

Bank Climate-Disclosure Mandates and Cross-Border Credit: Evidence from U.S. Corporate Loan Contracts^{*}

José Alves[†]

João Estevão[‡]

July 2026

Abstract

Mandatory climate-risk disclosure is central to the European Union's effort to align bank lending with the climate transition. We ask whether the rule changes cross-border corporate credit. The setting is international: the rule is European, the borrowers are U.S. public operating companies, and global banks connect the two. Using U.S. credit agreements filed with the Securities and Exchange Commission around the introduction of CRR Article 449a Pillar 3 ESG disclosure, we parse 1,002,805 filing-form records over 2018-2025, scan 2,000 operating-company filings, and retain 129 credit-agreement exhibits from 85 borrowers, including 34 in high-emission sectors and 48 from 2023 onward. Lender language appears in 94.6 percent of retained contracts and administrative-agent language in 81.4 percent. Validation yields 364 document-bank rows, 39 observed bank parents, and 12 Article 449a-treated EU/EEA parents matched to the EBA 2025 Transparency Exercise; we extract amounts, pricing, maturities, covenants, security, and sustainability-linked terms for 197 facilities. Aggregate comparisons mislead: an association between treated banks and post-2023 high-emission contracts positive under bank and year fixed effects turns negative once borrower-year and bank-year variation are absorbed. At the borrower-year level, the clearest movement is lender substitution: non-treated parent counts rise in post-2023 high-emission borrower-years while treated-parent presence barely moves. Facility-term estimates are imprecise but point to relevant margins for treated-bank participation in facility size, pricing, maturity, and sustainability-linked language. The evidence points to cross-border regulatory transmission through bank relationships, mainly via lender composition rather than a clean reduction in brown credit.

Keywords: bank lending; climate-risk disclosure; carbon risk; syndicated loans; ESG regulation; cross-border credit allocation; EU Taxonomy; Pillar 3.

JEL classification: G21, G28, G32, Q54, F36.

^{*} This work was supported by national funds through FCT - Fundação para a Ciência e a Tecnologia, I.P., in the framework of the unit UID/0097/2025 and UID/06522/2025, <https://doi.org/10.54499/UID/06522/2025>. The opinions expressed herein are those of the authors and do not necessarily reflect those of the authors' employers. Any remaining errors are the authors' sole responsibility.

[†] ISEG Research, ISEG - Lisbon School of Economics and Management, Universidade de Lisboa, Lisbon, Portugal; UECE - Research Unit on Complexity and Economics; CESifo (Center for Economic Studies and Ifo Institute). jalves@iseg.ulisboa.pt

[‡] CEGIST, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal; Advance/ISEG Research. Corresponding author: joaojestevao@tecnico.ulisboa.pt

1. Introduction

Climate policy has reached banks in a deceptively quiet form: disclosure. The rules do not order banks to abandon carbon-intensive borrowers. They ask banks to name, classify, measure, and report the climate-risk content of the assets they choose to hold. In banking, though, information is never a decorative appendix to the balance sheet. It enters credit committees, pricing grids, covenant packages, monitoring decisions, and the renewal of relationships that may have lasted for years. A disclosure mandate can therefore travel far beyond the Pillar 3 table in which it first appears. If it changes lending decisions, a rule written for banks reshapes the financing conditions of firms outside the regulator's jurisdiction. If it does not, disclosure makes bank balance sheets look more climate-aware while leaving the underlying flow of brown credit nearly unchanged.

The European Union offers an unusually clear setting for deciding between these views. Article 449a of the Capital Requirements Regulation requires large EU/EEA credit institutions to disclose detailed environmental, social, and governance risks in their Pillar 3 reports. Together with the EU Taxonomy and the Green Asset Ratio, the rule makes exposures to carbon-intensive counterparties visible to supervisors, investors, and the public. But the loan books of European banks do not stop at the edge of the single market. EU banking groups lend to U.S. corporations through syndicated and bilateral credit agreements, including borrowers in energy, utilities, chemicals, transportation, and metals, sectors in which transition risk is hard to ignore. The central question is therefore not whether European banks can produce greener disclosure tables but whether mandatory climate-risk disclosure changes who supplies credit to carbon-intensive foreign borrowers, and on what contractual terms.

Theory gives no reason to expect one clean response. Disclosure may support discipline: regulated banks stay in relationships with brown borrowers but charge higher spreads, shorten maturities, strengthen covenants, or tie pricing to transition performance. The bank does not exit; it uses the loan contract as a monitoring and incentive device. Disclosure may instead encourage retreat. A bank that wants to reduce visible carbon exposure can participate less often in brown loans, leaving the borrower to seek credit elsewhere. That response cleans the regulated balance sheet but leaves the real economy untouched if non-EU lenders replace the departing bank. A third possibility is thinner still: the contract acquires climate language while prices, quantities, covenants, lender composition, and emissions barely move. These outcomes can look almost identical in aggregate data. They mean very different things for climate policy.

We provide evidence on these possibilities by studying the contract through which international bank credit is actually supplied. The paper builds and validates a public panel of U.S. corporate credit agreements filed with the Securities and Exchange Commission, maps lenders and agents to consolidated bank parents, assigns Article 449a treatment from legal and supervisory sources, and extracts facility-level terms from the contract text. The main comparison is a triple difference: whether the lender belongs to an EU/EEA banking group subject to Article 449a, whether the credit agreement is signed after the disclosure regime becomes binding, and whether the borrower operates in a high-emission sector. We link that treatment margin to bank participation, facility amounts, pricing-spread language, maturity, covenants, security, sustainability-linked terms, and borrower-year lender composition.

