WORKING paper

Stress Testing European Banks' Transition Risk and Lending Behavior[1]

Olivier de Bandt,[2] David Nefzi,[3] Johannes-Gabriel Werner[4]

August 2026, WP #1059

ABSTRACT

This paper investigates whether supervisory climate stress tests change how euro-area banks price and allocate credit according to borrowers' transition risk. We combine loan-level data from the Eurosystem's AnaCredit database with a firm-level measure of greenhouse-gas (GHG) intensity and compare banks subject to the ECB's 2021-22 climate stress test with non-participating banks. A difference-in-differences-in-slopes design estimates whether the sensitivity of new-loan spreads to carbon intensity changed after the exercise. The average shift in this spread-intensity relationship is small and statistically insignificant. However, decomposing the aggregate effect reveals economically meaningful and offsetting adjustments. Participating banks reallocate new credit away from more carbon-intensive incumbent borrowers, while applying more carbon-sensitive pricing to relationships that enter or exit their portfolios; for new borrowers, a one-standard-deviation increase in carbon intensity is associated with about 10 basis points higher spreads. By contrast, repricing within continuing bank-firm relationships remains limited. The effects are strongest among the most carbon-intensive firms, and participation in the more demanding bottom-up module does not generate a clearly additional response. Overall, climate stress tests appear to operate mainly through portfolio recomposition and extensive-margin pricing, consistent with an information-and-supervisory-scrutiny channel rather than a broad repricing of existing loans.

Keywords: Banking, Climate Stress Tests, Transition Risk, Credit Register, Loan Pricing, Carbon Intensity

JEL classification: C23, E51, E58, G21, G28, G32, Q54

[1] The authors are grateful to Maria Alessia Aiello and Francesca Rinaldi (Banca d'Italia), as well as Fulvio Pegoraro (Autorité de Contrôle Prudentiel et de Résolution), for their insightful advice and constructive feedback. We thank for their valuable comments and engaging discussions the participants in: the Basel Committee on Banking Supervision (BCBS) Washington Meeting in May 2025, the Bundesbank Banking & Financial Research Seminar, the Research Conference on Climate- and Nature-Related Risks in the Economic and Financial System, held in Paris in March 2026, with special thanks to the discussant, Klaus Schaeck, and the French Economic Association Annual Congress (Nantes, 2026). Furthermore, we are indebted to Jean-Guillaume Sahuc for comments on an earlier draft of this article, as well as Valérie Mignon for her invaluable advice and continuous support.

[2] Banque de France, olivier.debandt@banque-france.fr

[3] Banque de France and Université Paris-Nanterre, david.nefzi@banque-france.fr;

[4] Deutsche Bundesbank, johannes-gabriel.werner@bundesbank.de

NON-TECHNICAL SUMMARY

Banks play a central role in the transition to a low-carbon economy. They finance investment in cleaner technologies, but they are also exposed to losses if climate policies, technological change, or shifts in consumer preferences reduce the value of carbon-intensive activities. Supervisors have therefore introduced climate stress tests to assess banks' preparedness, improve data and modelling, and identify vulnerabilities. Although these exercises were presented as learning tools rather than mechanisms that directly affect capital requirements, the information they generate and the accompanying supervisory scrutiny may nonetheless influence lending decisions.

The paper studies the ECB's 2021-22 climate stress test and asks whether participating banks changed the way they price and allocate new loans according to borrowers' greenhouse-gas intensity. The analysis combines more than two million observations on newly originated loans from the Eurosystem's AnaCredit register with a firm-level emissions-intensity measure. It compares significant banks that participated in the ECB exercise with smaller non-participating institutions, before and after the stress test, while accounting for the earlier French pilot exercise and controlling for differences across banks, firms, loan products, and time.

The headline result is that the average sensitivity of loan spreads to carbon intensity changes very little following the stress test. Taken in isolation, this near-zero effect could suggest that the exercise had little influence. However, the decomposition developed in the paper reveals a different picture. Banks subject to the stress test do not systematically reprice continuing bank-firm relationships. Instead, they adjust primarily through two channels that work in opposite but economically consistent directions. First, they reallocate new lending within their existing customer base away from more carbon-intensive firms. Second, they apply more carbon-sensitive pricing to relationships entering or leaving their portfolios. For new borrowers, a one-standard-deviation increase in carbon intensity is associated with loan spreads that are approximately 10 basis points higher. These effects are stronger among the highest-emitting firms.

The figure illustrates why the overall coefficient is close to zero. The negative contribution from reallocating credit among incumbent borrowers is largely offset by the positive contribution from entry-and-exit pricing, while repricing within continuing relationships and changes in borrower composition play only a minor role. Specifications that account for the different timing of the French and ECB exercises yield more moderate estimates but continue to identify entry-and-exit pricing as the main adjustment channel. Among significant banks, which were required to participate in the stress test, involvement in the most demanding bottom-up module does not generate a noticeably stronger response than participation in the lighter modules.

Overall, the results suggest that climate stress tests can influence bank behaviour even without an explicit 'brown' capital surcharge. Their main effect is not a uniform increase in borrowing costs for all carbon-intensive firms, but rather a gradual reallocation of lending portfolios combined with more risk-sensitive pricing at the margins of bank-firm relationships. This pattern is consistent with an information-and-scrutiny channel: the exercise makes transition risk more salient and encourages banks to reconsider which firms they finance. The findings should nevertheless be interpreted with caution because firms' emissions are estimated rather than directly observed. Although other climate-related supervisory measures were introduced over the same period, the analysis carefully controls for these concurrent policies. A wide range of robustness checks, based on state-of-the-art econometric method, confirms the main results.

For supervisors, the results imply that stress tests should provide sufficiently granular information to influence both loan pricing and portfolio allocation, and that their effects should be assessed using more than average loan spreads alone. The paper provides a clear empirical framework for assessing the impact of prudential regulation at the granular level.

How the climate stress test changes loan pricing and allocation

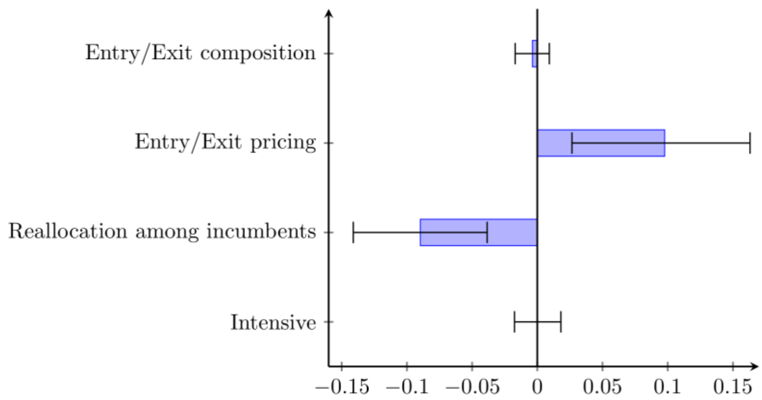

Note: The bars decompose the estimated post-test change in the sensitivity of loan spreads to borrowers' carbon intensity. The negative contribution from credit reallocation among incumbent borrowers is offset by the positive contribution from pricing at relationship entry and exit; while the remaining channels (credit allocation to new relationships, first row, and repricing of incumbent relationships, last row) are close to zero. The vertical lines denote the 95% confidence interval.

Tests de résistance climatique des banques européennes : risque de transition et comportement de prêt

Résumé

Cet article examine si les tests de résistance climatique modifient la tarification et l'allocation du crédit des banques de la zone euro en fonction du risque de transition des emprunteurs. Nous combinons les données de prêts de la base AnaCrédit, mise en place par l'Eurosystème, avec une mesure de l'intensité des émissions de gaz à effet de serre (GES) des entreprises non financières. Nous comparons les banques ayant participé au test de résistance climatique 2021-2022 de la BCE à des banques non participantes. Une approche de différences de différences appliquée aux pentes permet de mesurer l'évolution de la sensibilité des marges sur les nouveaux prêts à l'intensité carbone des emprunteurs. En moyenne, cette sensibilité évolue peu et la variation n'est pas statistiquement significative. La décomposition de l'effet agrégé révèle toutefois des ajustements importants, de sens opposé mais économiquement cohérents. Les banques participantes réallouent les nouveaux crédits au détriment des emprunteurs existants les plus intensifs en carbone et appliquent une tarification davantage sensible au carbone aux relations qui entrent dans leur portefeuille ou en sortent. Pour les nouveaux emprunteurs, une hausse d'un écart-type de l'intensité carbone est associée à une augmentation d'environ 10 points de base du spread de taux. En revanche, la révision des conditions tarifaires tarifaires au sein des relations bancaires existantes demeure limitée. Les effets sont particulièrement marqués pour les entreprises les plus émettrices. Les tests de résistance climatique semblent donc agir principalement par une recomposition des portefeuilles de crédit et une tarification plus sensible au risque de transition à la marge extensive, conformément à un canal d'information et de surveillance prudentielle.

Mots-clés : banques, tests de résistance climatique, risque de transition, registre de crédit, tarification des prêts, intensité carbone

---

1 Introduction

Banks play a central role in the low-carbon transition. They finance the large investment required in cleaner technologies, while their balance sheets remain exposed to transition and physical risks that may threaten financial stability. For central banks and supervisory authorities, the key question is not only how large these climate and environmental (C&E) risks are in the cross-section, but also whether existing prudential tools can steer banks' risk management and lending decisions in a more climate-sensitive direction.

Building on the work of the Network for Greening the Financial System (NGFS), it is now standard to distinguish between transition risks, physical risks and combinations thereof, depending on the timing and ambition of climate-policy action. These risks can disrupt global value chains, depress output, push up prices and ultimately feed back into the financial system. In Europe, supervisory authorities have for several years articulated expectations regarding the integration of C&E risks into banks' governance, risk-management frameworks and disclosures, although these expectations were initially principles-based and non-binding.

Centrally conducted climate stress tests (CST) have emerged as a key instrument to operationalise this agenda. Supervisors use them to assess banks' preparedness to manage C&E risks, to support capacity building in terms of data, modelling and governance, and to gauge the vulnerability of the banking system to transition and physical-risk scenarios. Since the Great Financial Crisis, stress tests have become a central tool for monitoring the resilience of banks to financial shocks in Europe and the United States (Borio, 2011; Flannery, Hirtle, and Kovner, 2017; IMF, 2021). They are now being extended to climate risks, notably by the Autorité de Contrôle Prudentiel et de Résolution (ACPR)[1], which ran a pilot CST in 2020–21, and by the European Central Bank (ECB), which conducted an EU-wide CST on significant institutions (SIs) in 2021–22 (Alogoskoufis, Dunz, Emambakhsh, Hennig, Kaijser, Kouratzoglou, Muñoz, Parisi, and Salleo, 2021). Jung, Santos, and Seltzer (2023) report similar developments in the United States, while Baudino and Svoronos (2021) and Acharya, Berner, Engle, Jung, Stroebel, Zeng, and Zhao (2023) survey climate stress-testing initiatives more broadly.

Formally, supervisors framed these exercises as stock-taking and learning tools. They aimed to map banks' exposures to C&E risks and to assess the adequacy of their risk-management practices, rather than to directly influence portfolio choices or regulatory capital. Whether CST nevertheless affect banks' behaviour remains an open question. By forcing banks to collect granular emissions data, to build forward-looking climate scenarios and to present quantitative results to supervisors, CST may change internal risk perceptions and the salience of climate risk in pricing and credit allocation. Banks may then adjust the way they price loans to more carbon-intensive borrowers or reallocate credit away from these borrowers, even in the absence of immediate capital consequences. While the ECB has reiterated that the CST was not intended to feed directly into capital requirements, it has also indicated that persistent non-compliance with supervisory expectations on C&E risks may trigger periodic penalty payments (Elderson, 2024, 2023), strengthening the incentive for banks to take the exercise seriously.[2]

[1] The French prudential supervisor.

[2] Recently, the ECB has imposed periodic penalty payments against a European bank for failure to comply with ECB C&E risk-management expectations; see: https://www.bankingsupervision.

This paper asks whether climate stress tests have in fact altered banks' lending behaviour along the transition-risk dimension. Using the Eurosystem's granular credit register AnaCredit, we study the impact of the ECB-wide Climate Stress Test (ECST), launched in October 2021, while controlling for the French Climate Stress Test (FCST) conducted in 2020, on how banks price loans as a function of borrowers' climate-related transition risk, proxied by firms' greenhouse-gas (GHG) emissions intensity.[3] Our main outcome variable is the credit spread on new loans granted by banks to non-financial corporations. Focusing on new originations is crucial, as loan pricing conditional on risk occurs at this stage.[4]

Our empirical analysis relies on a causal inference framework utilizing a difference-in-differences approach in slopes. While the CST was conducted concurrently with an array of climate initiatives led by the ECB and the Single Supervisory Mechanism (SSM), the unique characteristics of the stress test of 2022 (an information shock) validate our empirical design.[5] We estimate how the semi-elasticity of credit spreads with respect to borrower GHG intensity changes for CST-participating banks (treated SIs) relative to non-participating banks (LSIs), before and after the CST. We rely on Veron (2024) who argues that in many respects, LSIs can be seen of institutions that are very similar to SIs: one case in point are savings banks which are LSIs in Germany and Italy, while in France they belong to BPCE, a SI bank. We take avantage of the fact that only SI banks participated in the European Stress Test of 2022. In practice, this amounts to estimating a triple-difference specification in which the credit-spread–intensity relationship interacts with both treatment and post-CST indicators, while we control for rich sets of fixed effects at the bank, borrower, product and time dimensions.

Main findings. Our analysis yields three main empirical results. First, in the baseline non-staggered specification that treats the ECST as a single shock hitting all SIs in October 2021, we find that the average change in the spread–intensity slope for treated banks relative to controls is small and imprecisely estimated. However, a structural decomposition of this triple-difference estimand reveals economically meaningful offsetting margins: a negative reallocation among incumbent borrowers (treated banks tilt origination within their continuing relationships away from more carbon-intensive clients) is almost exactly offset by a positive entry/exit pricing component (treated banks charge more carbon-sensitive spreads on transitory relationships than on incumbents), while within-relationship repricing and entry/exit composition effects are negligible. In other words, the aggregate effect of the CST on the spread–intensity slope is close to zero, but it masks substantial portfolio rebalancing and extensive-margin pricing responses.

Second, when we move to a staggered design that respects the timing of the FCST and ECST waves, the qualitative picture is preserved but the magnitudes are more modest. An interaction-weighted estimand that averages cohort-specific effects over both CST

europa.eu/press/pr/date/2025/html/ssm.pr251110~3e0b6f579e.en.html.

[3]Since climate stress-test scenarios model transition risk primarily through higher carbon prices, GHG intensity is a natural proxy for firms' exposure to these risks. This abstraction ignores factors that may preserve the viability of some high-intensity business models, such as limited substitutability of their output, the ability to pass through costs, or rapid technological adjustment.

[4]Prolongations with renegotiated terms are recorded as new contracts in AnaCredit.

[5]The European stress test was compulsory, while the French stress test of 2021 , and the Dutch pilot stress test of 2021 was only run by the supervisor, without bank participation

waves indicates that the intensive margin remains small and statistically insignificant, the average reallocation among incumbents is close to zero, and the main contribution continues to come from the entry/exit pricing channel, which is positive but smaller than in the single-shock specification. This suggests that once we account for staggered adoption and dynamic cohort-adjustment paths, climate stress testing still promotes more carbon-sensitive pricing on the extensive margin, but the strength of the response is attenuated and spread out over time.

Third, within the universe of SIs we exploit variation in treatment intensity across modules of the 2022 ECST. Comparing banks that participated in the more demanding bottom-up Module 3 with those that only contributed to the lighter Modules 1 and 2, we find limited evidence of additional within-relationship repricing and only a modest incremental reallocation toward less carbon-intensive incumbents. This suggests that the main lending adjustments are associated with CST participation *per se* and the accompanying supervisory scrutiny, while the marginal effect of moving from a questionnaire-based to a fully modelled bottom-up exercise on loan-level pricing is relatively small over our sample horizon.

Contributions. Our contributions are both substantive and methodological. First, we use a very granular dataset of more than two million observations on new loans, constructed from the Eurosystem credit register. In contrast to most existing work, which focuses on outstanding exposures to large listed firms (either in AnaCredit (Altavilla, Boucinha, Pagano, and Polo, 2023) or in syndicated loan markets (Fuchs, Ngyuen, Nguyen, and Schaeck, 2024)), we study the pricing of newly originated bank credit across a much broader population of borrowers.

Second, we construct a firm-level GHG intensity metric explicitly designed to proxy exposure to transition risk at the country–sector–size–year level, which allows us to extend the analysis beyond a narrow set of large emitters and to document that climate stress tests affect the entire loan book rather than a small *brown tail* only.

Third, we develop a simple theoretical model in which a climate stress test acts as an information shock about firms' transition risks and supervisory scrutiny. The shock shifts banks' perceived distribution of climate-related losses and, through standard no-arbitrage and capital-constraint channels, the equilibrium bank lending rate as a function of borrowers' carbon intensity. The model clarifies under which conditions a CST should steepen, flatten, or leave unchanged the spread–intensity slope and guides the empirical interpretation of our coefficients.

Fourth, we take these predictions to the data with a difference-in-differences-in-slopes design and show, to the best of our knowledge, for the first time, how a standard triple-difference can be exactly decomposed into four economically meaningful channels. Using different samples and inverse-probability weighting, we separately identify (i) within-relationship repricing for continuing bank–firm pairs (intensive margin), (ii) reallocation of credit among incumbent borrowers, and two extensive margins—(iii) entry/exit pricing at a composition-adjusted portfolio and (iv) entry/exit composition in observables driven by differences between transitory and incumbent relationships. This decomposition makes it possible to distinguish a genuine steepening of the spread–intensity curve from portfolio recomposition and selection effects.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature, and Section 3 describes the institutional and supervisory setting. Section 4 presents a theoretical model highlighting information revelation through the stress test. Section 5 details the construction of the data. Section 6 sets out the baseline empirical strategy and main results. Section 7 reports robustness checks of the baseline strategy. Section 8 offers a further extension by examining the role of variation in treatment intensity across ECB-wide climate stress test modules. Section 9 concludes.

2 Related Literature

There is some evidence that banks' loan supply responds to climate change related transition risk (see Acharya et al. (2023) and de Bandt, Kuntz, Pankratz, Pegoraro, Solheim, Sutton, Takeyama, and Xia (2023) for surveys of the literature). Banks are shown to reduce the maturity of loans (Ivanov, Kruttli, and Watugala, 2024). Delis, Greiff, Iosifidi, and Ongena (2024) conclude that banks charge higher loan rates to fossil fuel firms. Jung et al. (2023) show that banks in the US which signed the Net-Zero Alliance reduced their exposures to the riskiest borrowers. However, de Bandt et al. (2023) conclude that the effect of climate change is relatively limited: on the basis of 10 academic papers concentrating on bank loan supply on a variety of loan portfolios, they report that most studies estimate credit spreads increase in the range of 25-50 bp following a transition risk shock, while only two studies identify a more significant impact, above 100bp. In addition, some of these estimates are based on small samples of observations, the transition shocks are not fully comparable across studies, some of them do not control for confounding factors, in particular borrower-specific credit risk. Furthermore they often fail to capture the most recent acceleration of climate change.

Since the 2007–2009 global financial crisis, stress testing has become a cornerstone of both microprudential and macroprudential regulatory frameworks, prompting a substantial expansion of the economic literature evaluating its impact on banking behavior and credit allocation. Stress tests aim to assess banks' resilience under adverse economic scenarios and enforce adjustments in risk-taking behavior, leading to a variety of observed effects on credit supply and financial stability.

The U.S. Supervisory Capital Assessment Program (SCAP) and the Comprehensive Capital Analysis and Review (CCAR) have played key roles in reshaping banks' lending behavior. Acharya, Berger, and Roman (2018) find that banks undergoing these stress tests reduce credit supply to relatively risky borrowers, prioritizing risk reduction over seeking higher spreads. Similarly, Connolly (2021) reports that the SCAP led to structural adjustments in syndicated loan markets, with tested banks reducing their participation. However, substitution by European and non-tested U.S. banks mitigated adverse credit effects, ensuring continued access to financing for most firms, except those highly dependent on banks that failed the stress tests.

Stress tests also influence portfolio composition. Bräuning and Fillat (2020) find that banks with poor stress-test results align their portfolios with better-performing banks, enhancing diversification across sectors and regions. While this reduces individual risk, it creates systemic concentration at the aggregate level. Moreover, weaker stress-test results lead banks to reduce lending, particularly for loans sensitive to adverse scenarios, causing credit and investment cutbacks among reliant firms. This highlights a macroprudential effect of stress testing: it restrains credit growth during expansions but strengthens banks' resilience during downturns.

In Europe, Konietschke, Ongena, and Marques (2022) show that stress tests prompt banks to reallocate credit from riskier to safer borrowers, particularly in the household sector. While this enhances financial stability, it reduces profitability, with the effect more pronounced among banks subject to undisclosed stress tests. Transparency emerges as a key factor: banks disclosing their capital requirements maintain stronger capital ratios, reinforcing market discipline and financial stability.

Stress testing also affects specific credit markets. Cortés, Demyanyk, Li, Loutskina, and Strahan (2020) document a shift in U.S. banks' credit supply away from riskier markets, particularly small business loans, following stress-test-induced capital requirements. Banks raise rates in markets where they possess local knowledge and withdraw where they do not, leaving room for smaller banks to increase their market share. Similarly, Calem, Correa, and Lee (2020) show that CCAR stress tests reduced jumbo mortgage originations and approval rates, particularly at banks with weaker capital positions. Additionally, the 2013 Supervisory Guidance on Leveraged Lending reduced speculative-grade loan originations at regulated banks, highlighting the broader influence of prudential policies on credit allocation and risk management.

Overall, the literature underscores the dual impact of stress testing: while it improves individual bank resilience and systemic stability, it can also constrain credit supply in certain markets, particularly for riskier borrowers. Building on the findings from traditional stress tests and prudential policies, recent research has increasingly focused on how supervisory interventions aimed at managing climate risks, such as CST, affect banks' credit allocation and shape their role in advancing green finance.

Against the backgound of the 2021 French CST, using a DiD framework, Fuchs et al. (2024), concentrating on syndicated loans, demonstrate that stress-tested banks increase lending to carbon-intensive borrowers but at higher interest rates, with some environmental improvements tied to borrowers' commitments to emission reduction, suggesting a causal link between climate stress tests and reduced transition risk. Similarly, Beyer and Schreiner (2024), also using a DiD framework, find that the most recent developments in ECB's climate-risk supervision measures significantly enhance banks' risk management and green finance activities, particularly among SIs, evidenced by greater green bond issuance and ESG assets under management, and green lending. Expanding on this, Aiello (2024) explores how the Single Supervisory Mechanism's (SSM) climate supervision leads to credit reallocation towards lower-emission firms, especially by banks with climate commitments. However, Aiello (2024) finds that high-emission borrowers' future reduction pledges have limited influence on banks' credit decisions, implying that banks prioritize current emissions over forward-looking information in their climate risk assessments. Along the same lines, Gschossmann, Kok, and De Cicco (2025) provide another empirical evaluation of how CST influence credit allocation within the euro area banking sector. Using loan-level data from AnaCredit and a DiD design exploiting the ECB CST, the authors find that participating banks significantly reduced lending to GHG–intensive firms relative to non-participants. The contraction is most pronounced for smaller borrowers, suggesting that transition risk aversion is mediated by firm size and relationship strength. Crucially, only banks that performed best in the stress test show a measurable reduction in brown lending ex post, implying that managerial sophistication conditions behavioral change. From a supervisory perspective, the study underscores that climate stress tests can serve as a soft instrument[6] for aligning bank portfolios with transition objectives, even absent binding regulatory consequences. However, the effect is heterogeneous: weaker-performing banks do not significantly adjust their exposures, highlighting the limits of voluntary or non-punitive supervisory frameworks. Our paper echoes and complements Gschossmann et al. (2025) approach. Altavilla et al. (2023) add a monetary policy perspective, identifying a climate risk-taking channel where banks charge higher interest rates to high-emission firms, particularly during tighter monetary conditions, which amplifies both credit and carbon premia. This results in a stronger contractionary impact on high-emission firms compared to those with lower emissions or emission-reduction commitments. In a different regulatory setting, Miguel, Pedraza, and Ruiz-Ortega (2024) examine Brazil's environmental capital requirements, showing that large banks reduce credit to environmentally exposed sectors, although smaller banks partially offset this contraction. This policy-driven reallocation does not significantly affect sectoral output or emissions.

3 Institutional Setting: Climate Change and European Banking Supervision

3.1 Significant and Less Significant Institutions

Within the European banking supervision framework, banks are designated as Significant Institutions (SIs) when they meet at least one of the criteria set out in the SSM Framework Regulation.[7] These criteria relate in particular to a bank's size, its importance for the national or EU economy, and the scale of its cross-border activities. Banks that do not meet these criteria are classified by default as Less Significant Institutions (LSIs).

This distinction determines the allocation of supervisory responsibilities. LSIs, generally small and medium-sized banks, remain directly supervised by national competent authorities (NCAs), under ECB oversight. By contrast, SIs are directly supervised by the ECB and account for the majority of banking assets in participating countries. The SI/LSI distinction is therefore essential for understanding the scope of climate-related supervision: some initiatives are addressed to the banking sector as a whole, whereas others are implemented only for institutions under direct ECB supervision. Nevertheless, as argued below, the distinction is to some extent institutional, and similar types of banks may be LSIs in some countries but SIs in others (e.g. savings banks in Germany are LSIs, while they belong to BPCE in France).

As direct supervisor of SIs, the ECB conducts several types of stress tests, including EU-wide exercises coordinated by the European Banking Authority (EBA), ECB-specific stress tests within the Supervisory Review and Evaluation Process (SREP), and thematic stress tests focused on particular risks.[8] In the absence of an EBA-wide exercise, thematic

[6] Similar to soft law mechanisms.

[7] For more details on the criteria, see: https://www.bankingsupervision.europa.eu/banking/list/criteria/html/index.en.html

[8] For more details on stress tests, see: https://www.bankingsupervision.europa.eu/banking/tasks/stresstests/html/index.en.htm

stress tests are used to assess SIs under common shock scenarios. This institutional framework explains why the European climate stress test was implemented on SIs rather than LSIs.

3.2 ECB Communication and Supervisory Guidance on Climate Risks

Since January 2020, the ECB has progressively integrated climate and environmental (C&E) risks into its supervisory framework. Its approach is based on two complementary components. The first is broad supervisory communication to banks, including guidance, reviews, feedback letters, speeches and supervisory dialogue. The second is the European climate stress test, a more targeted supervisory exercise applied to Significant Institutions.

The first component relies on communication and supervisory expectations addressed to the banking sector as a whole. The central reference is the *Guide on Climate-related and Environmental Risks*, jointly developed by the ECB and NCAs (ECB, 2020). The Guide sets out supervisory expectations on how banks should identify, assess, manage and disclose C&E risks. It requires institutions to conduct materiality assessments across their portfolios and to integrate climate risks into governance, strategy, risk management and capital adequacy processes.

Although the Guide plays a central role in the ECB's dialogue with SIs, it also provides a common benchmark for banks across the euro area. The objective was not to impose immediate quantitative capital requirements, but to clarify how climate risks should be understood and managed. In this sense, the Guide marked the starting point of the ECB's broader communication strategy on climate risks.

In early 2021, the ECB asked several SIs to assess themselves against the expectations in the Guide (ECB, 2021). Banks had to compare existing practices with supervisory expectations and submit implementation plans to close identified gaps. The ECB then issued individual feedback letters in September 2021, including qualitative requirements for some banks as part of the 2021 SREP. This process was intended to make banks engage with C&E risks in a structured way and to provide supervisors with information on their preparedness.

In 2022, the ECB extended this approach through a thematic review covering several SIs and LSIs. The review assessed whether banks were aligning their practices with the Guide and showed substantial deficiencies (ECB, 2022). In response, the ECB set institution-specific deadlines for full compliance by the end of 2024. Interim milestones were set until the end of 2024.

The ECB also reinforced these expectations through public communication. On November 14, 2023, Frank Elderson (2024), Member of the Executive Board of the ECB and Vice-Chair of the Supervisory Board, delivered a landmark speech reaffirming the ECB's commitment to the rigorous management of C&E risks.

3.3 The European Climate Stress Test

The second component of the ECB's approach is the European climate stress test (CST). Unlike the broad communication channel described above, the 2022 CST was implemented on Significant Institutions under the ECB's direct supervision and did not apply to LSIs.

This follows directly from the institutional distinction between SIs and LSIs: SIs are supervised directly by the ECB, whereas LSIs are supervised by NCAs under ECB oversight.

The 2022 CST was the first dedicated supervisory climate stress test conducted at the euro-area level. It was part of the broader effort to integrate C&E risks into the prudential framework, but it differed from general supervisory communication because it required SIs to assess climate-related exposures and potential losses under common scenarios. The exercise was explicitly framed as a learning and capacity-building tool rather than a capital-adequacy test. Its results fed into the SREP only in qualitative terms, and no automatic capital add-ons were imposed on the basis of projected losses (Alogoskoufis et al., 2021).

The CST had three main objectives: to assess the vulnerability of euro area SIs to climate-related risks, to evaluate banks' data and modelling capabilities, and to improve supervisors' and banks' understanding of how transition and physical risks affect balance sheets and income statements. To achieve these objectives, the exercise was organised into three modules. Module 1 consisted of a qualitative questionnaire on banks' climate stress-testing and risk-management capabilities, including governance, risk appetite, internal models and ICAAP processes. Module 2 collected harmonised climate-risk metrics at the exposure level, including sectoral and geographical exposures to high-emitting activities and financed emissions. Module 3 was a bottom-up stress test applied to a subset of 41 SIs, selected under a proportionality principle, requiring banks to project credit and market risk losses under common transition and physical risk scenarios (Alogoskoufis et al., 2021).

This architecture allowed supervisors to combine a system-wide view of climate-risk vulnerabilities with detailed information on banks' internal preparedness. The results pointed to non-negligible climate-related credit and market losses, estimated at around €70 billion across the short-term exercises, but also revealed that banks' data and modelling infrastructures remained at an early stage. Many institutions had limited risk coverage, insufficiently granular data and models that captured only part of the channels through which climate risks may affect financial risks.

The CST therefore functioned less as a traditional solvency test than as a supervisory experiment. Its purpose was to reveal where climate risks are concentrated on banks' balance sheets and to identify gaps in banks' capacity to measure and manage those risks. By requiring SIs to conduct detailed, exposure-level assessments, the CST generated information that could not be obtained through general supervisory communication alone.

Before the European CST, the French supervisory authority, the Autorité de Contrôle et de Résolution (ACPR), conducted a climate pilot exercise between July 2020 and April 2021. That exercise differed in scope because it included both banking groups and insurance companies and was conducted on a voluntary basis. However, it shared with the ECB CST a focus on long-term climate scenarios, sectoral analysis, and the assessment of both transition and physical risks (Clerc, Bontemps-Chanel, Diot, Overton, Soares de Albergaria, Vernet, and Louardi, 2021).

Given the intricate nature of the CST exercise, which allowed banks to implement detailed analyses of their risk exposures at the individual counterparty level, we posit that the CST functions as a significant regulatory shock on SIs, to be compared to LSIs. This shock, within the broader spectrum of the aforementioned regulatory interventions, has the potential to influence and modify the behavior of banking institutions. By compelling banks to scrutinize their portfolios more rigorously and integrate C&E risks into their strategic and operational frameworks, as well as given the publicity given to the exercise the CST can contribute to inducing a transformative effect on banks' risk management practices and decision-making processes, more than the other supervisory guidance and policy measures described above.

4 Model: Climate Stress Tests and Banks' Lending Behavior

In this section, we present a stylized model that illustrates how a supervisory regulatory shock—specifically, how a CST or a related Central Bank supervisory action can reveal borrowers' transition risk information—alters a bank's equilibrium lending conditions.

The objective of this model is to illustrate the effect of Central Bank's supervisory intervention on banks' lending behavior. Theoretically, it is crucial to distinguish two primary channels through which prudential regulation and supervision operate. First, prudential regulation imposes costs (e.g., penalties for non-compliance) and constraints on banks, including capital ratio requirements. Second, the supervisory process itself, including stress tests, drives data collection and enhanced transparency regarding the risk exposures of regulated entities. This second channel facilitates the discovery of private information about counterparties, reducing information asymmetry.[9] Such a reduction not only improves loan monitoring but can also lead to adjustments in credit terms for new loans.

In the following model,[10] we focus specifically on modeling the second channel. CST and, more broadly, prudential regulations regarding the management of C&E risks do not currently impose additional regulatory capital requirements. Instead, they serve to reveal previously unobserved risks, such as transition risks, which can reshape the lending conditions offered by banks. By doing so, this model captures how the process of Central Bank's supervisory intervention, by reducing information asymmetry, impacts equilibrium lending rates and banks monitoring decisions.

Therefore, the model highlights how transition risk affect both the bank's monitoring decisions and the equilibrium interest rate offered to borrowers. By doing so, we demonstrate that the equilibrium interest rate R is strictly increasing in the transition risk parameter τ as an outcome of an internal optimization and competitive equilibrium condition. For the sake of readability, proofs of intermediate derivation steps can be found in Appendix E.

Model Setup and Assumptions. Consider a two-period setting with dates $t = 0$ and $t = 1$. At date $t = 0$, the borrower's investment opportunity is a risky project which, at the initial date, requires an investment of I euros, yielding a stochastic payoff of X. A risk-neutral bank extends a loan of size $I > 0$ to the borrower.

[9]In a different setting, Flannery et al. (2017) findings suggest that disclosure of supervisory stress test results generates significant new information about stress tested U.S. bank holding companies.

[10]Our model is partly inspired by Ivashina (2009) model. Nevertheless, it is also conceptually aligned with the broader framework of credit market imperfections related to adverse selection and moral hazard, as first emphasized by Akerlof (1970) and later developed in the seminal work of Stiglitz and Weiss (1981).

We assume that $X > I > 0$ and that the risk-free rate is zero for simplicity. The project financed by this loan is characterized by stochastic outcomes that unfold in a single stage. At date $t = 1$ the project's outcome is random and pays X in the high state and zero in the low state. With probability $\theta \in [0, 1]$, the borrower repays the principal plus the interest r, totaling $I \cdot (1 + r)$. If the project fails, no repayment is made.

The probability θ is an endogenous function jointly determined by the bank's monitoring effort e_b and by the a borrower's choice variable that depends on the interest rate r. For simplicity:

$$\theta = \Theta(e_b, e_l(r))$$

where Θ increases both in e_b and e_l. However, e_l itself decreases in r, which can be interpreted as either moral hazard or adverse selection. A higher r depresses the borrower's incentive to choose or maintain a safer project. Only the bank knows its effort, and the financial contracts cannot be written contingent on monitoring effort or the state of nature.

Monitoring effort e_b is privately observed by the bank and entails a convex cost: $\frac{1}{2}\beta I e_b^2$

The resources required for monitoring increase not only with effort but also with the size of the investment. This may include more frequent audits, more detailed financial analysis, or the employment of additional staff to oversee the use of funds. Here, $\beta > 0$ parameterizes the firm's systematic characteristics as priced by lenders. A borrower that is easier to monitor would be associated with a lower β, as would a more transparent borrower, such as a publicly traded company, as compared to a privately held company. Similarly, a company providing collateral, having a better credit rating, or having a history of positive lending relationships, would have a lower β.

At $t = 0$, the bank offers a loan contract specifying interest rate r. The borrower accepts if it expects nonnegative utility net of any unmodeled effort costs. The bank chooses e_b, incurring the cost $\frac{1}{2}\beta I e_b^2$. At $t = 1$, the project succeeds with probability $\Theta(e_b, e_l(r))$. Success pays X, allowing repayment of $I(1+r)$. Failure yields zero repayment.

Under perfect competition, equilibrium requires that in any active market segment, the expected profit of the lending bank is driven to zero. Given that the project succeeds with probability θ, and assuming funding cost is zero, the bank's expected profit is:

$$\mathbb{E}[\Pi_b](r, e_b) = \Theta(e_b, e_l(r))I(1 + r) - I - \frac{1}{2}\beta I e_b^2 \tag{1}$$

This zero-profit condition, together with the borrower's endogenous response $e_l(r)$, can generate an interior solution for r^*.

Equilibrium. To illustrate concretely how an interior solution emerges, assume a simple parametric form for the success probability:

$$\Theta(e_b, e_l(r)) = e_b^\alpha \times (1 - \delta r)$$

where $0 < \alpha < 1$ captures decreasing marginal returns to the bank's own monitoring effort e_b, $0 < \delta < 1$ measures how strongly the borrower's effort or project quality declines with r (if r is too high, safer borrowers exit or the borrower exerts less effort), $1 - \delta r$ must remain positive to be meaningful, which implies an upper bound on feasible r.

Conditional on I, the bank then solves:

$$\max_{r,e_b} \left[e_b^\alpha(1-\delta r)(1+r) - 1 - \frac{1}{2}\beta e_b^2\right]$$

In equilibrium under competition, the optimal choice (r^*, e_b^*) will drive $\mathbb{E}[\Pi_b](r^*, e_b^*)$ to zero (or arbitrarily close to zero from above if the market is active).

After solving the FOC of the bank's optimization problem, we show in Appendix E that the equilibrium interest rate is a non linear function of the structural parameters only:

$$r^* = \frac{(\delta - 1) \pm \sqrt{(1-\delta)^2 + 4\delta(1-M^*)}}{2\delta} \tag{2}$$

with $M^* = M(\alpha, \beta)$.[11] The borrower's effective probability of success decreases as r increases. Beyond a certain point, raising r further reduces the borrower's quality so much that the bank's expected profit declines, causing the bank to set an interior r^*. The zero-profit condition pins down a unique positive root because once the interest rate becomes too high, $(1 - \delta r)$ shrinks sufficiently to reduce the bank's expected repayment and, thus, its expected profit. Even if some borrower were hypothetically willing to pay more than r^*, the bank's expected profit would decrease at higher rates, due to the endogenously lower probability of success. This directly leads to the outcome that the bank does not lend at any $r > r^*$.

