk for the financial sector: a framework for quantitative assessment. IMF Working Paper 18/143, International Monetary Fund.

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.

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.

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.

Financial Stability Board (2018). Cyber lexicon. Technical report, Financial Stability Board, Basel, Switzerland.

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.

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.

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.

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.

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.

Morris, S. and Shin, H. S. (2002). Social value of public information. *american economic review*, 92(5):1521–1534.

New York State Department of Financial Services (2020). Second annual report on an investigation of the new york insurance industry's cybersecurity programs. Technical report, New York State Department of Financial Services, New York, NY.

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.

Schelling, T. C. (1980). *The Strategy of Conflict: with a new Preface by the Author*. Harvard university press.

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.

Stein, J. C. (2012). Monetary policy as financial stability regulation. *The Quarterly Journal of Economics*, 127(1):57–95.

Figures

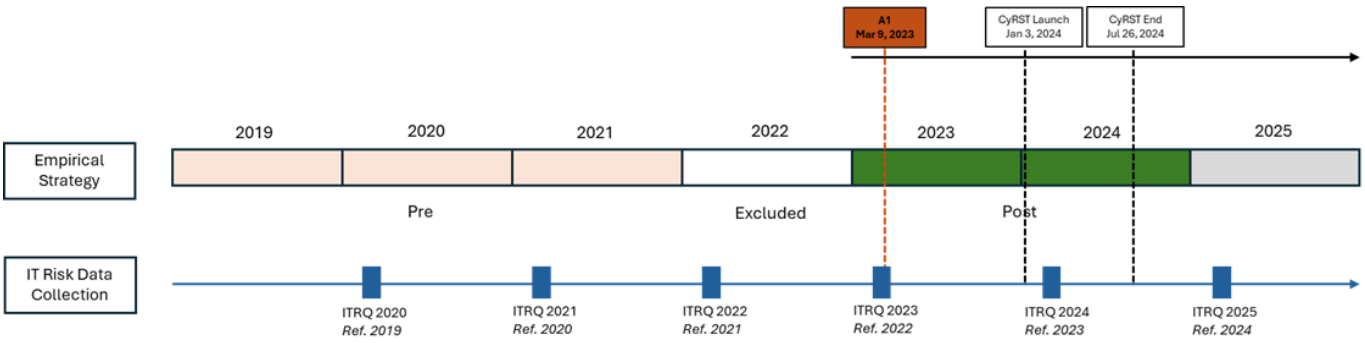

Figure 7: **Timeline.**
This figure illustrates the temporal structure of our identification strategy. The pre-treatment period (2019–2021) is followed by the exclusion of 2022, which coincided with the preparation of the stress test. The post-treatment period begins with the CyRST policy announcement in March 2023, continues with the confidential allocation of banks to high-intensity scrutiny in November 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.