The central result is not a single clean reduction in brown credit. In aggregate bank-parent comparisons, treated banks do not look less present in post-2023 high-emission contracts. Once the comparison is made within borrower-years and bank-years, however, treated-bank presence falls relative to the same borrowers' non-treated lenders. At the borrower-year level, the largest movement is on the substitution margin: non-treated bank-parent counts rise by 3.76 in post-2023 high-emission borrower-years, while treated-parent counts rise by only 0.40 and treated-parent shares move by less than one percentage point. The facility evidence is less precise, but it shows that the right contractual margins are measurable: facility size, spread language, maturity, covenants, security, and sustainability-linked clauses can all be observed in filed agreements. Together, the results point to a cross-border credit channel in which disclosure pressure changes lender composition more clearly than it contracts brown credit outright.

The comparison in the main bank panel is deliberately tight. We compare treated and non-treated lenders within the same borrower-year, so borrower-year fixed effects absorb credit demand, credit quality, investment opportunities, and borrower-level climate news. Bank-year fixed effects absorb lender funding costs, capital positions, country-year shocks, balance-sheet constraints, and time-varying risk appetite. The causal claim is correspondingly disciplined: within the public-contract sample, Article 449a exposure is associated with lower treated-bank relative participation in the post-2023 high-emission cell once borrower demand and bank-year conditions are held fixed. That is evidence of cross-border regulatory transmission through international bank credit. It is not, by itself, evidence that emissions have fallen or that brown borrowers have lost total access to finance. The design follows the credit-supply logic of Khwaja and Mian (2008), Jiménez

et al. (2014), Chodorow-Reich (2014), and Ivashina and Scharfstein (2010), and brings that logic to a climate-regulatory shock that crosses borders through bank balance sheets.

The question belongs naturally to international finance. The policy is written in Europe, the borrowers are U.S. corporations, and the transmission mechanism is the international banking system rather than trade flows, equity portfolios, or domestic environmental regulation. If European disclosure rules change which banks lend to carbon-intensive U.S. firms, prudential transparency has consequences beyond the jurisdiction that wrote the rule. If non-treated lenders replace the treated banks, the rule may clean regulated balance sheets while leaving real credit to brown activity largely intact. We therefore study Article 449a not as a narrow reporting requirement but as a test of how climate-financial regulation travels through cross-border credit markets.

A practical contribution of the paper is that the evidence is public from the start. Credit agreements come from SEC exhibits attached to Forms 8-K, 10-K, 10-Q, 20-F, and 6-K. The contracts identify borrowers, lenders, administrative agents, facility amounts, maturities, pricing grids, covenants, collateral provisions, and, in some cases, sustainability-linked terms. We parse 1,002,805 filing-form records, scan 2,000 operating-company filings, retain 129 credit-agreement exhibits, validate 364 document-bank rows, map them to 39 observed bank parents, match 12 treated parents to the EBA transparency metadata, and extract 197 facility rows. This architecture matters because many climate-finance questions are still studied through proprietary loan, ESG, and emissions data. Public data do not remove measurement judgment, but they leave every judgment open to inspection.

The paper contributes in five ways. First, it places banks, rather than firms or investors alone, at the center of the climate-finance mechanism: the treatment is a bank disclosure mandate, and the outcomes are bank participation, loan pricing, maturity, covenant intensity, sustainability-linked clauses, lender substitution, and borrower behaviour. Second, it treats climate regulation as an international credit shock: an EU rule matters for U.S. firms if regulated banks change their lending or if other lenders replace them. Third, it distinguishes real transition from balance-sheet cleaning and cross-border leakage, three mechanisms that policy debates often conflate. Fourth, it shows that public SEC filings can support a fully documented credit-agreement archive, built from full-year sampling, preserved source exhibits, and an explicit false-positive audit. Fifth, it turns document observability itself into an empirical object, so the reader can see where public contracts identify bank roles cleanly and where contract structure requires caution.

The rest of the paper proceeds as follows. Section 2 situates the argument in the literatures on climate risk, disclosure regulation, bank lending, credit-supply identification, and loan-contract design. Section 3 develops the conceptual framework and derives five hypotheses. Section 4 presents the triple-difference design, contract-term specifications, borrower-outcome tests, and lender-substitution tests. Section 5 describes the construction of the SEC credit-agreement archive and the stylized facts. Section 6 reports the document-level analysis that validates the contract archive. Section 7 adds filing-level selection tests, exact small-cell checks, influence diagnostics, functional-form checks, false-positive sensitivity, and randomization inference. Section 8 estimates the validated bank-parent schedule, the legal Article 449a crosswalk, the facility-term panel, and the borrower-year substitution tests. Section 9 sets out broader robustness checks, and Section 10 concludes with policy implications.

2. Literature Review

The climate-finance literature is now broad enough that the useful question is no longer whether climate risk matters. It does. Investors observe it, firms are valued differently when exposed to it, and several asset classes price at least part of it. For this paper, the relevant question is narrower and more institutional. Banks do not simply hold claims in a portfolio. They originate loans, screen borrowers, set spreads, write covenants, monitor firms, and decide whether old relationships remain worth renewing. A bank disclosure mandate must therefore be judged not only by whether climate risk becomes more visible, but by whether that visibility changes the allocation and governance of credit.