This outcome is in line with the type of credit-rationing result from Stiglitz and Weiss' model (see Appendix E for details). In standard competitive markets without informational frictions, one might expect the interest rate to adjust indefinitely to clear supply and demand. However, here the interest rate itself deteriorates the borrower's probability of success (through moral hazard or adverse selection), leading the bank's expected profit to fall if r is set too high. Rather than serve all borrowers willing to pay an even higher rate, the bank prefers to deny credit beyond r^*. In equilibrium, the market does not clear purely via price; instead, some borrowers who would accept $r > r^*$ remain rationed out of the market.

The bank does not simply keep increasing r to cover risk or monitoring costs; instead, it hits an optimal rate r^* and is unwilling to lend at rates above that, because higher rates engender unfavorable borrower behavior. Thus, some borrowers who would borrow at a higher r are rationed.

Impact of the Central Bank Intervention. Now suppose that the Central bank intervenes, via a stress test, wich allows discovering additional information indicating that the success probability θ is also affected by a transition risk τ. Specifically, the bank learns that its original assessment θ is now effectively $\theta|\tau$—a decreasing function in τ for any given monitoring effort. This new information shifts the perceived distribution of payoffs, making the project riskier when τ increases. Under these conditions, the bank re-optimizes its monitoring effort choice and updates the interest rate r accordingly.

From the point of view of welfare economics, the intervention of the Central bank is only justified if these policy measures make it possible to improve social welfare. We assume that there is a market failure warranting regulatory intervention. We posit that

[11] see Appendix E.

the banking system may collectively fail to incorporate transition risk into lending terms unless compelled by a Central bank. This may be explained, inter alia, by the cost to gather confidential information on firms' carbon exposure, to which only the Central bank has access.

Therefore, to explore the broader implications of a Central bank intervention across the banking system, we extend our analysis using a game theory approach to show the collective action problem banks are facing. We show in Supplementary Online Appendix Appendix J that a social optimal outcome can only be achieved by means of intervention of a benevolent social planner.

Equilibrium after Intervention. Let $\theta|\tau$ denote the success probability—a function of the bank's effort—conditional on the transition risk $\tau \in [0, 1]$.[12] The relationship between θ and $\theta|\tau$ is:

$$\theta|\tau = g(\theta, \tau),$$

where $g(\cdot)$ is a strictly decreasing function in τ, capturing the impact of transition risk on the success probability. Specifically, we assume a simple relationship:

$$\theta|\tau = e_b^{\alpha} \times (1 - \delta r)(1 - \tau).$$

The last term linearly and negatively incorporates transition risk: the larger τ is, the smaller $(1 - \tau)$, and thus the lower the overall probability of success. Since the success probability is decreasing in τ, it reflects the fact that the project is, on average, less likely to produce a favorable outcome. The fact that $\theta|\tau$ decreases with τ translates, in a simplified manner, the idea that the conditional distribution $f_{X|\tau}(x)$ places more mass on unfavorable realizations compared to the initial distribution $f_X(x)$ (For details see Appendix E). Consequently, the expected profit in presence of transition risk:

$$\mathbb{E}[\Pi_b](r, e_b) = e_b^{\alpha}(1 - \delta r)(1 - \tau)(1 + r) - 1 - \frac{1}{2}\beta e_b^2 \tag{3}$$

Notice that the assumption in Eq. (3) is that the bank's monitoring costs primarily reflect a technological or organizational dimension—such as internal procedures, staffing requirements, and analytic methodologies—which remains largely unaffected by the exogenous shock τ. In contrast, τ represents a systematic risk factor external to the bank's direct influence, altering the overall distribution of project outcomes rather than the marginal cost of monitoring itself. Thus, by separating these two components, Eq. (3) highlights that, while a higher τ reduces the ex-post success probability, it does not necessarily render monitoring more resource-intensive. Instead, the monitoring activity, which is undertaken ex-ante and designed to mitigate moral hazard and incentivize borrower discipline, retains the same underlying cost structure regardless of whether a transition risk materializes subsequently. This modeling choice thus captures the idea that the bank's cost of exerting effort depends on its internal monitoring technology and not on the particular nature of the external shock.

[12]The boundedness of τ is justified by its interpretation as a normalized measure of transition risk, where $\tau = 0$ signifies no transition risk and $\tau = 1$ represents the maximum conceivable level of transition risk affecting the project's outcome.

After solving the FOC of the bank's optimization problem, we show in Appendix E that the equilibrium interest rate with transition risk is:

$$r_{\tau}^{*} = r^{*} + \alpha \cdot \tau \tag{4}$$

Intuitively, the additional term in Eq. (4) captures how the bank adjusts its lending rate in response to the revealed transition risk. Because τ lowers the project's repayment, the interest rate r^* must rise to uphold the zero-profit condition. Specifically, β represents the bank's marginal cost of exerting monitoring effort, and thus a higher β indicates a more pronounced increase in the interest rate when transition risk rises. Similarly, δ measures the sensitivity of the project's success probability to changes in τ , so a larger δ leads to a steeper adjustment of lending terms. The factor I in the denominator reflects the loan size, implying that for a given rise in transition risk, the absolute rate adjustment is mitigated by a larger loan principal, as the bank can spread a similar profit-margin requirement over a greater volume of lending. Consequently, the additional term succinctly encapsulates how risk aversion, monitoring costs, and project scale jointly determine the incremental interest rate required to maintain the bank's expected profitability.

Discussion. This simple stylized model shows that when a Central Bank provides new information about transition risk that conditions the probability of the investment success, the bank perceives the project as riskier. The subsequent equilibrium response is a higher interest rate R. Furthermore, the model shows that not only does the equilibrium interest rate increase when the bank factors in transition risk, but it also increases monotonically with the level of this transition risk.

This theoretical framework suggests that policy measures enhancing transparency about transition risk exposures will likely affect credit conditions. Therefore, it sets the stage for empirical work. Assume two subsets of banks, a first subset receive a signal on the exposure of its counterparties to a transition risk while the second subset does not receive this signal. This can be reformulated in terms of potential outcome as:

$$\mathbb{E}[r^*|0 < \tau \leq 1] - \mathbb{E}[r^*|\tau = 0] = \alpha \cdot \tau \tag{5}$$

If the CST has the impact predicted by our model then the first group should set higher lending rates, all things being equal. In the subsequent Sections, we develop an empirical strategy to test this hypothesis. However, the empirical identification of the CST as a distinct regulatory shock presents significant challenges. This difficulty arises from the concurrent implementation of multiple successive regulatory initiatives aimed at addressing C&E risks within the European banking sector. These overlapping measures create a complex regulatory environment, making it arduous to isolate the specific impact of the CST. Additionally, variations in the timing, scope, and intensity of these interventions across different institutions further complicate the disentanglement of their individual effects. To address these challenges, our empirical analysis will employ robust methodological approaches to accurately identify the regulatory shock introduced by the CST. In the robustness analysis we also control for the influence of the staggered nature of the CST—first, the French CST, and second, the European CST.

5 Data

We use data on new loans from the Eurosystem's Anacredit database. After cleaning up as described in Annex D, we end up with 2 databases. Table **??** reports descriptive statistics for the two baseline datasets used in the econometric analysis. Panel A (FULL) comprises 2,214,833 loan observations for all bank-debtor pairs, while Panel B (INC)—our restricted incumbents sample—contains 1,187,416 observations. The INC sample retains only bank–firm pairs that are observed both before and after the CST methodology release, with at least two loan observations per pair. It therefore focuses on continuing relationships and mechanically drops purely transitory bank–firm matches, which explains the reduction in sample size relative to FULL.

Despite the sharp difference in sample size, the central moments of the key variables are remarkably similar across the two panels. Average interest rates are 3.55% in FULL and 3.48% in INC, and the corresponding means for credit spreads are 2.95% and 3.01%. Standard deviations and within-sample quantiles are likewise very close.

Two additional distributional features are salient for our analysis. First, even in logs, the firm-level GHG intensity variable remains right-skewed. In both samples, the median lies below the mean (2.77 vs 3.00 in FULL; 2.84 vs 3.07 in INC), and the interquartile range is relatively tight compared to the long right tail (p75 small and very similar in both samples but maxima significantly higher). Second, the loan outstanding amount variable displays a similar pattern: most exposures are of moderate size, but a small fraction of loans are very large and hence account for a large share of total credit volume.

Table 1. Descriptive Statistics SIs Only Samples

	N	Mean	St. Dev.	p10	Median	p90
A. All sample (FULL)						
Interest Rate (%)	2,214,833	3.55	2.46	0.85	3.03	6.86
Credit Spread (%)	2,214,833	2.95	2.17	1.07	2.55	5.85
GHG Intensity (log)	2,214,833	3.00	2.09	0.74	2.77	5.57
Out. Amount (log)	2,214,833	11.15	1.55	9.41	10.82	13.22
B. Restricted Sample (INC)						
Interest Rate (%)	1,187,416	3.48	2.49	0.81	3.00	6.85
Credit Spread (%)	1,187,416	3.01	2.23	1.10	2.55	6.11
GHG Intensity (log)	1,187,416	3.07	2.11	0.81	2.84	5.62
Out. Amount (log)	1,187,416	11.23	1.59	9.41	10.95	13.41

To protect the confidentiality of the underlying microdata, all reported values result from the aggregation of five observations around the relevant point of the respective distribution.

To assess firms' exposure to transition risk, prior studies (Altavilla et al., 2023; Fuchs et al., 2024) have primarily relied on carbon emission data, either reported or estimated. However, focusing exclusively on large firms that disclose their emissions would overlook micro, small, and medium-sized enterprises (MSMEs), which rarely provide such data. Aiello (2024) however extends the analysis to MSMEs by imputing firm-level GHG emissions. Along the same lines, but using an alternative approach, we construct an *ad hoc* firm-level GHG intensity metric by integrating both aggregate and firm-level data. First, we collect sector-specific data at the two-digit NACE level from Eurostat, including GHG emissions, the number of firms, and turnover disaggregated by firm size class. For each year, country, and NACE sector, we obtain the number of firms, their respective size classes, and the total sectoral turnover. To allocate GHG emissions within each NACE sector, we proportionally distribute total sectoral emissions across firm size categories based on their respective shares of total sectoral turnover. This yields GHG emissions estimates at the level of year-country-NACE sector-size class (henceforth referred to as a "cell"). We then compute the average GHG emissions per firm within each cell by dividing total estimated emissions by the number of firms. Finally, we normalize the cell by debtor turnover to obtain an estimate of firm-level GHG intensity.[13,14] We then obtain firm-level environmental intensity by normalizing this per-firm emissions measure with the debtor's turnover after capping turnover within each cell (year–country–sector–size) at the 1st and 99th percentiles in log space to limit outlier influence. Furthermore, we filter out firms with turnover below EUR 10,000.

Further details regarding the data description and sample construction procedure are available in Appendix Appendix C.

6 Empirical Strategy

Section 3 documents that the ECB's *green turn* has relied on a broad toolkit, ranging from general supervisory communication to targeted interventions such as the 2022 thematic review and the ECB's 2021–22 EU-wide climate stress test.[15] This multiplicity creates a standard identification challenge: many climate-related supervisory messages are economy-wide in scope, so they may affect the banking system as a whole through expectations, reputational pressure, or common learning. At the same time, a subset of measures is explicitly targeted at a well-defined population of institutions. Our empirical strategy leverages this institutional structure. In particular, we exploit the CST as a discrete supervisory intervention that applies to SIs but not to LSIs, thereby providing a *quasi-experimental*[16] source of cross-sectional variation in exposure to climate-focused supervisory scrutiny.

The CST is particularly useful for identification because it sharpens the distinction between broad and targeted supervisory actions. While ECB speeches, publications, and supervisory expectations on C&E risks are largely addressed to the banking system as

[13] See Appendix B for mathematical details.

[14] This methodology assumes that environmental efficiency varies at the cell level, with firms reporting lower turnover exhibiting higher estimated GHG intensity. Therefore, it implicitly assumes increasing environmental returns to scale.

[15] Jourde, Szczerbowicz, and van Dijk (2025) construct a Climate and Nature Communication index that captures the dynamic of ECB's speeches on climate-related topics.

[16] In economics, the concept *quasi-experiment* is typically used to describe a scenario in which treatment assignment is not explicitly randomized but contains an exogenous component that relies on institutional or policy thresholds that effectively mimic random assignment in certain respects; see Remler and Van Ryzin (2024) for an in-depth discussion.

a whole, the CST is operationally implemented as an intensive data, modelling, and governance exercise for the SI perimeter. LSIs may be indirectly exposed to the general communication environment, but they are not directly subject to the CST's reporting requirements, bottom-up templates, or supervisory challenge process. In a difference-in-differences framework, this means that common sector-wide climate communication is absorbed by time effects, whereas the incremental treatment intensity induced by the CST is concentrated among SIs. Put differently, the SI/LSI boundary isolates a supervisory *dose* difference that is central to our quasi-experimental design.

A natural concern is that SIs and LSIs differ mechanically in size, business models, and regulatory constraints. Two institutional features mitigate the force of this objection. First, the SI/LSI distinction is a legal-supervisory classification defined by the SSM, not a clean partition of fundamentally different banking technologies. In particular, many LSIs in countries such as Germany and Austria belong to savings-bank or cooperative networks that are partially mutualised through Institutional Protection Schemes (IPS). IPS arrangements provide ex-ante risk sharing, monitoring, and support mechanisms at the network level, which can make the relevant economic unit closer to a group than to a stand-alone local bank (Haselmann, Krahnen, Tröger, and Wahrenburg, 2022). Conversely, some large mutualist networks are organised as consolidated groups supervised as SIs (e.g., BPCE in France). These examples illustrate that the SI label is not synonymous with a unique economic model; rather, it reflects the supervisory perimeter and consolidation structure (Véron, 2024). This strengthens the case for using LSIs as an economically meaningful counterfactual group, while acknowledging that residual differences may remain.[17]

Second, even if levels differ across SIs and LSIs, our estimand is a difference-in-differences in slopes. Identification does not require treated and control banks to have the same average spreads, the same borrower mix, or the same baseline sensitivity to carbon intensity. What is required is that, absent the CST, the evolution of the spread–intensity semi-elasticity would have been parallel across the two groups once we condition on high-dimensional fixed effects and controls. Empirically, we implement this by saturating the specification with granular fixed effects that absorb time-varying bank heterogeneity and detailed borrower and loan characteristics, and by using reweighting strategies when necessary to improve overlap in observables. This approach is designed to ensure that remaining SI/LSI differences operate primarily through time-invariant (or smoothly evolving) components that are netted out by the triple-difference structure, rather than through discrete changes around the CST date. Therefore, the institutional threshold-based assignment to SI status forms the foundation of our quasi-experimental design, offering a credible avenue for isolating the causal impact of the CST on lending behavior as reflected in the spread–emissions sensitivity.

Formally, let $D_b = 1$ if bank b is an SI and thus directly subject to the CST, and $D_b = 0$ otherwise. Our causal object is the CST-induced change in the sensitivity of loan spreads to borrower carbon intensity. We estimate a triple-interaction specification in which the credit-spread–intensity relationship is interacted with both a post-CST indicator and the SI treatment indicator. The resulting coefficient identifies a DiD-in-slopes that captures whether the CST shifts the spread–intensity schedule for treated banks relative to controls.

[17] See also the following article: https://www.revue-banque.fr/archive/il-y-un-decalage-entre-analyse-economique-realite-YTRB16405

Let $y_{bdt}(d)$ denote the potential credit spread that bank b would charge debtor d at time t on a loan with carbon intensity x if the bank were, respectively, exposed or not exposed to the CST. For each period t, define the conditional pricing schedule

$$\mu_t^d(x) \equiv \mathbb{E}\left[y_{bdt}(d) \mid \log(x_{dt}) = x, \Phi_{FE}\right], \qquad \text{and} \qquad \theta_t^d \equiv \frac{\partial \mu_t^d(x)}{\partial x}.$$

where θ_t^d is the slope with respect to borrower carbon intensity. Our estimand is a difference-in-differences in slopes. Identification relies on a parallel-trends-in-slopes assumption: absent the CST, the evolution of θ_t^0 would have been the same for SIs and LSIs conditional on Φ_{FE}.

6.1 Identification

We detail and discuss how the main assumptions required for causal identification align with the institutional context of regulatory stress tests and how our empirical design is adapted to mitigate potential sources of bias.

Parallel Trend Assumption. A key identifying assumption in our DiD design is the parallel trend assumption: changes in outcomes for SIs would track changes in outcomes for LSIs if the policy intervention had not occurred. In other words, absent the CST, the slope of loan spreads with respect to debtor carbon intensity would have evolved similarly over time for SIs and LSIs, after conditioning on fixed effects. This assumption can be tested by evaluating pre-intervention trends in the spread–emissions slope, estimating event-time interactions for SIs relative to LSIs. If the pre-treatment slope coefficients are statistically indistinguishable from zero, any subsequent divergence in the spread–emissions sensitivity is more credibly interpreted as resulting from the CST. If systematic differences in outcomes between treated and untreated banks existed prior to treatment, and these differences vary over time in ways not fully captured by the model (e.g., due to time-varying confounders), residual selection bias could still emerge. In that regard, our fixed-effect specifications aim to absorb persistent traits at both the bank and borrower levels, thus reducing the scope for bias.

No-Anticipation Assumption. A condition for identification is that treated banks do not substantially alter their behavior before the policy intervention in response to anticipated changes in supervisory pressure. If banks subject to the CST had foreknowledge of the stringency or precise design of the tests, they could adjust their loan pricing or portfolio allocations in advance, contaminating the observed pre-intervention period and violating the parallel trends assumption. Even though the general CST hypotheses were disclosed by the ECB in November 2020, the exact methodology, scenarios, and implementation details were circulated to the stress-tested banks only in October 2021. We posit that detailed information regarding the exact scope and content of the CST was not public until October 2021 and that, without such details, banks had limited ability to implement targeted adjustments to the spread–emissions relationship. Empirically, we assess the plausibility of this assumption by examining pre-CST dynamics in the spread–emissions slope.

The Stable Unit Treatment Value Assumption (SUTVA). Under SUTVA, potential outcomes for each bank are assumed to be independent of the treatment status of other banks. This rules out unmodeled spillover or peer effects that could bias the estimation if a treated bank's behavior influences other banks' pricing strategies or risk assessments. In our baseline specification we assume that such spillovers are limited and do not materially change the spread–emissions relationship for LSIs. This assumption is admittedly strong in a supervisory environment where the European banking supervision *green turn* is manifest (see Section 3 for institutional details). The SSM CST, however, was primarily designed as a micro-prudential exercise for directly supervised SIs rather than as a broad sector-wide signaling tool. To the extent that LSIs nonetheless adjust partially in response to supervisors' public communication or to market learning from SIs, our estimates of the CST effect on SIs are likely to be conservative, as such spillovers would attenuate the difference between treated and control banks.

Endogeneity of Emissions Intensity. Our climate transition exposure proxy variable is an *ad hoc* intensity metric (see Appendix B). The emissions component is predetermined at the sector-country level and does not react to bank-level lending decisions at time t, and the allocation across size classes relies on sectoral aggregates rather than on banks' credit portfolios. As a result, most of the variation in carbon intensity is plausibly predetermined with respect to the CST. The only channel through which endogeneity may arise is the firm-specific turnover in the denominator, which can respond to credit supply and other shocks. However, our identification strategy does not require the carbon intensity to be fully exogenous in levels; it relies instead on the weaker assumption that any feedback from credit to firm turnover does not change differentially between SIs and LSIs exactly at the CST date, conditional on our high-dimensional fixed effects. Under this condition, the DiD slope coefficient can still be interpreted as the additional change in the spread–intensity sensitivity for SIs relative to LSIs attributable to the CST.

6.2 Baseline Model

Time is monthly. Event time is $\tau = t - T_0$ with $T_0 =$ Oct. 2021; we restrict to a symmetric window $\tau \in [-24, +24]$. The outcome y_{bdt} is the credit spread for bank b, debtor d, month t. Emissions intensity x_{dt} is defined in Appendix B. The treatment indicator TreatBank_b flags SIs; $\text{Post}_t = \mathbb{1}\{t \geq T_0\}$.

Our baseline specification allows isolating a change in slope on $\log(x_{dt})$ for treated banks in the post period:

$$\begin{aligned} y_{bdt} \;=\; & \beta_1 \log(x_{dt}) + \beta_2 \left[\text{Post}_t \times \log(x_{dt})\right] + \beta_3 \left[\text{TreatBank}_b \times \log(x_{dt})\right] \\ & + \theta \left[\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})\right] + \lambda_{cst} + \mu_{bt} + \phi_l + \varepsilon_{bdt}\,. \end{aligned} \tag{6}$$

Here, ϕ_l denotes loan-type fixed effects, λ_{cst} denotes country×size×time fixed effects, and μ_{bt} denotes bank×time fixed effects. The joint inclusion of λ_{cst} and μ_{bt} absorbs, respectively, time-varying *pseudo* demand and supply shocks.[18] Inference relies on standard

[18]See Appendix D for statistical tests that support the substitution of debtor×time fixed effects by country×time fixed effects.

errors clustered at the bank level [19]. We report four estimations: (i) FULL: all bank–debtor pairs within the window; (ii) INC: incumbent pairs present both before and after T_0; (iii) INC-IPW: the same set of incumbent pairs as in INC, but estimated using inverse-probability weights w_{bdt} that reweight incumbents so that, within year-by-treatment cells, their observable distribution of $\log(x_{dt})$ mirrors that of the FULL sample;[20] and (iv) INT: the INC sample augmented with bank×debtor fixed effects η_{bd} in Eq. (6), confining identification to within-relationship month-to-month variation.

We define the causal marginal impact of the shock on the spread–emissions slope as follows:

$$\text{Slope DiD} \equiv \Delta S_T - \Delta S_C,$$

$$\Delta S_T \equiv [\text{TreatBank}_b = 1, \text{Post}_t = 1] - [\text{TreatBank}_b = 1, \text{Post}_t = 0],$$

$$\Delta S_C \equiv [\text{TreatBank}_b = 0, \text{Post}_t = 1] - [\text{TreatBank}_b = 0, \text{Post}_t = 0].$$

Within model (6) framework, the four groups map to:

$$\begin{aligned}
&[\text{TreatBank}_b = 0, \text{Post}_t = 0] = \beta_1, \\
&[\text{TreatBank}_b = 0, \text{Post}_t = 1] = \beta_1 + \beta_2, \\
&[\text{TreatBank}_b = 1, \text{Post}_t = 0] = \beta_1 + \beta_3, \\
&[\text{TreatBank}_b = 1, \text{Post}_t = 1] = \beta_1 + \beta_2 + \beta_3 + \theta.
\end{aligned}$$

Therefore it is straightforward to see that Slope DiD $= \theta$. This parameter thus captures a change in slope in the post-shock period that is specific to the treated banks. By including $\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})$, the specification performs a DiD in slopes conditional on x_{dt}, so that the sensitivity of the spread to carbon intensity becomes contingent on period and treatment status.[21]

The dependent variable y_{bdt} is a credit spread expressed in percentage points (pp), so that a change of 0.01 percentage point corresponds to one basis point (bp). The explanatory variable of interest, carbon intensity, is included in logarithms, giving the baseline model (6) a level-log structure. This specification yields a *semi-elasticity*: a coefficient θ means a 1% increase in the level of intensity corresponds to an approximate 0.01 unit change in $\log(x_{dt})$. This, in turn, changes y_{bdt} by $(\theta \times 0.01)$ pp. Since 1 bp is 0.01 pp, this is equivalent to a change of exactly θ basis points. Therefore, the coefficients in Eq. (6) can be directly interpreted as the basis-point change in the spread for a 1% increase in x_{dt}.

The estimates in Table 2 point to a heterogeneous pattern in how banks price borrowers' carbon intensity. In Model (1), estimated on the FULL sample, the baseline pre–treatment slope of credit spreads with respect to carbon intensity for control banks, β_1, is positive and significant: a higher carbon footprint is associated with higher loan spreads even before the CST shock. The post–treatment difference in slope between treated and control banks, β_2, is essentially zero, and the pre–treatment change in the slope for the control group, β_3, is small and statistically insignificant. The triple–interaction

[19] The results are robust to two-way fixed effects distinguishing bank and time.

[20] See Subsection Appendix C.2 for more details.

[21] Note that the standard double-difference in means between the two groups is not tractable from model (6) because we include bank×time fixed effects, which are perfectly collinear with the double-difference coefficient.

coefficient, θ, is also small and not significant. As a result, the implied post–treatment slope for treated banks, $\beta_1 + \beta_2 + \beta_3 + \theta$, remains close to its pre–treatment counterpart. At the aggregate level, Model (1) therefore does not provide compelling evidence that the CST induced a discrete shift in carbon–risk sensitivity for treated banks.

Models (2) and (3), which restrict attention to incumbent borrower–bank relationships and, in the latter case, reweight observations using inverse probability weights to account for selection into the incumbent sample, reveal differential pattern. In both specifications, the triple–interaction coefficient θ is negative and statistically significant in both models, almost exactly offsetting the control-group post–treatment shift. These results suggest that, on incumbent exposures, the CST does not lead treated banks to strengthen the pricing of carbon risk relative to their own pre–shock behavior.

Model (4) isolates the intensive margin by including bank–debtor fixed effects and exploiting only within–relationship variation in both spreads and carbon intensity. In this specification, the baseline and interaction coefficients are small in magnitude, and the triple–interaction parameter θ is again close to zero and far from statistically significant. Once time–invariant relationship heterogeneity is fully absorbed, there is no evidence of a differential change in the within–relationship pricing of carbon risk for treated banks relative to comparable control relationships. Overall, across Models (1)–(4), the CST does not appear to generate a robust, uniform increase in treated banks' marginal pricing of carbon intensity. Instead, the main detectable adjustment in carbon–risk sensitivity arises for control banks in the incumbent sample, while treated banks' pricing of carbon risk remains broadly aligned with its pre–treatment pattern.

Table 2. Baseline Model

Dependent Variable:	Credit Spread			
Model:	(1) FULL	(2) INC	(3) INC (IPW)	(4) INT
Coefficients				
β_1	0.0682***	0.0204	0.0192	−0.0096**
	(0.0199)	(0.0181)	(0.0199)	(0.0032)
β_2	−0.0000	0.0382	0.0407	0.0120**
	(0.0256)	(0.0258)	(0.0275)	(0.0042)
β_3	0.0022	0.0933**	0.0966**	−0.0116
	(0.0315)	(0.0312)	(0.0318)	(0.0102)
θ	0.0044	−0.0896**	−0.0933**	0.0001
	(0.0296)	(0.0299)	(0.0306)	(0.0098)
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes
Fit statistics				
Observations	2,214,084	1,186,746	1,186,746	1,186,745
R^2	0.40091	0.42413	0.43793	0.80782
Within R^2	0.00527	0.00363	0.00385	0.00008

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

6.3 Decomposition of Extensive and Intensive Margins

To sharpen the economic interpretation of the slope DiD, we decompose θ into channels that correspond to intensive and extensive margins.[22] Let $Z \equiv \text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})$ denote the post-shock slope regressor. Denote by $\widehat{\theta}^F$ the OLS coefficient on Z in the FULL sample, by $\widehat{\theta}^R$ the corresponding coefficient in the INC sample, and by $\widehat{\theta}^I$ the coefficient in the INT specification. Although these coefficients are estimated on different samples and specifications, they can be used to construct a descriptive decomposition of

[22]See Appendix F for further details on the rationale for the channels decomposition.

the overall effect as follows:

$$\widehat{\theta}^F = \underbrace{\widehat{\theta}^I}_{\text{intensive margin}} + \underbrace{\underbrace{(\widehat{\theta}^R - \widehat{\theta}^I)}_{\text{reallocation among incumbents}} + \underbrace{(\widehat{\theta}^F - \widehat{\theta}^R)}_{\text{entry/exit (total)}}}_{\text{extensive margin}} . \tag{7}$$

While this identity is not algebraically exact in the Frisch–Waugh–Lovell sense—because $\widehat{\theta}^F$, $\widehat{\theta}^R$, and $\widehat{\theta}^I$ are estimated on different sets of observations—it provides an economically meaningful breakdown for interpreting the mechanisms through which the aggregate slope effect materializes.

Intensive Margin. With bank×debtor fixed effects η_{bd} in the specification, identification of $\widehat{\theta}^I$ comes solely from month-to-month movements within a given bank-borrower pair. Economically, $\widehat{\theta}^I$ answers whether, after the CST, a given bank changed the markup it applies to a given borrower as a function of that borrower's carbon intensity, holding the relationship fixed.

Reallocation Among Incumbents. The increment $\widehat{\theta}^R - \widehat{\theta}^I$ isolates between-relationship variation among continuing pairs within bank×month cells, after netting out within-pair repricing. This channel captures how treated banks reweighted their lending across existing clients. A negative value of $\widehat{\theta}^R - \widehat{\theta}^I$ indicates that treated banks reduced the share of new lending to more carbon-intensive incumbents relative to less carbon-intensive ones, thereby flattening the overall slope. Conversely, a positive value would indicate a shift toward more carbon-intensive incumbents. This is purely a loan origination reallocation effect: any within-relationship repricing is already captured by $\widehat{\theta}^I$.

Entry/Exit. The residual $\widehat{\theta}^F - \widehat{\theta}^R$ captures how relationship turnover contributes to the aggregate slope change. This component reflects both compositional changes in the pool of active relationships and potential pricing differences for new versus exiting relationships. The net entry/exit effect combines two distinct forces: (i) changes in the carbon intensity composition of entering versus exiting relationships, and (ii) differential pricing behavior for new versus departing relationships. To further disentangle the forces operating along the entry/exit channel, we exploit the fact that the INC sample is a selected subset of the FULL sample. For each year and treatment status, we estimate the probability that a bank-firm pair is an incumbent as a function of its emissions intensity and treatment status and construct inverse-probability weights.[23] Let $\widehat{\theta}^{R,w}$ denote the coefficient on Z estimated in the INC sample using these weights. By construction, the weighted incumbent sample approximates the FULL sample in terms of the observable distribution of $\log(x_{dt})$, within year-by-treatment cells.[24] Substituting $\widehat{\theta}^{R,w}$ into (7) yields the following equivalent decomposition:

$$\widehat{\theta}^F = \underbrace{\widehat{\theta}^I}_{\text{intensive margin}} + \underbrace{(\widehat{\theta}^R - \widehat{\theta}^I)}_{\text{reallocation among incumbents}} + \underbrace{\underbrace{(\widehat{\theta}^F - \widehat{\theta}^{R,w})}_{\text{entry/exit pricing}} + \underbrace{(\widehat{\theta}^{R,w} - \widehat{\theta}^R)}_{\text{entry/exit composition}}}_{\text{entry/exit (total)}} . \tag{8}$$

[23]See Subsection Appendix C.2 for more details.
[24]See Figure A3 compared to Figure A2 in Appendix A.

The last bracket decomposes the entry/exit margin[25] into a component that reflects changes in pricing conditional on the reweighted distribution of emissions intensity—the entry/exit pricing term $\widehat{\theta}^F - \widehat{\theta}^{R,w}$—and a residual component $\widehat{\theta}^{R,w} - \widehat{\theta}^R$ that captures how differences between the FULL and INC samples in the observable distribution of $\log(x_{dt})$ contribute to the aggregate slope change. In other words, the extensive margin itself can be disentangled as the sum of a composition effect (entry/exit composition in observables) and a pricing effect operating on the extensive margin.

The entry/exit pricing term $\widehat{\theta}^F - \widehat{\theta}^{R,w}$ captures whether treated banks price carbon risk differently for transitory relationships (entering or exiting) compared to incumbents. A positive value of $\widehat{\theta}^F - \widehat{\theta}^{R,w}$ indicates that, conditional on the same observable composition, the overall sensitivity of spreads to carbon intensity is higher in the FULL sample than what would be expected from incumbent pricing alone. This suggests that treated banks apply more carbon-sensitive pricing to transitory relationships than to incumbents. Conversely, a negative value indicates less carbon-sensitive pricing for transitory relationships relative to incumbents.

The sign of the entry/exit composition term in Eq. (8) can be interpreted in a straightforward way using the stylized illustration in Figure 1. Recall that the term $\widehat{\theta}^{R,w} - \widehat{\theta}^R$, where $\widehat{\theta}^R$ is the DiD in slope coefficient estimated on the INC sample, and $\widehat{\theta}^{R,w}$ is the same coefficient estimated on INC after reweighting observations so that, within year-by-treatment cells, the observable distribution of $\log(x_{dt})$ in the INC sample matches that of the FULL sample. In this sense, $\widehat{\theta}^{R,w}$ approximates the DiD in slope that would prevail if only incumbent relationships were priced, but their emissions intensities were distributed as in the FULL sample. A negative value of $\widehat{\theta}^{R,w} - \widehat{\theta}^R$ therefore means that, when we tilt the weights on stayers to reproduce the composition of the FULL sample in terms of $\log(x_{dt})$, the DiD in slope becomes more negative. Holding fixed the conditional pricing rule for incumbents, this can only occur if the FULL sample allocates relatively more weight to low-intensity borrowers and/or less weight to high-intensity borrowers than the INC sample.

The first Panel on the left-hand side of Figure 1 illustrates the case in which new relationships are concentrated at higher values of credit spreads for given level of $\log(x_{dt})$ (red dot): including these relationships rotates the fitted line upward relative to the baseline. Right-hand side panel illustrates the complementary case in which new loans to incumbent firms are concentrated at lower values of $\log(x_{dt})$ and lie below the baseline line, so that adding these relationships pulls the aggregate fitted line downward.

[25] This decomposition relies on the assumption that selection into the incumbent sample is based on observables $\log(x_{dt})$ and treatment status. Under this assumption, the IPW procedure creates a balanced sample where composition differences are eliminated, allowing us to isolate the pure pricing effect.

Figure 1. Composition Effect on Credit Spread

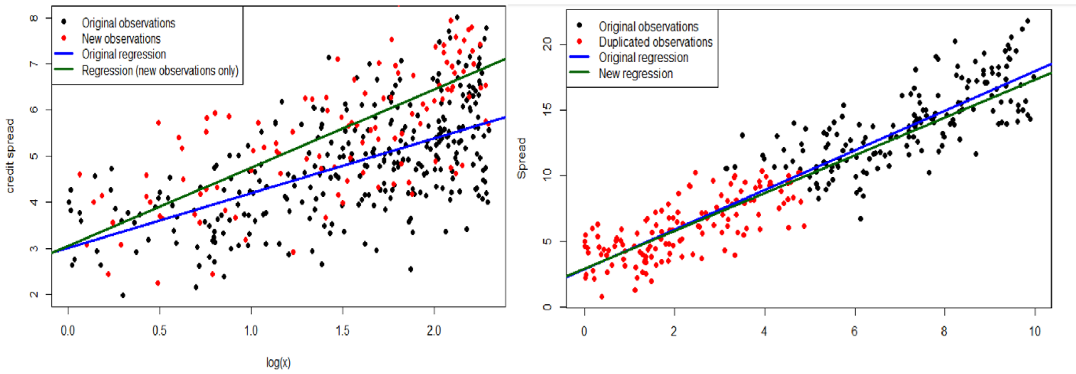

Notes: This figure illustrate the composition effect for the entry/exit extensive margin channel as well as for the reallocation among incumbents channel because in essence we highligth a pure between effect net of within effect.

The estimates in Figure 2 provide an empirical counterpart to the decomposition in Eqs. (7)–(8) and help clarify how the nearly zero aggregate effect of the CST on pricing in Model (1) of Table 2 masks sizable and offsetting adjustment margins. Recall that the triple–interaction coefficient on carbon intensity for treated banks, θ^F, is small and statistically insignificant, suggesting no clear net change in the overall sensitivity of loan spreads to borrowers' carbon intensity for treated institutions. The decomposition shows that this aggregate null is the result of two economically meaningful mechanisms that operate in opposite directions, while the other margins appear negligible.

Figure 2. Channels Decomposition

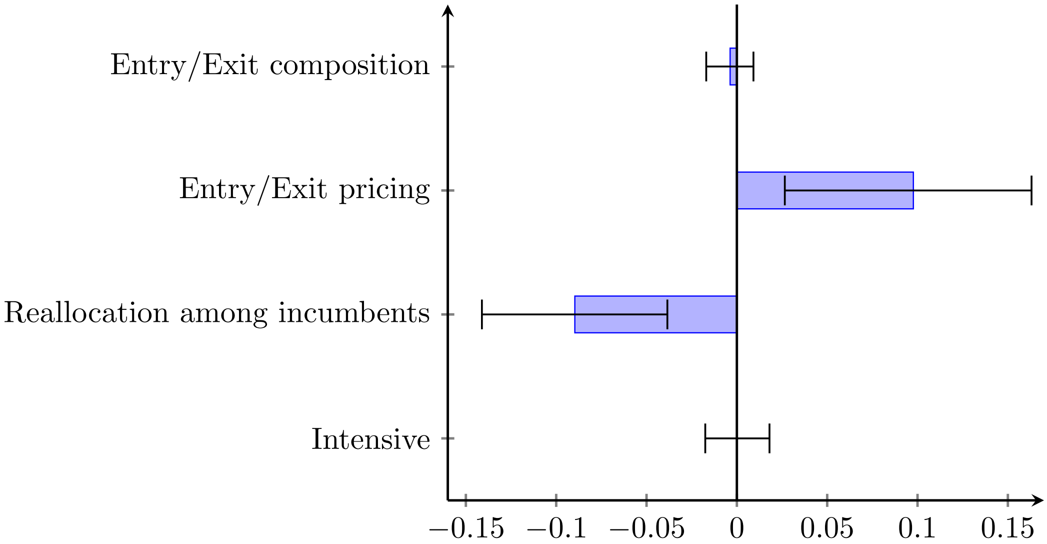

Notes: Point estimates derived from Table 2. Confidence intervals and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details on the bootstrap approach and see Table H2 for an alternative way to presenting the results).