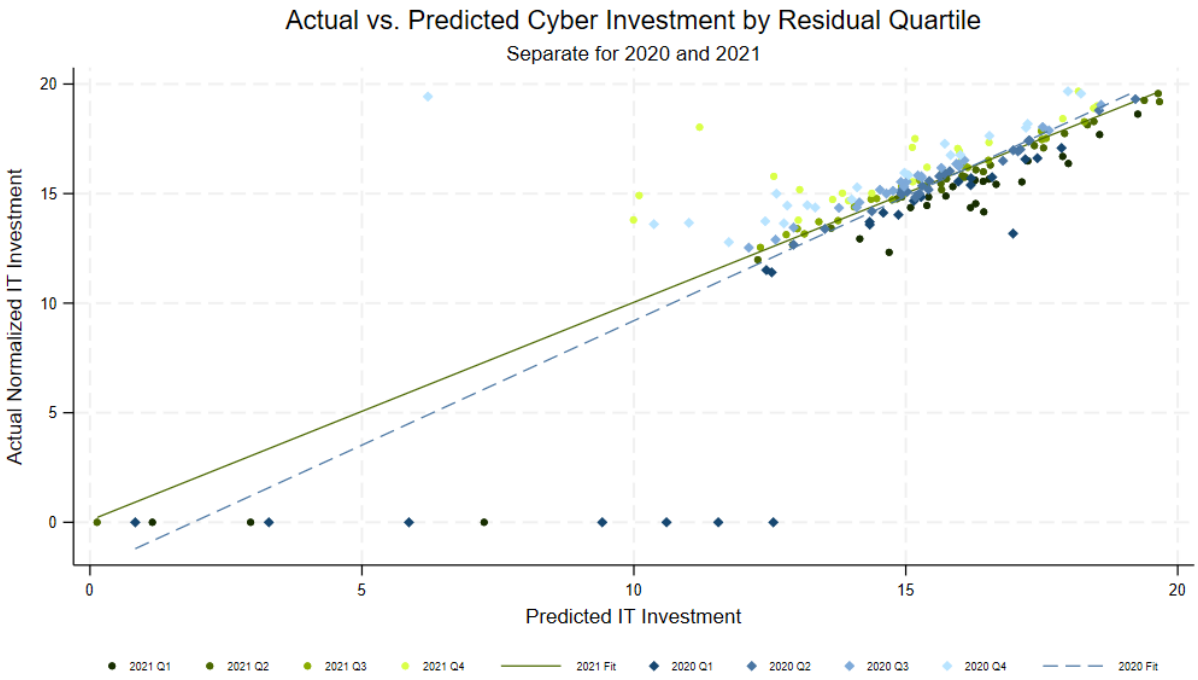

Figure 8: **Residuals normalized IT security investment and predicted IT security investment.**
Actual vs. predicted normalized IT investment for European banks, split by quarter in 2020 (blue, dashed fit) and 2021 (green, solid fit). Each point shows a bank-residual quartile; the fit lines indicate the model's predicted cyber investment accuracy by year, highlighting residual variation that defines under- or over-investing banks.

Appendix I: Variable Definitions

- **Construction of Severe Findings Dummy (`hsti_sevfind`)**: This indicator is constructed to identify banks subject to a high degree of supervisory concern, as evidenced by the substantive findings from the CyRST. The construction is a two-step process.

  Step 1: Weighted Severity Score. First, we create a continuous index of supervisory concern for each bank i. We use the confidential CyRST findings database, which assigns a severity score $s_{ij} \in \{1, 2, 3, 4\}$ to each finding j. We compute the weighted score for bank i as the linear sum of its findings' severities:

$$\text{wt\_severity\_finding}_i = \sum_{j \in \text{findings of bank } i} s_{ij}.$$

  This transparent specification ensures the score captures both the *frequency* of deficiencies (the number of terms in the sum) and their *criticality* (the value of each s_{ij}). For example, a bank with one minor (score=1) and one critical (score=4) finding receives a total score of 5.

  Step 2: Binary Classification. Second, we create the binary indicator by classifying banks based on a non-parametric median split of the scores across the entire sample. This method is robust to outliers in the distribution of the weighted scores. The indicator is formally defined as:

$$\texttt{hsti\_sevfind}_i = \mathbb{1}\left\{\text{wt\_severity\_finding}_i > \text{median}(\text{wt\_severity\_finding})\right\}.$$

  Thus, $\texttt{hsti\_sevfind}_i = 1$ identifies banks in the top half of the sample distribution of findings severity at baseline.

- **Construction of Severe Data-Quality Flags Dummy (`hsti_sevflag`)**:

  Step 1: Weighted Flag Score. The raw data for this measure are the supervisory flags (*Red* or *Amber*) assigned during the CyRST. We map these qualitative assessments to numerical severity weights, s_{iq}, setting *Red*=2 and *Amber*=1. The composite score for bank i is the aggregate of these weights:

$$\text{wt\_severity\_flag}_i = \sum_{q \in \text{flags of bank } i} s_{iq}.$$

  This index quantifies the overall severity of a bank's data governance weaknesses. For instance, a bank with one Red flag (weight=2) and three Amber flags (weight=1 each)

receives a total score of 5.

Step 2: Binary Classification. This continuous score is then dichotomized to create the final binary variable. A bank is classified as having high data-quality concerns if its score lies above the median of the cross-bank distribution:

$$\texttt{hsti\_sevflag}_i = \mathbb{1}\{\text{wt\_severity\_flag}_i > \text{median}(\text{wt\_severity\_flag})\}\,.$$

Accordingly, $\texttt{hsti\_sevflag}_i = 1$ denotes a bank with a baseline level of data-quality issues in the upper half of its peer group.