The prudential motivation for looking inside bank balance sheets is clear. Battiston et al. (2017) show that financial institutions are exposed, directly and through network linkages, to sectors that would be vulnerable in a climate-policy transition. Campiglio et al. (2018) bring that concern to central banks and supervisors, where the long horizon of climate risk sits uneasily beside the shorter horizon of contracts, capital rules, and supervisory review. Dafermos et al. (2018) make the macro-financial feedback explicit: climate damages weaken profits, raise default risk, depress asset values, and feed back into credit supply. These papers explain why regulators want banks to disclose carbon-intensive exposures. They leave open the harder question: whether disclosure changes the lending itself or only the language and tables through which lending is reported.

The asset-pricing evidence makes the financial relevance of transition risk hard to dismiss, but it does not settle the banking mechanism. Bolton and Kacperczyk (2021)

document a carbon premium in U.S. equity returns, and Bolton and Kacperczyk (2023) show that transition risk is priced across countries, especially where climate policy is stringent and investor attention strong. Ilhan et al. (2021) find that option markets charge for protection against carbon-related downside risk, while Engle et al. (2020) build hedge portfolios from climate-news exposure. Sautner et al. (2023) add a forward-looking firm-level measure based on climate language in earnings calls. Corporate debt points the same way: Seltzer et al. (2022) show that climate regulatory risk is reflected in bond ratings and yield spreads. This literature establishes a strong prior that markets care. It cannot tell us whether banks respond by repricing loans, tightening covenants, changing syndicate participation, or walking away from particular borrowers.

Work on physical climate risk teaches a similar lesson through different markets. Hong et al. (2019) show that equity prices underreact to drought trends in food industries. Bernstein et al. (2019) find discounts for homes exposed to projected sea-level rise, while Murfin and Spiegel (2020) find little evidence of capitalization once vertical land motion is used for identification. Giglio et al. (2021) use real-estate markets to infer very long-run discount rates, and Painter (2020) shows that sea-level exposure raises long-maturity municipal bond yields. These papers do not simply disagree; together they show that climate risk is priced unevenly, depending on salience, horizon, contract maturity, and who is on the margin of the trade. In bank lending, those same forces surface less as one clean market price than as changes in relationship lending, monitoring, and contractual control.

The sustainable-investing literature adds a warning about interpretation. Even when climate or ESG information matters, the observed price or flow can be ambiguous. Pástor et al. (2021) show that green assets may earn lower expected returns when investors value sustainability, even though unexpected shifts in preferences can generate green outperformance. Hartzmark and Sussman (2019) show how quickly fund flows react when sustainability rankings become salient, and Krueger et al. (2020) find that institutional investors see regulatory climate risk as especially important and often prefer engagement to divestment. Measurement remains a deep difficulty: Berg et al. (2022) document large disagreement across ESG ratings, much of it rooted in differences of scope and measurement. That evidence is one reason we measure borrower exposure with emissions, high-emission sectors, and taxonomy-relevant transition exposure rather than a single commercial ESG score.

The disclosure literature is the closest bridge to Article 449a, and it also shows why the banking setting deserves separate treatment. Christensen et al. (2021) argue that mandatory sustainability reporting can improve information while producing real effects and unintended consequences. Grewal et al. (2019) find negative equity-market reactions around the EU Non-Financial Reporting Directive, milder for firms that already had stronger non-financial disclosure and performance. Griffin et al. (2017) show that greenhouse-gas information is relevant for valuation even when firms do not disclose emissions directly, and Matsumura et al. (2014) document valuation penalties for both emissions and non-disclosure. Disclosure, in short, matters for valuation and governance. Article 449a poses a more intermediation-centred question: the entity that discloses is the bank, but the real adjustment may occur in the credit supplied to the borrowers on its balance sheet.

The bank-lending literature brings that adjustment into sharper focus. Goss and Roberts (2011) show that firms with weaker CSR profiles pay higher loan spreads, particularly on unsecured loans. Chava (2014) finds that environmental concerns raise both equity and bank-credit costs and reduce lender participation. Cheung et al. (2020) show that the loan pricing of CSR varies across national cultures, so institutions and social preferences shape how lenders act on this information. Houston and Shan (2022) go further: banks match with borrowers of similar ESG profiles and influence borrowers' later ESG behaviour, especially when firms depend heavily on bank finance. Earlier international evidence in Scholtens (2009) likewise shows meaningful cross-country variation in bank CSR practices. Banks, in other words, are not passive conduits for ESG signals. Through pricing, monitoring, and relationship renewal, they help decide which borrower characteristics are penalized, accommodated, or rewarded.

More recent work comes closer to the transition-risk margin in corporate loan markets. Delis et al. (2018) find that banks began pricing stranded-asset risk in fossil-fuel loans after the Paris Agreement. Ehlers et al. (2022) show that syndicated-loan markets price direct carbon emissions after Paris, though the premium is not confined to fossil-fuel firms and is not systematically stronger among green banks. Bruno and Lombini (2023) document non-linear, country-specific links between emissions and syndicated-loan terms, strongest where environmental policy is stringent. The closest precedent is Reghezza et al. (2022), who use euro-area credit-register data to show that banks reallocated lending away from more polluting firms after Paris, especially when well capitalized. Our setting differs in ways that matter. The shock is a bank disclosure

mandate rather than a global climate agreement; the borrowers are U.S. firms rather than domestic euro-area firms; and the key empirical margin is whether non-EU lenders replace EU banks when regulated banks reduce brown exposure.