First, in accordance with Model (4) of Table 2, the intensive margin is estimated to be essentially zero. The CST does not appear to have induced treated banks to systematically reprice carbon risk for a given borrower–bank pair, conditional on the relationship being maintained.

By contrast, the Reallocation among incumbents channel is negative (-0.09) and significant at the 1% level. Interpreted through the lens of the decomposition, this term captures how treated banks adjust the allocation of new lending across continuing relationships within bank–month cells, after netting out any within–pair repricing. The negative sign implies that, relative to the control group, treated institutions tilt origination away from more carbon–intensive incumbents and/or towards less carbon–intensive incumbents.

This negative reallocation effect is almost exactly offset by a positive (0.10) and significant Entry/Exit pricing component. The positive sign indicates that, conditional on the same composition in $\log(x_{dt})$, treated banks apply more carbon–sensitive pricing to transitory relationships than to the pool of incumbents. In other words, while the reallocation channel compresses the carbon premium among continuing clients, the pricing applied on the extensive margin (entry and exit) steepens the spread–intensity relationship relative to what would be implied by incumbents alone.

The remaining Entry/Exit composition term is small and statistically indistinguishable from zero. This suggests that, once we control for pricing behavior, compositional differences between stayers and transitory relationships along the carbon dimension do not materially contribute to the aggregate slope effect. Put differently, the extensive margin operates primarily through differential pricing of entering and exiting relationships, rather than through large shifts in the relative frequency of high- versus low–intensity borrowers.

Therefore, the decomposition framework indicates that the CST did not generate a uniform strengthening of carbon–risk pricing across all exposures. Instead, treated banks adjust along offsetting margins: they *green* the composition of new lending among incumbents while simultaneously applying more carbon–sensitive pricing on the margin of relationship turnover. The net effect at the portfolio level is an aggregate slope DiD that is small, but underpinned by economically meaningful reallocations and pricing responses.

6.4 Event-Study

We now turn to the estimation of a dynamic specification that allows capturing, month by month, how the sensitivity of credit spreads to carbon intensity changes for CST banks around the shock date. Let $k \in \{-24, \ldots, 24\}$ denote event time relative to $T_0 =$ Oct. 2021. The equation is:

$$\begin{aligned} y_{bdt} \;=\; & \beta_1 \log(x_{dt}) + \sum_{k \neq -1} \beta_{2,k} \Big[\mathbb{1}\{\tau = k\} \times \log(x_{dt}) \Big] + \beta_3 \Big[\text{TreatBank}_b \times \log(x_{dt}) \Big] \\ & + \sum_{k \neq -1} \theta_k \Big[\mathbb{1}\{\tau = k\} \times \text{TreatBank}_b \times \log(x_{dt}) \Big] + \lambda_{ilt} + \mu_{bt} + \varepsilon_{bdt} \,. \end{aligned} \tag{9}$$

where λ_{ilt} and μ_{bt} are fixed effects defined above. We normalize the month $k = -1$ as the reference, so each θ_k measures the change in the intensity slope for treated banks at event month k relative to $k = -1$, net of saturated confounders. Standard errors are clustered by bank. As for model (6), we report four samples.

Figure 3 allows us to assess the parallel–trends assumption in slope form. Recall that each coefficient θ_k measures, for event month k, the change in the carbon–intensity slope for treated banks relative to control banks, compared to the reference month $k = -1$, after saturating the fixed effects and lower–order interactions.

In the FULL sample (Figure 3, top–left panel), the estimated θ_k coefficients fluctuate moderately around zero both before and after the October 2021 onset. In the year immediately preceding the shock (roughly $k \in [-9, -2]$), the point estimates are small and their confidence intervals systematically include zero, consistent with the absence of strong differential pre-trends in the slope of credit spreads with respect to carbon intensity over this window. At longer leads (e.g., $k \leq -15$), some coefficients become positive and occasionally imprecisely estimated, which cautions against interpreting very long pre-periods as perfectly flat. After the shock ($k \geq 0$), the path of θ_k remains without a clear structural break relative to the late pre–shock months: most coefficients are modest in magnitude and not statistically distinguishable from zero. This dynamic pattern is fully in line with the static result in Model (1) of Table 2.

The incumbent sample without reweighting (Figure 3, top–right panel) exhibits a similar qualitative structure but with a clearer contrast between pre– and post–shock behavior. In the pre–treatment period, especially in the last year before the CST, the θ_k coefficients hover around zero and are generally imprecise, suggesting that treated and control banks had comparable carbon–slope dynamics among incumbents prior to the shock. After $k = 0$, however, the coefficients tend to drift into negative territory, particularly from about $k = 2$ onward, even though individual months are often only borderline significant given the sampling variability.

The IPW specification (Figure 3, bottom–left panel), shows a sharper dynamic response. Pre–treatment coefficients in the late pre–period remain close to zero, supporting the parallel–trends assumption in the reweighted incumbent sample over that window. From approximately two months after the shock, the θ_k path turns persistently negative, with several coefficients significantly below zero at conventional levels over a horizon of roughly one year. This indicates a sustained post–CST decline in the carbon–intensity slope for treated banks relative to matched controls once we control for compositional differences.

By contrast, the intensive–margin specification with bank–debtor fixed effects (Figure 3, bottom–right panel) shows θ_k coefficients that are uniformly small and very imprecise throughout the event window. Both before and after the shock, the confidence intervals for the dynamic intensive–margin coefficients overwhelmingly include zero, with no evident systematic drift in either direction.

The event–study results are coherent with the earlier static and decomposition analysis. Moreover, they show that over the late pre–CST period, slope differences between treated and control banks are broadly consistent with parallel trends in all four specifications.

Figure 3. Event-Study Coefficients (95% CI)

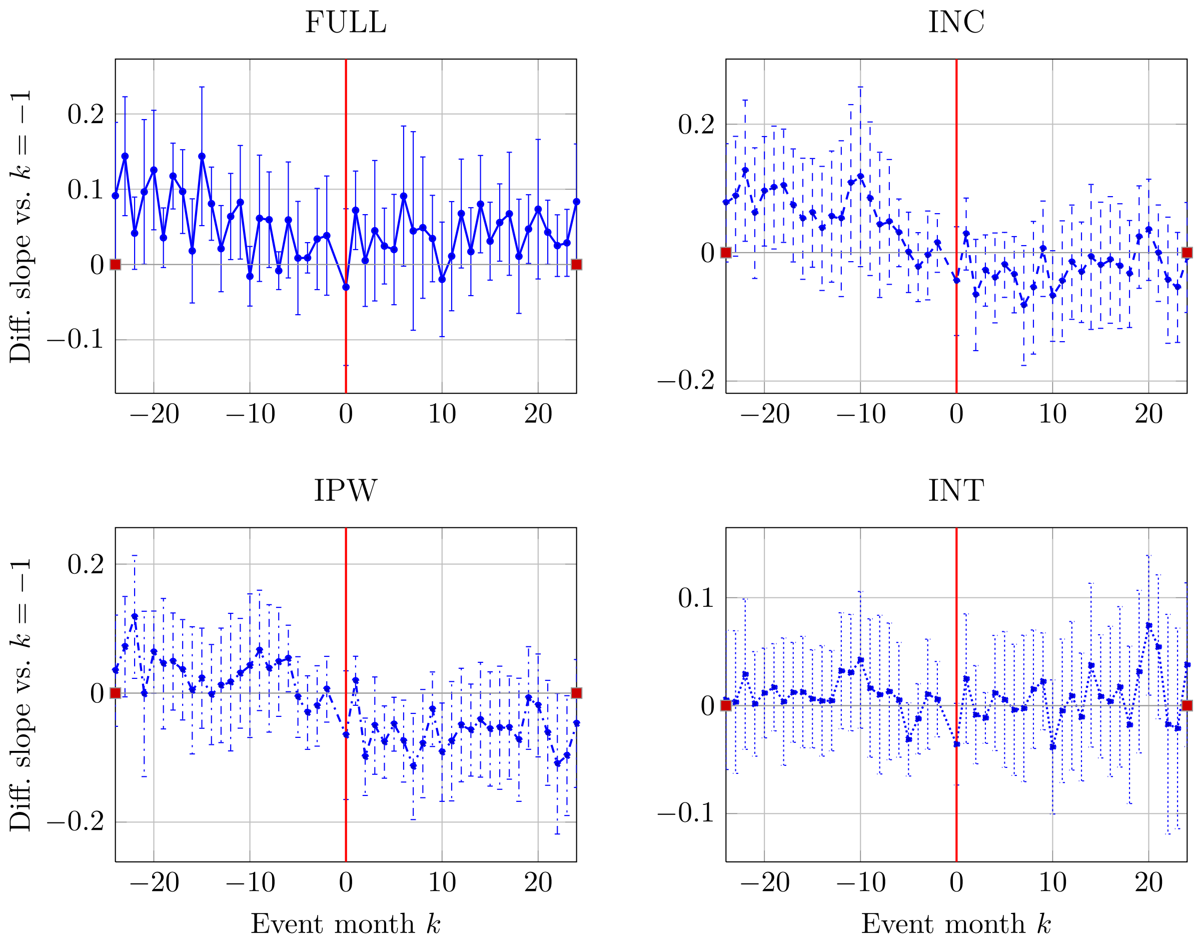

The dynamic channel decomposition in Figure 4 refines the interpretation of the event–study results by showing, month by month, how each margin contributes to the overall treatment–slope coefficient $\hat{\theta}_k$ in the FULL sample. Complementing the static decomposition in Figure 2, it confirms that the aggregate effect of the CST masks sizeable but offsetting adjustments along distinct margins of banks' lending behavior.

Figure 4. Monthly Channels Decomposition (95% CI)

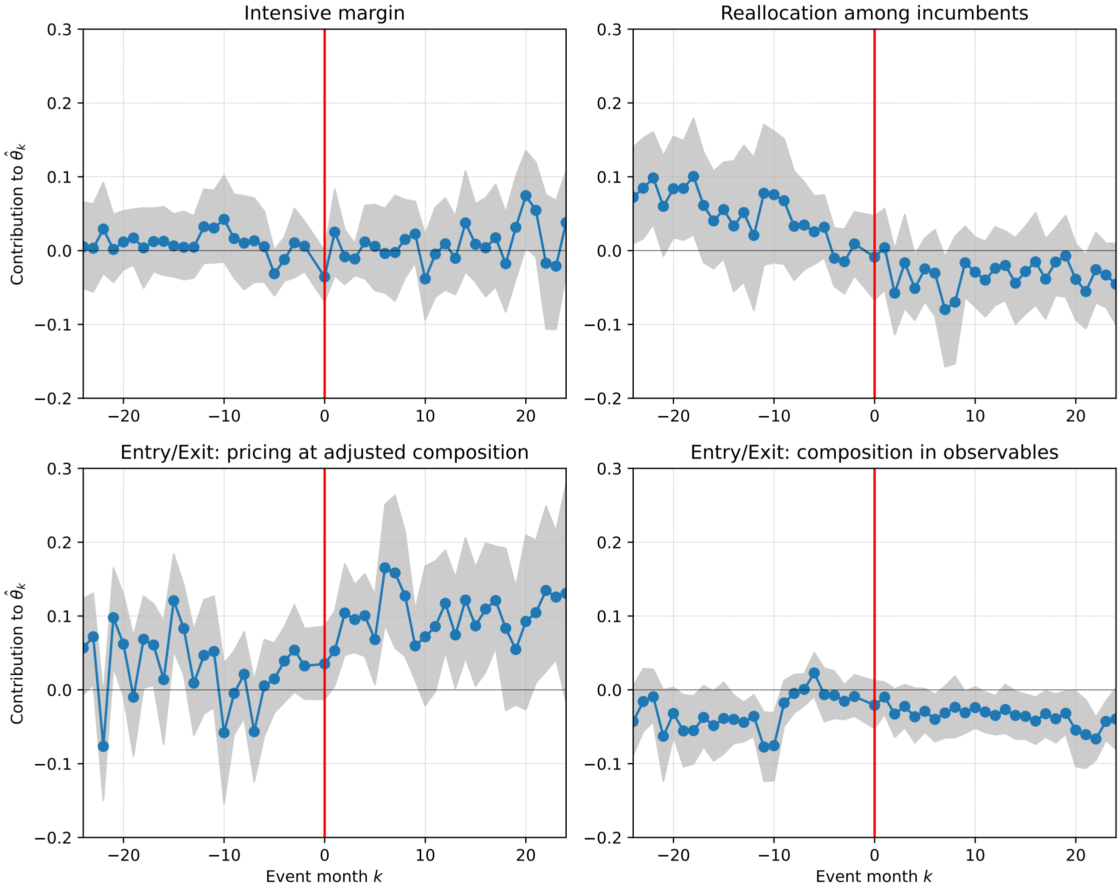

Notes: Point estimates derived from Figure 3. Confidence intervals and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details on the bootstrap approach).

The top–left panel shows that the contribution of within–relationship repricing to $\hat{\theta}_k$ is small and non significant throughout the event window.

By contrast, the top–right panel exhibits a clearer directional pattern. In the pre–CST period, contributions are mostly positive or near zero, suggesting no strong pre–trend relative to the reference month. After the shock, the incumbent–reallocation component gradually turns negative especially from the early post period onward, but remains mostly insignificant. This dynamic profile aligns with the negative and statistically significant reallocation estimate in the static decomposition. It indicates that treated banks progressively reweight lending away from more carbon–intensive incumbents toward less carbon–intensive continuing clients, thereby reducing the aggregate sensitivity of spreads to carbon intensity through a *green* rebalancing of their existing loan portfolio rather than through within–relationship repricing.

The bottom–left panel highlights the opposite force operating along the extensive margin. The entry/exit pricing contribution to $\hat{\theta}_k$ becomes clearly positive in the post–CST period. From shortly after $k = 0$ onwards, the point estimates are sizeable—often in the range of 0.08 to 0.17—and for many months the confidence intervals exclude zero. This pattern mirrors the positive and significant entry/exit pricing term in the static decomposition and confirms that treated banks apply more carbon–sensitive pricing to transitory relationships (new or exiting borrowers) than would be inferred from incumbent pricing alone. Dynamically, this channel appears both stronger and more persistent than suggested by a single post dummy, with a sustained positive contribution over much of the post–shock horizon.

Finally, the bottom–right panel shows that the composition component is predominantly negative, but relatively small, and most of the time non significant over the window. These results are in line with the insignificant composition effect in the static decomposition.

Similarly to the static case, the dynamic decomposition reconciles the apparently muted aggregate response of $\widehat{\theta}_k^F$ in the event–study. The monthly contributions from entry/exit pricing are persistently positive after the CST, indicating that treated banks charge more carbon–sensitive spreads on transitory relationships. At the same time, reallocation among incumbents and, to a lesser extent, entry/exit composition contribute negatively, reflecting a gradual shift of credit toward less carbon–intensive borrowers. These opposing forces largely offset each other in the aggregate, even though the underlying margins are actively adjusting. In this sense, the dynamic channel decomposition shows that the CST did not leave pricing behavior unchanged, but rather induced a combination of more stringent *brown* pricing responses on transitory exposures and *green* reallocation on the core portfolio, whose net effect at the portfolio level is close to neutral.

6.5 Reallocation Across Intensity Bins

To complement the DiD and event–study evidence on extensive margins, we quantify directly how treated banks reshaped the within–bank–month composition of new originations across the distribution of emissions intensity. We discretize x_{dt} into deciles $q \in \{1, \ldots, 10\}$ using pre–treatment loans of treated banks to fix the cutoffs.[26] Let N_{bqt} denote the number of new loans originated by bank b in month t to borrowers in bin q. Define $\mathrm{TP}_{bt} \equiv \mathrm{TreatBank}_b \times \mathrm{Post}_t$ and take the cleanest decile $q = 1$ as the reference category.

We estimate Poisson pseudo–maximum likelihood (PPML) models with bank–month fixed effects μ_{bt} and bin fixed effects δ_q for loan counts:

$$\mathbb{E}\left[N_{bqt} \mid \cdot\right] = \exp\left(\mu_{bt} + \delta_q + \sum_{q \geq 2} \beta_q^N \mathbb{1}\left\{\mathrm{TP}_{bt} = 1\right\} \mathbb{1}\left\{q\right\}\right),$$

[26]Fixing the bins in the pre–treatment period prevents the thresholds from moving endogenously with treatment.

By construction, each coefficient $\beta_q^{(\cdot)}$ captures, for treated banks in the post–CST period, the within–bank–month log change in the expected mass of new lending allocated to decile q relative to the cleanest decile $q = 1$, net of common time shocks and bank–specific time shocks.

The fixed–effects structure and the interaction terms imply a multinomial allocation across bins. Let $\hat{\delta}_q$ denote the estimated bin fixed effects and $\widehat{\beta}_q^{(\cdot)}$ the estimated treatment–post interaction for decile q in the number of new loans originated specification. The corresponding predicted within–bank–month shares are given by

$$s_q^{\text{pre}} = \frac{\exp(\hat{\delta}_q)}{\sum_{j=1}^{10} \exp(\hat{\delta}_j)}, \tag{10}$$

$$s_q^{\text{post}} = \frac{\exp(\hat{\delta}_q + \hat{\beta}_q)}{\sum_{j=1}^{10} \exp(\hat{\delta}_j + \hat{\beta}_j)}. \tag{11}$$

for the post–treatment allocation under treatment. Eqs. (10)–(11) correspond to the standard softmax mapping: a negative $\widehat{\beta}_q^{(\cdot)}$ shifts mass away from decile q (relative to $q = 1$) in the treated–post regime, while renormalization ensures that shares sum to one across all deciles.

Table 3 reports the PPML estimates. Model (1), based on number of new loans originated, indicates that treated banks substantially reweight new originations away from the most carbon–intensive decile after the CST: the coefficient for $q = 10$ is large and negative, with additional declines for several upper and middle deciles.

Translating these estimates into predicted shares using (10)–(11), Table 3 shows that the share of the most carbon–intensive decile $q = 10$ falls by 6.49 percentage points in terms of the number of new loans, with offsetting increases in cleaner deciles. These patterns are fully consistent with our earlier decomposition results: the cross–sectional spread–intensity slope becomes more negative mainly because treated banks reallocate new originations across clients (reallocation among incumbents and entry/exit composition channels), while changes in the within–relationship slope among incumbents remain comparatively limited.

The PPML evidence makes the mechanism behind the aggregate estimates transparent. Following the European CST, treated banks intensify climate–risk differentiation at the portfolio level by redirecting new originations toward less carbon–intensive borrowers within each bank–month cell. This reallocation is pronounced in terms of the number of newly granted loans. These findings reinforce the interpretation that the observed adjustments in the spread–intensity slope primarily reflect between–client reallocation rather than systematic within–pair repricing.

Table 3. PPML Reallocation and Predicted Shares by Intensity Decile

	PPML Model	Predicted Shares		
Decile q	Coefficient	Pre (%)	Post (%)	Δ pp
1	*Ref.*	9.89	12.36	+2.47
2	-0.1626^{**}	10.15	10.78	+0.63
	(0.0589)			
3	-0.1452^{*}	10.40	11.24	+0.84
	(0.0740)			
4	-0.1206	10.20	11.30	+1.10
	(0.0822)			
5	-0.0725	9.57	11.12	+1.55
	(0.0581)			
6	-0.1483^{*}	9.80	10.56	+0.76
	(0.0596)			
7	-0.2255^{*}	10.25	10.22	-0.03
	(0.0892)			
8	-0.2068^{*}	10.15	10.31	+0.16
	(0.0950)			
9	-0.3217^{*}	10.61	9.61	-1.00
	(0.1447)			
10	-1.502^{***}	8.99	2.50	-6.49
	(0.2813)			
Fit statistics & Fixed-effects				
Bank-Time & Bin FEs			Yes	
Observations			37,490	
Pseudo R^2			0.9338	

Notes: The table reports the PPML reallocation coefficients and the predicted within-bank-month shares. Decile 1 is the omitted reference category. Clustered (Bank) standard errors are in parentheses. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$. Shares are computed via the softmax mapping s_q^{pre} in (10) and s_q^{post} in (11), renormalized to sum to 100%. Deciles are defined using pre-treatment loans of treated banks.

Lastly, in Appendix G we also document asymmetric and non-linear effects of the EU-wide CST.

7 Robustness Checks

In this section we conduct a series of robustness checks to ensure that our results are robust to the particular design of the CST. Recall that the CST hypotheses were disclosed by the ECB in November 2020, but the exact methodology was circulated to the stress-tested banks in October 2021. In our previous specifications, we posited as a condition for identification that treated banks do not alter their behavior before the policy intervention in response to anticipated changes. Therefore, we assumed that detailed information regarding the exact scope and content of the CST was not public until October 2021. This is a somewhat strong assumption. Also, as aforementioned some French banks had already participated in a national ACPR pilot exercise in 2020, which might be interpreted as a staggered treatment setting. Therefore, we will also control for this particular configuration. In the following, we consider (i) alternative shock dates as well as a Placebo exercice, and (ii) a staggered design of the stress test. Differences in treatment intensity are investigated in the following section.

7.1 Alternative Climate Stress Test Shock

We first assess the sensitivity of our findings to the timing of the CST shock. In the baseline specification, the treatment onset T_0 is set to October 2021, when the detailed methodology and templates of the ECB climate stress test were circulated to participating banks. This implicitly assumes that banks did not adjust their lending behaviour in anticipation of the exercise based solely on the earlier publication of the CST hypotheses in November 2020. To relax this assumption, we re–estimate the baseline model and the associated channel decomposition using November 2020 as the alternative shock date, thereby treating the disclosure of the broad CST design as the relevant information arrival.

Table 4 reports the resulting decomposition of the alternative full–sample slope effect $\widehat{\theta}^F_{\text{alt}}$ into intensive and extensive margins. The intensive–margin component remains essentially zero. This confirms that redefining the shock date does not uncover any hidden within–relationship repricing of carbon risk. The pattern for the extensive margin is also qualitatively stable. The reallocation among incumbents channel is still negative and statistically significant at the one–percent level, only slightly smaller in magnitude than the baseline estimate. Hence, even when the post–CST period is expanded to start in late 2020, treated banks continue to tilt new lending away from more carbon–intensive incumbent borrowers relative to cleaner incumbents. The entry/exit pricing component remains positive, although it is now not significant, while the entry/exit is stil insignificant. Compared with the baseline decomposition in Table H2, the alternative–shock results suggest some attenuation and redistribution of the extensive–margin contribution across pricing and composition, but no reversal of sign.

Table 4. Channel Decomposition Alternative Shock

Channels	*Estimands*
Intensive margin	0.0002
	(0.0086)
Reallocation among incumbents	-0.0788^{***}
	(0.0276)
Entry/Exit pricing	0.0343
	(0.0230)
Entry/Exit composition	0.0186
	(0.0184)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates derived from Table A3. Standard errors (in parentheses) and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details on the bootstrap approach).

From an econometric perspective, redefining the shock date is closely related to a placebo exercise, but it does not constitute a standard placebo design with a period of pure non-treatment. In many DiD settings, moving the treatment date strictly before the actual policy intervention serves as a pure falsification test: any sizeable effect estimated around the placebo date would point to spurious pre-trends or confounding shocks. In our case, however, November 2020 is not a period of pure non-treatment: it corresponds to the public disclosure of the CST hypotheses, which constitutes an earlier and weaker information event preceding the full methodological release of October 2021. The November 2020 specification should therefore be interpreted as an alternative treatment definition that captures the possibility of anticipatory adjustments, rather than as a standard placebo with no economic content.

Viewed in this light, the comparison between Tables H2 and 4 is informative. If our baseline October 2021 results were mainly driven by pre-existing trends between treated and control banks, then, under reasonably smooth pre-trends, one would expect the November 2020 dating to generate effects of similar magnitude—or even larger—simply because more of the pre-trend would be mechanically absorbed into the post period. Instead, the core qualitative structure of the decomposition remains intact but somewhat attenuated. The intensive margin continues to play no role in either specification. The reallocation among incumbents channel stays negative and statistically significant, though slightly smaller in absolute value under the alternative shock, and the entry/exit margin remains positive but becomes less precisely estimated and partially reallocated from pricing to composition. This pattern is consistent with a gradual build-up of the CST effect—under an interpretation in which the main portfolio rebalancing is triggered by the detailed methodological implementation in late 2021, while the initial hypothesis disclosure in November 2020 induces only limited anticipatory adjustment.

Put differently, using November 2020 as the onset date can be viewed as a conservative robustness check: it both allows for early anticipatory responses and, by construction, increases the risk that pre-trends contaminate the post period. The fact that we still recover the same sign pattern across channels—no intensive-margin effect, a persistent negative reallocation among incumbents, and a positive but noisier entry/exit component—is reassuring and, under the maintained assumptions about pre-trends, suggests that our baseline findings are not merely an artefact of an arbitrary choice of October 2021. Rather, they appear qualitatively robust to a broader window that treats the earlier CST communication as part of the treatment process, and this broader definition does not reveal any large, spurious effects that would be clearly inconsistent with the DiD identification strategy.

Figure 5. Alternative Monthly Channels Decomposition (95% CI)

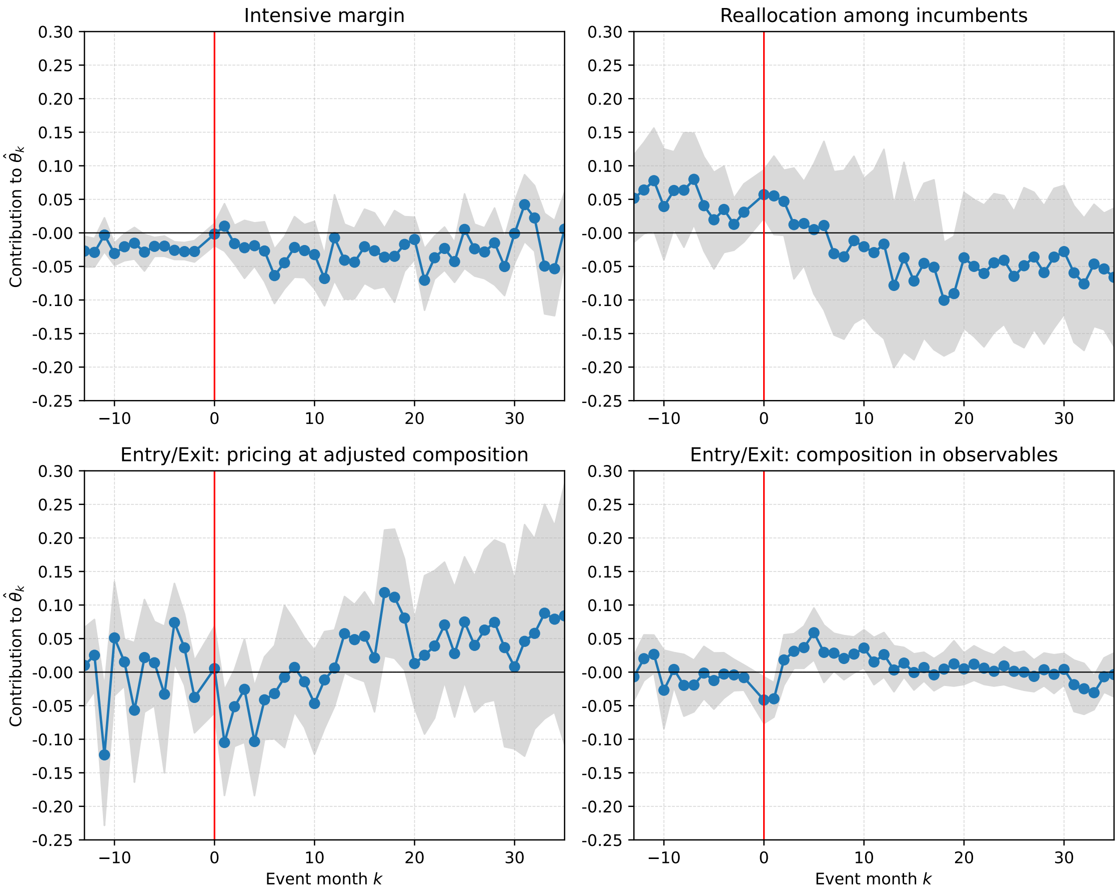

Notes: Point estimates are derived from Eq. (9) with an alternative shock. Confidence intervals and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details on the bootstrap approach).

As an additional falsification exercise, we re-estimate the baseline non-staggered specification using a placebo shock date set to November 2019, with a symmetric 12-month window around this date (see Appendix I). This placebo predates the methodological release of the ECB climate stress test in October 2021 and also occurs before the main phase of climate-stress-test implementation. The θ coefficient of equation 6 is not statistically significant, and the channel decomposition does not provide significant results with the correct sign. Overall, the placebo exercise corroborates the identifying assumptions underlying the main estimates.

7.2 Staggered Design

We also address the staggered nature of banks' exposure to CST: the French exercise (FCST) in 2020 preceded the ECB-wide exercise (ECST) in 2022. Under such staggered adoption recent literature showed that the canonical estimators can be biased.[27] Therefore, we use a staggered DiD design robust to these issues. We follow the modern DiD literature in separating identification, estimation and aggregation. Our parameter of interest is the change, for treated banks, in the conditional slope of the credit spread with respect to carbon intensity relative to a control group, which gives each loan the same weight and correspond to the effect for a typical loan in the sample, in the spirit of Callaway and Sant'Anna (2021) and Borusyak et al. (2024).

Identification. Identification relies on a parallel trends in slopes assumption: absent exposure to climate stress testing, the evolution over time of this slope would have been the same for treated cohorts and for never–treated or not–yet–treated banks, conditional on our rich set of fixed effects. We also maintain standard no-anticipation and no-spillover assumptions.

Estimation. For each bank b we define a bank–specific onset month T_b as the first month in which it is exposed to either the FCST or the ECST, and set $T_b = +\infty$ for institutions that are never treated. We then adopt an interaction–weighted event–study specification that compares treated banks to either never–treated or not–yet–treated banks and allows treatment effects to vary across cohorts and over time. Let the event time be $\tau_{bt} \equiv t - T_b$ for treated banks, and introduce the treatment indicator $D_{bt} \equiv \mathbb{1}[t \geq T_b]$ for those banks, while $D_{bt} \equiv 0$ for never–treated institutions, so that τ_{bt} is never in the post–treatment region for the latter. As before, x_{dt} denotes the carbon intensity, y_{bdt} the credit spread on the loan between bank b and debtor d at time t, and λ_{cst}, μ_{bt}, ϕ_l and η_{bd} are country–size–time, bank–time, loan–type and bank–debtor fixed effects, respectively.

We partition treated banks into cohorts $g \in G$ according to their onset month T_b, and let B_g be the set of banks in cohort g. Following Sun and Abraham (2021), for each cohort g, we conceptually construct a stacked sample that contains banks in B_g together with comparison units that are either never treated or not yet treated at time t. On this g–specific stack, we estimate the following interaction–weighted specification:

$$y_{bdt} = \beta_1 x_{dt} + \sum_{g \in G} \beta_g \Big(\mathbb{1}[b \in B_g] \cdot x_{dt} \Big) + \sum_{g \in G} \sum_{\substack{k \in \mathcal{K} \\ k \neq -1}} \theta_{g,k} \Big(\mathbb{1}[b \in B_g] \cdot \mathbb{1}[\tau_{bt} = k] \cdot x_{dt} \Big) + \lambda_{cst} + \mu_{bt} + \phi_l + \varepsilon_{bdt}, \tag{12}$$

where event month $k = -1$ is omitted and thus serves as the reference period. By construction, each coefficient $\theta_{g,k}$ measures, for cohort g, the change at event time k in the slope of the credit spread with respect to climate intensity x_{dt}, relative to the pre–treatment month $k = -1$, , while the comparison banks (never–treated or not–yet–treated) provide the overall baseline level of this slope. Under the assumptions discussed above, $\theta_{g,k}$ can thus be

[27] See, for instance, Strezhnev (2018); de Chaisemartin and d'Haultfœuille (2020); Callaway and Sant'Anna (2021); Goodman-Bacon (2021); Sun and Abraham (2021); Sun and Shapiro (2022); Borusyak, Jaravel, and Spiess (2024); Wooldridge (2021, 2023) For recent surveys, see de Chaisemartin and d'Haultfœuille (2023); Roth, Sant'Anna, Bilinski, and Poe (2023) and Zhu (2025).

interpreted as an average treatment–on–the–treated effect on the spread–carbon–intensity slope for cohort g at event time k. This interacted specification corresponds to the first step of Sun and Abraham (2021) interaction–weighted estimator applied to a DDD object (the slope with respect to x_{dt}): it recovers a collection of cohort–by–event–time effects $\{\theta_{g,k}\}_{g,k}$that play the same role as their cohort–specific CATTs.[28]

Aggregation. We estimate specification (12) on several samples or weighting schemes $s \in \mathcal{S}$. For each sample s, this delivers a collection of coefficients $\widehat{\theta}^{(s)}_{g,k}$ and cell counts $n^{(s)}_{gk}$ for cohort g at event time k. To construct scalar estimands that summarize the post-onset effects of the French and European CST, we partition the set of treated cohorts into two mutually exclusive subsets G^F and G^E according to whether their onset month T_b falls in the FCST wave or in the subsequent ECB-wide ECST wave. For $m \in \{F, E\}$, let G^m denote the corresponding subset of cohorts, and let $\mathcal{K}_m$ denote the set of post-treatment event months over which we wish to summarize the effect (e.g., all $k \geq 0$ within the estimation window). In the spirit of Sun and Abraham (2021), we interpret the aggregated parameters as interaction–weighted (IW) averages.

We consider event-time averages across cohorts (i.e., dynamic IW effects). Let $G^{\text{CST}} \equiv G^F \cup G^E$ denote the set of all treated cohorts and let $\mathcal{K}_{\text{CST}}$ denote a common post-treatment event-time window. For each sample s and each event month $k \in \mathcal{K}_{\text{CST}}$, we compute an event-time–specific, observation-weighted IW average of the cohort-specific effects at that relative time:

$$\widehat{\nu}^{(s)}_k = \sum_{g \in G^{\text{CST}}} \omega^{(s)}_{gk|\text{dyn}} \, \widehat{\theta}^{(s)}_{g,k}, \qquad \sum_{g \in G^{\text{CST}}} \omega^{(s)}_{gk|\text{dyn}} = 1,$$

where the weights $\omega^{(s)}_{gk|\text{dyn}}$ are proportional to the number of observations in each (g, k) cell of sample s at event time k,

$$\omega^{(s)}_{gk|\text{dyn}} \;\propto\; n^{(s)}_{gk} \quad \text{for } g \in G^{\text{CST}}.$$

The resulting path $\{\widehat{\nu}^{(s)}_k : k \in \mathcal{K}_{\text{CST}}\}$ is the DDD analogue of the interaction–weighted dynamic treatment effects of Sun and Abraham (2021), and should be interpreted as event–time–specific, observation–weighted averages of cohort–level effects for a typical loan in sample s.[29]

We also define an overall post-treatment IW estimand that averages across all treated cohorts and all post-onset event months in $\mathcal{K}_{\text{CST}}$:

$$\widehat{\nu}^{(s)}_{\text{CST}} = \sum_{g \in G^{\text{CST}}} \sum_{k \in \mathcal{K}_{\text{CST}}} \omega^{(s)}_{gk|\text{CST}} \, \widehat{\theta}^{(s)}_{g,k}, \qquad \sum_{g \in G^{\text{CST}}} \sum_{k \in \mathcal{K}_{\text{CST}}} \omega^{(s)}_{gk|\text{CST}} = 1,$$

with weights $\omega^{(s)}_{gk|\text{CST}}$ proportional to the number of observations in cell (g, k) of sample s. This estimand provides a global post-onset average treatment-on-the-treated effect on the spread–carbon–intensity slope for a typical loan across both the FCST and ECST waves.[30]

[28]Figure A7 presents the coefficients for each sample s, event-time k, and cohort g.
[29]Figure A8 presents the results for different underlying populations.
[30]Table A4 presents the results for different underlying populations.

Inference. For all aggregated estimands, we compute standard errors via the delta method applied to the clustered covariance matrix of the regression coefficients and treating aggregation weights as fixed.

Channels Decomposition. Lastly, we use the framework laid out in Section 6.3 to decompose the estimands and we present the results below.

In the non-staggered specification, where we treat the CST as a single shock hitting all SIs in October 2021, the channel decomposition in Table H2 delivers a very sharp narrative. The aggregate DiD in the spread–intensity slope is close to zero because a sizeable negative reallocation among incumbents is almost exactly offset by a positive entry/exit pricing component of similar magnitude, while the intensive margin and the entry/exit composition term are essentially negligible. In that framework, the CST appears to induce banks to *green* the composition of new lending within incumbent relationships (tilting origination away from more carbon-intensive incumbents) but, at the same time, to apply more carbon-sensitive pricing on the extensive margin (entry and exit), so that the overall portfolio-level slope effect is small but masks large offsetting mechanisms.

Table 5 revisits this decomposition using the staggered estimand $\widehat{\nu}^{(s)}_{\mathrm{CST}}$, which averages cohort-specific effects over both the French pilot and the ECB-wide CST and over post-onset event months. Once exposure timing is taken into account, the pattern of channels becomes somewhat different. The intensive margin is now slightly negative (-0.011) but remains statistically insignificant, so there is still no robust evidence of systematic within-relationship repricing of carbon risk. The reallocation among incumbents channel, which was large and negative in the non-staggered design, shrinks to a small positive and insignificant value. This suggests that the strong *greening among incumbents* effect in the single-shock specification was at least partly an artefact of forcing all CST-related adjustments into a single break date: when we allow for staggered adoption and dynamic cohort-specific responses, the average post-onset rebalancing within the incumbent portfolio is much closer to zero. By contrast, the entry/exit pricing channel remains positive and statistically significant, although its magnitude is substantially smaller than in the non-staggered case. This indicates that the extensive margin continues to be the main source of carbon-sensitive adjustment once we aggregate across CST waves and event times: treated banks price transitory relationships more steeply in carbon intensity than incumbents, but the strength of this effect is more modest and spread out over time and cohorts. Finally, the entry/exit composition term becomes small, negative, and statistically significant, but its magnitude is very small relative to the pricing channel, implying that differences in the observable composition of entering and exiting borrowers play a secondary role.