- **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 5) over the 2020–2021 pre-treatment period, as formally defined in Equation 6.
- **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, insurancecontractsdedu
- *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: Theoretical Model: Complete Framework and Derivations

This appendix provides the complete theoretical framework, including all assumptions, formal derivations, and proofs for the conceptual model presented in Section 4.

II.1 The Economic Environment and Model Primitives

Agents and Actions. The economy consists of a continuum of risk-neutral banks, indexed by $i \in [0, 1]$, and a single, risk-neutral Supervisor. Banks simultaneously choose an investment action, $e_i \in \{0, 1\}$, where $e_i = 1$ denotes investing in a critical cybersecurity measure and $e_i = 0$ denotes not investing. The Supervisor chooses a policy to implement the full-investment equilibrium at minimum inspection cost. This objective is aligned with maximizing social welfare—defined as $W = B \cdot \Omega - \int c \cdot e(c) g(c) dc -$(Total Inspection Costs)—provided the public benefit B is sufficiently large relative to the average investment cost $\mathbb{E}[c]$. Each inspection incurs a cost $K > 0$.

Information Structure. Each bank i is characterized by a private cost of investment, c_i, which is its "type." The cost c_i is private information to bank i and is drawn independently and identically from a common knowledge distribution with a continuous probability density function (PDF) $g(c) > 0$ and a cumulative distribution function (CDF) $G(c)$ on the support $[c_L, c_H]$, where $0 < c_L < c_H$.

Systemic Externality and Payoffs. System-wide cybersecurity integrity, Ω, is modeled as a "weakest-link" public good. This is a standard approach for capturing acute systemic risk where the failure of one critical component can compromise the entire system (e.g., in clearing houses or interbank networks). While stylized, this assumption captures the essence of a regulator's concern about contagion from a single point of failure. Formally:

$$\Omega = \min_{i \in [0,1]} \{e_i\}. \tag{10}$$

A secure system ($\Omega = 1$) generates a public benefit $B > 0$ that accrues to all banks, regardless of their individual investment choice. An insecure system ($\Omega = 0$) provides no such benefit. Due to the weakest-link structure, if any single bank i chooses $e_i = 0$, then $\Omega = 0$. A bank i that shirks ($e_i = 0$) also faces a regulatory penalty $P > 0$ if it is inspected. Let $\mathbb{I}_i \in \{0, 1\}$ be an indicator variable for an inspection of bank i. The ex-post utility for a bank of type

c_i is given by:

$$U_B(e_i, \Omega, \mathbb{I}_i \mid c_i) = \Omega \cdot B - e_i \cdot c_i - (1 - e_i) \cdot \mathbb{I}_i \cdot P. \tag{11}$$

Parametric Assumptions. To ensure a well-defined pre-policy coordination failure and a non-trivial policy problem, we impose the following standard assumptions:

Assumption A1 (Insufficient Private Benefit) $B < c_L$. The social benefit B is insufficient to motivate even the lowest-cost bank to invest unilaterally.

Assumption A2 (Sufficiently High Penalty) $P > c_H$. The penalty for non-compliance is large enough to deter any bank type, provided inspection is certain.

Assumption A3 (Pre-Policy Coordination Failure) $B + q_0 P < c_L$. Under the baseline supervisory regime with a low, uniform inspection probability q_0, the net utility from investing is negative even for the lowest-cost bank, regardless of its beliefs about being pivotal.

II.2 Equilibrium Analysis and Proofs

The solution concept is Bayesian Nash Equilibrium (BNE). A bank's pure strategy is a mapping from its type to its action, $e : [c_L, c_H] \to \{0, 1\}$.

II.2.1 Pre-Policy Equilibrium (Proof of Proposition 1)

Proposition 3 (Pre-Policy Free-Rider Equilibrium)**.** *Under Assumptions A1–A3, the unique Bayesian Nash Equilibrium is for all banks to shirk, i.e., $e^*(c) = 0$ for all $c \in [c_L, c_H]$.*

Proof. A bank of type c chooses $e \in \{0, 1\}$ to maximize its expected utility. A bank's action is pivotal if and only if all other banks invest; let this event be $\mathcal{E}_{-i}$.