That replacement margin is precisely why the credit-supply literature matters. Khwaja and Mian (2008) show that comparing banks' lending to the same firm separates bank supply shocks from borrower demand. Jiménez et al. (2014) use firm-time and bank-time fixed effects to study monetary policy and bank risk taking in credit-register data. Chodorow-Reich (2014) shows that bank health affected firm employment during the financial crisis, and Ivashina and Scharfstein (2010) document the 2008 contraction in bank lending and the role of credit-line drawdowns. The discipline of this work applies directly here. If EU and non-EU lenders to the same borrower move differently after Article 449a, borrower demand is a weak explanation. If bank-year fixed effects absorb time-varying lender conditions, the remaining variation speaks to how banks exposed to the disclosure mandate treat brown borrowers differently.

Loan contracts are where that treatment should become visible if the mandate changes behaviour rather than portfolio labels. Sufi (2007) shows that information asymmetry shapes syndicate structure and the allocation of lending roles. Chava and Roberts (2008) show that covenant violations transfer control rights to lenders and reduce investment. Nini et al. (2009) document that explicit investment restrictions in private credit agreements constrain borrower policy and precede changes in operating performance. Murfin (2012) shows that shocks to lenders tighten the contracts written for later borrowers. For a climate-disclosure mandate, the lender list is therefore only part of the evidence. Engagement may appear in spreads, maturity, collateral, covenants, or sustainability-linked pricing grids. Exit appears instead as the treated bank's disappearance from the syndicate.

The green-bond literature warns against overinterpreting labels. Flammer (2021) finds that corporate green bonds signal environmental commitment and are followed by improvements in environmental performance and ownership. Tang and Zhang (2020) find positive announcement effects but limited evidence of a durable reduction in the cost of debt, while Zerbib (2019) estimates a small green-bond premium. The lesson is not that sustainable-finance labels are empty but that they persuade only when tied to prices, quantities, contractual terms, or real outcomes. The same caution applies to sustainability-linked loan language: greener or transition wording matters only if it travels with changes in credit allocation, contractual discipline, borrower emissions, or lender substitution.

The gap this literature leaves is narrow but economically important. We know that climate risk can be priced, that disclosure can move markets, that banks condition credit on environmental characteristics, and that credit-supply shocks can be isolated with within-borrower designs. We know less about whether a disclosure rule imposed on banks changes cross-border corporate credit. This paper addresses that question by combining a bank-level regulatory shock, public loan-contract data, objective borrower brownness measures, harmonized lender identities, and a triple-difference design that can distinguish real transition from divestment and leakage.

3. Theoretical Framework

The framework starts from a simple intuition. A climate-disclosure mandate matters for credit only if it changes the economics of keeping a loan on a bank's books, or if it changes the contract a bank is willing to offer. Consider a loan market with EU-regulated banks, $b \in \mathcal{B}^{EU}$, and non-EU banks, $b \in \mathcal{B}^N$. Borrowers differ in transition exposure. Brown firms, $f \in \mathcal{F}^B$, have higher emissions intensity $e_f \geq 0$ than green firms, $f \in \mathcal{F}^G$. A loan between bank b and borrower f in year t is summarized by $c_{bft} = (s_{bft}, m_{bft}, k_{bft}, q_{bft}, \sigma_{bft})$: the spread s , maturity m , covenant intensity k , collateral intensity q , and an indicator σ for sustainability-linked terms. The bank first decides whether to join the loan and, if it does, what contract to write. Separating those two decisions is important. A disclosure rule may make credit more expensive, shorter, or more restrictive while the relationship survives. It may also push the borrower toward a different lender, leaving the underlying activity much the same.

Let the expected payoff from a commitment L_{bft} be

$$\pi_{bft}(c_{bft}, L_{bft}) = [s_{bft} - r_{bt} - \mathbb{E}_t(\ell_{bft}) - \kappa_{bt}(e_f)]L_{bft} + R_{bf} - \phi(c_{bft}), \quad (1)$$

where r_{bt} is the bank's funding and balance-sheet cost, $\mathbb{E}_t(\ell_{bft})$ is expected credit loss, R_{bf} is the private value of the lending relationship, and $\phi(c_{bft})$ collects the practical costs of writing and managing the contract: screening, monitoring, collateral administration, covenant enforcement, and verification of sustainability-linked provisions. The object doing the main work is $\kappa_{bt}(e_f)$, the bank-specific shadow cost of carrying a climate-exposed loan. It is kept separate from expected default loss because it represents a different burden: the extra cost of holding a brown exposure once that

exposure becomes visible to supervisors, investors, wholesale funders, or the public. The cost should rise with borrower emissions intensity, but the same borrower need not be equally costly for every lender.