Table 5. Staggered Channels Decomposition by Cohort

Channels	*Estimands*
Intensive margin	-0.0110
	(0.0106)
Reallocation among incumbents	0.0147
	(0.0098)
Entry/Exit pricing	0.0264^{**}
	(0.0130)
Entry/Exit composition	-0.0038^{***}
	(0.0016)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates are derived from Table A4. Standard errors (in parentheses) and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details).

However, these staggered estimands should be interpreted with caution because they rely on a strong modelling simplification: once a bank is exposed to a stress test, it is treated forever (absorbing treatment), and we compress treatment into a single binary indicator. In practice, banks move through a sequence of treatment states—FCST only, then FCST+ECST for French SIs, and ECST-only for non-French SIs—so we are closer to a multi-valued treatment setting than to a pure one-time binary shock. The recent literature on staggered designs with multi-valued or sequential treatments emphasises that collapsing such states into a single treated dummy generally identifies a complicated mixture of effects, weighting together transitions such as never $\rightarrow$ FCST $\rightarrow$ FCST+ECST and never $\rightarrow$ ECST-only in a non-trivial way. In our context, this means that the staggered CST channels reported in Table 5—and the underlying estimand in Table A4—should be read as average effects along the observed treatment paths, not as clean marginal effects of adding the ECST on top of the FCST or of the CST in isolation. Hence these patterns should be interpreted as suggestive rather than fully structural: they highlight where the adjustment margins are likely to be, but a richer multi-valued staggered framework would be needed to cleanly separate the incremental effects of the FCST and the ECST.

Figure 6. Staggered Channels Decomposition (95% CI)

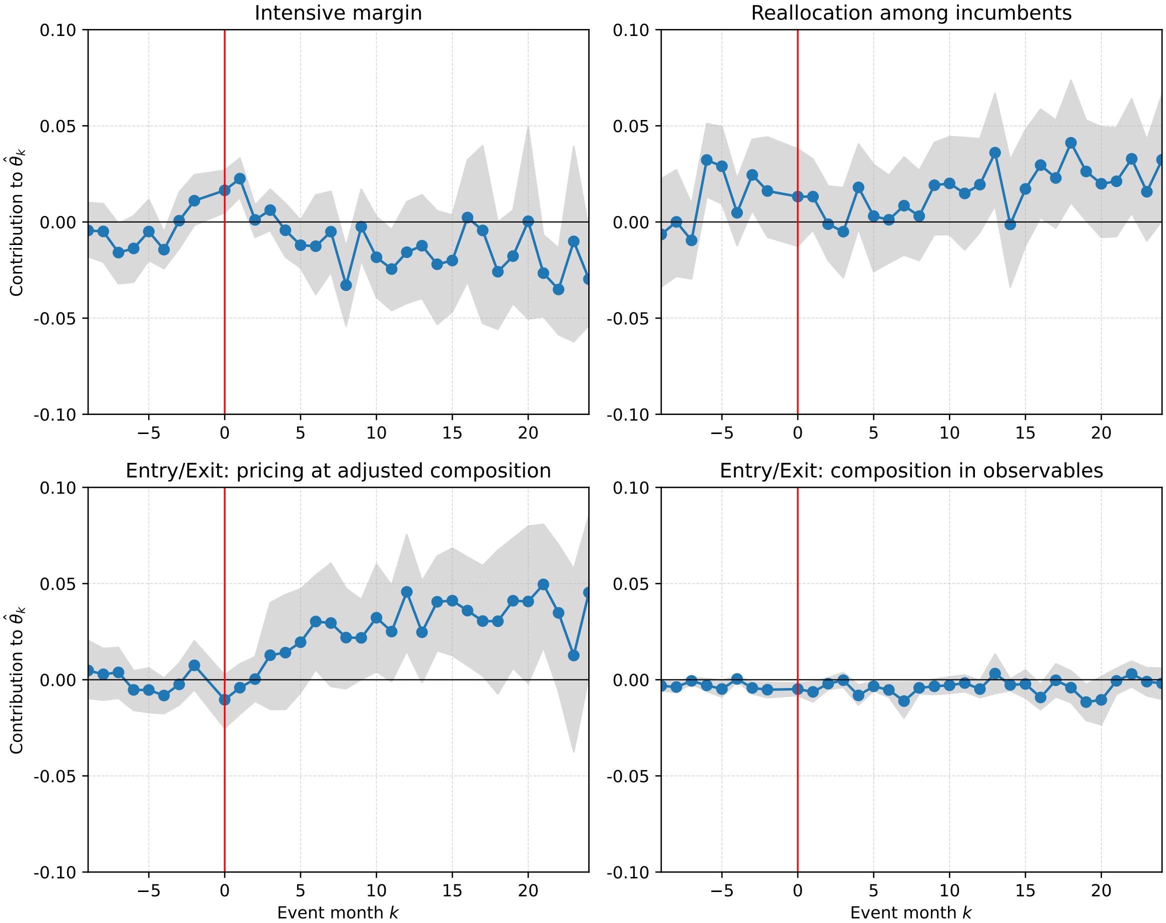

Notes: Point estimates are derived from Eq. (12). Confidence intervals and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details).

Finally, Figure 6 decomposes the dynamic estimand into the four channels over event time. Overall, the profiles confirm that the typical-loan response of the spread–intensity slope is not driven by systematic within-relationship repricing. This channel remains small throughout: estimates hover close to zero, with only a modest uptick around the onset of climate stress testing and then drifting slightly into negative territory.

The pattern of the reallocation–among–incumbents channel in the staggered specification is quite different from the single-shock baseline. In the standard event-study design, reallocation turns clearly negative after the CST date and remains so for most post-treatment months, suggesting a persistent shift of lending away from more carbon-intensive incumbent borrowers. In the staggered event–study, the reallocation series instead hovers close to zero throughout the window, and there is no pronounced break in the neighbourhood of the onset date. This pattern implies that once we allow treatment to arrive in different waves and aggregate across cohorts, there is little evidence of a systematic *greening* or *browning* of the incumbent portfolio. The main adjustment margins in the staggered design thus come from the extensive margin—captured by the entry/exit pricing channel—rather than from large, persistent reallocations of credit among continuing bank–firm relationships.

Indeed, the entry/exit pricing channel is close to zero before the shock, then becomes increasingly positive a few months after treatment and stays elevated over most of the post-onset window. This points to a gradual strengthening of carbon-sensitive pricing on the extensive margin: new and terminating relationships are priced with a steeper spread–intensity slope than what would be implied by incumbents alone. In contrast, the entry/exit composition component remains quantitatively small and statistically indistinguishable from zero at almost all horizons, indicating that differences in observable characteristics between stayers and non-stayers play only a minor role. Taken together, the dynamic decomposition underscores that, in the staggered setting, climate stress testing primarily operates through portfolio recomposition and differential pricing of entry and exit, rather than through large, sustained adjustments in within-relationship semi-elasticities.

8 Difference in Treatment Intensity

In this section, we retain the econometric framework used in the previous sections but address a different question. Whereas our baseline analysis focused on identifying the average effect of the CST on banks' lending behaviour—by comparing SIs (treated) with LSIs (control)—we now ask whether the intensity of the treatment matters. Recall that the theoretical model in Section 4 treats the CST as an information shock. The present empirical exercise can therefore be viewed as probing the role of signal quality: we test whether banks subject to a more demanding and informative version of the CST adjust their pricing more strongly.

To do so, we exploit an additional institutional feature of the ECB's 2022 CST, which is organised in three complementary modules. All SIs were required to participate in Modules 1 and 2, which collect information on banks' internal climate stress testing capabilities, their exposures to carbon-intensive sectors, and financed emissions. By contrast, only a subset of 41 SIs were selected—under a principle of proportionality—to participate in Module 3, a bottom-up stress test requiring banks to project credit losses under transition and physical risk scenarios. This latter module entails substantially greater supervisory scrutiny and a deeper quantitative engagement with climate risks than in Modules 1 and 2. In our setting, this naturally induces a second treatment dimension: banks participating in Module 3 are classified as treated institutions, while the remaining SIs that only contribute to Modules 1 and 2 form the control group.

We estimate a variant of Eq. (6) in which the key quantity of interest is the triple interaction between climate intensity, Module 3 participation and the post-CST period. We compare the evolution of credit spreads for more and less climate-exposed loans granted by Module 3 banks to the corresponding evolution for otherwise similar loans extended by banks that only participated in Modules 1 and 2. The triple-interaction coefficient θ then measures the differential change, after the CST, in the sensitivity of credit spreads to climate intensity for Module 3 banks relative to other SIs.

The choice of Module 3 participation as a treatment indicator is tightly linked to the institutional logic of the exercise. Selection into the bottom-up module is not random: these banks are typically larger or more advanced in terms of data infrastructure and climate risk modelling. From our perspective, this selection reflects the supervisor's view that these institutions are sufficiently prepared to internalise and operationalise detailed climate stress test results. The parameter θ is therefore interpreted as the incremental effect of a more intrusive and quantitatively demanding CST on the pricing of credit risk, over and above the common information and governance channel that arises from participation in Modules 1 and 2, which applies to the full SI universe.

The channel decomposition in Table 6 suggests that the answer to the question: does undergoing the more demanding Module 3 generate an additional shift in carbon–risk pricing? is, at least in the short run and on average, essentially no.

The estimates indicate that Module 3 marginal tightening of supervision does not translate into a detectable amplification of carbon-sensitive pricing along any of the three margins. The intensive margin coefficient is very close to zero, indicating no systematic incremental repricing of a given bank–borrower relationship for Module 3 banks relative to other SIs. The reallocation channel is positive in magnitude, which is suggestive of stronger portfolio tilting among incumbents for Module 3 participants, but insignificant. Likewise, the entry/exit component is small and statistically indistinguishable from zero, pointing to no clear additional adjustment on the extensive margin within the SI universe.

Table 6. Channel Decomposition Modules 1&2 vs. Module 3

Channels	*Estimands*
Intensive margin	0.0035
	(0.0105)
Reallocation among incumbents	0.0330
	(0.0208)
Entry/Exit (total)	0.0023
	(0.0167)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates derived from Table A3. Standard errors (in parentheses) and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details on the bootstrap approach).

From the perspective of the hypothesis that CST corresponds to an information shock, this pattern is informative. The main action seems to come from crossing the threshold into being stress-tested as an SI at all: once banks are exposed to Modules 1&2, any extra information or discipline delivered by Module 3 does not show up as a further, robust strengthening of the spread–carbon–intensity slope. This can be read in two ways. One interpretation is that the baseline CST common to all SIs is already sufficient to move pricing and portfolios, so that the marginal supervisory pressure from Module 3 operates more on internal risk processes, disclosure and governance than on observable loan rates. Another, complementary view is that organisational frictions, capital and business-model constraints, or expectations about the permanence of the exercise limit how far banks can adjust their lending policies in response to additional stress-test modules. In either case, the comparison between SIs and LSIs captures a first-order shift in behaviour, whereas the within-SI contrast between Modules 1&2 and Module 3 points to the absence of any incremental effects on the pricing of carbon risk.

9 Conclusion

This paper has studied how the ECB's 2022 CST is reflected in banks' pricing and allocation of credit across borrowers with different carbon intensities. Using AnaCredit loan-level data matched to an *ad hoc* firm-level emissions intensity measure, we estimate a difference-in-differences-in-differences design in which the triple interaction between CST exposure, a post-CST dummy, and log carbon intensity identifies changes in the spread–emissions semi-elasticity for treated banks. By nesting two increasingly restrictive samples, we decompose the aggregate slope effect into four economically distinct channels: within-relationship repricing (intensive margin), reallocation of lending among incumbents, entry/exit pricing at a given observable composition, and entry/exit composition effects.

The baseline specification delivers a small and statistically insignificant net effect of the CST on the overall spread–intensity slope. However, the decomposition shows that this aggregate masks two sizeable and offsetting adjustments. Treated banks do not systematically reprice existing relationships as a function of carbon intensity once we control for bank–debtor fixed effects, indicating little evidence of intensive-margin repricing. Instead, they reallocate new lending away from more carbon-intensive incumbent borrowers, consistent with a tightening of credit supply for *brown* firms relative to *greener* incumbent firms. At the same time, treated banks apply more carbon-sensitive spreads on the extensive margin—entry and exit of relationships—than on the pool of incumbents, which steepens the spread–intensity profile on transitory relationships. The net portfolio effect is therefore close to zero in the aggregate, but it is underpinned by economically meaningful reallocations and differential pricing across margins.

Extending the analysis to a staggered setting that aligns banks by their first exposure to either the French pilot or the ECB-wide CST and aggregates across cohorts delivers a more nuanced picture. Once timing of adoption is taken into account, the reallocation among incumbents channel is statistically insignificant, while the main contribution continues to come from entry/exit pricing, which remains positive and significant but is substantially smaller than in the single-shock design. In the staggered framework, the typical-loan effect of climate stress testing therefore arises primarily through more carbon-sensitive pricing of entering and exiting relationships, with only limited evidence of systematic portfolio rebalancing among continuing bank–firm pairs or of large changes in observable composition.

These staggered results should nonetheless be interpreted cautiously. The empirical implementation adopts an absorbing binary treatment indicator: once a bank has participated in any climate stress test, it is defined as treated for all subsequent periods. In practice, banks move through a sequence of treatment states—no CST, French pilot only, French pilot plus ECB-wide CST, or ECB-wide CST only—so the data generating process is probably closer to a multi-valued, sequential treatment setting than to a single one-time intervention. The recent literature on staggered designs with multi-valued treatments shows that collapsing such states into a single treated dummy recovers mixtures of transition effects rather than clean marginal impacts of each step.

A complementary exercise exploits heterogeneity in treatment intensity within the set of SIs by comparing banks that participated in the CST's bottom-up Module 3 with SIs that only contributed to the more qualitative Modules 1 and 2. Here the evidence for additional adjustment at higher treatment intensity is limited. The triple-interaction coefficient for Module-3 banks indicates only a modest and statistically insignificant increase in the spread–intensity slope relative to other SIs, and the associated channel decomposition does not reveal a clearly distinguishable pattern on the intensive or extensive margins. One interpretation is that the main supervisory signal is already delivered by the CST as a whole—through common scenarios, data collection and supervisory dialogue—and that the incremental effect of the more demanding bottom-up module on pricing behaviour is modest relative to the heterogeneity in banks' pre-existing capabilities and business models.

We also document that adjustments are asymmetric with respect to the direction of change in borrower emissions intensity. When we allow the spread–intensity semi-elasticity to differ between months in which a borrower's intensity is rising and months in which it is falling, treated banks display more negative post-CST slopes in both regimes in the incumbents sample, with the effect being more pronounced when intensity decreases. Once bank–debtor fixed effects are added, the asymmetric triple-interaction terms become statistically insignificant, and the associated channel decomposition attributes essentially none of the asymmetry to within-relationship repricing. Instead, the entire asymmetric response loads on reallocation among incumbents: treated banks tilt new lending away from more carbon-intensive borrowers both when their intensity increases and when it decreases, but the reweighting is stronger in periods of decarbonisation. This pattern is consistent with banks conditioning their portfolio choices on both the level and the trajectory of emissions intensity, rewarding transitions away from carbon dependence more strongly than penalising temporary setbacks.

Taken together, these findings are compatible with the hypothesis that climate stress tests influence lending behaviour in a way that is reminiscent of the impact of traditional solvency stress tests on risk-taking. Acharya et al. (2018) show that banks facing tighter supervisory scrutiny reduce credit supply to relatively risky borrowers rather than compensating risk with higher spreads. In our context, SIs exposed to the CST do not uniformly increase the carbon premium across all loans. Instead, they reallocate credit away from carbon-intensive incumbents and use the extensive margin to apply more carbon-sensitive pricing to marginal relationships. In that sense, CSTs appear to operate primarily as portfolio-rebalancing devices—shifting the composition of lending towards relatively greener firms—rather than as broad price-level shocks on the intensive margin.

Nevertheless, because of data limitations and concurrent regulatory developments, we remain cautious in making strong causal claims. Our firm-level intensity measure is imputed from sector–country–size emissions and turnover rather than directly observed at the borrower level, and other climate-related supervisory initiatives and disclosure requirements were rolled out over the same period. These factors imply that our estimates could capture the combined effect of the CST and the broader supervisory *green turn* rather than a pure CST treatment effect. Nevertheless, the consistent sign pattern across specifications, robustness to alternative shock timings, and stability of the channel decomposition suggest that the CST played a substantive role in reshaping the way euro-area SIs allocate and price carbon-intensive credit risk.

From a policy perspective, our results underscore that climate stress tests can affect banks' behaviour even in the absence of explicit *brown penalising* capital charges. By revealing portfolio vulnerabilities and forcing banks to quantify sectoral and borrower-level exposures, the CST appears to encourage a gradual rebalancing away from carbon-intensive incumbents and a sharper pricing of carbon risk on the margin, while leaving average spreads broadly unchanged. Future research could explore the real-economy consequences of this reallocation—on investment, emissions trajectories, and firm entry and exit—as well as the interaction between stress tests and other climate policies. On the supervisory side, our evidence points to the importance of designing stress-test exercises that not only generate aggregate risk metrics but also provide sufficiently granular information to influence bank-level portfolio decisions along both the intensive and extensive margins.

References

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

Acharya, V. V., R. Berner, R. Engle, H. Jung, J. Stroebel, X. Zeng, and Y. Zhao (2023). Climate stress testing. *Annual Review of Financial Economics 15*(1), 291–326.

Agarwal, I., R. Duttagupta, and A. F. Presbitero (2020). Commodity prices and bank lending. *Economic inquiry 58*(2), 953–979.

Aiello, M. A. (2024). Climate supervisory shocks and bank lending: Empirical evidence from microdata. *Available at SSRN 4756516*.

Akerlof, G. A. (1970). The market for 'lemons': quality uncertainty and the market mechanism. *Market Failure or Success 66*.

Allen, F. and D. Gale (2000). Financial contagion. *Journal of political economy 108*(1), 1–33.

Alogoskoufis, S., N. Dunz, T. Emambakhsh, T. Hennig, M. Kaijser, C. Kouratzoglou, M. A. Muñoz, L. Parisi, and C. Salleo (2021). *ECB economy-wide climate stress test: Methodology and results*. Number 281. ECB Occasional Paper.

Altavilla, C., M. Boucinha, M. Pagano, and A. Polo (2023). Climate risk, bank lending and monetary policy. *European Corporate Governance Institute–Finance Working Paper* (936).

Altavilla, C., M. Boucinha, F. Smets, et al. (2020). Banking supervision, monetary policy and risk-taking: Big data evidence from 15 credit registers. Technical report, CEPR Discussion Papers.

Amiti, M. and D. E. Weinstein (2018). How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy 126*(2), 525–587.

Baudino, P. and J.-P. Svoronos (2021). *Stress-testing banks for climate change: a comparison of practices*. Bank for International Settlements, Financial Stability Institute.

Beyer, A. and L. Schreiner (2024). The impact of ecb banking supervision on climate risk and sustainable finance. Technical report.

Borio, C. (2011). Implementing a macroprudential framework: Blending boldness and realism. *Capitalism and society 6*(1).

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*.

Bräuning, F. and J. L. Fillat (2020). The impact of regulatory stress tests on bank lending and its macroeconomic consequences.

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

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The review of economics and statistics 90*(3), 414–427.

Clerc, L., A.-L. Bontemps-Chanel, S. Diot, G. Overton, S. Soares de Albergaria, L. Vernet, and M. Louardi (2021). *A first assessment of financial risks stemming from climate change: The main results of the 2020 climate pilot exercise.* Number 122-2021. ACPR Analyses et synthèses.

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

Cooper, R. and A. John (1988). Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics 103*(3), 441–463.

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

de Bandt, O., L.-C. Kuntz, N. Pankratz, F. Pegoraro, H. Solheim, G. Sutton, A. Takeyama, and F. D. Xia (2023). The effects of climate change-related risks on banks: A literature review. *Journal of Economic Surveys*.

de Chaisemartin, C. and X. d'Haultfœuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–2996.

de Chaisemartin, C. and X. d'Haultfœuille (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *The Econometrics Journal 26*(3), C1–C30.

Degryse, H., O. De Jonghe, S. Jakovljević, K. Mulier, and G. Schepens (2019). Identifying credit supply shocks with bank-firm data: Methods and applications. *Journal of Financial Intermediation 40*, 100813.

Delis, M. D., K. d. Greiff, M. Iosifidi, and S. Ongena (2024). Being stranded with fossil fuel reserves? climate policy risk and the pricing of bank loans. *Financial markets, institutions & instruments 33*(3), 239–265.

Durrani, A., J. Metzler, N. Michail, and J.-G. Werner (2020). *Bank lending rates and the remuneration for risk: evidence from portfolio and loan level data.* Number 2753. ECB Working Paper.

ECB (2020). *Guide on climate-related and environmental risks.* European Central Bank.

ECB (2021). *The state of climate and environmental risk management in the banking sector.* European Central Bank.

ECB (2022). *Walking the talk.* European Central Bank.

Eisenberg, L. and T. H. Noe (2001). Systemic risk in financial systems. *Management Science 47*(2), 236–249.

Elderson, F. (2023, November). Making finance fit for paris: achieving negative splits. Keynote speech by Frank Elderson, Member of the Executive Board of the ECB and Vice-Chair of the Supervisory Board of the ECB, at the conference on "The decade of sustainable finance: half-time evaluation", Brussels.

Elderson, F. (2024, March). Making banks resilient to climate and environmental risks – good practices to overcome the remaining stumbling blocks. Speech by Frank Elderson, Member of the Executive Board of the ECB and Vice-Chair of the Supervisory Board of the ECB, at the 331st European Banking Federation Executive Committee meeting, Frankfurt am Main.

Flannery, M., B. Hirtle, and A. Kovner (2017). Evaluating the information in the federal reserve stress tests. *Journal of Financial Intermediation 29*, 1–18.

Fuchs, L., H. Ngyuen, T. Nguyen, and K. Schaeck (2024). Climate stress tests, bank lending, and the transition to the carbon-neutral economy. Technical report, IWH Discussion Papers.

Gan, J. (2007). The real effects of asset market bubbles: Loan-and firm-level evidence of a lending channel. *The Review of Financial Studies 20*(6), 1941–1973.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Gschossmann, I., C. Kok, and V. De Cicco (2025). *Bank lending implications of climate stress tests.* Number 3088. ECB Working Paper.

Haselmann, R., J. P. Krahnen, T. Tröger, and M. Wahrenburg (2022). *Institutional protection schemes: What are their differences, strengths, weaknesses, and track records?* Number 88. SAFE White Paper.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, 153–161.

IMF (2021). 2021 comprehensive surveillance review. *Policy Paper No. 2021/033*, 1–32.

Ivanov, I. T., M. S. Kruttli, and S. W. Watugala (2024). Banking on carbon: Corporate lending and cap-and-trade policy. *The Review of Financial Studies 37*(5), 1640–1684.

Ivashina, V. (2009). Asymmetric information effects on loan spreads. *Journal of financial Economics 92*(2), 300–319.

Jakovljević, S., H. Degryse, and S. Ongena (2020). Introduction to the symposium on contemporary banking research: the use of fixed effects to disentangle loan demand from loan supply. *Economic Inquiry 58*(2), 917–920.

Jourde, T., U. A. Szczerbowicz, and F. van Dijk (2025). The ecb's green put: From cheap talk to priced action. *Available at SSRN*.

Jung, H., J. A. Santos, and L. Seltzer (2023). Us banks' exposures to climate transition risks. *FRB of New York Staff Report* (1058).

Khwaja, A. I. and A. Mian (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review 98*(4), 1413–1442.

Konietschke, P., S. Ongena, and A. P. Marques (2022). *Stress tests and capital requirement disclosures: Do they impact banks' lending and risk-taking decisions?* Number 2679. ECB Working Paper.

MacKinnon, J. G. and M. D. Webb (2017). Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics 32*(2), 233–254.

Miguel, F., A. Pedraza, and C. Ruiz-Ortega (2024). Climate-change regulations: Bank lending and real effects. *Journal of Financial Stability 70*, 101212.

Milgrom, P. and J. Roberts (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society*, 1255–1277.

Paligorova, T. and J. A. Santos (2017). Monetary policy and bank risk-taking: Evidence from the corporate loan market. *Journal of Financial Intermediation 30*, 35–49.

Pigou, A. (2017). *The economics of welfare*. Routledge.

Remler, D. K. and G. G. Van Ryzin (2024). Control, exogeneity, and directness: Understanding and designing quasi-and natural experiments. *American Journal of Evaluation*, 10982140241265708.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *The American economic review 71*(3), 393–410.

Strezhnev, A. (2018). Semiparametric weighting estimators for multi-period difference-in-differences designs. In *Annual Conference of the American Political Science Association, August*, Volume 30.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics 225*(2), 175–199.

Sun, L. and J. M. Shapiro (2022). A linear panel model with heterogeneous coefficients and variation in exposure. *Journal of Economic Perspectives 36*(4), 193–204.

Tielens, J. and J. Van Hove (2017). The amiti–weinstein estimator: an equivalence result. *Economics Letters 151*, 19–22.

Véron, N. (2024). Europe's banking union at ten unfinished yet transformative. *Peterson Institute for International Economics Working Paper* (24-15).

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. *Available at SSRN 3906345*.

Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal 26*(3), C31–C66.

Zhu, S. (2025). Recent developments in event study estimation.

Appendix (and Supplementary online Material

Appendices

Appendix A Additional Tables and Figures

Figure A1. Kernel Density of log(intensity) - FULL sample

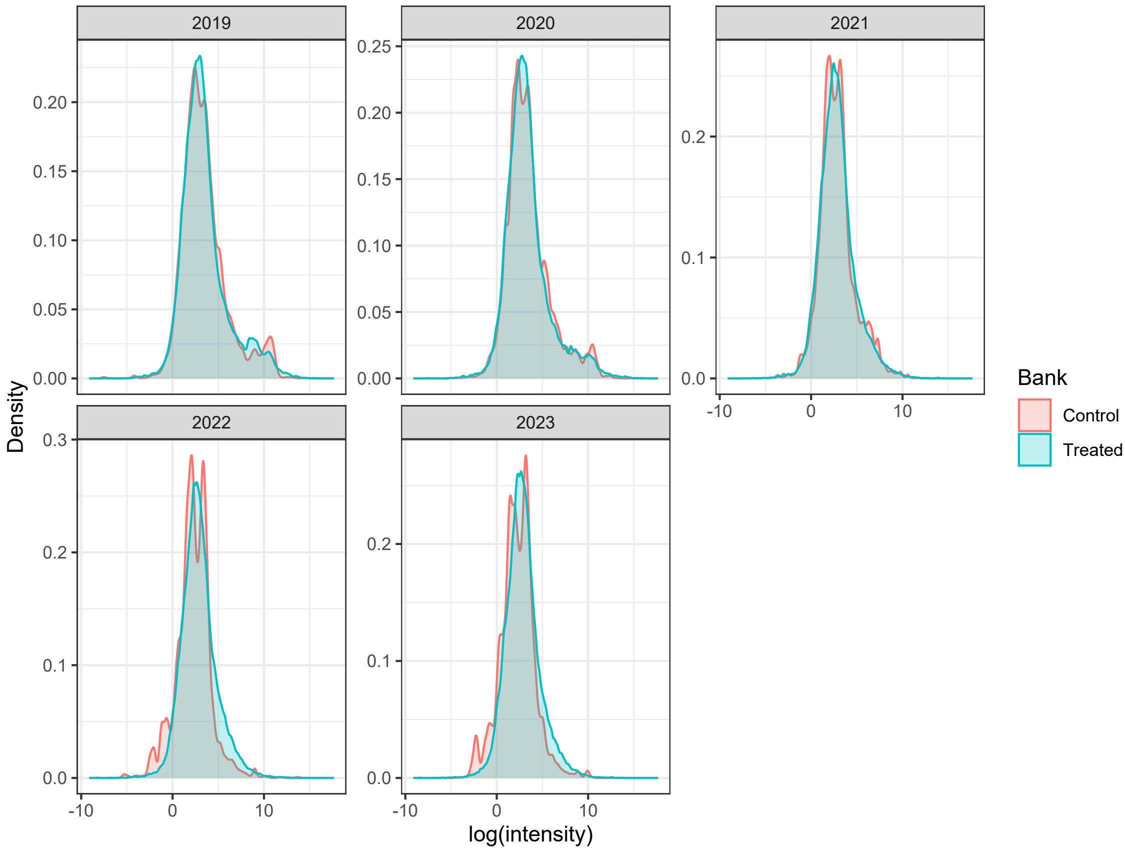

Sources: AnaCredit data and Authors' calculations.

Figure A2. Kernel Density of log(intensity) - INC sample

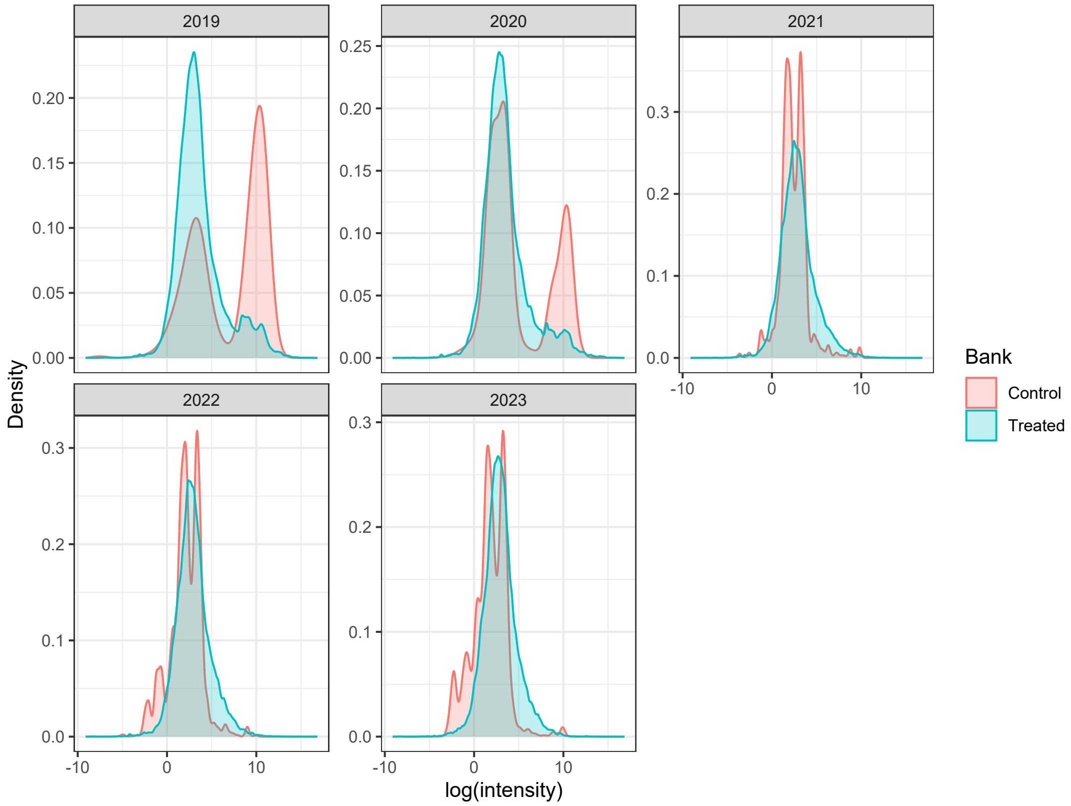

Sources: AnaCredit data and Authors' calculations.

Figure A3. Kernel Density of log(intensity) - INC sample (IPW corrected)

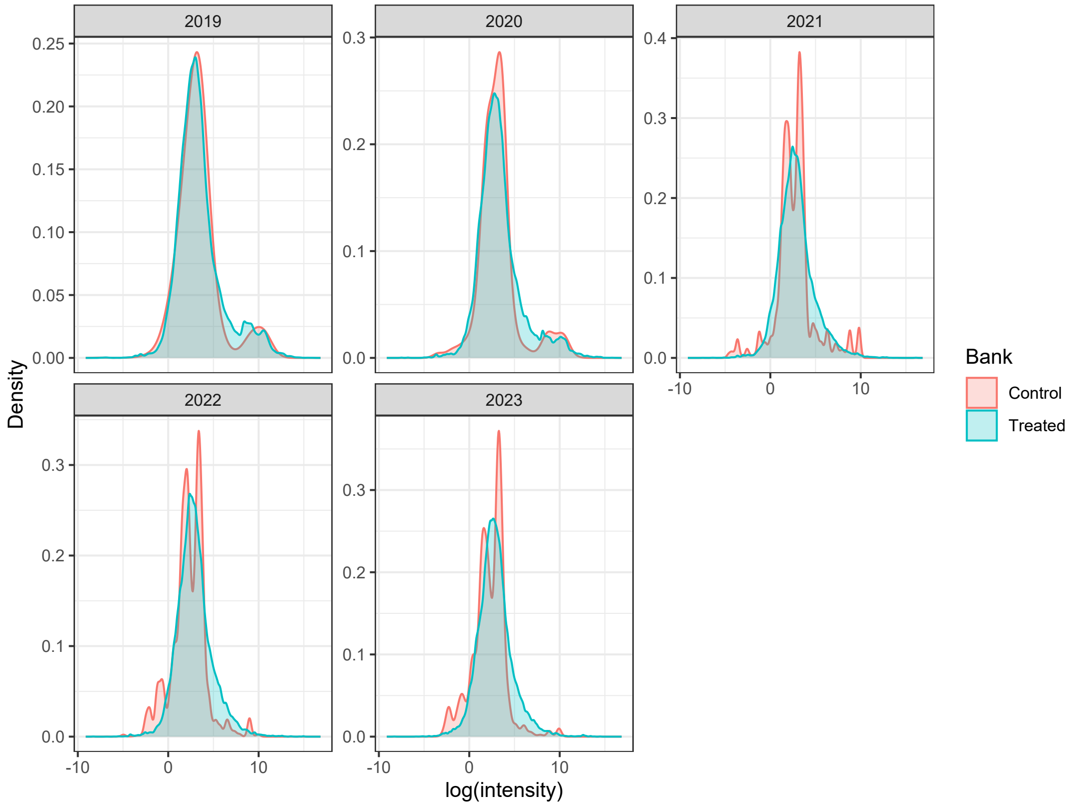

Sources: AnaCredit data and Authors' calculations.

Figure A4. Kernel Density of log(intensity) - FULL sample (alternative)

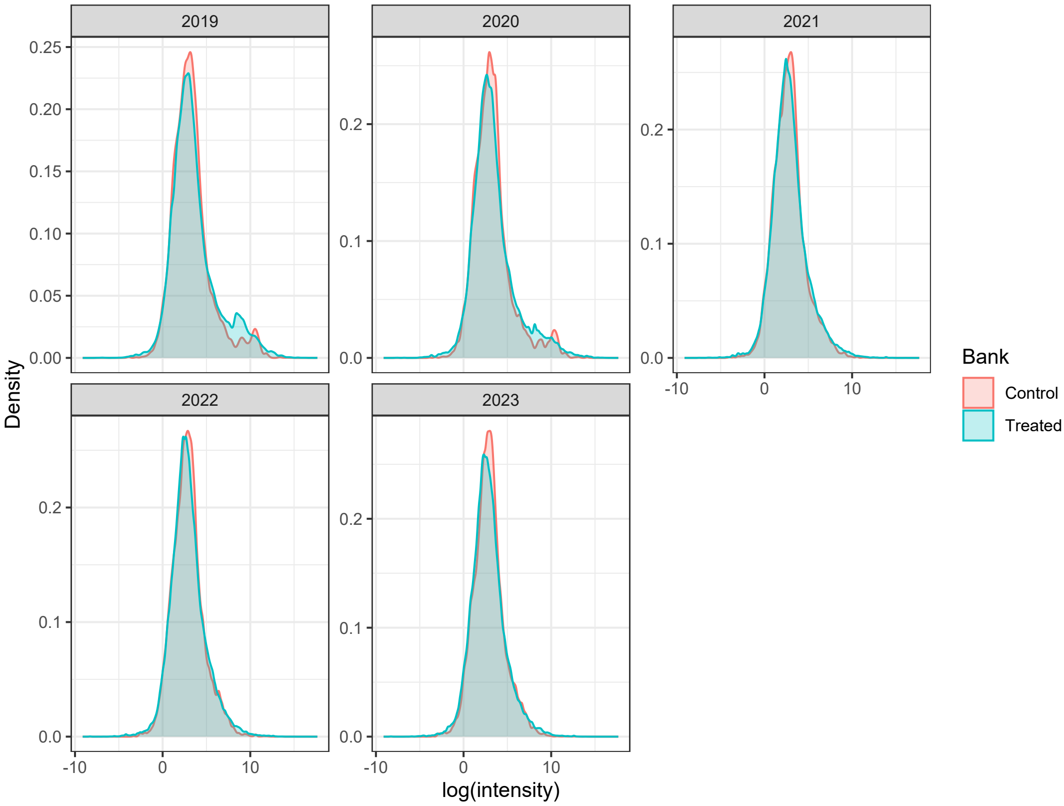

Sources: AnaCredit data and Authors' calculations.