The expected utility from investing ($e = 1$) is:

$$\mathbb{E}[U_B(e = 1) \mid c] = \mathbb{E}[\Omega \cdot B - c] = \Pr(\mathcal{E}_{-i}) \cdot B - c. \tag{12}$$

The expected utility from not investing ($e = 0$), which implies an insecure system ($\Omega = 0$), is based solely on the expected penalty from being found non-compliant:

$$\mathbb{E}[U_B(e = 0) \mid c] = \mathbb{E}[-\mathbb{I} \cdot P] = -q_0 P. \tag{13}$$

The net utility of investing is the difference between these two expected utilities:

$$\Delta U(c) = \mathbb{E}[U_B(e = 1) \mid c] - \mathbb{E}[U_B(e = 0) \mid c] = \Pr(\mathcal{E}_{-i}) \cdot B - c + q_0 P. \tag{14}$$

To find a dominant strategy, we consider the most favorable belief for investing, which is that the bank is certain to be pivotal, i.e., $\Pr(\mathcal{E}_{-i}) = 1$. In this case, the net utility is:

$$\Delta U_{\text{pivotal}}(c) = B + q_0 P - c. \tag{15}$$

By Assumption A3, $B + q_0 P < c_L$. Since $c_L \leq c$ for all types in the support, it follows that $\Delta U_{\text{pivotal}}(c) < 0$ for all $c \in [c_L, c_H]$. As the net utility of investing is strictly negative even under the most optimistic beliefs about being pivotal, not investing ($e = 0$) is the strictly dominant strategy for every bank type. Thus, the unique BNE is $e^*(c) = 0$ for all c. □

II.2.2 Post-Policy Equilibrium (Proof of Proposition 2)

Proposition 4 (Policy-Induced Disciplined Equilibrium)**.** *A supervisory policy $\mathcal{P}^* = (s^*, q^*_{target})$ that credibly increases the inspection probability for high-cost banks can uniquely implement the Pareto-superior equilibrium where all banks invest ($e^*(c) = 1$ for all types c).*

Proof. The proof proceeds by construction. We define the Supervisor's policy instruments and optimization problem, solve for the optimal policy, and show that it uniquely implements the full-investment equilibrium.

1. Type-Dependent Inspection Probability. The Supervisor commits to a policy $\mathcal{P} = (s, q_{\text{target}})$ based on a noisy signal of each bank's type, $s_i^{\text{sup}} = c_i + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$. The ex-ante probability of inspection for a bank of type c is:

$$q(c; s, q_{\text{target}}) = q_0 \cdot \Pr(s_i^{\text{sup}} \leq s \mid c_i = c) + q_{\text{target}} \cdot \Pr(s_i^{\text{sup}} > s \mid c_i = c) \tag{16}$$

$$= q_0 \cdot \Pr(\varepsilon_i \leq s - c) + q_{\text{target}} \cdot [1 - \Pr(\varepsilon_i \leq s - c)] \tag{17}$$

Letting $\rho \equiv 1/\sigma$ denote the signal's precision and $\Phi(\cdot)$ be the standard normal CDF:

$$q(c; s, q_{\text{target}}) = q_0 \Phi(\rho(s - c)) + q_{\text{target}}[1 - \Phi(\rho(s - c))] \tag{18}$$

$$= q_{\text{target}} - (q_{\text{target}} - q_0)\Phi(\rho(s - c)). \tag{19}$$

2. Supervisor's Optimization Problem. The Supervisor solves the following cost-minimization problem:

$$\min_{s, q_{\text{target}}} C(s, q_{\text{target}}) = K \int_{c_L}^{c_H} q(c; s, q_{\text{target}}) g(c) dc, \tag{20}$$

subject to two constraints:

1. **Incentive Compatibility (IC):** To induce full investment, the net utility from investing must be non-negative for all types, assuming each bank believes it is pivotal:

$$\Delta U_{\text{post}}(c) = B - c + q(c; s, q_{\text{target}})P \geq 0 \quad \forall c \in [c_L, c_H]. \tag{21}$$

2. **Uniqueness:** To ensure the full-investment equilibrium is unique, the net utility function $\Delta U_{\text{post}}(c)$ must be strictly decreasing in c. Taking the derivative with respect to c:

$$\begin{aligned}
\frac{d}{dc}\Delta U_{\text{post}}(c) &= -1 + P\frac{d}{dc}q(c) && (22)\\
&= -1 + P\frac{d}{dc}\left[q_{\text{target}} - (q_{\text{target}} - q_0)\Phi(\rho(s-c))\right] && (23)\\
&= -1 + P\left[-(q_{\text{target}} - q_0)\phi(\rho(s-c))(-\rho)\right] && (24)\\
&= -1 + P\rho(q_{\text{target}} - q_0)\phi(\rho(s-c)). && (25)
\end{aligned}$$

For uniqueness, we require $\frac{d}{dc}\Delta U_{\text{post}}(c) < 0$, which implies $P\rho(q_{\text{target}} - q_0)\phi(\rho(s-c)) < 1$.