The participation decision follows from the same payoff. A bank lends only if the best feasible contract delivers at least as much value as the next best use of its balance-sheet capacity:

$$P_{bft} = \mathbb{1}\left\{\max_{c_{bft}, L_{bft}} \pi_{bft}(c_{bft}, L_{bft}) \geq \bar{\pi}_{bt}\right\}, \quad (2)$$

where P_{bft} equals one when the bank joins the loan, and $\bar{\pi}_{bt}$ is the bank-year return available from deploying the same scarce risk-bearing capacity elsewhere. Equation (2) describes the extensive margin. Participation falls only when the disclosure shock pushes the best attainable payoff below the outside option. If the bank can preserve the relationship surplus by repricing the loan, shortening maturity, adding covenants, requiring collateral, or writing transition-linked clauses, the same regulation changes the contract rather than the lender list.

Article 449a enters the model by raising $\kappa_{bt}(e_f)$ for treated EU banks after the disclosure regime becomes effective. Let D_t denote the post-2023 period. The shadow cost can be written as

$$\kappa_{bt}(e_f) = \kappa_0(e_f) + \mathbb{1}\{b \in \mathcal{B}^{EU}\} D_t [\kappa_b^R(e_f) + \kappa_b^P(e_f) + \kappa_b^C(e_f)], \quad (3)$$

where $\kappa_0(e_f)$ is the baseline transition-risk cost that lenders may face even without the mandate. The bracketed terms are the costs that become more salient for treated EU banks once disclosure begins. The regulatory component, κ_b^R , captures supervisory attention to Green Asset Ratios and transition-risk templates. The public component, κ_b^P , captures the reputational cost of reporting brown exposures in a comparable format. The capital and funding component, κ_b^C , allows markets, supervisors, and internal capital allocation to attach a higher price to those exposures over time. Evidence in Krueger et al. (2020) supports the relevance of regulatory and capital-cost channels, while Grewal et al. (2019) and Matsumura et al. (2014) motivate the role of public scrutiny. For a non-EU bank, the

mandate-specific bracketed term is absent. The same brown borrower can therefore be a manageable exposure for one lender and an unattractive exposure for another.

The same logic generates the model's heterogeneity. The disclosure shock is larger when the borrower is more carbon intensive, and larger for banks that are more exposed to climate-sensitive assets, more visible to outside audiences, or less able to absorb the added cost through capital and funding capacity:

$$\frac{\partial \tau_b(e_f)}{\partial e_f} > 0, \frac{\partial \tau_b(e_f)}{\partial GAR_b} < 0, \frac{\partial \tau_b(e_f)}{\partial CarbonExposure_b} > 0, \frac{\partial \tau_b(e_f)}{\partial Capital_b} < 0, \frac{\partial \tau_b(e_f)}{\partial Scrutiny_b} > 0, \quad (4)$$

where $\tau_b(e_f) \equiv \kappa_b^R(e_f) + \kappa_b^P(e_f) + \kappa_b^C(e_f)$. Equation (4) makes clear that the mandate is not economically uniform, even if the legal rule is common. It should bind more strongly for banks with lower Green Asset Ratios, larger carbon-related exposures, thinner capital buffers, or greater public visibility. It should also bind more strongly when the borrower has emissions, activities, or transition plans that are difficult to justify. The empirical implication is that average effects are likely to hide meaningful variation on both sides of the lending relationship.

If the bank stays in the relationship, the adjustment has to appear in the contract. Let c_{bft}^* denote the contract that maximizes (1), taking the participation constraint into account. A higher disclosure-driven shadow cost implies the following comparative statics:

$$\frac{\partial s_{bft}^*}{\partial \tau_b(e_f)} > 0, \quad \frac{\partial m_{bft}^*}{\partial \tau_b(e_f)} < 0, \quad \frac{\partial k_{bft}^*}{\partial \tau_b(e_f)} > 0, \quad \frac{\partial q_{bft}^*}{\partial \tau_b(e_f)} > 0, \quad \frac{\partial \Pr(\sigma_{bft}^*=1)}{\partial \tau_b(e_f)} > 0. \quad (5)$$

These signs read naturally in contractual terms. A lender that keeps financing a brown borrower must either be paid for the additional shadow cost or write a contract that makes the exposure easier to monitor, discipline, or exit. Higher spreads compensate the lender. Shorter maturities cut the time over which transition risk can build. Stronger covenants and collateral give the bank control if the borrower weakens. Sustainability-linked terms tie the cost of credit to observable transition performance. In this version of the mechanism, the mandate works through engagement rather than retreat: the bank stays in the relationship, but the loan becomes more demanding.

Exit is the case in which contractual adjustment cannot preserve the surplus. From (2), an increase in $\tau_b(e_f)$ lowers the probability that a treated EU bank participates in a loan to a brown borrower:

$$\frac{\partial \Pr(P_{bft}=1)}{\partial (\mathbb{1}\{b \in \mathcal{B}^{EU}\} D_t e_f)} < 0. \quad (6)$$

This is the credit-reallocation prediction. It stops well short of claiming that EU banks abandon brown borrowers wholesale. What it does claim is that, holding borrower demand and bank-year conditions fixed, EU participation falls most where borrower transition exposure makes the disclosed position hardest to carry.