Figure A5. Kernel Density of log(intensity) - INC sample (alternative)

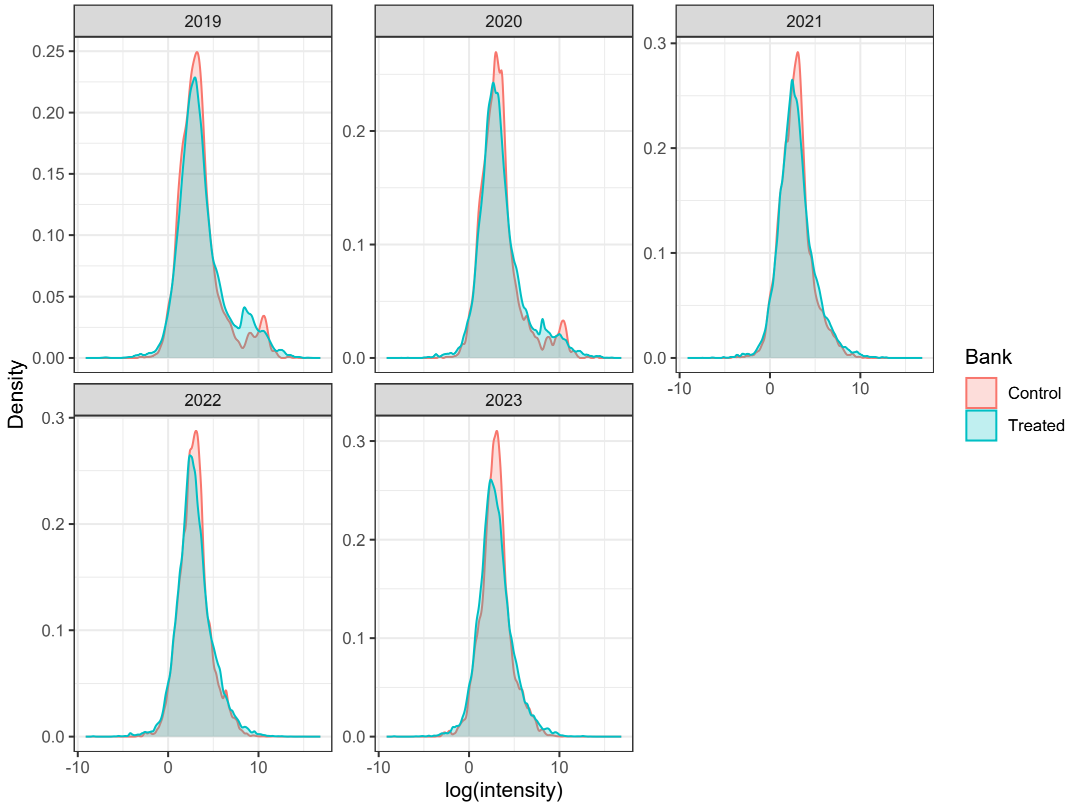

Sources: AnaCredit data and Authors' calculations.

Table A1. Channel Decomposition

Channels	*Estimands*
Intensive margin	0.0000
	(0.0095)
Reallocation among incumbents	-0.0897^{***}
	(0.0273)
Entry/Exit pricing	0.0977^{***}
	(0.0365)
Entry/Exit composition	-0.0037
	(0.0067)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates derived from Table 2. Standard errors (in parentheses) are computed via $B = 499$ wild bootstrap with Rademacher weights (See Appendix F for details on the bootstrap approach).

Table A2. Quantile Model

Panel A ($q = 0.10$)

Dependent Variable:	Credit Spread			
Model:	(1)	(2)	(3)	(4)
	FULL	INC	INC (IPW)	INT
Coefficients				
θ	0.0552	-0.1838	-0.1818	-0.0213
	(0.1130)	(0.1268)	(0.1272)	(0.0828)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel B ($q = 0.25$)

Model:	(1)	(2)	(3)	(4)
Coefficients				
θ	0.1232	$-0.1389^{\cdot}$	-0.1896^{*}	0.0408
	(0.0915)	(0.0749)	(0.0896)	(0.0649)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel C ($q = 0.50$)

Model:	(1)	(2)	(3)	(4)
Coefficients				
θ	$0.1334^{\cdot}$	-0.1097^{*}	-0.2269^{**}	$0.0848^{\cdot}$
	(0.0775)	(0.0543)	(0.0693)	(0.0474)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel D ($q = 0.75$)

Model:	(1)	(2)	(3)	(4)
Coefficients				
θ	-0.0090	-0.4116^{**}	-0.4650^{***}	-0.1860^{**}
	(0.1270)	(0.1337)	(0.1281)	(0.0690)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel E ($q = 0.90$)

Model:	(1)	(2)	(3)	(4)
Coefficients				
θ	-0.1483	-0.6127^{**}	-0.5802^{**}	-0.3708^{***}
	(0.1357)	(0.1797)	(0.1731)	(0.0880)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes
Fit statistics				
Observations	2,214,084	1,186,746	1,186,746	1,186,745
R^2	0.39838	0.42240	0.43616	0.80781
Within R^2	0.00107	0.00064	0.00071	0.00006

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$, $^{\cdot}p < 0.15$.

Table A3. Alternative Shock

Dependent Variable:	Credit Spread			
Model:	(1)	(2)	(3)	(4)
	FULL	INC	INC (IPW)	INT
Coefficients				
θ	-0.0257	-0.0786^{***}	-0.0600^{***}	0.0002
	(0.0257)	(0.0286)	(0.0224)	(0.0085)
Controls	Yes	Yes	Yes	Yes
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes
Fit statistics				
Observations	2,214,084	1,186,746	1,186,746	1,186,745
R^2	0.40091	0.42417	0.43793	0.80782
Within R^2	0.00527	0.00370	0.00385	0.00008

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure A6. Alternative Event-Study Coefficients (95% CI)

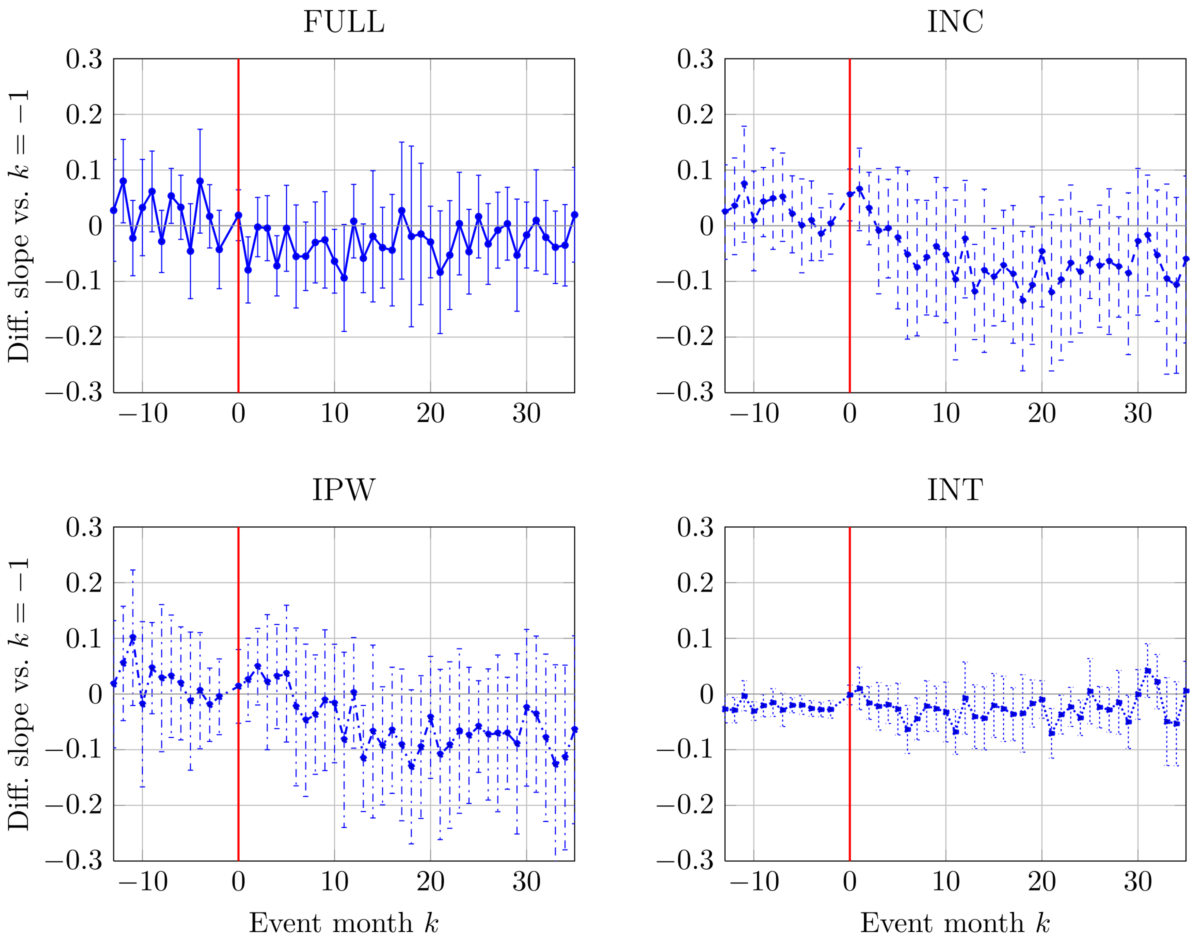

Table A4. Staggered by Cohort

Dependent Variable:	Credit Spread			
Model:	(1)	(2)	(3)	(4)
	FULL	INC	INC (IPW)	INT
Estimands				
ν_{CST}	0.0283**	0.0055	0.0017	−0.0113
	(0.0134)	(0.0141)	(0.0135)	(0.0096)
Fixed-effects				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes
Fit statistics				
Observations	1,628,037	899,241	899,241	858,042
R^2	0.40359	0.41778	0.38582	0.82627
Within R^2	0.00658	0.00575	0.00592	0.00048

Point estimates correspond to the estimands described in the Subsection 7.2.
Clustered (Bank) standard errors in parentheses, via delta method.
Fit statistics are taken from the underlying regression (12).
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Figure A7. Monthly Staggered Coefficients by Cohort (95% CI)

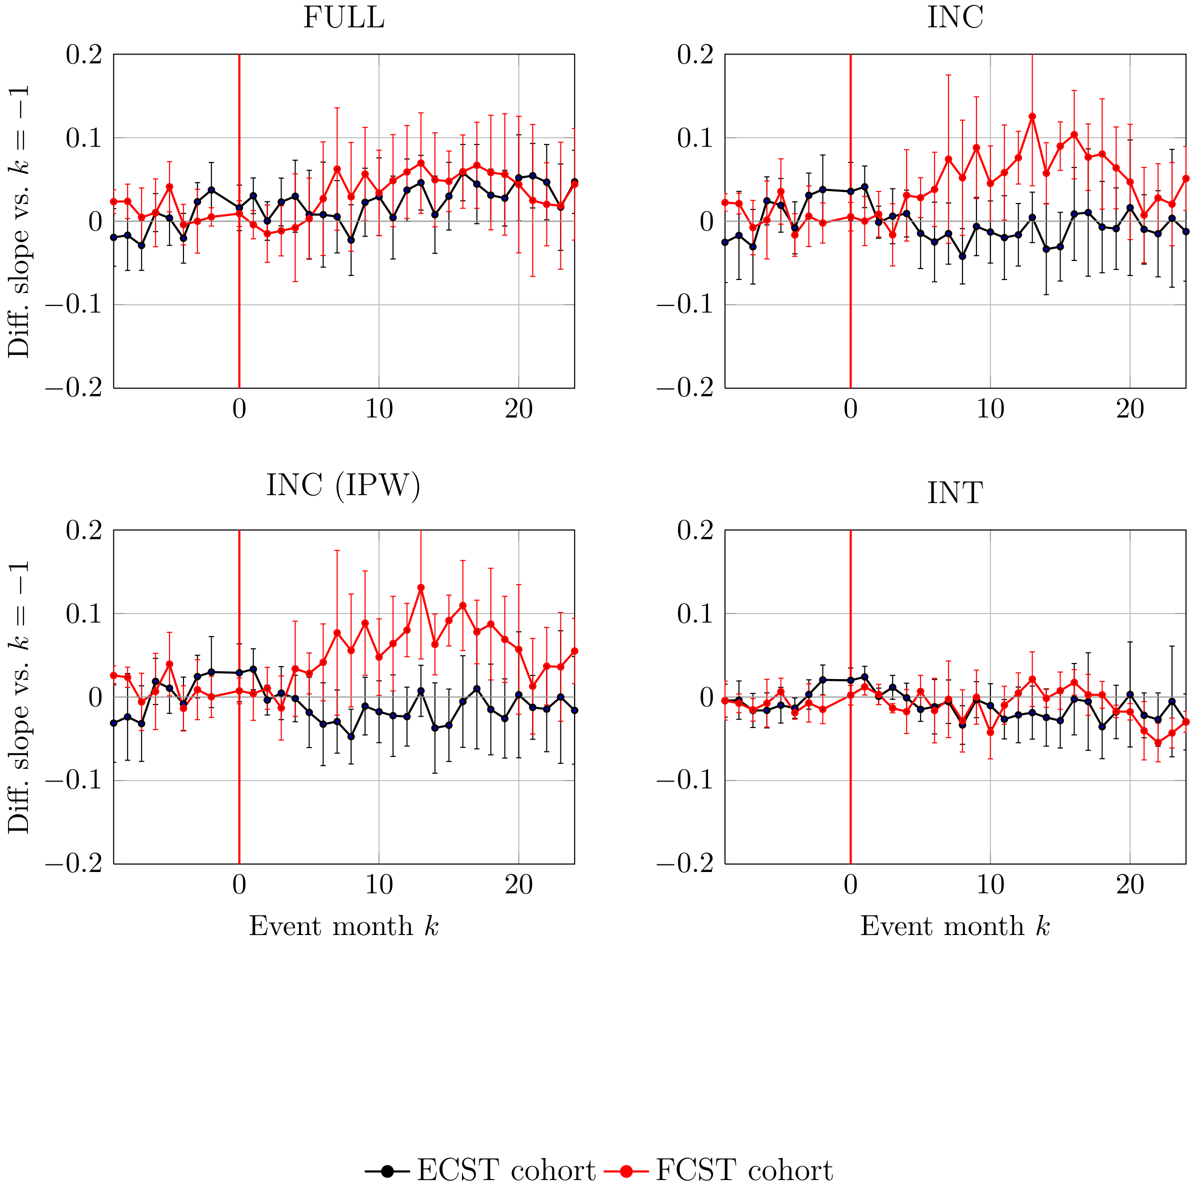

Figure A8. Monthly Staggered Estimands (95% CI)

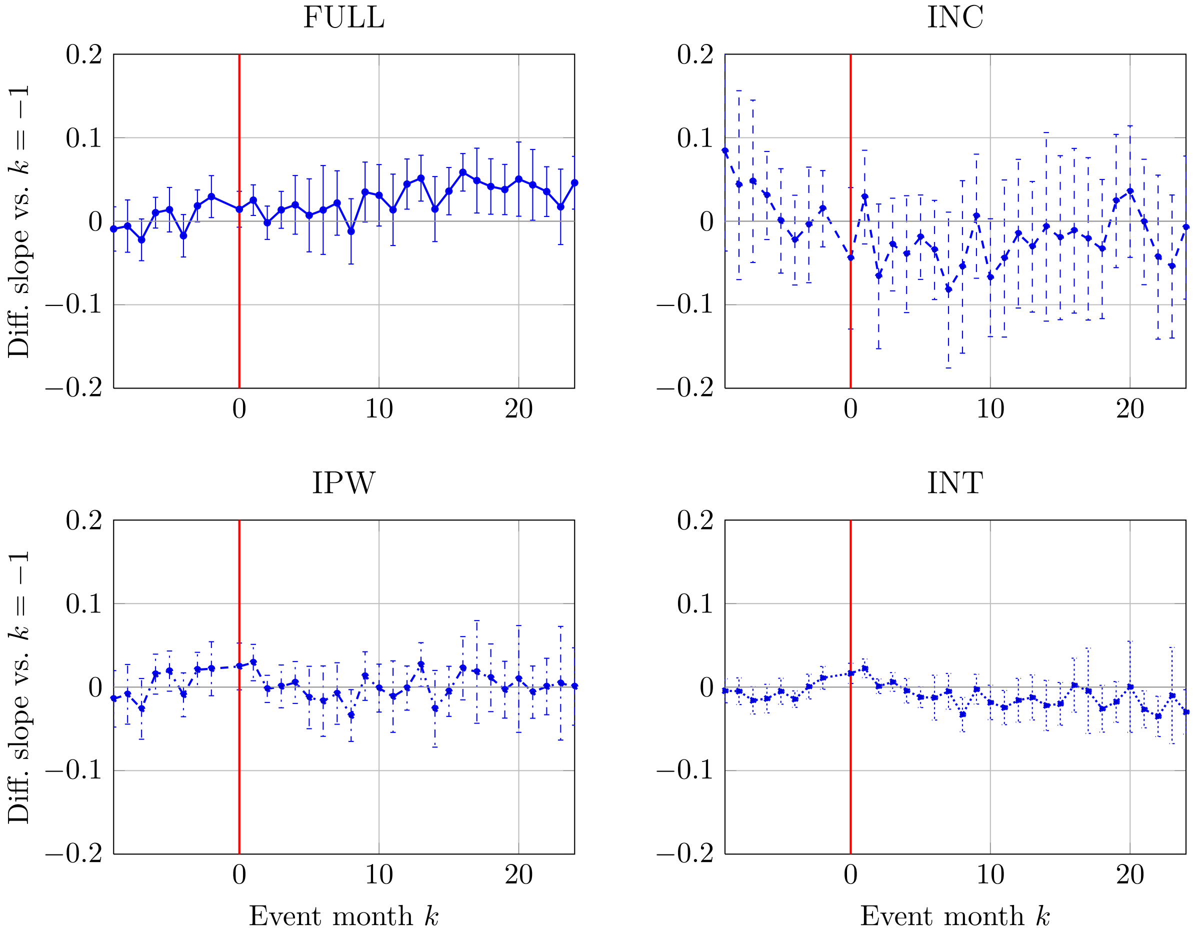

Notes: Point estimates correspond to the group-time event estimands *à la* Sun and Abraham described in the Subsection 7.2. Clustered (Bank) standard errors are derived via delta method.

Table A5. Modules 1&2 vs. Module 3

Dependent Variable:	Credit Spread		
Model:	(1)	(2)	(3)
	FULL	INC	INT
Coefficients			
θ	0.0388*	0.0365	0.0035
	(0.0232)	(0.0260)	(0.0110)
Controls	Yes	Yes	Yes
Fixed-effects			
Loan Type	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes
Bank-Debtor	No	No	Yes
Fit statistics			
Observations	2,084,975	1,149,188	1,149,187
R^2	0.37518	0.37067	0.79589
Within R^2	0.00541	0.00359	0.00009

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Appendix B Firm-Level GHG Intensity Metric

We define our firm-level GHG intensity metric by following a stepwise allocation process. Let t denote the year, c denote the country, s denote the NACE sector (at the two-digit level), k denote the firm size class, $N_{t,c,s,k}$ be the number of firms in size class k for a given (t, c, s), $T_{t,c,s,k}$ be the total turnover of firms in size class k in sector s, year t, and country c, $E_{t-1,c,s}$ be the total GHG emissions of sector s in country c and year t.

We assume that GHG emissions within a sector are distributed proportionally to the share of turnover of each firm size category. Thus, the emissions attributed to size class k in sector, s, year t, and country c are:

$$E_{t,c,s,k} = E_{t-1,c,s} \times \frac{T_{t,c,s,k}}{\sum_k T_{t,c,s,k}}$$

where $\sum_k T_{t,c,s,k}$ is the total turnover of sector s in country, c and year t.

We compute the average emissions per firm in each size class as:

$$\bar{E}_{t,c,s,k} = \frac{E_{t,c,s,k}}{N_{t,c,s,k}}$$

where $\bar{E}_{t,c,s,k}$ represents the estimated emissions per firm in size class k, sector s, country c, and year t.

For a given firm i operating in sector s, country c, and year t, with observed turnover $T_{i,t}$, we define its estimated GHG intensity as:

$$I_{i,t} = \frac{\bar{E}_{t,c,s,k}}{T_{i,t}}$$

where $I_{i,t}$ represents the estimated GHG emissions per unit of turnover for firm i.

Appendix C Data Description

Appendix C.1 Sample Construction

We construct our dataset by combining multiple data sources, starting with the Eurosystem's AnaCredit (Analytical Credit datasets).[31] AnaCredit is a harmonized credit registry for the euro area, which offers granular information, provided by the banks on a monthly basis via a standardised set of templates, on individual bank loans to corporates. The first observation is available for November 2018.

For each instrument, we consider information on outstanding nominal and off-balance-sheet amounts, on the type of instruments and interest rate from November 2018 to August 2024. However, given its nature as a administrative microdata credit register, several data cleaning steps were undertaken to ensure accuracy and consistency. Without any data cleaning, the dataset initially includes approximately 20 million new loans across the entire time span, representing the full population of new loans since end 2018 at monthly frequency.

First, loans without valid maturity values are excluded to preserve the reliability of maturity-related computations. Interest rates are then matched with the ECB yield curve data,[32] using AAA-rated euro area central government bond yields. Across the entire time span, we use spot rates (SR) as our baseline measure to capture the risk-free rate for each maturity due to their established role in representing the comprehensive term structure of interest rates. Spot rates provide a stable and consistent foundation, aligning with prevalent methodologies in the existing literature, which enhances the comparability and reliability of our findings. By utilizing spot rates, we ensure that our primary analysis remains grounded in a widely accepted framework, facilitating a clear interpretation of credit spreads across varying maturities. For fixed-rate loans, the risk-free rate is aligned with the loan's maturity at its inception date. For variable-rate loans, the risk-free rate is matched to the maturity reset date. Credit spread is subsequently calculated by subtracting the relevant risk-free rates from the observed loan interest rates.[33]

Second, to distinguish stand-alone banks from banking groups, bank-level information is consolidated at the parent level when applicable. This consolidation is paramount, as the European CST primarily targeted SIs, which are typically part of larger banking groups. The consolidation permits the accurate identification of participating banks. For geographical consistency, we keep only loans granted by banks headquartered within EU member states. Furthermore, for consistency we manually checked mergers and acquisitions to reconstitute accurate bank corporate structures.[34]

After excluding loans—mainly arising from credit lines—under EUR 5,000 and a few outliers with loans above EUR 500 million, additional preprocessing steps are taken to

[31] For more details, see: https://www.ecb.europa.eu/stats/ecb_statistics/anacredit/html/index.en.html

[32] https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/euro_area_yield_curves/html/index.en.html

[33] As a robustness check, we also use both Instantaneous Forward Rate and Interest Rate Swap data. These alternative measures do not significantly change the distribution of the credit spreads over the all time span. Data are available upon request.

[34] These steps are conducted based on banking group information provided by the SSM: https://www.bankingsupervision.europa.eu/framework/supervised-banks/html/index.en.html

address potential data inconsistencies and outliers. Loans with interest rates outside the plausible range (0.1% to 25%) are excluded to eliminate errors. Given the nature of this microdata as an administrative credit register, it is reasonable to assume that errors in interest rate reporting are random. Observations with "Mixed" interest rate types are also removed to ensure analytical consistency.

We also control for loans covered by public guarantee schemes during the COVID-19 pandemic. Similarly to Durrani, Metzler, Michail, and Werner (2020), we identify these loans via the information on instrument-protection received and we exclude state guaranteed loans since banks' loan granting and pricing decisions for this kind of loans might not correctly reflect debtors' risks.

However, an unrestricted unbalanced panel may not be well suited for assessing the impact of the CST on banks. We therefore construct two different samples around a single shock date (October 2021, when the CST methodology was transmitted to banks). We then restrict the data to a symmetric window of ± 24 months around the shock, which yields the FULL sample. Within this window, we flag bank–borrower pairs that are observed both before and after the shock and have at least two loan observations; these pairs are classified as incumbents, and the corresponding subsample defines our INC sample.

We also construct two additional samples to study heterogeneity in treatment intensity within the set of SIs only. Using supervisory information on participation in the Module 3, we classify each SI as either participating or not participating in Module 3. On this restricted universe, we define (i) a FULL–SI sample, which coincides with the FULL sample, and (ii) an INC–SI sample, obtained by intersecting this FULL–SI sample with the incumbent pairs defined above. These samples are used when estimating the triple–difference specifications in Section 8.

After conducting all these data cleaning steps, our final dataset comprises over a million observations. Given that we exploit granular data at the bank-borrower-time level it is crucial to exhaustively control for banks and debtors characteristics to avoid endogeneity. However, since we do not have systematically granular banks and debtors information, to exhaustively control for multiple sources of unobserved heterogeneity we use fixed-effects. Following Altavilla, Boucinha, Smets, et al. (2020), we notably use firm-time fixed-effects that control for all time-varying unobservable characteristics at the firm level, including demand shifts, firm balance-sheet structure, or risk. However, to avoid perfect multicollinearity this approach requires a sample restriction, namely that firms borrow from more than one bank. Furthermore, these multiple-bank firms should also have a loan granted by a bank in the treatment group and in the control group in each period. That is because firm-time fixed-effects control for demand shocks at the firm level, thus one needs to compare how the same firm allocates its borrowing across different banks at a given point in time.[35]

In our case, around $16-20\%$[36] of our observed firms have multiple borrowing relationships in each time period.[37] Yet, single-bank firms dominate our sample, it is then

[35] We develop more extensively our econometric approach, including the rationale for fixed-effects in Section 6.

[36] $16,5\%$ for the sample FULL and $20,3\%$ for the sample INC.

[37] This figure is close to Altavilla et al. (2020) findings. Using AnaCredit, they find that the share of firms with more than one borrowing relationship in the total number of firms may vary between 10% and

problematic to discard these observations, especially as Degryse, De Jonghe, Jakovljević, Mulier, and Schepens (2019) find, on Belgium data, that single-bank firms exhibit systematically different features.[38] An alternative to firm-time fixed effect is then to follow a strategy *à la* Degryse et al. by clustering firms into industry-location-size-time bins which allows incorporating single-bank firms into the analysis of the effects of supply shocks on borrowing firms. However, it comes at the cost of imposing stronger assumption, namely a homogeneous loan demand within each cluster. Following Degryse et al.'s procedure, we check the coherence of the alternative *pseudo* demand control (see Appendix D for details).

Appendix C.2 Data Description

The kernel densities in Figure D1 confirm that the entire distribution of spreads overlaps strongly across the two samples: the incumbent restriction does not mechanically isolate a segment with systematically higher or lower pricing, but rather trims the extensive margin of short-lived relationships. This is important for the subsequent analyses: differences between FULL and INC estimates can be interpreted as differences in margins of adjustment (transitory vs continuing relationships), rather than artefacts of very different pricing distributions.

Figure D1. Credit Spread (Kernel Density)

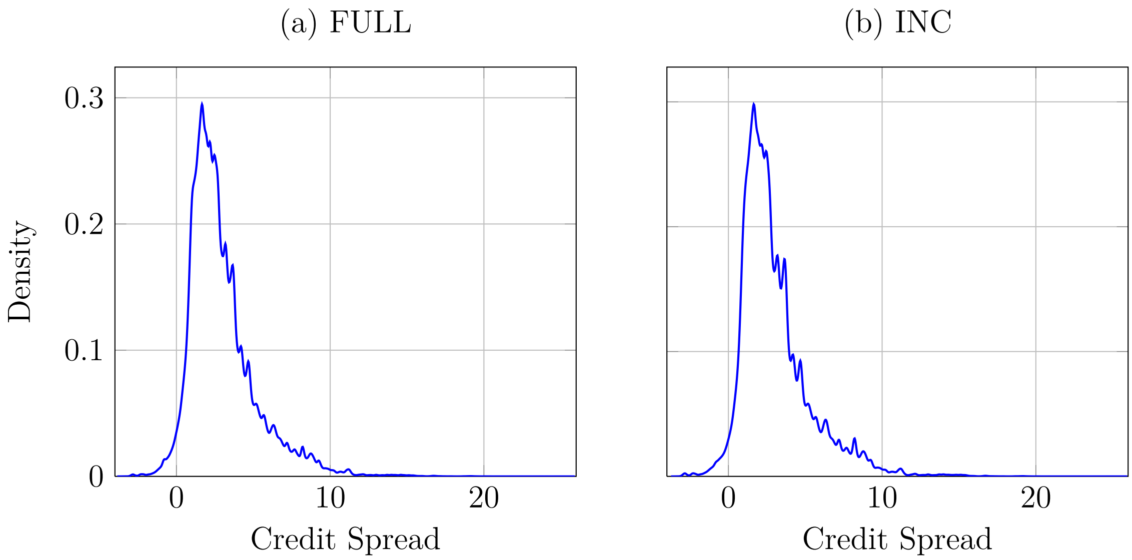

Sources: AnaCredit data and Authors' calculations.

46% depending on EU countries.

[38]They are on average younger, smaller, they hold more tangible assets and borrow smaller loan amounts.

Figure D2. log(intensity) (Kernel Density)

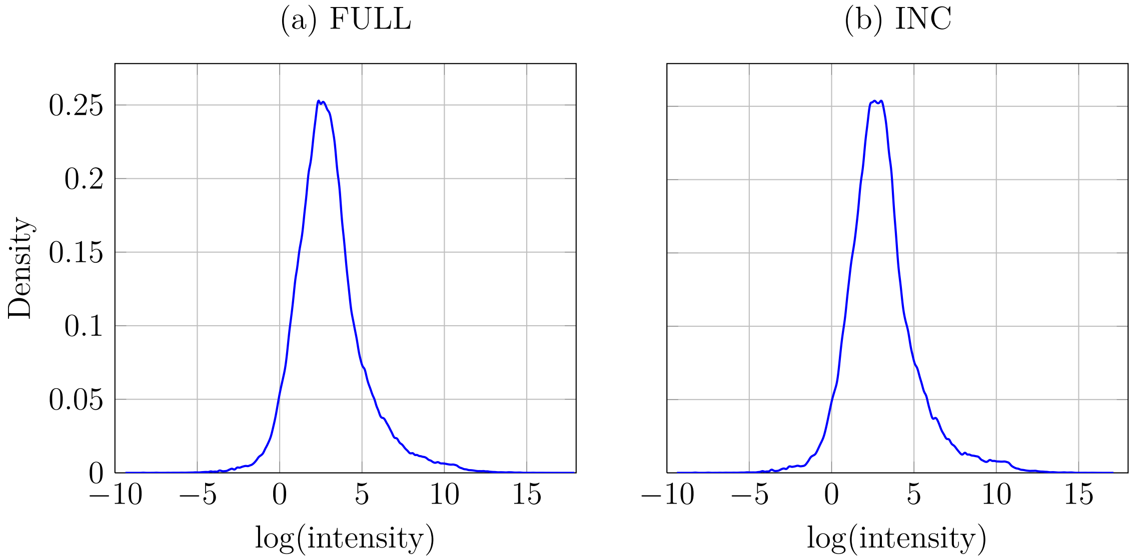

Sources: AnaCredit data and Authors' calculations.

These heavy-tailed distributions are consistent with earlier evidence on the joint distribution of bank credit and emissions across European firms (e.g. Altavilla et al., 2023): a relatively small number of highly carbon-intensive and/or large borrowers coexist with a broad mass of firms in a more moderate range. For our purposes, this has two implications. First, the level–log specification for spreads is naturally adapted to this structure: most of the identifying variation comes from the dense central part of the intensity distribution, while extreme *super-brown* exposures contribute relatively little in terms of new loans originated, even if they may be important in terms of risk concentration. Second, the close alignment of the GHG intensity and outstanding-amount distributions between FULL and INC indicates that the incumbent sample is not disproportionately composed of either very green or very brown borrowers, nor of only small or only large loans. As a result, the incumbents-based specifications can be viewed as zooming in on a stable subset of relationships that is broadly representative of the overall portfolio in terms of climate and size characteristics.

Figure D3. Firm Size Composition by Carbon Intensity Quartile

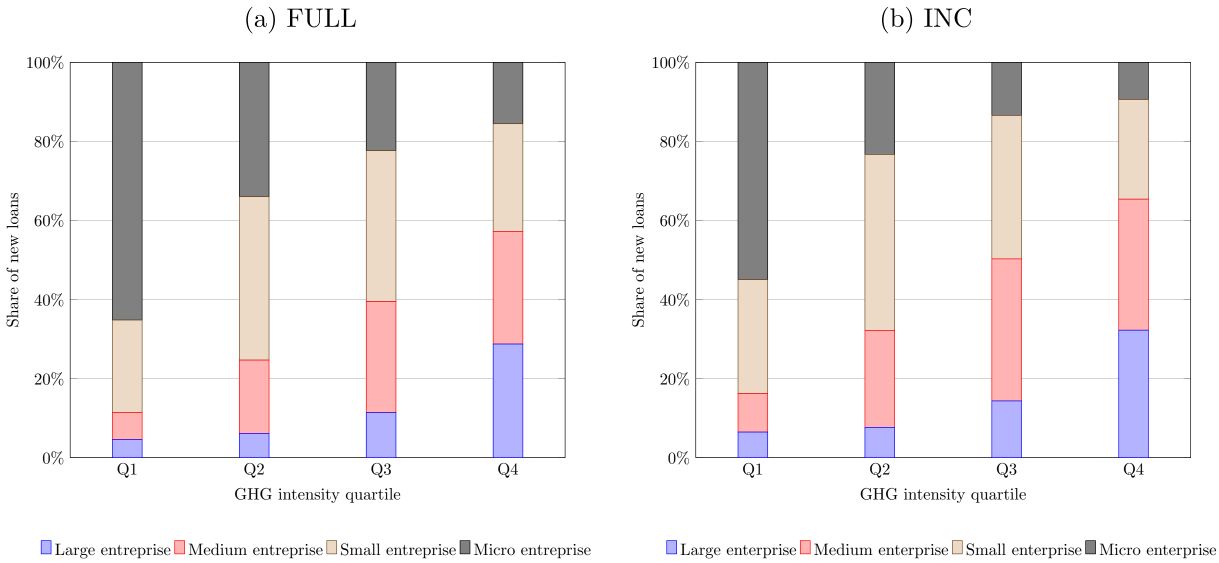

Sources: AnaCredit data and Authors' calculations.

Figure D3 further documents how the composition of borrowers varies along the carbon-intensity distribution. In both FULL and INC, moving from the first to the fourth quartile of GHG intensity is associated with a pronounced shift from micro towards larger firms. Low-intensity loans (Q1) are predominantly granted to micro and small enterprises, whereas the share of medium and especially large firms rises steadily with intensity and becomes sizeable in Q4. This monotonic pattern is very similar across the FULL and INC samples, with only a mild tilt towards larger firms in INC, reflecting the fact that long-standing bank–firm relationships are more common among bigger borrowers. Econometrically, this is reassuring: restricting to incumbents does not fundamentally alter the mapping between carbon intensity and firm size, so differences between FULL and INC estimates are unlikely to be driven by a radically different firm-size mix across intensity bins.

A remaining concern is that the incumbent sample could be subject to non-random selection and hence to survivor bias. By construction, the INC sample only includes bank–firm pairs that are observed both before and after the CST window. The probability that a given relationship survives is plausibly related to the debtor's carbon intensity and to the bank's treatment status. Indeed, when we compare unconditional distributions, we find that in the early years of the sample LSIs in the INC sample are disproportionately exposed to very carbon-intensive borrowers relative to the FULL sample, while this discrepancy largely disappears in later years.[39] As a consequence, naïve comparisons of coefficients estimated on FULL and INC would conflate intensive-margin repricing with changes in the composition of relationships along the carbon dimension.

To mitigate this compositional survivor bias, we reweight observations in the INC sample using an inverse-probability weighting (IPW) scheme. For each calendar year, we estimate a logit model for the probability of being an incumbent relationship, $\mathbb{P}\Big(\mathrm{INC}_{bdt} =$

[39]See Figures A1, A2 and A3 in Appendix A for details.

$1 \mid \log(\text{intensity}_{dt}), \text{TreatBank}_b\Big)$, allowing for a flexible third degree polynomial dependence on the debtor's log-intensity and interactions with treatment status. The resulting fitted probabilities are then inverted to construct observation-specific weights for the INC sample, which are normalised to have unit mean. This procedure leaves the within-bank, within-year relationship between spreads and carbon intensity untouched, but adjusts the relative contribution of different parts of the intensity distribution to match more closely the composition observed in the FULL sample.

In practice, the IPW correction brings the weighted mean of log-intensity in the INC sample very close to the corresponding mean in the FULL sample for each year and treatment group, while the unweighted incumbent means can differ substantially in the early years. This reweighting therefore removes a mechanical source of divergence between FULL and INC that is purely due to selective survival of high-emission relationships.

This correction is not needed when we study the heterogeneity in treatment intensity within the set of SIs only. Figures A4, A5 in Appendix A show that the unconditional distributions of the carbon intensity by treatment status are similar and do not call for any IPW correction.

Appendix D Multiple Banks vs Single Bank Test

Distinguishing loan demand from supply is a central challenge in banking research, particularly given the policy relevance of identifying supply-driven shocks (See Jakovljević, Degryse, and Ongena (2020) for a short review). Early contributions (Gan, 2007; Khwaja and Mian, 2008) showed that demand effects could be accounted for by using firm fixed effects in regressions where bank-firm exposures were observed. By comparing different banks lending to the same firm over time, one can interpret any residual variation as coming from bank-specific supply dynamics. For instance, Agarwal, Duttagupta, and Presbitero (2020) use this method to establish how reductions of commodity prices affect lending in developing countries, or more specifically, which bank characteristics lead to a stronger transmission of the global price shock through loan supply.

As richer credit registers datasets emerged, researchers extended this approach to include firm-time fixed effects, thereby controlling for time-varying demand at the firm level. This method has proven especially useful in Difference-in-Differences (DiD) settings where loan supply is hypothesized to change following an exogenous shock, allowing researchers to trace the causal impact of such shocks on credit growth. Yet it relies on firms borrowing from multiple banks; when a large share of firms maintain only one banking relationship, these strategies risk excluding a significant part of the data and potentially biasing the resulting inferences.