3. Solving the Supervisor's Problem. To solve for the cost-minimizing policy parameters, we first satisfy the constraints. The IC constraint (21) must hold for all types, so it must hold for the highest-cost type, c_H. To minimize cost (and thus q_{target}), the Supervisor sets the policy such that this constraint binds for c_H:

$$B - c_H + q(c_H; s, q_{\text{target}})P = 0 \implies q(c_H; s, q_{\text{target}}) = \frac{c_H - B}{P}. \tag{26}$$

Substituting the expression for $q(c)$ from (19) into (26) yields the cost-minimizing targeted inspection rate, q^*_{target}, as a function of the threshold s:

$$q^*_{\text{target}}(s) = \frac{(c_H - B)/P - q_0\Phi(\rho(s - c_H))}{1 - \Phi(\rho(s - c_H))}. \tag{27}$$

Next, consider the uniqueness constraint. The term $\phi(\rho(s-c))$ is maximized when its argument is zero, i.e., at $c = s$, where the standard normal PDF $\phi(0) = 1/\sqrt{2\pi}$. The condition is therefore tightest at this point and becomes:

$$P\rho(q^*_{\text{target}}(s) - q_0)\frac{1}{\sqrt{2\pi}} < 1 \implies \rho < \frac{\sqrt{2\pi}}{P(q^*_{\text{target}}(s) - q_0)} \equiv \rho^*. \tag{28}$$

This "Precision Paradox" condition requires that signal precision ρ be bounded from above. If monitoring is too precise, the incentive gradient becomes too steep around the threshold s,

which can create non-monotonicities in a bank's net utility of investing and lead to multiple equilibria. Bounding precision ensures the incentive to invest is well-behaved. The Supervisor chooses the optimal threshold s^* by substituting $q^*_{\text{target}}(s)$ into the cost function (20) and solving $\frac{dC}{ds} = 0$.

4. Conclusion of Proof. An optimal policy $\mathcal{P}^* = (s^*, q^*_{\text{target}}(s^*))$ solves this cost-minimization problem. By construction, it satisfies the IC constraint for the highest-cost type c_H, such that $\Delta U_{\text{post}}(c_H) = 0$. The uniqueness condition ($\rho < \rho^*$) ensures that $\Delta U_{\text{post}}(c)$ is strictly decreasing in c. Therefore, for any bank with cost $c_i < c_H$, its incentive to invest is strictly positive: $\Delta U_{\text{post}}(c_i) > \Delta U_{\text{post}}(c_H) = 0$. Since $\Delta U_{\text{post}}(c) \geq 0$ for all $c \in [c_L, c_H]$ and the equilibrium is unique, investing ($e = 1$) is the optimal action for every bank type. □

II.2.3 Derivation of the Heterogeneous Treatment Effect

The model directly motivates the paper's focus on heterogeneous treatment effects. To illustrate this, we can relax Assumption A3 such that $B + q_0 P > c_L$. In this case, an equilibrium with partial investment can exist, where low-cost banks successfully coordinate.

Illustrative Case: A Pre-Policy Equilibrium with Partial Investment. In such an equilibrium, banks that invest must believe they are pivotal. The net utility from investing is $\Delta U_{\text{pivotal}}(c) = B + q_0 P - c$, which defines an investment threshold $c^*_{\text{pre}} = B + q_0 P$. The pre-policy strategy profile is:

$$e_{\text{pre}}(c) = \begin{cases} 1 & \text{if } c \leq c^*_{\text{pre}} \quad \text{(Compliant Banks)} \\ 0 & \text{if } c > c^*_{\text{pre}} \quad \text{(Free-Riding Banks)} \end{cases}$$

Post-Policy Equilibrium and Treatment Effect. As proven above, the optimal policy $\mathcal{P}^*$ induces $e_{\text{post}}(c) = 1$ for all $c \in [c_L, c_H]$. The change in investment strategy for a bank of type c is $\Delta e(c) = e_{\text{post}}(c) - e_{\text{pre}}(c)$. This change is:

$$\Delta e(c) = \begin{cases} 1 - 1 = 0 & \text{if } c \leq c^*_{\text{pre}} \\ 1 - 0 = 1 & \text{if } c > c^*_{\text{pre}} \end{cases}$$

This confirms the core hypothesis: the policy's causal effect on investment behavior, $\Delta e(c)$, is concentrated entirely among the high-cost banks ($c > c^*_{\text{pre}}$) that were free-riding pre-policy.