The economic meaning of reallocation depends on what happens next. Let $L_{ft}^{EU,B}$ be credit supplied by EU banks to brown borrower f , $L_{ft}^{N,B}$ credit supplied by non-EU banks, and L_{ft}^B total credit to the same brown borrower. The accounting identity is

$$\Delta L_{ft}^B = \Delta L_{ft}^{EU,B} + \Delta L_{ft}^{N,B}. \quad (7)$$

If EU lending falls and no one replaces it, then $\Delta L_{ft}^{EU,B} < 0$, $\Delta L_{ft}^{N,B} \approx 0$, and total brown credit falls. That is divestment without substitution. If EU lending falls but non-EU lending rises, then $\Delta L_{ft}^{EU,B} < 0$, $\Delta L_{ft}^{N,B} > 0$, and ΔL_{ft}^B may remain close to zero. That is leakage. The distinction is central: in the first case, the borrower faces a genuine credit contraction; in the second, the EU banking sector looks cleaner, but the borrower may continue financing the same carbon-intensive activity through someone else's balance sheet.

Borrower behaviour provides the final test of whether the mandate reaches the real economy. Let $\Delta e_{f,t+h}$ denote the change in borrower emissions intensity over a horizon $h > 0$. Let $Disc_{ft}$ summarize the discipline embedded in loan terms, including spreads, maturity, covenants, collateral, and sustainability-linked provisions. The real-transition channel can be written as

$$\Delta e_{f,t+h} = \psi_1 \Delta Disc_{ft} + \psi_2 \Delta L_{ft}^B + u_{f,t+h}, \quad \psi_1 < 0, \quad \psi_2 > 0. \quad (8)$$

The restriction $\psi_1 < 0$ says that stronger contractual discipline should lower future emissions intensity when it pushes the borrower toward transition. The restriction $\psi_2 > 0$ says that more total credit to brown activity relaxes financing constraints and, all else equal, should make emissions reductions less likely. The empirical distinction is concrete. If EU banks tighten loan terms and borrower emissions fall afterward, the evidence points toward real transition. If EU banks leave, non-EU lenders replace them, and emissions do not move, the evidence points toward leakage.

A weaker possibility also exists. The language of transition may become more visible while the economics of the loan remain largely unchanged. In the notation above, this cosmetic-compliance case appears as

$$\Delta \Pr(\sigma_{bft} = 1) > 0, \quad \Delta s_{bft} \approx \Delta m_{bft} \approx \Delta k_{bft} \approx \Delta q_{bft} \approx \Delta L_{ft}^B \approx \Delta e_{f,t+h} \approx 0. \quad (9)$$

This case matters because sustainability-linked language can look like progress even when credit allocation, contractual discipline, and borrower emissions are unchanged. We therefore treat such language as economically meaningful only when it comes with movement in prices, quantities, control rights, lender composition, or borrower outcomes.

The empirical hypotheses follow directly from the model. H1 is the participation prediction in (6): after 2023, EU-regulated banks reduce participation in loans to high-transition-risk U.S. borrowers relative to non-EU banks' lending to the same borrowers. H2 is the contract-term prediction in (5): conditional on lending, treated EU banks should write tighter contracts for brown borrowers. H3 combines (7) and (8): real transition requires discipline or credit contraction to be followed by lower emissions, whereas leakage implies replacement by non-EU lenders and little emissions response. H4 is the bank-heterogeneity prediction in (4). H5 is the borrower-heterogeneity prediction that the response should be strongest when e_f and transition salience are highest.

The framework exists to prevent one specific mistake: treating cleaner bank disclosures, by themselves, as evidence of cleaner real activity. The same mandate can produce engagement, divestment, leakage, or cosmetic adjustment. Evidence for H1 and H2, followed by lower borrower emissions under H3, is consistent with real transition through bank engagement. A strong H1 effect with no compensating non-EU expansion looks like divestment without substitution. A strong H1 effect offset by non-EU expansion, with no emissions response, points to leakage. The empirical design is

therefore organized around one question: does the disclosure shock change the allocation and governance of credit, or does it only change which lender holds the same brown exposure?

4. Empirical Strategy

The empirical strategy begins where the theoretical framework says the regulatory pressure should first become visible: inside the lending relationship. If Article 449a makes brown credit more costly to carry, or more costly to explain, for EU-regulated banks, then the most immediate margin is not an aggregate loan volume but the identity and role of the banks willing to sit inside a borrower syndicate. Let Y_{bft} denote bank b 's participation in borrower f 's credit agreement in year t . We measure participation in several ways, because SEC credit agreements do not disclose syndicate information with the same detail in every filing. The preferred measure is an indicator equal to one when bank b is named as a lender in Lenders_{ft} . We also record whether the bank serves as administrative agent or arranger, its commitment share when a lender schedule is available, and the number of treated-bank seats in the syndicate. These margins capture different parts of the lending relationship. Agent and arranger roles identify banks with real influence over the loan; named-lender indicators capture the breadth of participation; and commitment shares, when disclosed, come closest to the bank's balance-sheet exposure. The baseline specification is

$$Y_{bft} = \beta (\text{EU}_b \times \text{Post}_t \times \text{Brown}_f) + \alpha_{bt} + \alpha_{ft} + \varepsilon_{bft}, \quad (10)$$

where EU_b is an indicator for an EU/EEA banking group subject to CRR Article 449a, Post_t equals one after the disclosure regime becomes effective, and Brown_f measures the borrower's transition exposure. The baseline post period is $t \geq 2023$, with 2022 and 2024 cutoffs used in robustness exercises to allow for anticipation and gradual implementation. The bank-year fixed effect, α_{bt} , absorbs the many reasons a bank might lend more or less in a given year: capital, liquidity, risk appetite, home-country macroeconomic conditions, supervisory pressure, and changes in the bank's overall U.S. lending strategy. The borrower-year fixed effect, α_{ft} , absorbs the corresponding borrower-side forces in that same year: credit demand, default risk, investment opportunities, refinancing needs, commodity-price exposure, and sector shocks as they are felt by that firm. Standard errors

are clustered by bank and by borrower, allowing lending decisions to be correlated within a bank’s portfolio and within a borrower’s lending relationships over time.