Recent work has sought to mitigate this problem by devising alternative proxies for time-varying demand that do not require multi-bank borrowers. Degryse et al. (2019) propose clustering firms into industry-location-size bins, arguing that firms in the same cluster share similar demand conditions, thus allowing single-bank firms to remain in the sample. Their results suggest that this industry-location-size-time approach yields estimates of supply shocks that closely resemble those obtained with firm-time fixed effects among the subset of multi-bank firms. Other enhancements include weighting procedures reflecting the importance of each firm in a bank's portfolio (Amiti and Weinstein, 2018; Tielens and Van Hove, 2017) and incorporating bank-firm fixed effects (Paligorova and Santos, 2017; Altavilla et al., 2020) to absorb unobserved match-specific heterogeneity. Although these methods improve the precision of supply shock measurement, they still hinge on observing established bank-firm relationships, overlooking the role of loan application decisions and rejections in the initial stages of credit allocation.

We follow Degryse et al.'s procedure.[40] In the first step, we estimate, in the multiple-bank firms sample, the following model:

$$y_{bdt} = \mu_{bt}^{FT} + \gamma_{dt} + \eta_{bd} + \varepsilon_{bft} \tag{13}$$

where δ_{bt} is the bank-time fixed effect, γ_{dt} is the firm-time fixed effect and η_{bd} is the bank-firm fixed effect.

We also estimate the following model using the same sample:

$$y_{bdt} = \mu_{bt}^{CST} + \lambda_{cst} + \eta_{bd} + \varepsilon_{bdt} \tag{14}$$

[40]Note that we do not follow the *tetxbook* procedure. Initially, Degryse et al. use industry×country×size×time fixed effects as a *pseudo* demand-side controls. However, in our application we cannot use such a grannular interaction because it will absorb all the variation from the numerator of our carbon intensity proxy (See Appendix D for details).

where λ_{cst} is the country×size×time fixed effects.

Based on results from Eqs. (13) and (14), we compare the supply-side shocks $\widehat{\mu}_{bt}$. We find that supply-side shocks estimates do not vary substantially. The Pearson correlation coefficient is higher than 0.80 over the entire timespan for the FULL sample, and around 0.52 for the INC sample, implying that the *pseudo* demand-side control has not led to a misspecification (See Figure C1).

As an alternative check, we regress $\widehat{\mu}_{bt}^{FT}$ from Eq. (13) on $\widehat{\mu}_{bt}^{CST}$ from Eq. (14):

$$\widehat{\mu}_{bt}^{FT} = \widehat{\mu}_{bt}^{CST} + \delta_b + \theta_t + \varepsilon_{bt}$$

where δ_b are bank fixed effects and θ_t time fixed effects.

We find that the proportion of the variance that is explained by the supply-side shocks from Eq. (14) is around 0.42 and 0.51 (See Within R^2 in Table C1.A). These results suggest that the demand-side controls are similar.

In the second step, we compare the results of Eq. (13), based on multiple-bank firms sample, and Eq. (14), based on the full sample, to the extent that potential differences in the two supply-side shocks estimates can therefore be attributed to variation in supply-side shocks faced by multiple-bank firms and single-bank firms. Figure C2 presents the month-by-month correlation between $\widehat{\mu}_{bt}^{FT}$ and $\widehat{\mu}_{bt}^{CST}$. The Pearson coefficients are lower on average than in the first step (around 0.80) but they also vary more significantly. These results corroborate Degryse et al. (2019) findings and support the assumption that excluding single-bank firms could lead to biased estimates of the strenght of supply-side shocks. The regression procedure also support this assumption (See Table C1.B).

Table C1. Comparing Supply Shocks

	A. First Step	
Outcome:	$\widehat{\mu}_{bt}^{FT}$	$\widehat{\mu}_{bt}^{FT}$
Sample:	FULL	INC
Variables		
$\widehat{\mu}_{bt}^{CST}$	0.837***	0.744***
	(0.049)	(0.059)
Fixed-effects		
Bank	Yes	Yes
Year–Month	Yes	Yes
Fit statistics		
Observations	2,144	1,751
Within R^2	0.459	0.517
	B. Second Step	
Outcome:	$\widehat{\mu}_{bt}^{FT}$	$\widehat{\mu}_{bt}^{FT}$
Sample:	FULL	INC
Variables		
$\widehat{\delta}_{bt}^{CST}$	0.756***	0.686***
	(0.079)	(0.108)
Fixed-effects		
Bank	Yes	Yes
Year–Month	Yes	Yes
Fit statistics		
Observations	2,144	1,751
Within R^2	0.256	0.262

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Panel (a) presents the regression of $\widehat{\mu}_{bt}$ from (13) on $\widehat{\mu}_{bt}$ from (14), using only the multiple-bank firms sample. Panel (b) presents the regression of $\widehat{\mu}_{bt}$ from (13), using the multiple-bank firms sample, on $\widehat{\mu}_{bt}$ from (14), using the single-bank sample.

Figure C1. Monthly correlation between supply shocks (multiple-banks sample)

(a) FULL sample

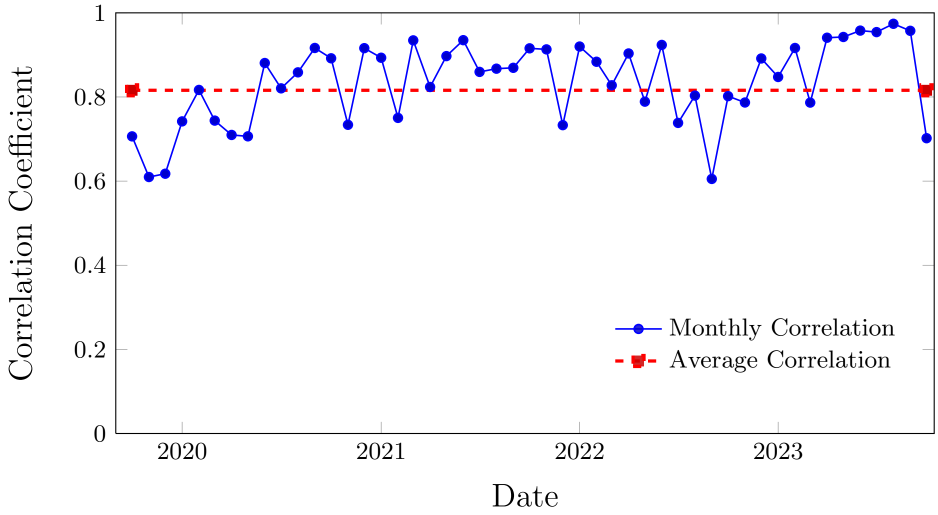

(b) INC sample

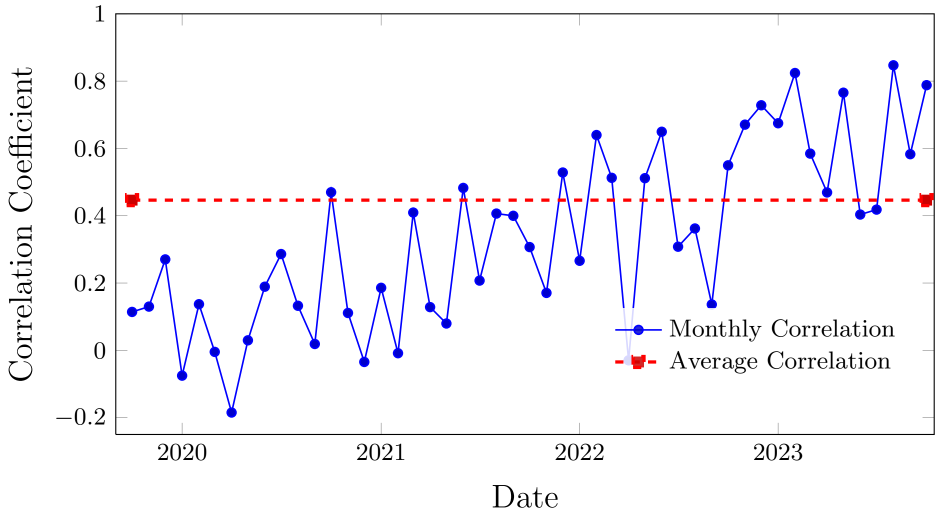

Source: AnaCredit data and Authors' calculations.
Notes: Both panels shows the monthly correlation of the $\widehat{\mu}_{bt}$ from (13) with the $\widehat{\mu}_{bt}$ from (14) using the multiple-bank firms sample with the Credit Spread as the dependant variable. Panel (a) focuses on the FULL sample while Panel (b) focuses on the INC sample.

Figure C2. Monthly correlation between supply shocks (multiple banks vs single bank)

(a) FULL sample

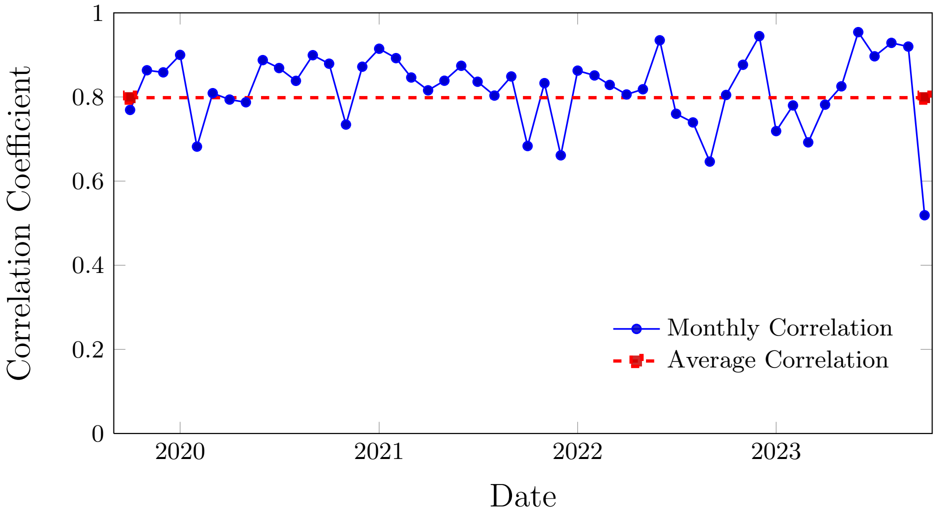

(b) INC sample

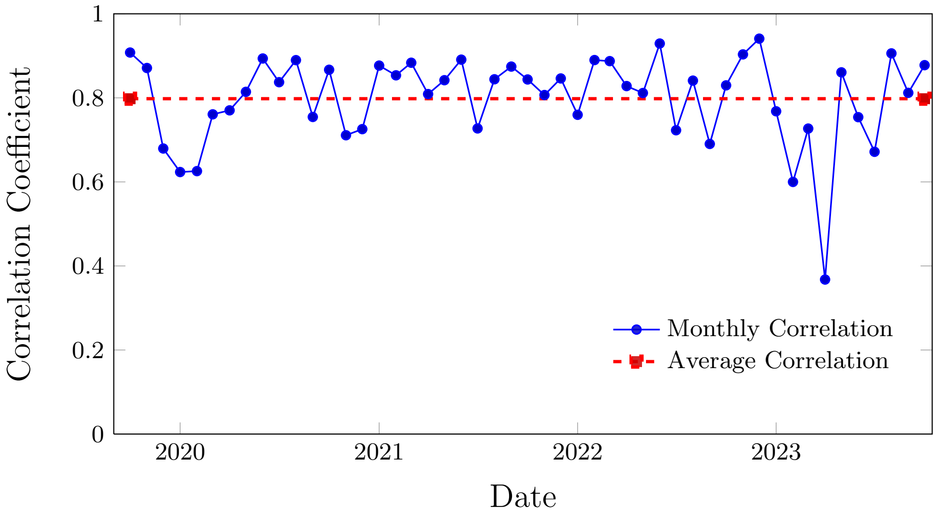

Source: AnaCredit data and Authors' calculations.
Notes: Both panels show the monthly correlation of the $\widehat{\mu}_{bt}$ from (13), using the multiple-bank firms sample, with the $\widehat{\mu}_{bt}$ from (14) using the single-bank sample with Credit Spread as the endogenous variable. Panel (a) focuses on the FULL sample while Panel (b) focuses on the INC sample..

Appendix E Proofs of Section 4

This section comprises various technical foundations. Beginning with a derivation of a social optimal intervention by the Central Bank in order to establish banks' implementation of C&E risks into their risk management.

Proof Eq. (2)

For any fixed r, the bank chooses e_b to maximize:

$$\max_{e_b} \ \Pi_b(r, e_b) = \left[e_b^{\alpha}(1 - \delta r)(1 + r) - 1 - \frac{1}{2}\beta e_b^2\right]$$

The first-order condition (FOC) w.r.t. e_b is:

$$\frac{\partial \Pi_b}{\partial e_b} = \alpha e_b^{\alpha-1}(1 - \delta r)(1 + r) - \beta e_b = 0 \quad \Rightarrow \quad \beta e_b = \alpha e_b^{\alpha-1}(1 - \delta r)(1 + r)$$

Dividing both sides by $e_b^{\alpha-1}$ and solving for e_b yields the bank's optimal monitoring effort as a function of r:

$$e_b^*(r) = \left[\frac{\alpha(1 - \delta r)(1 + r)}{\beta}\right]^{\frac{1}{2-\alpha}}$$

Now substituting $e_b^*(r)$ back into the profit $\Pi_b(r, e_b)$ and imposing the zero-profit equilibrium condition gives:

$$\Pi_b(r, e_b^*(r)) = e_b^*(r)^{\alpha}(1 - \delta r)(1 + r) - 1 - \frac{1}{2}\beta(e_b^*(r))^2 = 0$$

Define:

$$e_b^*(r) = K\left[(1 + r)(1 - \delta r)\right]^{\frac{1}{2-\alpha}}, \quad \text{where} \quad K \equiv \left(\frac{\alpha}{\beta}\right)^{\frac{1}{2-\alpha}}$$

Therefore, we have:

$$(e_b^*(r))^{\alpha} = K^{\alpha}\left[(1 + r)(1 - \delta r)\right]^{\frac{\alpha}{2-\alpha}}, \quad (e_b^*(r))^{2} = K^{2}\left[(1 + r)(1 - \delta r)\right]^{\frac{2}{2-\alpha}}$$

Define:

$$M \equiv (1 + r)(1 - \delta r)$$

Plug these back into $\Pi_b(r, e_b^*(r)) = 0$ yields:

$$
\begin{aligned}
\Pi_b(r, e_b^*(r)) &= K^\alpha M^{\frac{\alpha}{2-\alpha}} M - \frac{1}{2}\beta K^2 M^{\frac{2}{2-\alpha}} - 1 \\
&= K^\alpha M^{\frac{\alpha}{2-\alpha}+1} - \frac{1}{2}\beta K^2 M^{\frac{2}{2-\alpha}} - 1 \\
&= K^\alpha M^{\frac{2}{2-\alpha}} - \frac{1}{2}\beta K^2 M^{\frac{2}{2-\alpha}} - 1 = 0
\end{aligned}
$$

The zero-profit condition can be rewritten:

$$
\left[K^\alpha - \frac{1}{2}\beta K^2\right] M^{\frac{2}{2-\alpha}} = 1
$$

Solving for M:

$$
M^{\frac{2}{2-\alpha}} = \frac{1}{K^\alpha - \frac{1}{2}\beta K^2} \quad \Rightarrow \quad M^* = \left[\frac{1}{K^\alpha - \frac{1}{2}\beta K^2}\right]^{\frac{2-\alpha}{2}}
$$

Recall that $M \equiv (1+r)(1-\delta r)$. Hence, r^* must satisfy:

$$
1 + r^* - \delta r^* - \delta r^{*2} = M^*
$$

This is a quadratic equation in r^*, that can be solved explicitly using the quadratic formula:

$$
r^* = \frac{(1-\delta) \pm \sqrt{(1-\delta)^2 + 4\delta(1-M^*)}}{2\delta}
$$

Since $r^* \in \mathbb{R}^+$ for an economically meaningful solution, the bank selects the positive root of the quadratic equation and imposes $(1+r^*)(1-\delta r^*) > 0$ so that the success probability and expected returns are well-defined. Denoting by M^* the equilibrium value of $M(r) \equiv (1+r)(1-\delta r)$, the condition $\Pi_b(r^*) = 0$ pins down M^* and hence r^* via $M(r^*) = M^*$.

To see that r^* is finite and corresponds to an interior solution, it is useful to study how the bank's reduced-form profit varies with r. Substituting the optimal monitoring effort $e_b^*(r)$ into the profit function yields

$$
\Pi_b(r) \equiv \Pi_b(r, e_b^*(r)) = \left[K^\alpha - \frac{1}{2}\beta K^2\right] M(r)^{\frac{2}{2-\alpha}} - 1,
$$

Define

$$
L \equiv K^\alpha - \frac{1}{2}\beta K^2 \quad \Longrightarrow \quad L = K^\alpha\left(1 - \frac{\alpha}{2}\right),
$$

note that $L > 0$ whenever $0 < \alpha < 2$.[41]

[41] In our setting, α measures the sensitivity of project success to monitoring effort, while β scales the quadratic cost of effort. The condition $0 < \alpha < 2$ therefore simply requires that monitoring is productive but not too disproportionate relative to its cost. Economically, $L > 0$ ensures that, at the optimal effort choice, the bank can earn strictly positive expected profits for some range of interest rates. This is necessary for a competitive zero-profit equilibrium in which rates are bid down from a profitable level.

We rewrite the profit function as:

$$\Pi_b(r) = L\, M(r)^{\frac{2}{2-\alpha}} - 1.$$

Differentiating w.r.t. r gives

$$\frac{\partial \Pi_b(r)}{\partial r} = L\frac{2}{2-\alpha} M(r)^{\frac{2}{2-\alpha}-1}\, \frac{\partial M(r)}{\partial r}.$$

Because

$$\frac{\partial M(r)}{\partial r} = \frac{\partial}{\partial r}\Big[(1+r)(1-\delta r)\Big] = 1 - \delta - 2\delta r,$$

we obtain

$$\frac{\partial \Pi_b(r)}{\partial r} = L\frac{2}{2-\alpha} M(r)^{\frac{2}{2-\alpha}-1}(1-\delta-2\delta r).$$

We restrict attention to values of r such that $M(r) = (1+r)(1-\delta r) > 0$ and $r \geq 0$. This implies $r \in (0, 1/\delta)$ when $\delta > 0$, which ensures that the borrower's success probability remains nonnegative.[42] On this interval, $M(r) > 0$, so $M(r)^{\frac{2}{2-\alpha}-1} > 0$. Under $L > 0$, the sign of $\partial\Pi_b(r)/\partial r$ is therefore entirely determined by $(1 - \delta - 2\delta r)$.

Suppose $\delta < 1$. Then $1 - \delta > 0$ and

$$\left.\frac{\partial \Pi_b(r)}{\partial r}\right|_{r=0} = L\frac{2}{2-\alpha} M(0)^{\frac{2}{2-\alpha}-1}(1-\delta) > 0,$$

so profits initially increase with r for r close to zero. Moreover, the derivative of $M(r)$ vanishes at

$$r_T \equiv \frac{1-\delta}{2\delta},$$

which lies strictly between 0 and $1/\delta$ when $\delta \in (0, 1)$. For all $r > r_T$ we have

$$1 - \delta - 2\delta r < 0,$$

and thus

$$\frac{\partial \Pi_b(r)}{\partial r} < 0 \quad \text{for all } r \in (r_T, 1/\delta),$$

since $L > 0$ and $M(r)^{\frac{2}{2-\alpha}-1} > 0$ on this domain.

In other words, $\Pi_b(r)$ increases for small r (when the marginal gain from a higher interest rate dominates the loss in success probability), but eventually decreases for larger r as the term $(1 - \delta r)$ in the success probability shrinks. Hence there exists a local maximum $r^\dagger \in (0, 1/\delta)$ at which

$$\frac{\partial \Pi_b(r^\dagger)}{\partial r} = 0 \quad \text{and} \quad \frac{\partial \Pi_b(r)}{\partial r} < 0 \text{ for all } r > r^\dagger.$$

If the maximum value $\Pi_b(r^\dagger)$ is positive, competition among banks drives expected

[42] Recall that we assume $\delta \in (0, 1)$, so that the adverse effect of the interest rate on project success, captured by $(1 - \delta r)$, is moderate. Starting from a very low rate, a small increase in r then raises the bank's profit: the higher margin more than offsets the slight drop in success probability.

profits down to zero by slightly undercutting $r^\dagger$. The unique equilibrium rate $r^* > 0$ that satisfies $\Pi_b(r^*) = 0$ must then lie (weakly) below $r^\dagger$. Consequently, the equilibrium interest rate is finite: beyond $r^\dagger$ the slope $\partial\Pi_b(r)/\partial r$ is negative, so any further increase in r would reduce the bank's expected profit.

Economically, even if the borrower would be willing to pay more than $r^\dagger$, the bank would reduce its expected payoff by charging such a high rate, because the associated drop in success probability more than offsets the higher margin. The bank instead chooses to ration credit at or below r^*. This mechanism is the standard credit-rationing logic in Stiglitz and Weiss (1981).

Proof Eq. (4)

Under perfect competition, the bank takes r as determined in equiibirum, the bank's optimization problem upon receiving the signal τ becomes:

$$\max_{e_b} \Pi_b(r, e_b) = \left[e_b^\alpha (1-\delta r)(1-\tau)(1+r) - 1 - \frac{1}{2}\beta e_b^2 \right]$$

The first-order condition (FOC) w.r.t. e_b is:

$$\frac{\partial \Pi_b}{\partial e_b} = \alpha e_b^{\alpha-1}(1-\delta r)(1-\tau)(1+r) - \beta e_b \quad \Rightarrow \quad \beta e_b = \alpha e_b^{\alpha-1}(1-\delta r)(1-\tau)(1+r)$$

Dividing both sides by $e_b^{\alpha-1}$ and solving for e_b yields the bank's optimal monitoring effort as a function of r:

$$e_b^*(r) = \left[\frac{\alpha(1-\delta r)(1-\tau)(1+r)}{\beta} \right]^{\frac{1}{2-\alpha}}$$

Now substituting $e_b^*(r)$ back into the profit $\Pi_b(r, e_b)$ and imposing the zero-profit equilibrium condition gives:

$$\Pi_b(r, e_b^*(r)) = e_b^*(r)^\alpha (1-\delta r)(1-\tau)(1+r) - 1 - \frac{1}{2}\beta (e_b^*(r))^2 = 0$$

Define:

$$e_b^*(r,\tau) = K \left[(1+r)(1-\tau)(1-\delta r) \right]^{\frac{1}{2-\alpha}}, \quad \text{where} \quad K \equiv \left(\frac{\alpha}{\beta} \right)^{\frac{1}{2-\alpha}}$$

Therefore, we have:

$$(e_b^*(r,\tau))^\alpha = K^\alpha \left[(1+r)(1-\tau)(1-\delta r) \right]^{\frac{\alpha}{2-\alpha}}, \quad (e_b^*(r,\tau))^2 = K^2 \left[(1+r)(1-\tau)(1-\delta r) \right]^{\frac{2}{2-\alpha}}$$

Define:

$$M_\tau \equiv (1+r)(1-\tau)(1-\delta r)$$

Plug these back into $\Pi_b(r, e_b^*; \tau) = 0$ yields:

$$\begin{aligned}\Pi_b(r,\tau) \equiv \Pi_b(r, e_b^*(r,\tau)) &= K^\alpha M_\tau^{\frac{\alpha}{2-\alpha}} M_\tau - \frac{1}{2}\beta K^2 M_\tau^{\frac{2}{2-\alpha}} - 1 \\ &= K^\alpha M_\tau^{\frac{\alpha}{2-\alpha}+1} - \frac{1}{2}\beta K^2 M_\tau^{\frac{2}{2-\alpha}} - 1 \\ &= K^\alpha M_\tau^{\frac{2}{2-\alpha}} - \frac{1}{2}\beta K^2 M_\tau^{\frac{2}{2-\alpha}} - 1 = 0\end{aligned}$$

The zero-profit condition can be rewritten:

$$\left[K^\alpha - \frac{1}{2}\beta K^2\right] M_\tau^{\frac{2}{2-\alpha}} = 1$$

Solving for M_τ:

$$M_\tau^{\frac{2}{2-\alpha}} = \frac{1}{K^\alpha - \frac{1}{2}\beta K^2} \quad \Rightarrow \quad M_\tau = \left[\frac{1}{K^\alpha - \frac{1}{2}\beta K^2}\right]^{\frac{2-\alpha}{2}}$$

Recall that $M_\tau \equiv (1+r)(1-\tau)(1-\delta r)$. Hence, r^* must satisfy:

$$(1-\tau)(1+r^* - \delta r^* - \delta r^{*2}) = M_\tau \quad \Rightarrow \quad -\delta(1-\tau)r^{*2} + (1-\delta)(1-\tau)r + (1-\tau) - M_\tau = 0$$

Solving using the quadratic formula yields:

$$r_\tau^* = \frac{-(1-\delta)(1-\tau) \pm \sqrt{[(1-\delta)(1-\tau)]^2 - 4\left[-\delta(1-\tau)\right]\left[(1-\tau) - M_\tau\right]}}{2\left[-\delta(1-\tau)\right]}$$

Since $r \in \mathbb{R}^+$ for an economically meaningful solution, only the positive root ($(1+r)(1-\tau)(1-\delta r) > 0$) is relevant.

To simplify the expression for r_τ^* we perform a first-order Taylor approximation. Suppose the bank's equilibrium interest rate is denoted by r_τ^* when the climate risk is $\tau \in [0,1]$. In the absence of τ, we denote the climate-risk-free rate by $r_0^* := r_{\tau=0}^*$. A general reduced-form zero-profit condition for the bank may be written as

$$F\Big(r_\tau^*, \tau\Big) \;=\; 0,$$

where $F(\cdot,\cdot)$ is a continuously differentiable function capturing, for instance, the expected revenue net of costs. In our framework, $F(r_\tau^*, \tau)$ is factored into the following canonical form:

$$F(r_\tau^*, \tau) \;=\; (1+r)\Big(1 - \delta\, r\Big)\Big(1-\tau\Big) \;-\; M$$

We treat $F(r_\tau^*, \tau) = 0$ as an implicit function of (r, τ). Differentiating totally with respect to τ, we obtain

$$\frac{\partial F}{\partial r}(r_\tau^*, \tau)\,\frac{dr_\tau^*}{d\tau} \;+\; \frac{\partial F}{\partial \tau}(r_\tau^*, \tau) \;=\; 0$$

We shall evaluate this at $\tau = 0$ and $r = r_0$, i.e. $(r_0, 0)$ satisfies $F(r_0, 0) = 0$. Solving

$F(r^*_\tau, \tau)$ for $\left.\frac{dr^*_\tau}{d\tau}\right|_{\tau=0}$ yields

$$\left.\frac{dr^*_\tau}{d\tau}\right|_{\tau=0} \;=\; -\,\frac{\left.\frac{\partial F}{\partial \tau}\right|_{(r_0,0)}}{\left.\frac{\partial F}{\partial r}\right|_{(r_0,0)}}$$

We compute the partial derivatives $\frac{\partial F}{\partial r}$ and $\frac{\partial F}{\partial \tau}$, and we evaluate them at $(r_0, 0)$:

$$\frac{\partial F}{\partial r}(r,\tau) \;=\; (1-\tau)\,\frac{d}{dr}\Big[(1+r)\,(1-\delta r)\Big] \;=\; (1-\tau)\Big(1 \;-\; \delta \;-\; 2\delta\, r\Big)$$

At $\tau = 0$, $r = r_0$:

$$\left.\frac{\partial F}{\partial r}\right|_{(r_0,0)} = 1 - \delta - 2\delta r_0$$

$$\frac{\partial F}{\partial \tau}(r,\tau) \;=\; \frac{\partial}{\partial \tau}\Big[(1+r)\,(1-\delta r)\,(1-\tau)\Big] \;=\; -(1+r)(1-\delta r)$$

At $(r_0, 0)$, note that $(1 + r_0)(1 - \delta\, r_0) = M$ from $F(r_0, 0) = 0$. Hence

$$\left.\frac{\partial F}{\partial \tau}\right|_{(r_0,0)} \;=\; -\,(1+r_0)\,\Big(1-\delta\, r_0\Big) \;=\; -\,M$$

It follows that:

$$\left.\frac{dr^*_\tau}{d\tau}\right|_{\tau=0} \;=\; \frac{M}{1 \;-\; \delta \;-\; 2\,\delta\, r_0}$$

The first-order Taylor approximation of r^*_τ around $\tau = 0$ is thus:

$$r^*_\tau \;\approx\; r^*_0 \;+\; \left.\frac{dr^*_\tau}{d\tau}\right|_{\tau=0} \tau \;=\; r^*_0 \;+\; \frac{M}{1-\delta-2\,\delta\, r_0}\,\tau \;=\; r^* + \alpha\cdot\tau$$

Finally, we see that introducing the transition risk parameter τ into the model leads to a higher equilibrium interest rate. Also, note that r^*_τ is strictly increasing in τ.

Stochastic Dominance

The central bank's intervention provides a signal about the transition risk $\tau \in [0,1]$, which leads the bank to update its beliefs about the project payoff X. The project without transition risk yields a payoff $X > 0$ in the high state and 0 in the low state, with success probability $\theta = \theta(e_b, r)$ determined by the bank's monitoring effort and the loan rate. In the main text, we specify

$$\theta(e_b, r) = e_b^\alpha(1-\delta r),$$

so that, conditional on transition risk τ, the success probability becomes

$$\theta|\tau = \theta(e_b, r)\,(1-\tau) = e_b^\alpha(1-\delta r)(1-\tau),$$

and for any $\tau > 0$ the project is strictly less likely to succeed. The following proposition formalizes this as a first-order stochastic dominance relation between the unconditional and conditional payoff distributions.

Proposition E.1 (First-order stochastic dominance under transition risk)**.** Let X denote the project payoff at $t = 1$. Without conditioning on τ, X takes value $X > 0$ with probability $\theta \in (0, 1)$ and 0 with probability $1 - \theta$, yielding the unconditional cumulative distribution function

$$F_X(x) = \mathbb{P}(X \leq x).$$

Conditional on transition risk $\tau > 0$, the success probability is $\theta|\tau \in (0, \theta)$, and the conditional CDF is

$$F_{X|\tau}(x) = \mathbb{P}(X \leq x \mid \tau).$$

Then, for any $\tau > 0$,

$$F_X(x) \;\leq\; F_{X|\tau}(x) \quad \text{for all } x \in \mathbb{R},$$

with strict inequality on $[0, X)$. Hence, the unconditional payoff distribution X first-order stochastically dominates the conditional distribution $X|\tau$.

Proof. Given the binary structure of outcomes at $t = 1$, there are two possible states:

- High state: X euros with probability θ (unconditional) or $\theta|\tau$ (conditional),
- Low state: 0 euros with probability $1 - \theta$ (unconditional) or $1 - \theta|\tau$ (conditional).

The unconditional CDF $F_X(x)$ is therefore

$$F_X(x) = \begin{cases} 0 & \text{if } x < 0, \\ 1 - \theta & \text{if } 0 \leq x < X, \\ 1 & \text{if } x \geq X, \end{cases}$$

while the conditional CDF $F_{X|\tau}(x)$, given transition risk $\tau > 0$, is

$$F_{X|\tau}(x) = \begin{cases} 0 & \text{if } x < 0, \\ 1 - \theta|\tau & \text{if } 0 \leq x < X, \\ 1 & \text{if } x \geq X. \end{cases}$$

By construction of the main model,

$$\theta|\tau = \theta(e_b, r)(1 - \tau) < \theta(e_b, r) = \theta \quad \text{for all } \tau \in (0, 1],$$

so that $1 - \theta|\tau > 1 - \theta$.

We now compare $F_X(x)$ and $F_{X|\tau}(x)$ case by case:

1. For $x < 0$:
$$F_X(x) = 0 = F_{X|\tau}(x).$$

2. For $0 \leq x < X$:
$$F_X(x) = 1 - \theta, \qquad F_{X|\tau}(x) = 1 - \theta|\tau.$$

Since $\theta|\tau < \theta$, it follows that

$$F_X(x) = 1 - \theta \; < \; 1 - \theta|\tau = F_{X|\tau}(x).$$

3. For $x \geq X$:

$$F_X(x) = 1 = F_{X|\tau}(x).$$

Thus, for all $x \in \mathbb{R}$,

$$F_X(x) \leq F_{X|\tau}(x),$$

with strict inequality on the interval $[0, X)$ whenever $\tau > 0$. This is exactly the definition of (strict) first-order stochastic dominance: the unconditional payoff distribution X first-order stochastically dominates the conditional distribution $X|\tau$. $\square$

Appendix F Disentangling Price and Composition Effects

Our empirical specification conditions on loan origination, which means the estimated coefficient on the triple interaction in Eqs. (6) & (9) combines two distinct economic mechanisms: (i) how banks adjust spreads for loans they actually make (intensive margin), and (ii) which loans get made in the first place (extensive margin); both are modulated by the borrower's carbon intensity. The following decomposition formalizes this intuition. Note that we proceed under a static specification; the argument extends to the dynamic specification case. We formalize this setting using a two-equation system *à la* Heckman:

$$\begin{aligned}
\text{Latent price:} \quad & y^*_{bdt} = \theta\Big[\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})\Big] + \mu(x) + u_{bdt}, && (1)\\
\text{Selection:} \quad & G_{bdt} = \mathbb{1}\left\{V_{bdt} > 0\right\}, && (2)\\
& V_{bdt} = \delta\Big[\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})\Big] + \mu(x) + \eta_{bdt}.
\end{aligned}$$

Here y^*_{bdt} denotes the latent spread on a loan from bank b to debtor d in month t, V_{bdt} is the latent selection index, and $G_{bdt} \in \{0,1\}$ indicates origination. By construction, the observed spread is

$$y^{\text{obs}}_{bdt} = y^*_{bdt}\,\mathbb{1}\{G_{bdt} = 1\},$$

so that spreads are observed only if an origination occurs for the pair (bd) at t. All our empirical regressions are thus implicitly estimated conditional on $G_{bdt} = 1$. This conditioning implies that θ loads on two components: a pure price effect (intensive margin) among granted loans, and a composition effect operating through the origination probability (extensive margin). For notational economy, let $d_{bt} \equiv \mathbb{1}\{\text{TreatBank}_b = 1, \text{Post}_t = 1\}$, $z \equiv \log(x_{dt})$, and denote the fixed effects and other controls compactly by $\mu(x)$. Dropping indices for clarity, we then write

$$y^* = d\,z + \mu(x) + u, \quad d \equiv \theta d_{bt}, \qquad V = c\,z + \mu(x) + \eta, \quad c \equiv \delta d_{bt}, \qquad G = \mathbb{1}\{V > 0\}$$

In the granted sample, the object of interest is the conditional mean

$$m(z) \;\equiv\; \mathbb{E}\Big[y^* \mid G = 1,\, x,\, z\Big].$$

In this framework, the empirical slope on intensity corresponds to the derivative

$$\frac{\partial}{\partial z}\,\mathbb{E}\Big[y^* \mid G = 1,\, x, z\Big] \;=\; m'(z).$$

Define the conditional grant probability

$$\pi(z) \equiv \mathbb{P}(G = 1 \mid x, z) = \mathbb{E}\Big[\mathbb{1}\{V > 0\} \mid x, z\Big],$$

and the joint numerator

$$\nu(z) \;\equiv\; \mathbb{E}\Big[y^*\mathbb{1}\{G = 1\} \mid x, z\Big].$$

Then $m(z) = \nu(z)/\pi(z)$ and, by the quotient rule,

$$m'(z) \;=\; \frac{\nu'(z)}{\pi(z)} \;-\; m(z)\,\frac{\pi'(z)}{\pi(z)}.$$

We now assume the usual probit/logit case: the selection shock η has a smooth density f_η independent of z (e.g., standard normal in probit). Writing $a(z) \equiv -cz - \mu(x)$, we have

$$\pi(z) = \mathbb{P}(\eta > a(z)) = \int_{a(z)}^{+\infty} f_\eta(e)\,de, \quad \Rightarrow \quad \pi'(z) = c\,f_\eta\big(a(z)\big),$$

by the Leibniz rule for integrals with a moving lower bound. Likewise,

$$\nu(z) = \int_{a(z)}^{+\infty} m_s(e,z)\,f_\eta(e)\,de, \qquad m_s(e,z) \equiv \mathbb{E}\big[y^* \mid x,z,\eta=e\big] = d\,z + \mu(x) + \mathbb{E}\big[u \mid x,z,\eta=e\big].$$

Differentiating under the integral sign (Leibniz),

$$\nu'(z) = c\,f_\eta\big(a(z)\big)\,m_s\big(a(z),z\big) + d\,\pi(z) + \int_{a(z)}^{+\infty} \frac{\partial}{\partial z}\mathbb{E}\big[u \mid x,z,\eta=e\big] f_\eta(e)\,de.$$

Combining these pieces,

$$m'(z) = d + \frac{1}{\pi(z)} \int_{a(z)}^{+\infty} \frac{\partial}{\partial z}\mathbb{E}\big[u \mid x,z,\eta=e\big] f_\eta(e)\,de + \frac{c\,f_\eta\big(a(z)\big)}{\pi(z)}\Big(m_s\big(a(z),z\big) - m(z)\Big). \quad (4)$$

Note that $\pi(z)$ is a probability; the indicator notation $\mathbb{1}\{V > 0\}$ is used inside conditional expectations to emphasize that $\pi(z) = \mathbb{E}[\mathbb{1}\{V > 0\} \mid x, z]$. In our empirical application, we adopt the benchmark restriction that the internal Missing Not at Random (MNAR) term vanishes, namely

$$\frac{\partial}{\partial z}\mathbb{E}\big[u \mid x,z,\eta=e\big] = 0 \quad \text{for all } e.$$

This condition is consistent with (i) orthogonality between the pricing shock and the selection shock, $u \perp \eta$ conditional on (x, z), and (ii) mean independence of u with respect to z once x is held fixed. Economically, it means that unobserved determinants of pricing (e.g., idiosyncratic bargaining, transitory markups) are independent of unobserved drivers of the grant decision (e.g., screening noise) once fixed effects and observables are controlled for, and do not load systematically on borrower intensity z. Under this benchmark, the internal MNAR term drops out, so that the decomposition reveals that the observed slope $m'(z)$ consists of: (i) direct pricing effect (d); (ii) boundary term capturing a composition effect from selection at the approval margin. Intuitively, the more mass of potential loans sits near the cut-off, the more a small change in z tilts the composition of granted loans and, therefore, the conditional mean spread. This formalizes why our empirical strategy must separate intensive and extensive margins to identify the causal pricing effect.