Appendix III: Alternative Results and Figures

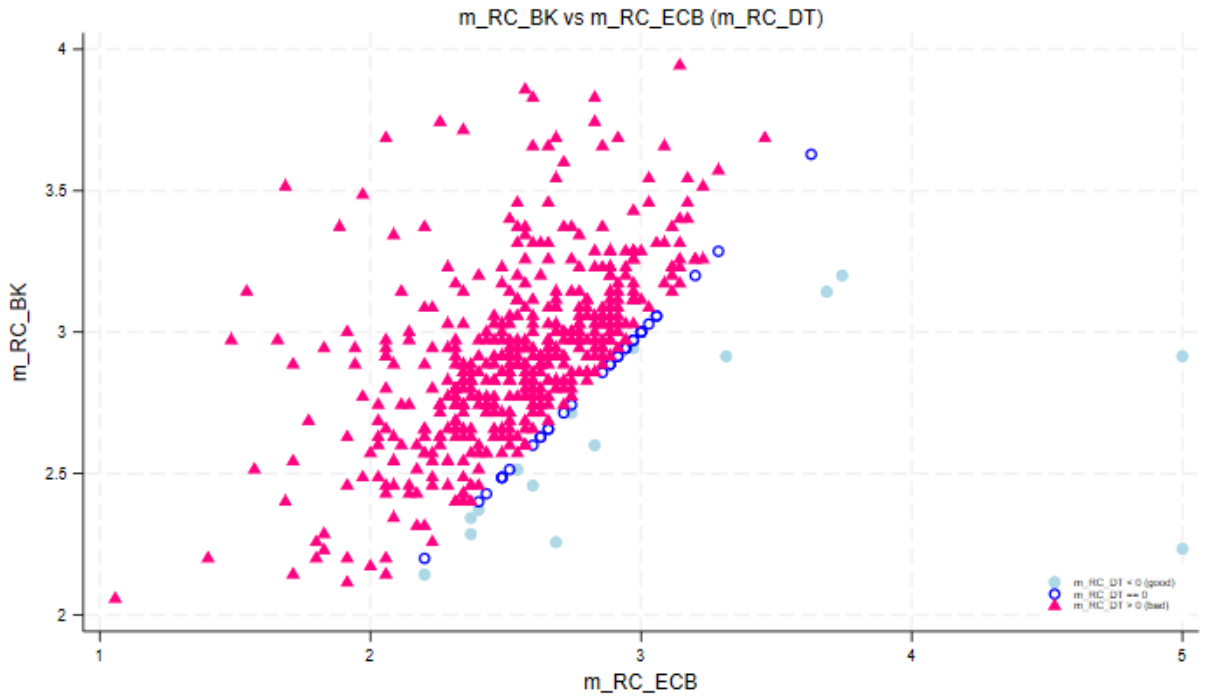

Figure 9: **Supervisory Misalignment in Risk Control Perception: Risk Control Misalignment.**
This scatterplot compares banks' average self-assessed cyber risk control scores (m_RC_BK) with those assessed by the ECB (m_RC_ECB). Both scores are constructed by reversing original risk control indicators (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.