The identifying variation in equation (10) is narrow. The coefficient β compares EU and non-EU banks that lend to the same borrower in the same year, and asks whether the EU bank’s participation falls differentially when that borrower is brown after the Article 449a disclosure shock. Under H1, $\beta < 0$: EU-regulated banks should become less willing to participate in high-transition-risk borrowers relative to otherwise comparable non-EU lenders. The fixed effects impose much of the discipline in this comparison. The bank-year fixed effect absorbs the bank-level exposure to the post period, including $EU_b \times Post_t$; the borrower-year fixed effect absorbs the borrower-side post-period exposure, including $Brown_f \times Post_t$. Time-invariant differences across banks and borrowers are likewise swept out by the relevant comparisons. What remains is the part of the relationship where all three features meet: the treated lender, the treated period, and the borrower’s transition exposure.

The appeal of this design is that it reduces a cross-border regulatory question to a comparison within the same credit market. Many U.S. borrowers borrow from syndicates that include both EU and non-EU global banks. The same firm can therefore be matched, in the same year and often inside the same credit agreement, to lenders facing different Pillar 3 climate-disclosure obligations. This is the bank-side analogue of the within-borrower logic in Khwaja and Mian (2008): borrower demand is held fixed by comparing lenders that face the same firm at the same time. The bank-year fixed effect completes the symmetry by holding fixed each bank’s annual lending environment. The estimate therefore does not simply pick up brown firms borrowing less, EU banks cutting back across the board, or a borrower-specific deterioration that all lenders observe together.

The main identifying assumption is parallel trends in this narrow comparison. Absent Article 449a, EU and non-EU banks should not have changed their relative participation in brown-borrower credit at the same time, and in the same direction, as the new disclosure regime. We examine that condition with a dynamic event-study specification,

$$Y_{bft} = \sum_{\tau \neq -1} \beta_{\tau} (EU_b \times \mathbb{1}\{t = 2023 + \tau\} \times Brown_f) + \alpha_{bt} + \alpha_{ft} + \varepsilon_{bft}, \quad (11)$$

normalizing the year immediately before treatment to zero. The pre-treatment coefficients, $\{\beta_{-k}\}_{k \geq 2}$, are the key diagnostic. If EU and non-EU banks were not already

parting ways in brown credit before the mandate, these coefficients should be small and statistically indistinguishable from zero. The post-treatment coefficients then show when the response arrives. A break concentrated around 2023 would fit the disclosure channel most closely; a movement that begins well before 2023 would instead suggest anticipation or a broader climate-finance trend that is not specific to Article 449a.

We take two threats seriously. The first is anticipation. The EBA Implementing Technical Standards were available before the first full post-regime observations, and large banks may have started adjusting portfolios before the disclosure templates became formally visible to the market. We therefore re-estimate the baseline around alternative post-period cutoffs in 2022, 2023, and 2024. The second threat is spillover to the control group. Non-EU banks may respond to the same climate politics, investor pressure, or supervisory convergence that shaped EU banks; U.S. climate-disclosure rulemaking is the obvious example. If these forces push non-EU banks in the same direction as EU banks, the estimated treatment effect is attenuated. If non-EU banks instead step into relationships that EU banks leave, the treatment effect coexists with lender substitution. Non-EU lender behaviour is therefore not only part of the control comparison but an outcome to be studied in its own right.

Participation is only one way a bank can adjust. A bank that stays in the relationship can still demand a higher spread, a shorter maturity, more collateral, tighter covenants, or language that ties pricing to sustainability performance. For facility ℓ to borrower f at time t , let $T_{\ell ft}$ denote a contract term such as the all-in drawn spread, maturity, secured status, covenant intensity, or the presence of sustainability-linked pricing language. The contract-term specification is

$$T_{\ell ft} = \gamma (\text{EUExp}_{\ell} \times \text{Post}_t \times \text{Brown}_f) + \alpha_{ft} + \alpha_{jt} + X'_{\ell} \delta + u_{\ell ft}, \quad (12)$$

where EUExp_{ℓ} is the share of EU-regulated banks in the facility syndicate. When commitment shares are reported, we use them directly; otherwise, we rely on lead-bank, arranger, or agent counts. The borrower-year fixed effect α_{ft} keeps the comparison close to the same borrower at the same point in time, while α_{jt} absorbs industry-year movements at the NAICS-3 level. The vector X_{ℓ} includes facility size, loan type, benchmark rate, loan purpose, and other contract features that shape pricing and maturity. H2 predicts $\gamma > 0$ for spreads, $\gamma < 0$ for maturity, and $\gamma > 0$ for secured status, covenant

intensity, and the probability that the loan contains sustainability-linked terms. EU banks do not dictate every term in a syndicated loan. But when EU exposure carries more weight inside the syndicate, brown borrowers should face contracts that ask for more compensation, more control, or more explicit transition discipline once the disclosure regime begins. This is the contract-design logic in Murfin (2012), Sufi (2007), and Nini et al. (2009), where loan terms reveal how lenders allocate control rights and protect themselves against downside risk.