Two thought experiments help fix ideas. In a cross-section at fixed (b, t) across firms with different intensities, a regression for new-loan spreads primarily reflects *between-debtor* differences among those who do obtain a loan at t. Conversely, along the time dimension for a single pair (b, d) observed around the shock, regressing y^*_{bdt} on intensity identifies a *pure price* effect because composition is held fixed. Our unbalanced panel—with one-off and repeated relationships—mechanically mixes these objects, and the channels decomposition is designed precisely to separate within-price, reallocation-among-incumbents, and entry/exit effects.

Appendix G Heterogeneous Effects of the Climate Stress Test

We now extend the baseline model in (6) by allowing for heterogeneous effects of the CST along two distinct dimensions. First, we investigate whether banks respond asymmetrically to increases versus decreases in borrowers' carbon intensity. Second, we examine whether the impact of the CST on the sensitivity of credit spreads to carbon intensity is nonlinear across the distribution of ex-ante emissions, for instance if pricing adjustments are disproportionately concentrated among the most carbon-intensive borrowers. These two exercises provide complementary tests of whether the CST induces state-dependent reallocation and repricing of credit risk along the carbon dimension.

Appendix G.1 Asymmetric Responses to Increases and Decreases in Carbon Intensity

A natural concern is that the post-shock change in the spread–emissions semi-elasticity may be state dependent: banks may react differently when a debtor's emissions intensity is increasing or decreasing. To test for such asymmetry, we exploit the month-to-month variation in debtor emissions intensity. Let $\Delta \log(x_{dt}) \equiv \log(x_{dt}) - \log(x_{d,t-1})$, we then split the intensity regressor into two mutually exclusive components:

$$\log(x_{dt})^{+} \equiv \mathbb{1}\{\Delta \log(x_{dt}) \geq 0\} \log(x_{dt}), \qquad \log(x_{dt})^{-} \equiv \mathbb{1}\{\Delta \log(x_{dt}) \leq 0\} \log(x_{dt}),$$

and replace $\log(x_{dt})$ in Eq. (6) by $\log(x_{dt})^{+}$ and $\log(x_{dt})^{-}$, allowing all interactions with Post_t and TreatBank_b to differ across the two regimes. The key coefficients of interest are the two triple interactions,

$$\theta^{+} [\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})^{+}] \quad \text{and} \quad \theta^{-} [\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})^{-}],$$

which identify a slope DiD separately for months in which the debtor's intensity is rising versus falling. As in the baseline level–log structure, the magnitudes can be read as basis-point changes in the spread for a 1% increase in x_{dt}, conditional on being in the *increasing* or *decreasing* regime.

Before turning to the estimates, the raw shares of $\mathbb{1}\{\Delta \log(x_{dt}) \geq 0\}$ indicate that the distribution of intensity changes differs across treatment status and periods: for treated banks the share of *increasing* months declines from 0.84 pre-shock to 0.78 post-shock, whereas for control banks it increases from 0.79 to 0.84. This motivates an explicit asymmetric specification rather than imposing a common post-treatment slope shift regardless of the sign of $\Delta \log(x_{dt})$.

Table H1. Asymmetric Model

Dependent Variable:	Credit Spread	
Model:	(1)	(2)
	INC	INT
Coefficients		
θ^+	-0.0815^{*}	-0.0000
	(0.0330)	(0.0103)
θ^-	-0.1201^{***}	-0.0005
	(0.0261)	(0.0123)
Controls	Yes	Yes
Fixed-effects		
Loan Type	Yes	Yes
Country-Size-Time	Yes	Yes
Bank-Time	Yes	Yes
Bank-Debtor	No	Yes
Fit statistics		
Observations	1,186,746	1,186,745
R^2	0.42423	0.80782
Within R^2	0.00381	0.00008

Clustered (Bank) standard errors in parentheses.
Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table H1 reports the results on the INC sample[43] and the within-relationship specification. In INC, we find economically meaningful and statistically significant post-treatment slope shifts in both regimes, with a clear asymmetry. Thus, the post-shock change in the spread–intensity semi-elasticity for treated banks is more negative when the debtor's intensity is *decreasing* than when it is increasing. Interpreted through the lens of transition-risk integration, this pattern is consistent with treated banks conditioning their slope response not only on the level of emissions intensity, but also on its *direction of change*, reacting more strongly when intensity improvements materialize.

As for the previous results, once we add bank×debtor fixed effects, both asymmetric triple-interaction coefficients become insignificant. These results implied channel decomposition on the incumbent sample therefore attributes none of the effect to within-relationship repricing, while essentially all of the asymmetric response loads on reallocation among incumbents. This channel is negative both when borrower intensity increases and when it decreases, implying that, in both cases, treated banks tilt origination away from more carbon-intensive incumbent borrowers. Nevertheless, the effect is significantly

[43] We focus solely on this sample to compare pre- versus post-shock variation.

more negative when intensity decreases (down vs up difference: -0.038 bp per 1% intensity), indicating an asymmetric reweighting that is stronger for decarbonizing borrowers.

Beyond documenting an asymmetric response to increases and decreases in debtor intensity, these patterns also reveal the underlying shape of the portfolio adjustment. The fact that the reallocation channel is negative in both regimes, and more negative when carbon intensity is decreasing, suggests that the reshaping of treated banks' incumbent portfolios is not simply a passive reflection of firms' own *greening* path. Instead, banks appear to use information about both the level and the trajectory of emissions to actively reweight their exposures within the existing client base, systematically tilting credit away from relatively more carbon-intensive incumbents even in periods when aggregate intensity is increasing. In this interpretation, improvements in borrowers' carbon intensity do not merely reduce transition risk mechanically; they also seem to trigger a more pronounced reallocation of lending toward the cleaner end of the incumbent distribution. The CST thus appears to operate less as a one-off price adjustment mechanism and more as a catalyst for continuous, forward-looking portfolio reshaping, with banks exploiting directional changes in intensity to accelerate the rebalancing of their books toward lower-carbon counterparties.

Table H2. Channel Decomposition (up vs. down)

Channels	*Estimands*
Intensive margin (up)	0.0000
	(0.0100)
Intensive margin (down)	-0.0005
	(0.0121)
Difference intensive (down $-$ up)	-0.0004
	(0.0104)
Reallocation among incumbents (up)	-0.0815^{***}
	(0.0286)
Reallocation among incumbents (down)	-0.1196^{***}
	(0.0220)
Difference reallocation (down $-$ up)	-0.0381^{**}
	(0.0182)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates and bootstrap distributions are obtained from the asymmetric baseline specification (see corresponding table in the main text). Standard errors (in parentheses) and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights (see Appendix F for details). "Up" and "down" refer to positive vs. negative movements in log(intensity).

Appendix G.2 Nonlinear Effects across the Distribution of Carbon Intensity

We also examine whether the CST shock generates non-linearities in the spread–carbon relationship. Instead of interacting the treatment indicators with the continuous variable $\log(x_{dt})$, we replace it with a high–intensity dummy constructed from ex-ante intensity quantiles. Specifically, let $q \in \{0.10, 0.25, 0.50, 0.75, 0.90\}$, and denote by c_q the q-th quantile of the pre-CST distribution of $\log(x_{dt})$. For each q, we define $I_{dt}^{(q)} \equiv \mathbb{1}\{\log(x_{dt}) \geq c_q\}$, and we estimate the following specification separately for each quantile q:

$$
\begin{aligned}
y_{bdt} \;=\; & \beta_1^{(q)} I_{dt}^{(q)} + \beta_2^{(q)} \Big[\text{Post}_t \times I_{dt}^{(q)}\Big] + \beta_3^{(q)} \Big[\text{TreatBank}_b \times I_{dt}^{(q)}\Big] \\
& + \theta^{(q)} \Big[\text{Post}_t \times \text{TreatBank}_b \times I_{dt}^{(q)}\Big] + \lambda_{cst} + \mu_{bt} + \phi_l + \varepsilon_{bdt}.
\end{aligned} \tag{15}
$$

In this setup, $\theta^{(q)}$ measures how the post-CST slope differential between treated and control banks changes for loans to borrowers whose carbon intensity lies above the q-th quantile, relative to otherwise similar low-intensity loans.[44] Table H3 extends the static channel decomposition from Eq. (8) to the quantile-based Eq. (15). The results suggest that the CST-induced adjustment in carbon-risk pricing operates primarily through extensive margins, with marked heterogeneity across the emissions distribution.

Table H3. Channel Decomposition by Quantiles

	Intensity quantile q				
Channel	0.10	0.25	0.50	0.75	0.90
Intensive margin	−0.021	0.041	0.085*	−0.186***	−0.371***
	(0.083)	(0.066)	(0.046)	(0.068)	(0.085)
Reallocation among incumbents	−0.162	−0.180***	−0.194***	−0.226**	−0.242
	(0.114)	(0.027)	(0.038)	(0.117)	(0.221)
Entry/Exit pricing	0.237	0.313***	0.360***	0.456***	0.432**
	(0.150)	(0.082)	(0.087)	(0.174)	(0.235)
Entry/Exit composition	0.002	−0.051	−0.117***	−0.053**	0.033
	(0.036)	(0.037)	(0.032)	(0.029)	(0.033)

Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Notes: Point estimates are obtained from Table A2. Standard errors (in parentheses) and p-values are computed from $B = 499$ wild bootstrap replications with Rademacher weights, clustered at the bank level (see Appendix F for details).

At low and intermediate thresholds ($q = 0.10$ and $q = 0.25$), the intensive margin contribution is small and insignificant, while the reallocation and entry/exit terms display

[44] Table A2 presents the results of Eq. (15) for each sample.

sizeable but offsetting effects. At the median threshold ($q = 0.50$), the intensive margin becomes modestly positive and borderline significant, suggesting some post-CST steepening of the spread–carbon intensity slope within existing relationships for borrowers above the median intensity. However, for higher thresholds ($q = 0.75$ and $q = 0.90$) the intensive component turns strongly negative and highly significant, indicating that treated banks do not systematically increase, and may even attenuate, the within-relationship sensitivity of spreads to carbon intensity for the most carbon-intensive clients. This pattern is consistent with the static decomposition, where the intensive margin was already found to be limited on average, and suggests that relationship frictions or strategic considerations may dampen repricing for the most exposed borrowers.

Conversely, the extensive margins display both larger magnitudes and clearer monotonicity in q. The reallocation among incumbents channel is negative at all quantiles and statistically significant from $q = 0.25$ to $q = 0.75$, with estimates around -0.18 to -0.23. This points to a systematic post-CST shift in treated banks' origination towards less carbon-intensive incumbents relative to more intensive ones, particularly once the threshold moves above the lower part of the distribution. The entry/exit pricing channel is positive and highly significant from $q = 0.25$ onwards, and its magnitude increases with q, reaching about 0.46 at the 75-th percentile and remaining large at $q = 0.90$. Conditional on composition, treated banks thus apply markedly more carbon-sensitive pricing to transitory relationships than would be implied by incumbent pricing alone, especially for borrowers in the upper tail of the emissions distribution. Finally, the entry/exit composition term is small and insignificant at the 10-th and 25-th quantiles, but becomes significantly negative at $q = 0.50$ and $q = 0.75$, consistent with a *green* composition effect whereby the pattern of entry and exit reallocates credit toward relatively cleaner borrowers among those above the median intensity.

Overall, the quantile decomposition refines the picture obtained from the continuous-slope and static analyses. On average, the intensive margin is close to zero, but this masks substantial within-relationship adjustments concentrated among the most carbon-intensive exposures. At the same time, extensive-margin forces—reallocation across incumbents and more carbon-sensitive pricing of transitory relationships—remain the dominant and more systematically monotone channels through which the CST reshapes the distribution of credit along the carbon dimension.

These patterns are also consistent with an interpretation in which treated banks differentiate sharply between long-standing and more transitory borrower relationships when pricing transition risk. Among repeat borrowers, the negative intensive-margin contributions at the upper quantiles, together with the systematically negative reallocation terms, suggest that banks both reduce the share of new lending to the most carbon-intensive incumbents and, within the subset of high-intensity relationships they maintain, do not impose especially steep spread premia relative to the control group. One plausible mechanism is enhanced screening among these high-emitting incumbents: treated banks may continue relationships primarily with firms that credibly signal transition plans or emissions-reduction commitments, or for which bank funding is explicitly tied to decarbonisation investments. Under this reading, the comparatively mild within-relationship penalisation of high-intensity incumbents is consistent with the idea that relationship lenders condition not only on current emissions levels but also on forward-looking transition prospects, in line with recent evidence on the interaction between climate risk and relationship lending (see, for example, Fuchs et al., 2024).

By contrast, the sizeable and increasing entry/exit pricing effects at higher quantiles indicate that, for more transitory relationships where information about firms' transition trajectories is likely thinner, treated banks price carbon risk much more aggressively. Together, the quantile-specific intensive and extensive components therefore point to a screening-based mechanism: treated banks appear to protect, and relatively soften the pricing of, those high-emission borrowers for whom they acquire rich information and credible transition signals, while imposing steeper carbon premia on high-intensity borrowers with whom they have weaker informational ties.

Appendix H Channels Decomposition

Appendix H.1 Algebraic Decomposition within the Incumbent Sample

This subsection provides a formal derivation of the relationship between the three coefficients $\widehat{\theta}^F$, $\widehat{\theta}^R$, and $\widehat{\theta}^I$ used in the channels decomposition in Section 6.3, focusing on the two incumbent-sample coefficients $\widehat{\theta}^R$ (INC) and $\widehat{\theta}^I$ (INT). We show that $\widehat{\theta}^I$ is the "pure within-relationship" slope, whereas $\widehat{\theta}^R$ combines within- and between-relationship variation in a transparent way. In particular, we show that the "reallocation among incumbents" term $\widehat{\theta}^R - \widehat{\theta}^I$ is entirely driven by between-relationship variation among incumbent pairs.

For the purposes of this subsection, it is convenient to abstract from time dynamics and from the extensive margin, and to focus exclusively on the *incumbent* sample (INC). We collect all incumbent observations (b, d, t) in a vector of residualized outcomes y and a vector of residualized regressors Z, where Z is the triple interaction

$$Z_{bdt} \equiv \text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt}),$$

and both y and Z have been partialled out of the fixed effects λ_{cst}, μ_{bt} and ϕ_l that are common to both the INC and INT specifications. Formally, let A denote the matrix of these common fixed effects, and define the "residual maker"

$$M_A \equiv I - A(A'A)^{-1}A'.$$

We then work with the transformed variables

$$\tilde{y} \equiv M_A y, \qquad \tilde{Z} \equiv M_A Z.$$

In this notation, the INC regression (without relationship fixed effects) can be written as

$$\widehat{\theta}^R = \frac{\tilde{Z}'\tilde{y}}{\tilde{Z}'\tilde{Z}}. \tag{16}$$

In the INT specification, we add bank×debtor fixed effects, collected in a matrix D of relationship dummies. Let

$$M_D \equiv I - D(D'D)^{-1}D'$$

be the corresponding residual maker. The INT regression is equivalent (by the Frisch–Waugh–Lovell theorem) to regressing $\tilde{y}$ on $\tilde{Z}$ after projecting both variables orthogonally to D, i.e.,

$$\widehat{\theta}^I = \frac{(M_D\tilde{Z})'(M_D\tilde{y})}{(M_D\tilde{Z})'(M_D\tilde{Z})}.$$

We therefore define the *within-relationship* components:

$$Z^W \equiv M_D\tilde{Z}, \qquad y^W \equiv M_D\tilde{y},$$

so that

$$\widehat{\theta}^I = \frac{(Z^W)'y^W}{(Z^W)'Z^W}. \tag{17}$$

We now decompose the residualized regressor and outcome into their within- and between-relationship components. Define the projection onto the space spanned by relationship dummies:

$$P_D \equiv D(D'D)^{-1}D' = I - M_D,$$

and set

$$Z^B \equiv P_D\tilde{Z} = (I - M_D)\tilde{Z}, \qquad y^B \equiv P_D\tilde{y} = (I - M_D)\tilde{y}.$$

By construction,

$$\tilde{Z} = Z^W + Z^B, \qquad \tilde{y} = y^W + y^B.$$

The following lemma records a standard orthogonality property of these components.

Lemma H.1 (Within–between orthogonality)**.** *The within- and between-relationship components are orthogonal:*

$$(Z^W)'Z^B = 0, \qquad (y^W)'y^B = 0.$$

Proof. By construction, M_D and P_D are symmetric idempotent matrices: $M_D' = M_D$, $M_D^2 = M_D$, $P_D' = P_D$, $P_D^2 = P_D$, and $M_D P_D = 0$. We have

$$Z^W = M_D\tilde{Z}, \qquad Z^B = P_D\tilde{Z},$$

hence

$$(Z^W)'Z^B = \tilde{Z}'M_D'P_D\tilde{Z} = \tilde{Z}'M_D P_D\tilde{Z}.$$

But $P_D = I - M_D$ and $M_D^2 = M_D$, so

$$M_D P_D = M_D(I - M_D) = M_D - M_D^2 = M_D - M_D = 0.$$

Therefore $(Z^W)'Z^B = \tilde{Z}'(0)\tilde{Z} = 0$. The same argument applies to y, since $y^W = M_D\tilde{y}$ and $y^B = P_D\tilde{y}$, yielding $(y^W)'y^B = 0$. □

We can now express $\widehat{\theta}^R$ in terms of within and between components.

Proposition H.2 (Decomposition of $\widehat{\theta}^R$ into within and between slopes)**.** Let $\widehat{\theta}^R$ and $\widehat{\theta}^I$ be defined as in (16)–(17), and define the "pure between" slope

$$\theta^B \equiv \frac{(Z^B)'y^B}{(Z^B)'Z^B}$$

whenever $(Z^B)'Z^B > 0$. Then:

1. $\widehat{\theta}^I$ coincides with the within-relationship slope,

$$\widehat{\theta}^I = \frac{(Z^W)'y^W}{(Z^W)'Z^W}.$$

2. $\widehat{\theta}^R$ is a convex combination of the within- and between-relationship slopes:

$$\widehat{\theta}^R = \omega\,\widehat{\theta}^I + (1-\omega)\,\theta^B,$$

where

$$\omega \equiv \frac{(Z^W)'Z^W}{(Z^W)'Z^W + (Z^B)'Z^B} \in [0,1].$$

3. The difference between the two incumbent-sample coefficients satisfies

$$\widehat{\theta}^R - \widehat{\theta}^I = \frac{(Z^B)'Z^B}{(Z^W)'Z^W + (Z^B)'Z^B}\left(\theta^B - \widehat{\theta}^I\right).$$

Proof. Part (i) is just (17). For (ii), start from the definition (16) and use the decomposition $\tilde{Z} = Z^W + Z^B$, $\tilde{y} = y^W + y^B$:

$$\widehat{\theta}^R = \frac{\tilde{Z}'\tilde{y}}{\tilde{Z}'\tilde{Z}} = \frac{(Z^W + Z^B)'(y^W + y^B)}{(Z^W + Z^B)'(Z^W + Z^B)}.$$

Expand the numerator and denominator:

$$\tilde{Z}'\tilde{y} = (Z^W)'y^W + (Z^W)'y^B + (Z^B)'y^W + (Z^B)'y^B,$$

$$\tilde{Z}'\tilde{Z} = (Z^W)'Z^W + (Z^W)'Z^B + (Z^B)'Z^W + (Z^B)'Z^B.$$

By Lemma H.1, we have $(Z^W)'Z^B = (Z^B)'Z^W = 0$. Moreover, y^W is orthogonal to y^B and both are linear combinations of the columns of $\tilde{y}$; while this orthogonality does not imply that $(Z^W)'y^B$ and $(Z^B)'y^W$ vanish individually in general, we can proceed by writing the numerator in terms of the within and between regressions.

Define the within-relationship slope

$$\theta^W \equiv \frac{(Z^W)'y^W}{(Z^W)'Z^W} = \widehat{\theta}^I,$$

and the between-relationship slope

$$\theta^B \equiv \frac{(Z^B)'y^B}{(Z^B)'Z^B}.$$

We can always decompose y^W and y^B into the part explained by Z^W or Z^B and a residual. Specifically, there exist residual vectors r^W and r^B such that

$$y^W = \theta^W Z^W + r^W, \qquad (Z^W)'r^W = 0,$$

$$y^B = \theta^B Z^B + r^B, \qquad (Z^B)'r^B = 0,$$

because the OLS slopes θ^W and θ^B are linear projections of y^W on Z^W and of y^B on Z^B, respectively.

Substituting into the numerator:

$$\begin{aligned}\tilde{Z}'\tilde{y} &= (Z^W)'y^W + (Z^B)'y^B + (Z^W)'y^B + (Z^B)'y^W \\ &= (Z^W)'(\theta^W Z^W + r^W) + (Z^B)'(\theta^B Z^B + r^B) + (Z^W)'y^B + (Z^B)'y^W \\ &= \theta^W (Z^W)'Z^W + \theta^B (Z^B)'Z^B + (Z^W)'r^W + (Z^B)'r^B + (Z^W)'y^B + (Z^B)'y^W.\end{aligned}$$

By construction of r^W and r^B, we have $(Z^W)'r^W = 0$ and $(Z^B)'r^B = 0$. Moreover, both y^W and y^B lie in the column space of D plus the column space of M_D, while Z^W and Z^B decompose similarly; using $M_D P_D = 0$ and $P_D M_D = 0$, one can verify that the cross terms $(Z^W)'y^B$ and $(Z^B)'y^W$ vanish.[45] Thus

$$\tilde{Z}'\tilde{y} = \theta^W (Z^W)'Z^W + \theta^B (Z^B)'Z^B.$$

For the denominator, Lemma H.1 yields

$$\tilde{Z}'\tilde{Z} = (Z^W)'Z^W + (Z^B)'Z^B.$$

Combining these expressions, we obtain

$$\widehat{\theta}^R = \frac{\tilde{Z}'\tilde{y}}{\tilde{Z}'\tilde{Z}} = \frac{\theta^W (Z^W)'Z^W + \theta^B (Z^B)'Z^B}{(Z^W)'Z^W + (Z^B)'Z^B}.$$

Define

$$\omega \equiv \frac{(Z^W)'Z^W}{(Z^W)'Z^W + (Z^B)'Z^B} \in [0,1].$$

Then we can rewrite

$$\widehat{\theta}^R = \omega\,\theta^W + (1-\omega)\,\theta^B.$$

Finally, since $\theta^W = \widehat{\theta}^I$, we obtain part (ii). Part (iii) follows immediately:

$$\begin{aligned}\widehat{\theta}^R - \widehat{\theta}^I &= \omega\,\widehat{\theta}^I + (1-\omega)\,\theta^B - \widehat{\theta}^I \\ &= (1-\omega)\,\theta^B - (1-\omega)\,\widehat{\theta}^I \\ &= (1-\omega)\left(\theta^B - \widehat{\theta}^I\right),\end{aligned}$$

with $1-\omega = (Z^B)'Z^B/[(Z^W)'Z^W + (Z^B)'Z^B]$. □

Proposition H.2 shows that the coefficient $\widehat{\theta}^I$ is the within-relationship slope, whereas $\widehat{\theta}^R$ is a variance-weighted average of the within- and between-relationship slopes. Consequently, the "reallocation among incumbents" term

$$\widehat{\theta}^R - \widehat{\theta}^I$$

used in Eq. (7) is entirely driven by between-relationship variation among incumbents: it vanishes if and only if the pure between-relationship slope θ^B coincides with the within-relationship slope $\widehat{\theta}^I$, and its magnitude is scaled by the share of total variation in Z

[45] Formally, $y^B = P_D\tilde{y}$ and $Z^W = M_D\tilde{Z}$, so $(Z^W)'y^B = \tilde{Z}'M_D'P_D\tilde{y} = \tilde{Z}'M_DP_D\tilde{y} = 0$ because $M_DP_D = 0$. Similarly, $(Z^B)'y^W = \tilde{Z}'P_D'M_D\tilde{y} = \tilde{Z}'P_DM_D\tilde{y} = 0$.

that comes from between-relationship differences. In this precise sense, $\widehat{\theta}^R - \widehat{\theta}^I$ captures reallocation of credit across incumbent relationships within bank×time cells, over and above within-relationship repricing.

Appendix H.2 Decomposition of the Entry/Exit Margin

This subsection formalizes the "entry/exit" component of the slope DiD, i.e. the difference between the FULL and INC coefficients, $\widehat{\theta}^F - \widehat{\theta}^R$. We show that, after residualizing out the common fixed effects, the FULL-sample slope is a variance-weighted average of the slopes in the incumbent and transitory (entering or exiting) subsamples. As a result, the entry/exit margin is entirely driven by the difference in slopes between transitory and incumbent relationships, scaled by the relative contribution of transitory relationships to the variance of the slope regressor.

Setup. As in Appendix Appendix H.1, let Z_{bdt} denote the triple interaction

$$Z_{bdt} \equiv \text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt}),$$

and y_{bdt} denote the spread. Let A denote the matrix of all fixed effects and controls that are common to the FULL and INC specifications (e.g. λ_{cst}, μ_{bt}, and ϕ_l), and define the residualized variables

$$\tilde{y} \equiv M_A y, \qquad \tilde{Z} \equiv M_A Z, \qquad M_A \equiv I - A(A'A)^{-1}A'.$$

All regressions in this subsection are run on $(\tilde{y}, \tilde{Z})$, so we omit the tildes for notational simplicity and write (y, Z).

We focus on the FULL sample, which contains all granted loans in the time window. Partition this sample into two disjoint groups:

- $g = I$ (incumbents): bank–debtor pairs that are present both before and after the shock,
- $g = E$ (transitory): bank–debtor pairs that either enter or exit within the window (entrants and exiters).

Let $\mathcal{I}_I$ and $\mathcal{I}_E$ denote the index sets for incumbents and transitory relationships, respectively, so that the FULL sample is $\mathcal{I}_F = \mathcal{I}_I \cup \mathcal{I}_E$ with $\mathcal{I}_I \cap \mathcal{I}_E = \emptyset$.

We collect the observations (Z_i, y_i) for $i \in \mathcal{I}_I$ into vectors (Z_I, y_I) and those for $i \in \mathcal{I}_E$ into vectors (Z_E, y_E). Stacking these, the FULL-sample regressor and outcome vectors can be written as

$$Z = \begin{pmatrix} Z_I \\ Z_E \end{pmatrix}, \qquad y = \begin{pmatrix} y_I \\ y_E \end{pmatrix}.$$

Group-specific slopes. Define the group-specific OLS slopes:

$$\theta_I \equiv \frac{Z_I' y_I}{Z_I' Z_I}, \qquad \theta_E \equiv \frac{Z_E' y_E}{Z_E' Z_E},$$

whenever $Z_I'Z_I > 0$ and $Z_E'Z_E > 0$. These are the linear projection slopes of y on Z when the regression is run separately in each group. By definition, the INC coefficient $\widehat{\theta}^R$ is the OLS slope computed on the incumbent subsample only:

$$\widehat{\theta}^R = \frac{Z_I'y_I}{Z_I'Z_I} = \theta_I.$$

The FULL-sample OLS slope, denoted $\widehat{\theta}^F$, is

$$\widehat{\theta}^F = \frac{Z'y}{Z'Z} = \frac{Z_I'y_I + Z_E'y_E}{Z_I'Z_I + Z_E'Z_E}. \tag{18}$$

The following proposition provides a convenient decomposition of $\widehat{\theta}^F$.

Proposition H.3 (Mixture decomposition of the FULL-sample slope)**.** Let $\widehat{\theta}^F$ and $\widehat{\theta}^R$ be defined as in (18), and define the group-specific slopes θ_I, θ_E as above. Then:

1. The FULL-sample slope is a variance-weighted average of the group-specific slopes:

   $$\widehat{\theta}^F = \omega_I\,\theta_I + (1-\omega_I)\,\theta_E,$$

   where

   $$\omega_I \equiv \frac{Z_I'Z_I}{Z_I'Z_I + Z_E'Z_E} \in [0,1].$$

2. The difference between the FULL and INC slopes (the entry/exit margin) can be written as

   $$\widehat{\theta}^F - \widehat{\theta}^R = (1-\omega_I)\Big(\theta_E - \theta_I\Big).$$

Proof. We start from the definition (18):

$$\widehat{\theta}^F = \frac{Z_I'y_I + Z_E'y_E}{Z_I'Z_I + Z_E'Z_E}.$$

By construction of the group-specific slopes,

$$\theta_I = \frac{Z_I'y_I}{Z_I'Z_I}, \qquad \theta_E = \frac{Z_E'y_E}{Z_E'Z_E},$$

so that we can write

$$Z_I'y_I = \theta_I\,Z_I'Z_I, \qquad Z_E'y_E = \theta_E\,Z_E'Z_E.$$

Substituting back into the numerator of $\widehat{\theta}^F$ yields

$$\widehat{\theta}^F = \frac{\theta_I\,Z_I'Z_I + \theta_E\,Z_E'Z_E}{Z_I'Z_I + Z_E'Z_E}.$$

Define

$$\omega_I \equiv \frac{Z_I'Z_I}{Z_I'Z_I + Z_E'Z_E}, \qquad 1-\omega_I = \frac{Z_E'Z_E}{Z_I'Z_I + Z_E'Z_E}.$$

Then we can rewrite

$$\widehat{\theta}^F = \omega_I \, \theta_I + (1 - \omega_I) \, \theta_E,$$

which establishes part (i).

For part (ii), recall that $\widehat{\theta}^R = \theta_I$ by definition (INC sample is precisely the incumbent subsample). Therefore,

$$\begin{aligned} \widehat{\theta}^F - \widehat{\theta}^R &= \Big[\omega_I \, \theta_I + (1 - \omega_I) \, \theta_E\Big] - \theta_I \\ &= (1 - \omega_I) \, \theta_E - (1 - \omega_I) \, \theta_I \\ &= (1 - \omega_I) \, \Big(\theta_E - \theta_I\Big). \end{aligned}$$

This completes the proof. □

Proposition H.3 implies that the entry/exit term $\widehat{\theta}^F - \widehat{\theta}^R$ is entirely driven by the difference between the slope for transitory relationships (θ_E) and the slope for incumbents (θ_I), scaled by the relative contribution of transitory relationships to the variance of the slope regressor Z in the FULL sample.

In population notation, we can write the analogue of Proposition H.3 as

$$\theta^F = \omega_I \, \theta^I + (1 - \omega_I) \, \theta^E, \qquad \theta^F - \theta^R = (1 - \omega_I) \, \Big(\theta^E - \theta^I\Big),$$

where θ^F is the population FULL-sample slope, θ^R is the population incumbent slope, and θ^I, θ^E are the group-specific population slopes. The sample estimators $\widehat{\theta}^F$, $\widehat{\theta}^R$, $\widehat{\theta}_I$, and $\widehat{\theta}_E$ are consistent for their population counterparts under standard regularity conditions. In particular, the sign of the entry/exit term reflects the sign of $\theta^E - \theta^I$ whenever the variance share $1 - \omega_I$ is non-negligible.

Appendix H.3 Sub-Decomposition of the Entry/Exit Margin into Pricing and Composition

We now refine the entry/exit margin $\widehat{\theta}^F - \widehat{\theta}^R$ by splitting it into a *pricing* component and a *composition* component, as in Eq. (8). This subsection provides a detailed formal justification for this sub-decomposition, using the inverse-probability weighting (IPW) scheme that maps the INC sample to the FULL-sample composition in observables.

Weighted least squares representation. For any sample k (FULL, INC, or INC-IPW), possibly with observation-specific weights $w_i^k > 0$, the OLS (or WLS) slope on Z in a regression of y on Z (assuming a zero intercept after demeaning) can be written as

$$\widehat{\theta}^k = \frac{\sum_{i \in \mathcal{I}_k} w_i^k Z_i y_i}{\sum_{i \in \mathcal{I}_k} w_i^k Z_i^2}, \tag{19}$$

with the convention $w_i^k \equiv 1$ for unweighted OLS and where $\mathcal{I}_k$ denotes the index set of observations in sample k.

We consider three samples:

- **FULL sample** ($k = F$): all granted loans; weights $w_i^F \equiv 1$.

- **INC sample** ($k = R$): only incumbents; weights $w_i^R \equiv 1$ for $i \in \mathcal{I}_I$ and no observations outside $\mathcal{I}_I$.
- **INC-IPW sample** ($k = R, w$): same incumbents as in INC, but each observation $i \in \mathcal{I}_I$ is assigned a positive weight $w_i^{R,w}$ constructed as

$$w_i^{R,w} = \frac{1}{\widehat{\mathbb{P}}(\text{INC}_i = 1 \mid Z_i, \text{Treat}_i)},$$

normalized to have mean one among incumbents. These weights are estimated separately within each year-by-treatment cell.

Under standard regularity conditions, the estimators $\widehat{\theta}^k$ converge to population slopes of the form

$$\theta^k = \frac{\mathbb{E}_k[Zy]}{\mathbb{E}_k[Z^2]}, \tag{20}$$

where $\mathbb{E}_k$ denotes an expectation under a corresponding population measure μ_k that incorporates both the sampling scheme and the weighting. For concision, we use the population notation in what follows, understanding that all expressions have sample analogues given by (19).

Pricing functions and distributions of Z. For each measure μ_k, let $F_{Z,k}$ denote the marginal distribution of Z and define the conditional mean function

$$m_k(z) \equiv \mathbb{E}_k[y \mid Z = z].$$

Writing $y = m_k(Z) + \varepsilon_k$ with $\mathbb{E}_k[\varepsilon_k \mid Z] = 0$ and assuming that Z is centered under μ_k (as ensured by the inclusion of fixed effects and within-cell demeaning), we can express the slope θ^k as

$$\theta^k = \frac{\mathbb{E}_k[Zy]}{\mathbb{E}_k[Z^2]} = \frac{\displaystyle\int z\, m_k(z)\, dF_{Z,k}(z)}{\displaystyle\int z^2\, dF_{Z,k}(z)}. \tag{21}$$

Hence each θ^k depends on two components:

- the *pricing function* $m_k(z)$, and
- the *distribution of* Z given by $F_{Z,k}$.

Assumptions for the IPW scheme. We formalize the role of the IPW weights with the following assumption, which parallels the selection-on-observables condition discussed in Section 6.3.

Assumption H.4 (IPW balancing and incumbent pricing)**.**

1. **Balancing in Z within year-by-treatment cells.** Within each year-by-treatment cell, the IPW weights $w_i^{R,w}$ are such that the marginal distribution of Z in the INC-IPW sample matches that of the FULL sample:

$$F_{Z,R,w} = F_{Z,F} \quad \text{(within year-by-treatment cells)}.$$

2. **Invariance of the incumbent pricing function.** The conditional mean of y given Z among incumbents is the same under the INC and INC-IPW measures:

$$m_R(z) = m_{R,w}(z) \equiv m_I(z) \quad \text{for all } z.$$

This reflects the fact that reweighting incumbents changes their relative frequencies across values of Z but does not change their pricing behavior conditional on Z.

3. **Mixture representation for FULL.** Under the FULL measure, the conditional mean of y given $Z = z$ combines incumbents and transitory relationships:

$$m_F(z) = p_I(z)\, m_I(z) + p_E(z)\, m_E(z),$$

where $p_I(z)$ and $p_E(z)$ are the conditional shares of incumbents and transitory relationships among granted loans with $Z = z$, and $m_E(z)$ is the pricing function for transitory relationships.

By (21), we can write

$$\theta^F = \theta[F_{Z,F}, m_F], \qquad \theta^R = \theta[F_{Z,R}, m_I], \qquad \theta^{R,w} = \theta[F_{Z,F}, m_I].$$

Two generic lemmas. We first prove two simple lemmas that separate the roles of the pricing function and the distribution of Z.

Lemma H.5 (Same composition, different pricing rule). *Let $(F_{Z,1}, m_1)$ and $(F_{Z,2}, m_2)$ be two pairs of marginal distributions and pricing functions. If $F_{Z,1} = F_{Z,2} \equiv F_Z$, then the difference in slopes*

$$\theta[F_{Z,1}, m_1] - \theta[F_{Z,2}, m_2]$$

is entirely due to the difference between m_1 and m_2. In particular, if $m_1(z) = m_2(z)$ for all z, then $\theta[F_{Z,1}, m_1] = \theta[F_{Z,2}, m_2]$.

Proof. By definition,

$$\theta[F_{Z,1}, m_1] = \frac{\displaystyle\int z\, m_1(z)\, dF_{Z,1}(z)}{\displaystyle\int z^2\, dF_{Z,1}(z)}, \qquad \theta[F_{Z,2}, m_2] = \frac{\displaystyle\int z\, m_2(z)\, dF_{Z,2}(z)}{\displaystyle\int z^2\, dF_{Z,2}(z)}.$$

If $F_{Z,1} = F_{Z,2} = F_Z$, the denominators coincide:

$$\int z^2\, dF_{Z,1}(z) = \int z^2\, dF_{Z,2}(z) = \int z^2\, dF_Z(z).$$

Thus

$$\theta[F_{Z,1}, m_1] - \theta[F_{Z,2}, m_2] = \frac{\displaystyle\int z\Big(m_1(z) - m_2(z)\Big)\, dF_Z(z)}{\displaystyle\int z^2\, dF_Z(z)}.$$

If $m_1(z) = m_2(z)$ for all z, the numerator vanishes and the two slopes are equal. Conversely, any difference in slopes must come from the difference $m_1 - m_2$, since F_Z is held fixed. $\square$

Lemma H.6 (Same pricing rule, different composition). *Let* $(F_{Z,1}, m)$ *and* $(F_{Z,2}, m)$ *share the same pricing function* $m(z)$ *but have different distributions* $F_{Z,1}$ *and* $F_{Z,2}$. *Then the difference in slopes*

$$\theta[F_{Z,1}, m] - \theta[F_{Z,2}, m]$$

is entirely due to the difference between $F_{Z,1}$ *and* $F_{Z,2}$. *In particular, if* $F_{Z,1} = F_{Z,2}$, *then* $\theta[F_{Z,1}, m] = \theta[F_{Z,2}, m]$.