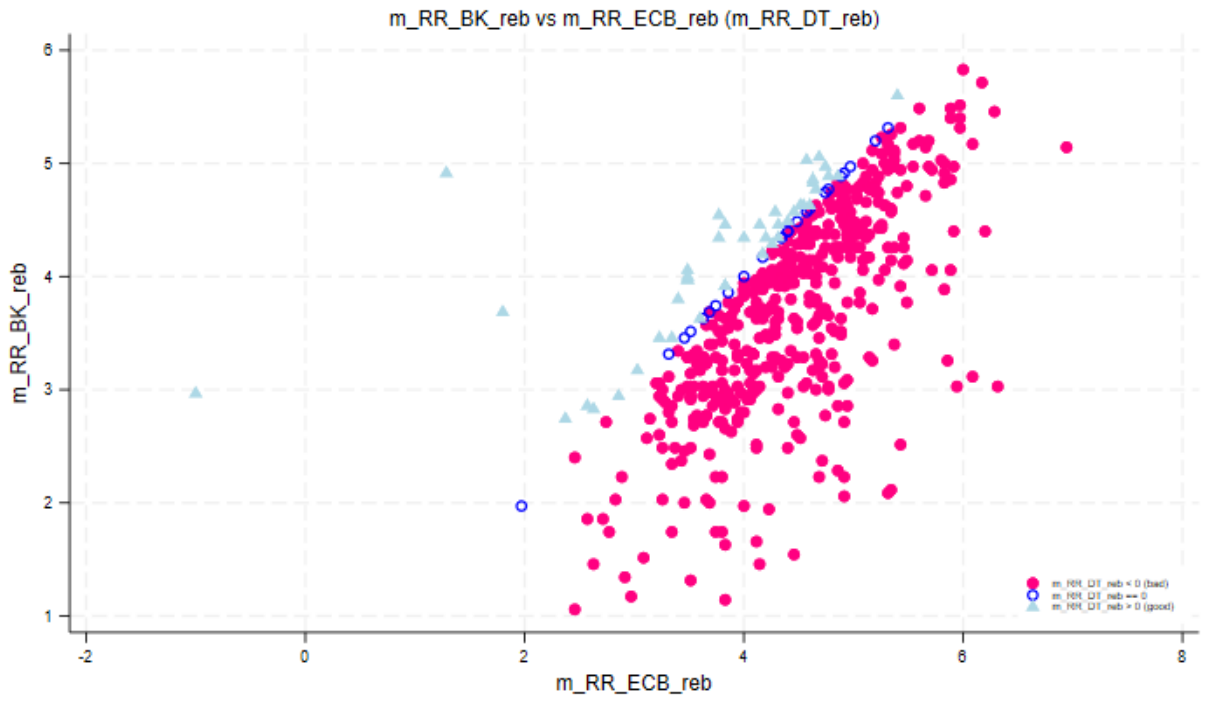

Figure 10: **Supervisory Misalignment in Risk Level Perception: Risk Level Misalignment.**
This scatterplot depicts banks' perceived residual IT risk (*m_RR_BK_reb*) against the ECB's assessment (*m_RR_ECB_reb*). Residual risk is defined as the difference between perceived risk level and risk control, 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.

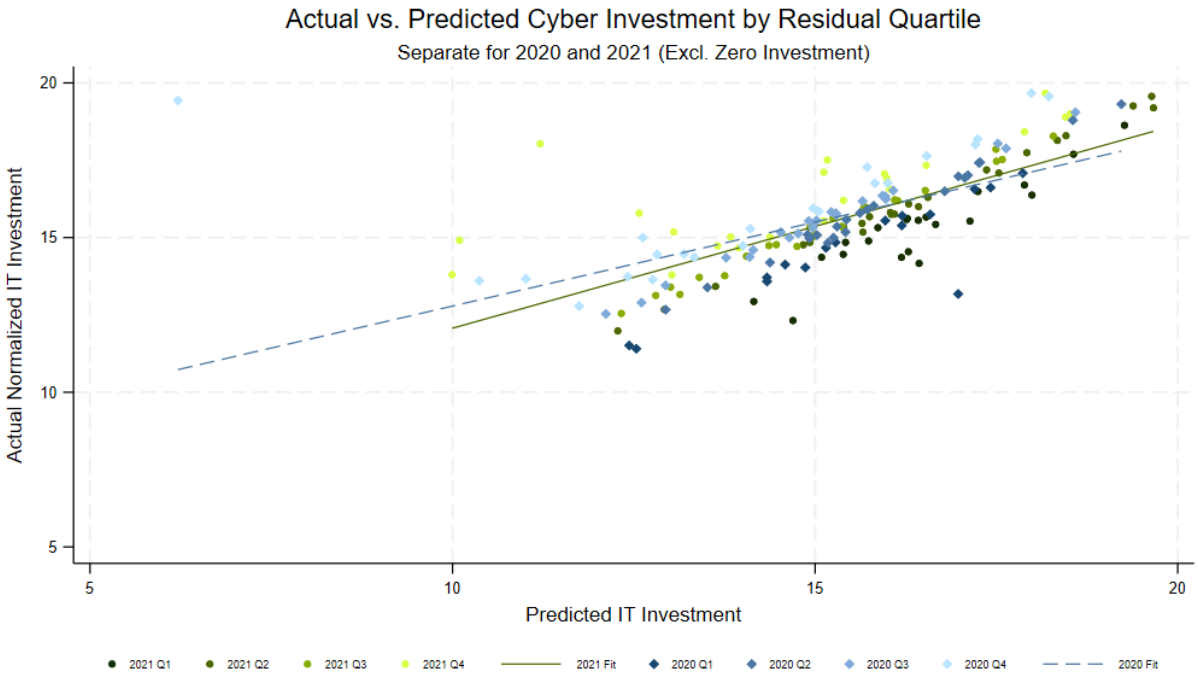

Figure 11: **Actual and Predicted Cyber Security Investment by Residual Quartile.**

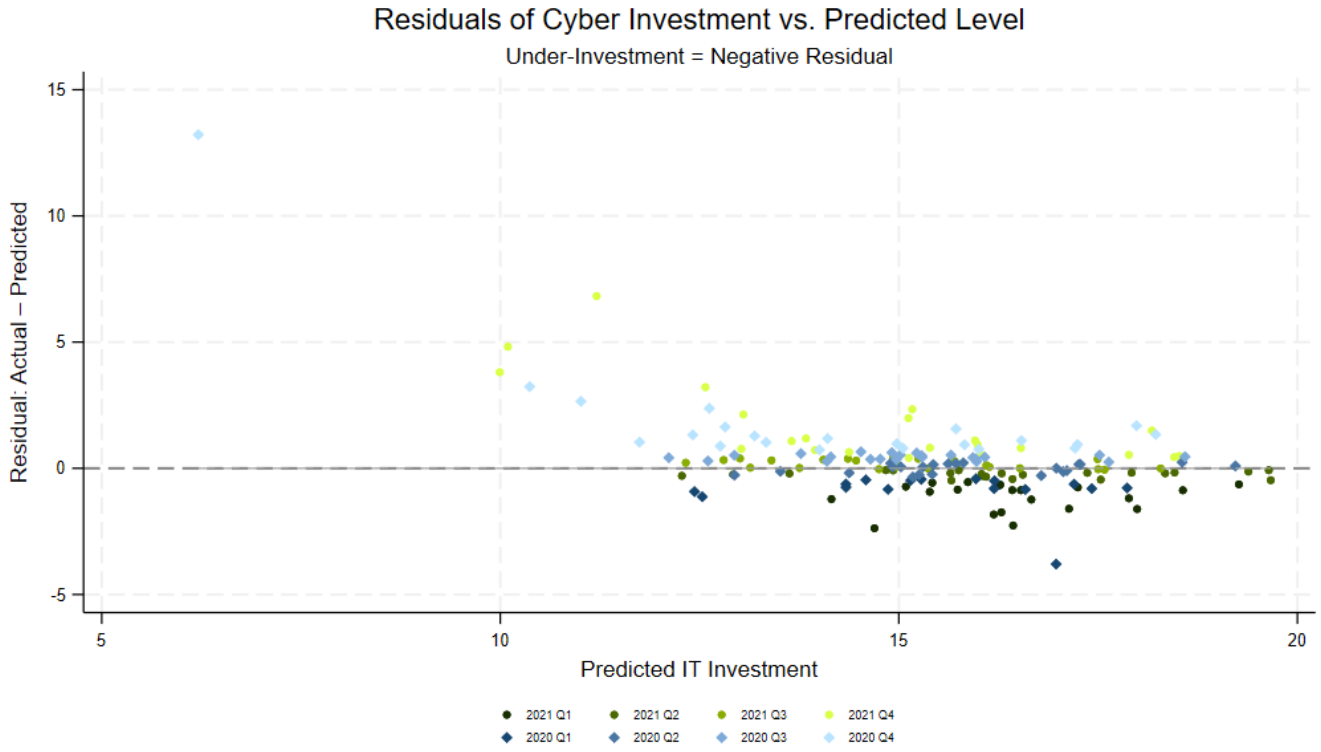

Figure 12: **Residuals of Cyber Investment vs. Predicted Level.**