Proof. From (21),

$$\theta[F_{Z,1}, m] = \frac{\int z\, m(z)\, dF_{Z,1}(z)}{\int z^2\, dF_{Z,1}(z)}, \qquad \theta[F_{Z,2}, m] = \frac{\int z\, m(z)\, dF_{Z,2}(z)}{\int z^2\, dF_{Z,2}(z)}.$$

If $F_{Z,1} = F_{Z,2}$, both numerator and denominator are identical, hence the slopes coincide. More generally, any difference in slopes must come from the change in the integrals with respect to $F_{Z,1}$ and $F_{Z,2}$, since the function m is common. □

Application to FULL, INC, and INC-IPW. We now apply Lemmas H.5 and H.6 to the triple $(\theta^F, \theta^R, \theta^{R,w})$.

By Assumption H.4(i), the marginal distribution of Z in the INC-IPW sample matches that of the FULL sample:

$$F_{Z,R,w} = F_{Z,F}.$$

By Assumption H.4(ii), the pricing function in INC and INC-IPW coincides:

$$m_R(z) = m_{R,w}(z) = m_I(z).$$

Thus we can write

$$\theta^F = \theta[F_{Z,F}, m_F], \qquad \theta^R = \theta[F_{Z,R}, m_I], \qquad \theta^{R,w} = \theta[F_{Z,F}, m_I].$$

We first interpret $\theta^F - \theta^{R,w}$. The pair $(F_{Z,F}, m_F)$ corresponds to the FULL sample, while $(F_{Z,F}, m_I)$ corresponds to a hypothetical world in which the distribution of Z is the same as in the FULL sample but all loans are priced according to the incumbent pricing rule m_I. Since both arguments share the same distribution $F_{Z,F}$, Lemma H.5 implies that

$$\theta^F - \theta^{R,w} = \theta[F_{Z,F}, m_F] - \theta[F_{Z,F}, m_I]$$

is entirely due to the difference between m_F and m_I, i.e. to the fact that transitory relationships may be priced differently from incumbents. Economically, $\theta^F - \theta^{R,w}$ is an *entry/exit pricing effect*: it measures how the slope changes when, holding the composition of Z fixed at the FULL-sample distribution, we replace the incumbent pricing rule by the actual mixture of incumbent and transitory pricing.

We next interpret $\theta^{R,w} - \theta^R$. The pair $(F_{Z,R}, m_I)$ corresponds to the actual INC sample, while $(F_{Z,F}, m_I)$ corresponds to the reweighted INC sample that reproduces the FULL distribution of Z but keeps the same pricing function m_I. Since both arguments share the same pricing function m_I, Lemma H.6 implies that

$$\theta^{R,w} - \theta^R = \theta[F_{Z,F}, m_I] - \theta[F_{Z,R}, m_I]$$

is entirely due to the difference between $F_{Z,F}$ and $F_{Z,R}$, i.e. to differences in the observable composition of emissions intensity between the FULL and INC samples. Economically, $\theta^{R,w} - \theta^R$ is an *entry/exit composition effect*: it measures how the slope for incumbents would change if the distribution of Z among incumbents were tilted to match the FULL-sample distribution.

Combining these two pieces yields the formal decomposition.

Proposition H.7 (Pricing vs composition in the entry/exit margin)**.** Under Assumption H.4, the population entry/exit margin satisfies

$$\theta^F - \theta^R = \underbrace{\left(\theta^F - \theta^{R,w}\right)}_{\text{entry/exit pricing}} + \underbrace{\left(\theta^{R,w} - \theta^R\right)}_{\text{entry/exit composition}},$$

where:

1. $\theta^F - \theta^{R,w}$ isolates differences in pricing behavior between transitory and incumbent relationships for a fixed composition of Z (the FULL distribution).
2. $\theta^{R,w} - \theta^R$ isolates differences in the observable distribution of Z between the FULL and INC samples, holding the incumbent pricing rule m_I fixed.

Proof. The identity

$$\theta^F - \theta^R = \left(\theta^F - \theta^{R,w}\right) + \left(\theta^{R,w} - \theta^R\right)$$

is algebraic. Under Assumption H.4, we have

$$\theta^F = \theta[F_{Z,F}, m_F], \quad \theta^{R,w} = \theta[F_{Z,F}, m_I], \quad \theta^R = \theta[F_{Z,R}, m_I].$$

By Lemma H.5, the difference

$$\theta^F - \theta^{R,w} = \theta[F_{Z,F}, m_F] - \theta[F_{Z,F}, m_I]$$

is entirely due to the difference between m_F and m_I for a fixed distribution $F_{Z,F}$, which yields part (i). By Lemma H.6, the difference

$$\theta^{R,w} - \theta^R = \theta[F_{Z,F}, m_I] - \theta[F_{Z,R}, m_I]$$

is entirely due to the difference between $F_{Z,F}$ and $F_{Z,R}$ for a fixed pricing function m_I, which yields part (ii). □

In sample notation, Proposition H.7 justifies the empirical sub-decomposition

$$\widehat{\theta}^F - \widehat{\theta}^R = \underbrace{\left(\widehat{\theta}^F - \widehat{\theta}^{R,w}\right)}_{\text{entry/exit pricing}} + \underbrace{\left(\widehat{\theta}^{R,w} - \widehat{\theta}^R\right)}_{\text{entry/exit composition}},$$

and, together with Proposition H.3, underpins the full channels decomposition used in the main text and in Eq. (8).

Appendix I Placebo Shock

We provide here the results of a baseline non-staggered specification using a placebo shock date set to November 2019, with a symmetric 12-month window around this date, i. e. earlier than the release of the methodology of the climate stress test in October 2021 and before the main phase of climate-stress-test implementation. It therefore provides a more stringent falsification test than alternative treatment timings based on early CST communication, because no comparable supervisory information shock should be expected to affect the relative carbon-sensitivity of loan pricing around this date.

Table G1 reports the results. Panel A shows the four baseline specifications. The coefficient of interest, the triple interaction between borrower carbon intensity, treated-bank status, and the post-placebo indicator, is negative in all specifications but small and statistically insignificant. In the FULL sample, the estimated placebo effect is -0.0246 with a clustered standard error of 0.0225. The corresponding estimates are even closer to zero in the incumbent specifications: -0.0047 in INC, -0.0053 in INC-IPW, and -0.0018 in the within-relationship INT specification. Thus, the placebo exercise does not reveal a significant differential change in the spread–carbon-intensity slope for treated banks around November 2019.

Panel B reports the associated channel decomposition with wild-cluster bootstrap inference. The total placebo effect is statistically insignificant: the point estimate is -0.0246, with a bootstrap standard error of 0.0220 and a two-sided bootstrap p-value of 0.285. The intensive margin is economically negligible and statistically insignificant, as are reallocation among incumbents and the observable-composition component of entry and exit. The only statistically significant component is the entry/exit pricing term, which is negative. This pattern does not reproduce the main baseline findings, where the economically relevant post-CST adjustment operates through portfolio rebalancing and extensive-margin pricing responses associated with the actual supervisory intervention. Instead, the placebo decomposition points to no robust aggregate pre-trend in the carbon-pricing semi-elasticity and no evidence of systematic within-relationship repricing before the CST.

Overall, the placebo exercise corroborates the identifying assumptions underlying the main estimates. If the baseline findings were driven by mechanical differences in pre-existing trends between SIs and LSIs, one would expect a sizeable and similarly structured effect to appear around the pre-treatment placebo date. The data reject this hypothesis. The aggregate placebo effect is small and not statistically significant; the incumbent and intensive-margin effects are close to zero; and the only significant component (entry/exit pricing) has the opposite sign to the extensive-margin pricing response emphasized in the main results. These findings are therefore consistent with the interpretation that the main effects documented in the paper are attributable to the climate stress-test information and scrutiny channel, rather than to spurious pre-existing differences in carbon-sensitive loan pricing.

Table G1. Placebo shock: November 2019

	FULL	INC	INC-IPW	INT
Panel A. Triple-difference estimates				
log(GHG intensity)	0.0538**	0.0624***	0.0684***	−0.0098
	(0.0162)	(0.0149)	(0.0151)	(0.0204)
log(GHG intensity) × TreatBank	0.0452*	0.0280	0.0234	0.0138
	(0.0216)	(0.0217)	(0.0212)	(0.0209)
log(GHG intensity) × Post	−0.0032	0.0046	0.0053	0.0070
	(0.0203)	(0.0165)	(0.0173)	(0.0084)
log(GHG intensity) × TreatBank × Post	−0.0246	−0.0047	−0.0053	−0.0018
	(0.0225)	(0.0198)	(0.0203)	(0.0094)
Country–size–month FE	Yes	Yes	Yes	Yes
Bank–month FE	Yes	Yes	Yes	Yes
Instrument FE	Yes	Yes	Yes	Yes
Bank–debtor FE	No	No	No	Yes
Observations	966,863	539,384	539,384	539,383
R^2	0.3246	0.3324	0.3341	0.8326
Within R^2	0.0071	0.0082	0.0084	0.0001

Panel B. Channel decomposition and wild-cluster bootstrap inference

Channel	Estimate	Share	Bootstrap s.e.	p-value
Intensive margin	−0.0018	0.072	0.0092	0.846
Reallocation among incumbents	−0.0029	0.119	0.0110	0.768
Entry/exit pricing at adjusted composition	−0.0193**	0.785	0.0084	0.016
Entry/exit composition in observables	−0.0006	0.024	0.0012	0.645
Total effect, FULL	−0.0246	1.000	0.0220	0.285

Notes: The table reports a placebo specification using November 2019 as the shock date and a symmetric 12-month event window. Panel A reports the baseline triple-difference-in-slopes estimates. Standard errors in parentheses are clustered at the parent-bank level. Panel B reports the decomposition of the FULL triple-difference coefficient into the intensive margin, reallocation among incumbent bank–debtor relationships, entry/exit pricing at adjusted composition, and entry/exit composition in observables. Bootstrap standard errors and p-values are based on 499 Rademacher wild-cluster bootstrap draws clustered by parent bank, with 121 clusters. Significance levels: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

Online Supplementary Material

Appendix J Collective Action Problem in the Banking Sector and Climate Regulation

We consider a competitive banking market in which each bank decides whether to incorporate transition risk into its lending decisions. We use a simple binary-action coordination game with strategic complementarities, which is a special case of the noncooperative supermodular games analyzed by Milgrom and Roberts (1990). We model transition-risk integration as a coordination problem with strategic complementarities[46]: when only a few banks integrate transition risks into their lending behavior, (i) discipline on carbon-intensive borrowers is weak because credit demand can reallocate to non-adopting banks, and (ii) the private effectiveness of screening and pricing is limited because firms' incentives to adjust remain weak. When many banks integrate transition risks into their lending behavior, leakage shrinks and the price signal becomes stringent economy-wide, increasing both private and social returns. This structure can generate multiple Nash equilibria (see, Cooper and John, 1988; Milgrom and Roberts, 1990), which in turn motivates regulatory intervention to select the socially desirable equilibrium. The framework is deliberately stylized and should be interpreted as a proof-of-concept model: it abstracts from richer dynamics and institutional details, but it provides formal micro-founded substance for the argument that uncoordinated market outcomes may justify public intervention.

Model setup. There are N competing banks $i \in \{1, \dots, N\}$ and aggregate loan demand $D > 0$. Each bank chooses a binary action $a_i \in \{C, NC\}$, where C denotes integrating transition risk (e.g. enhanced data collection, pricing and/or limits for transition-exposed borrowers), and NC denotes not doing so. Let

$$x \equiv \frac{1}{N} \sum_{j=1}^{N} \mathbb{1}\{a_j = C\} \in [0, 1]$$

denote the fraction of banks adopting C. We consider a highly competitive environment in which baseline lending conditions are similar across banks. The key competitive margin is the allocation of credit to transition-exposed (*brown*) borrowers: banks choosing C increase interest rates for such borrowers, which induces a reallocation of loan demand towards NC banks whenever they remain available as outside options.

Large-N (nonatomic) approximation. When analyzing unilateral deviations, the relevant state is the adoption rate among other banks, $x_{-i} \equiv \frac{1}{N-1} \sum_{j \neq i} \mathbb{1}\{a_j = C\}$. We work with a standard large-N (mean-field) approximation in which each bank treats x as exogenous, and $|x - x_{-i}| = O(1/N)$ is negligible for incentive comparisons.

[46]Strategic complementarities arise when the optimal strategy of an agent depends positively upon the strategies of the other agents (see, Cooper and John, 1988).

Demand reallocation. We capture demand reallocation in reduced form by allowing per-bank loan demand to depend on x:

$$q_C(x) = \frac{D}{N}\Big(1 - \lambda(1-x)\Big), \tag{1}$$

$$q_{NC}(x) = \frac{D}{N}\Big(1 + \lambda x\Big), \tag{2}$$

where $\lambda \in (0,1)$ measures how strongly borrowers can substitute away from lenders more sensitive to transition risks. When x is low (few adopters), $1-x$ is large and a bank adopting C loses more market shares to NC banks. When x is high, few NC banks remain and demand reallocation is limited. Note that the aggregate loan volume allocated to non-adopters equals $(1-x)Nq_{NC}(x)$ and vanishes as $x \to 1$, even though the per-bank demand wedge $q_{NC}(x) - q_C(x)$ is constant. Hence, demand reallocation shrinks on the extensive margin: there are fewer non-adopting outside options.

Transition-loss technology and strategic complementarity. Let $m \equiv R - c$ denote the per-unit interest margin net of operating costs, where R is the interest rate and c is marginal cost. Let $\ell_{NC}(x)$ be the expected transition-related monetary loss per unit of credit for a non-adopting bank, with $\ell'_{NC}(x) < 0$, reflecting that more widespread integration of transition risk reduces aggregate vulnerability: as more banks adopt strategy C, economy-wide pricing on transition-exposed activities lowers expected losses even for NC banks. A bank adopting C can reduce its expected loss by an amount $\eta(x) \geq 0$:

$$\ell_C(x) = \ell_{NC}(x) - \eta(x), \qquad \eta(0) = 0, \quad \eta'(x) > 0.$$

The key condition $\eta'(x) > 0$ corresponds to strategic complementarity: the private effectiveness of integrating transition risk increases with x because (i) fewer non-adopting banks remain to refinance exposed borrowers, and (ii) the resulting economy-wide price signal induces broader emissions adjustment by firms.

Payoffs. A bank adopting C pays a fixed private cost $\sigma > 0$ (e.g. for data, modelling and governance). The profit functions of both strategies are

$$\Pi_i(C;x) = q_C(x)\Big[m - \ell_C(x)\Big] - \sigma, \tag{3}$$

$$\Pi_i(NC;x) = q_{NC}(x)\Big[m - \ell_{NC}(x)\Big]. \tag{4}$$

Best response and payoff differential. Given x (and taking N large so that each bank is approximately a price-taker in x), bank i prefers C if and only if

$$\Delta\Pi(x) \equiv \Pi_i(C;x) - \Pi_i(NC;x) \geq 0.$$

Substituting from Eqs. (3)–(4) and using $\ell_C(x) = \ell_{NC}(x) - \eta(x)$ yields

$$\begin{aligned}\Delta\Pi(x) &= q_C(x)\Big[m - \ell_{NC}(x) + \eta(x)\Big] - \sigma - q_{NC}(x)\Big[m - \ell_{NC}(x)\Big] \\ &= \underbrace{\Big(q_C(x) - q_{NC}(x)\Big)}_{\text{demand differential}}\Big[m - \ell_{NC}(x)\Big] + \underbrace{q_C(x)\eta(x)}_{\text{risk reduction}} - \sigma.\end{aligned}$$

From Eqs. (1)–(2),

$$\begin{aligned}q_C(x) - q_{NC}(x) &= \frac{D}{N}\Big(1 - \lambda + \lambda x - (1 + \lambda x)\Big) = -\frac{\lambda D}{N}, \\ q_C(x)\eta(x) &= \frac{D}{N}\Big(1 - \lambda + \lambda x\Big)\eta(x),\end{aligned}$$

so that

$$\Delta\Pi(x) = -\frac{\lambda D}{N}\Big[m - \ell_{NC}(x)\Big] + \frac{D}{N}\Big(1 - \lambda + \lambda x\Big)\eta(x) - \sigma. \tag{5}$$

The first term in (5) is strictly negative and reflects the competitive disadvantage of adopting C: each adopting bank serves less loan demand than each non-adopting bank by an amount $\lambda D/N$. The second term is the private risk-reduction benefit from integrating transition risk, which grows with x through both $\eta(x)$ and $q_C(x)$, and can also increase via the indirect effect of x on $\ell_{NC}(x)$.

Differentiating (5) with respect to x yields

$$\frac{\partial}{\partial x}\Delta\Pi(x) = \frac{D}{N}\Big[\lambda \ell'_{NC}(x) + \lambda\eta(x) + \Big(1 - \lambda + \lambda x\Big)\eta'(x)\Big]. \tag{6}$$

The first term in brackets is negative because $\ell'_{NC}(x) < 0$ and $\lambda > 0$: as x rises, expected losses $\ell_{NC}(x)$ fall, raising the effective margin $m - \ell_{NC}(x)$ and thereby amplifying the competitive loss associated with serving a smaller loan volume ($q_C < q_{NC}$). The linear combination of the remaining terms is positive because $\eta(x) \geq 0$, $\eta'(x) > 0$ and $1 - \lambda + \lambda x > 0$ for all $x \in [0, 1]$. This captures the positive complementarity effect: as x increases, the risk reduction $\eta(x)$ becomes larger and applies to a loan volume $q_C(x)$ that itself increases with x.

Assumption J.1 (Transition losses and complementarities)**.** The functions $\ell_{NC} : [0, 1] \to \mathbb{R}$ and $\eta : [0, 1] \to \mathbb{R}_+$ satisfy

$$\ell'_{NC}(x) < 0, \qquad \eta(x) \geq 0, \qquad \eta'(x) > 0, \quad \text{for all } x \in [0, 1],$$

and are continuous on $[0, 1]$. The demand system satisfies

$$q_C(x) = \frac{D}{N}\Big(1 - \lambda + \lambda x\Big), \qquad q_{NC}(x) = \frac{D}{N}\Big(1 + \lambda x\Big), \qquad \lambda \in (0, 1).$$

Proposition J.2 (Increasing differences)**.** Under Assumption J.1, the payoff differential

$$\Delta\Pi(x) \equiv \Pi_i(C; x) - \Pi_i(NC; x)$$

is continuous and strictly increasing in x on $[0, 1]$. In particular, the game exhibits increasing differences in the sense of Milgrom and Roberts (1990): the private incentive to adopt C is higher when a larger fraction of other banks adopt C.

Proof. Continuity follows from the continuity of $q_C(\cdot)$, $q_{NC}(\cdot)$, $\ell_{NC}(\cdot)$ and $\eta(\cdot)$. Differentiating (5) with respect to x yields

$$\frac{\partial}{\partial x}\Delta\Pi(x) = \frac{D}{N}\Big[\lambda\ell'_{NC}(x) + \lambda\eta(x) + \big(1-\lambda+\lambda x\big)\eta'(x)\Big],$$

see Eq. (6). Under Assumption J.1, we have $\ell'_{NC}(x) < 0$, $\eta(x) \geq 0$, $\eta'(x) > 0$, and $1-\lambda+\lambda x > 0$ for all $x \in [0,1]$. Hence the term in brackets is strictly positive, implying $\partial\Delta\Pi(x)/\partial x > 0$ for all $x \in [0,1]$. This establishes that $\Delta\Pi(x)$ is strictly increasing. $\square$

Assumption J.3 (Stag-hunt boundary conditions)**.** The payoff differential satisfies

$$\Delta\Pi(0) < 0 \quad \text{and} \quad \Delta\Pi(1) > 0.$$

Proposition J.4 (Existence and uniqueness of the adoption threshold)**.** Suppose Assumptions J.1 and J.3 hold. Then there exists a unique $\bar{x} \in (0,1)$ such that $\Delta\Pi(\bar{x}) = 0$. Moreover,

$$\Delta\Pi(x) < 0 \quad \text{for } x < \bar{x}, \qquad \Delta\Pi(x) > 0 \quad \text{for } x > \bar{x}.$$

In the large-N approximation where each bank treats x as given, the individual best response is NC whenever $x < \bar{x}$, C whenever $x > \bar{x}$, and both actions are optimal when $x = \bar{x}$.

Proof. By Proposition J.2, $\Delta\Pi(x)$ is continuous and strictly increasing on $[0,1]$. Assumption J.3 implies $\Delta\Pi(0) < 0$ and $\Delta\Pi(1) > 0$. By the Intermediate Value Theorem, there exists at least one $\bar{x} \in (0,1)$ such that $\Delta\Pi(\bar{x}) = 0$. Strict monotonicity rules out multiple roots, so $\bar{x}$ is unique. Monotonicity further implies $\Delta\Pi(x) < 0$ for all $x < \bar{x}$ and $\Delta\Pi(x) > 0$ for all $x > \bar{x}$. The characterization of best responses follows immediately from the sign of $\Delta\Pi(x)$. $\square$

Given Proposition J.4, the symmetric best-response correspondence can be written as

$$\mathrm{BR}(x) = \begin{cases} \{NC\} & \text{if } x < \bar{x}, \\ \{C, NC\} & \text{if } x = \bar{x}, \\ \{C\} & \text{if } x > \bar{x}. \end{cases}$$

Proposition J.5 (Nash equilibria and coordination structure)**.** Under Assumptions J.1 and J.3, the following statements hold:

(i) The profile in which all banks choose NC (i.e. $x = 0$) is a symmetric pure-strategy Nash equilibrium.

(ii) The profile in which all banks choose C (i.e. $x = 1$) is also a symmetric pure-strategy Nash equilibrium.

(iii) In a finite N-player game, there exists a symmetric mixed-strategy equilibrium in which each bank adopts C with probability $p \in (0,1)$ such that $\Delta\Pi(p) = 0$. In the large-N limit, this probability converges to the unique threshold $\bar{x}$ characterized in Proposition J.4. Under standard adjustment dynamics, $\bar{x}$ acts as an unstable tipping point separating the basins of attraction of $x = 0$ and $x = 1$.

Proof. (i) If all banks choose NC, then $x = 0$ and any unilateral deviation to C faces an adoption rate $x_{-i} \approx 0$. By Assumption J.3, $\Delta\Pi(0) < 0$, so deviating to C strictly lowers the deviator's payoff. Hence $x = 0$ is a pure-strategy Nash equilibrium.

(ii) If all banks choose C, then $x = 1$ and any unilateral deviation to NC faces $x_{-i} \approx 1$. Assumption J.3 implies $\Delta\Pi(1) > 0$, so adopting C yields a strictly higher payoff than NC. Deviating to NC is therefore unprofitable, and $x = 1$ is a pure-strategy Nash equilibrium.

(iii) In a finite N-player game, the usual arguments for two-action coordination games imply the existence of a symmetric mixed-strategy equilibrium in which each bank adopts C with probability $p \in (0, 1)$ that solves the indifference condition $\Delta\Pi(p) = 0$, interpreted as the adoption rate among other banks. As $N \to \infty$, this adoption probability converges to the unique threshold $\bar{x}$ in Proposition J.4, which separates regions where C or NC are strictly preferred. Under standard myopic best-response or replicator dynamics, $\bar{x}$ is unstable, while $x = 0$ and $x = 1$ are locally stable, yielding the familiar stag-hunt coordination structure. □

Consequently, the game features two robust pure-strategy Nash equilibria:

$$x = 0 \quad (\text{all } NC), \qquad x = 1 \quad (\text{all } C),$$

which correspond to a standard stag-hunt coordination structure. If others do not integrate transition risk (x close to 0), it is privately unattractive to be an early adopter; if others do integrate (x close to 1), deviating to non-adoption becomes unattractive because leakage disappears and the private benefits of risk integration are high.

Social welfare and systemic-risk externality. Let $E(x)$ denote the social benefit from lower systemic risk when the fraction of adopters is x, with $E'(x) > 0$ and $E''(x) \geq 0$, allowing for convexity when systemic resilience improves sharply once climate risk pricing becomes sufficiently widespread. We allow $E(x)$ to be an aggregate externality term that can scale with market size N (e.g. $E(x) = N\,\tilde{E}(x)$), so that $E'(x)/N$ remains $O(1)$ under the large-N approximation. We interpret $\ell_{NC}(x)$ and $\ell_C(x)$ as privately internalized expected credit losses. Thus, the welfare term $E(x)$ captures residual systemic externalities not priced by individual banks, such as fire-sale spillovers, network default cascades, and other real-economy costs of crises (Allen and Gale, 2000; Eisenberg and Noe, 2001).

Define social welfare as the sum of banks' profits and a residual systemic externality term:

$$W(x) = \sum_{i=1}^{N} \Pi_i(a_i; x) + E(x). \tag{7}$$

Under symmetry, let $\Pi_C(x) \equiv \Pi_i(C; x)$ and $\Pi_{NC}(x) \equiv \Pi_i(NC; x)$ for any bank i. Then welfare can be written as

$$W(x) = xN\,\Pi_C(x) + (1 - x)N\,\Pi_{NC}(x) + E(x).$$

Differentiating $W(x)$ w.r.t. x yields

$$\begin{aligned} W'(x) &= \frac{d}{dx}\Big(xN\Pi_C(x) + (1-x)N\Pi_{NC}(x) + E(x)\Big) \\ &= N\Pi_C(x) + xN\Pi_C'(x) - N\Pi_{NC}(x) + (1-x)N\Pi_{NC}'(x) + E'(x) \\ &= N\Big[\Pi_C(x) - \Pi_{NC}(x)\Big] + N\Big[x\Pi_C'(x) + (1-x)\Pi_{NC}'(x)\Big] + E'(x). \end{aligned} \tag{8}$$

Rewrite Eq. (8) more compactly as

$$W'(x) = N\,\Delta\Pi(x) + N\Big[x\Pi_C'(x) + (1-x)\Pi_{NC}'(x)\Big] + E'(x). \tag{9}$$

Dividing by N yields an expression for the marginal social gain per bank of an increase in the adoption rate:

$$\frac{1}{N}W'(x) = \Delta\Pi(x) + \underbrace{x\Pi_C'(x) + (1-x)\Pi_{NC}'(x)}_{\text{endogenous payoff effects as } x \text{ changes}} + \underbrace{E'(x)/N}_{\text{systemic-risk externality}}. \tag{10}$$

Eq. (10) highlights the wedge between private and social incentives (Pigou, 2017). An individual bank, treated as atomistic in a large-N environment, compares $\Pi_C(x)$ and $\Pi_{NC}(x)$ taking x as given. Its decision rule is therefore characterized by the sign of $\Delta\Pi(x)$ alone. In contrast, a social planner internalizes not only the direct payoff differential $\Delta\Pi(x)$, but also (i) the impact of a higher x on all banks' payoffs, captured by $x\Pi_C'(x) + (1-x)\Pi_{NC}'(x)$, and (ii) the marginal reduction in systemic risk, $E'(x)/N$.

To make (10) more explicit, recall that under symmetry

$$\Pi_C(x) = q_C(x)\Big[m - \ell_C(x)\Big] - \sigma, \qquad \Pi_{NC}(x) = q_{NC}(x)\Big[m - \ell_{NC}(x)\Big],$$

with

$$\ell_C(x) = \ell_{NC}(x) - \eta(x), \quad \ell_{NC}'(x) < 0, \quad \eta(x) \geq 0, \quad \eta'(x) > 0,$$

and

$$q_C(x) = \frac{D}{N}\Big(1 - \lambda + \lambda x\Big), \qquad q_{NC}(x) = \frac{D}{N}\Big(1 + \lambda x\Big), \qquad \lambda \in (0,1).$$

Differentiating $\Pi_C(x)$ and $\Pi_{NC}(x)$ w.r.t. x gives

$$\begin{aligned} \Pi_C'(x) &= q_C'(x)\Big[m - \ell_C(x)\Big] + q_C(x)\Big[-\ell_C'(x)\Big] \\ &= q_C'(x)\Big[m - \ell_{NC}(x) + \eta(x)\Big] - q_C(x)\Big[\ell_{NC}'(x) - \eta'(x)\Big], \end{aligned} \tag{11}$$

$$\Pi_{NC}'(x) = q_{NC}'(x)\Big[m - \ell_{NC}(x)\Big] + q_{NC}(x)\Big[-\ell_{NC}'(x)\Big]. \tag{12}$$

Given $q_C'(x) = q_{NC}'(x) = \frac{D}{N}\lambda > 0$, $\ell_{NC}'(x) < 0$, $\eta(x) \geq 0$ and $\eta'(x) > 0$, Eqs. (11)–(12) highlight two effects of an increase in x on private payoffs: (i) changes in loan volumes $q_C(x)$ and $q_{NC}(x)$ through demand reallocation, and (ii) changes in expected climate transition monetary losses through $\ell_{NC}(x)$ and $\eta(x)$.

Substituting (11)–(12) into (9) yields an explicit, micro-founded expression for $W'(x)$ in terms of $\ell_{NC}(\cdot)$, $\eta(\cdot)$, and the demand system $q_C(\cdot)$, $q_{NC}(\cdot)$. Although the exact closed form is not required for the results that follow, equations (11)–(12) make it immediate that $W'(x)$ includes terms proportional to $\ell'_{NC}(x)$ and $\eta'(x)$ that are ignored by individual banks when they take x as exogenous, in addition to the purely systemic-risk component $E'(x)$.

Proposition J.6 (Private versus social incentives)**.** Under symmetry, social welfare is given by

$$W(x) = xN\,\Pi_C(x) + (1-x)N\,\Pi_{NC}(x) + E(x),$$

and the marginal social gain per bank of a marginal increase in the adoption rate is

$$\frac{1}{N}W'(x) = \Delta\Pi(x) + \Big[x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)\Big] + \frac{E'(x)}{N}.$$

Any interior solution $x^\star \in (0,1)$ to the planner's problem

$$x^\star \in \arg\max_{x\in[0,1]} W(x)$$

must satisfy

$$\Delta\Pi(x^\star) = -\Big[x^\star\Pi'_C(x^\star) + \big(1-x^\star\big)\Pi'_{NC}(x^\star)\Big] - \frac{E'(x^\star)}{N}. \tag{13}$$

In particular, at any interior social optimum $x^\star$, the private payoff differential $\Delta\Pi(x^\star)$ is strictly negative.

Proof. The expression for $W(x)$ follows directly from the definition of social welfare and symmetry. Differentiating with respect to x yields

$$W'(x) = N\,\Delta\Pi(x) + N\Big[x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)\Big] + E'(x),$$

which implies the stated formula for $W'(x)/N$. An interior optimum $x^\star \in (0,1)$ for the planner's problem must satisfy the first-order condition $W'(x^\star) = 0$, which can be rewritten as (13). Under Assumption J.1, the term in brackets is strictly positive, and $E'(x^\star)/N > 0$ by construction. Hence the right-hand side of (13) is strictly negative, implying $\Delta\Pi(x^\star) < 0$. □

Private vs. social optima. A symmetric decentralized equilibrium adoption rate is characterized, in the large-N approximation, by the fixed-point condition that each bank best responds to x based solely on the sign of $\Delta\Pi(x)$, which generates the low-adoption equilibrium $x = 0$ and the high-adoption equilibrium $x = 1$ described above. By contrast, a social planner chooses x to maximize $W(x)$ and internalizes both the impact of higher adoption on all banks' payoffs and the marginal reduction in systemic risk. As a result, there exist parameter constellations under which the decentralized economy is trapped in the low-adoption equilibrium $x = 0$ even though the planner's problem is maximized at a strictly higher $x^\star$, potentially at $x^\star = 1$ if $E(\cdot)$ is sufficiently large and/or convex.

Appendix K Wild-Cluster Bootstrap Inference

This appendix describes the wild-cluster bootstrap procedure used throughout the paper for inference on all channel decompositions, in the static, event-study, and staggered specifications.

General Setup

All empirical specifications in the paper are linear fixed-effects models of the form

$$Y = X\beta + D\gamma + F\lambda + u,$$

where Y is the outcome (credit spread), X collects the moderator(s) of interest (typically log-emissions intensity and its interactions with treatment and time indicators), D denotes additional controls, F are high-dimensional fixed effects, and u is the disturbance term. Estimation is conducted by OLS or WLS (when using IPW). Let $\widehat{\beta}$ denote the estimated coefficient vector from a given specification, and let $\widehat{\Theta} = f(\widehat{\beta})$ be a linear functional of $\widehat{\beta}$:

$$\widehat{\Theta} = A\widehat{\beta},$$

for some known matrix A. In all applications, $f(\cdot)$ maps regression coefficients into a set of four channels based on post-shock (or event-time-specific) slope coefficients measured in four samples. Let $\widehat{\theta}_F$, $\widehat{\theta}_R$, $\widehat{\theta}_I$, and $\widehat{\theta}_R^{\text{IPW}}$ denote the relevant slope coefficients in these four specifications (for a given post-shock indicator, event time, or staggered treatment-slope term). The four channels are always defined by the same linear mapping as in Eq. (8)

In the static specification, each $\widehat{\theta}.$ is a single post-shock slope coefficient. In the event-study specification, we obtain collections $\{\widehat{\theta}_{F,k}\}$, $\{\widehat{\theta}_{R,k}\}$, $\{\widehat{\theta}_{I,k}\}$, and $\{\widehat{\theta}_{R,k}^{\text{IPW}}\}$ indexed by event time k, and the same mapping is applied for each k. In the staggered design, the objects aggregated into FCST or ECST estimands are again linear combinations of underlying quantities, so that the channel decomposition remains linear in $\widehat{\beta}$.

Wild-Cluster Bootstrap Algorithm

Inference relies on the wild-cluster bootstrap at the bank level, following Cameron, Gelbach, and Miller (2008) and MacKinnon and Webb (2017). The same algorithm is used for all specifications; only the mapping $f(\cdot)$ changes depending on whether we work with a single post-shock coefficient, an event-time profile, or a staggered aggregate. In all settings—static, dynamic, and staggered—the wild-cluster bootstrap therefore provides a unified, cluster-robust inference procedure for the four-channel decomposition reported in the main text.

Algorithm 1 Wild-Cluster Bootstrap for Linear Channel Functionals

1: **Inputs:** outcome Y, moderator(s) X, treatment and timing indicators, bank-level clusters $c(i)$, fixed effects F, weight matrix W (identity or IPW), number of replications B, wild multiplier law $\mathcal{G}$ (Rademacher in the baseline), confidence level $1-\alpha$.

2: **Baseline estimation:**

1. Estimate the relevant linear model (e.g., Eq. (6), (9), or (12)) on each sample, using the same set of regressors and fixed effects as in the main text.
2. Recover the slope coefficients $\widehat{\theta}_F$, $\widehat{\theta}_R$, $\widehat{\theta}_I$, and $\widehat{\theta}_R^{\mathrm{IPW}}$, as well as fitted values $\widehat{Y}$ and residuals $\widehat{u}$ on the effective estimation support of each regression.
3. Map $(\widehat{\theta}_F, \widehat{\theta}_R, \widehat{\theta}_I, \widehat{\theta}_R^{\mathrm{IPW}})$ into channel estimates $\widehat{\Theta} = f(\widehat{\beta})$ using the four linear combinations defined above.

3: **Cluster structure:** construct the union of all bank identifiers appearing in the four estimation samples and index them by $c = 1, \ldots, C$. This ensures that the same multiplier is drawn for a given bank across all four regressions in each replication.

4: **for** $b = 1, \ldots, B$ **do**

5: Draw i.i.d. cluster multipliers $\{g_c^{(b)}\}_{c=1}^{C}$ from $\mathcal{G}$ and set $g_{it}^{(b)} = g_{c(i)}^{(b)}$ for every observation (i,t) in each of the four samples.

6: For each sample $s \in \{F, R, I, R^{\mathrm{IPW}}\}$, generate bootstrap outcomes on the original estimation support:

$$Y_{it}^{s,*(b)} = \widehat{Y}_{it}^{s} + \widehat{u}_{it}^{s}\, g_{it}^{(b)}.$$

7: Re-estimate the four regressions on $(Y^{s,*(b)}, X, F, W)$ with exactly the same fixed effects and weighting scheme, obtaining bootstrap slopes $\widehat{\theta}_F^{*(b)}$, $\widehat{\theta}_R^{*(b)}$, $\widehat{\theta}_I^{*(b)}$, and $\widehat{\theta}_R^{\mathrm{IPW},*(b)}$ (or their event-time/staggered analogues).

8: Map these into a bootstrap draw of the channels:

$$\Theta^{*(b)} = f\left(\widehat{\theta}_F^{*(b)}, \widehat{\theta}_R^{*(b)}, \widehat{\theta}_I^{*(b)}, \widehat{\theta}_R^{\mathrm{IPW},*(b)}\right),$$

applied either to the post-shock coefficients, to the whole set $\{\widehat{\theta}_{\cdot,k}^{*(b)}\}_{k\in\mathcal{K}}$, or to the staggered aggregates.

9: **end for**

10: **Inference:** for each channel (and, in the dynamic case, for each k), compute:

- the bootstrap standard error as the sample standard deviation of $\{\Theta^{*(b)}\}_{b=1}^{B}$;
- percentile (or basic/reversed-percentile) confidence intervals as the empirical $\alpha/2$ and $1-\alpha/2$ quantiles of the bootstrap distribution.