The next question is whether these credit-market adjustments show up in the borrower's real behaviour. A cleaner EU bank portfolio is not, by itself, evidence of transition. If the policy channel is real, borrowers more exposed to EU banks before the mandate should later display lower emissions intensity, stronger transition investment, or more concrete transition language in their public filings. We therefore estimate borrower-level outcome regressions that relate future emissions, emissions intensity, transition-investment language, and green-capex indicators to pre-treatment EU-bank exposure. Let $O_{f,t+1}$ denote a borrower outcome measured after the lending decision:

$$O_{f,t+1} = \theta \left(\text{Pre2023EUExp}_f \times \text{Post}_t \times \text{Brown}_f \right) + \alpha_f + \alpha_{jt} + Z'_{ft}\phi + v_{ft}, \quad (13)$$

where Pre2023EUExp_f is the borrower's average exposure to EU-regulated banks over 2018–2022. Using pre-treatment exposure is crucial. It anchors the design in the borrower's inherited lending structure rather than in the lender choices it makes after the regime begins, when those choices may already reflect decarbonization plans. The borrower fixed effect α_f absorbs persistent differences across firms, while the industry-year fixed effect α_{jt} absorbs sectoral shocks to emissions, disclosure, and investment. The control vector Z_{ft} includes time-varying firm characteristics such as size, leverage, profitability, capital expenditure, cash holdings, and credit quality where available. Text-based outcome measures are constructed from 10-K risk-factor and business-description sections using the approach of Sautner et al. (2023). A negative θ for future emissions intensity would be consistent with real transition if it appears together with tighter loan terms or a contraction in total brown credit. A null emissions response, especially when EU credit is replaced by non-EU credit, would instead point toward leakage.

Because a decline in EU lending can carry more than one economic meaning, we estimate lender substitution directly. The central question is whether non-EU banks fill the space left by treated EU banks. For each borrower-year, we estimate

$$\text{NonEUExp}_{ft} = \rho \left(\text{Pre2023EUExp}_f \times \text{Post}_t \times \text{Brown}_f \right) + \alpha_f + \alpha_{jt} + w_{ft}. \quad (14)$$

A positive ρ means that non-EU banks expand precisely among brown borrowers that were more exposed to EU banks before the mandate. In the framework above, that pattern is evidence of leakage. A negative H1 effect combined with a null ρ is closer to divestment without substitution: EU banks leave, and other global lenders do not appear to replace them. A positive ρ with no subsequent emissions decline is the sharpest empirical signature of balance-sheet relocation rather than real transition.

Borrower brownness is measured in several ways because no single public proxy captures transition exposure cleanly, and commercial ESG ratings would import a separate measurement problem into the design. The document-level implementation uses high-emission SIC sectors and separates utilities, whose contracts are institutionally distinctive. The broader bank-level implementation adds facility-level Scope 1 emissions from the EPA Greenhouse Gas Reporting Program, linked to corporate parents and scaled by assets, sales, or output proxies; EU Taxonomy-relevant transition sectors mapped from NACE to NAICS; and text-based fossil-fuel and transition-risk vocabulary from 10-K risk-factor and business-description sections in the spirit of Sautner et al. (2023). The aim is not to force all proxies to tell the same story but to see whether any estimated disclosure effect is tied to measured emissions, sectoral transition exposure, taxonomy exposure, or the way firms describe climate risk in their filings. The preference for objective emissions and sectoral measures over commercial ESG scores follows the concern, documented by Berg et al. (2022), that ESG ratings often disagree about what they measure.

Treatment in the causal panel is assigned at the consolidated bank-parent level. A bank is treated when its parent banking group is subject to EU/EEA Pillar 3 ESG disclosure under CRR Article 449a. The classification draws on EBA Pillar 3 templates, the legal scope of Article 449a, bank Pillar 3 reports, and consolidated parent information. U.S. broker-dealer, lending, and branch affiliates of EU banking groups count as part of the same parent, because the climate-disclosure obligation attaches to the consolidated group rather than to the local legal name that happens to appear in a credit agreement. Non-EU

banks include U.S., U.K., Swiss, Canadian, Japanese, and other global lenders not subject to the Article 449a requirement during the sample period. We implement that assignment for the observed bank parents in the contract archive: recurring lender and administrative-agent names are mapped to consolidated parents, the contract language surrounding each name is checked for a credit role, mention-only rows are excluded, and observed EU/EEA parents are matched to the EBA 2025 Transparency Exercise under the Article 449a legal framework. GLEIF parent linkages and subsidiary-level checks for every legal entity in the commitment schedules would refine the mapping further, but treatment is not assigned by country alone.

The empirical setting consists of U.S. public non-financial firms borrowing from global banks between 2018 and 2025. Loan agreements come from SEC filings, including Forms 8-K, 10-K, 10-Q, 20-F, and 6-K. The start date provides pre-treatment history before the disclosure shock; the post-2023 years capture the first period in which the new Pillar 3 climate templates could plausibly affect lending. Financial firms are excluded from the baseline because their balance sheets, regulation, and emissions profiles make them poor comparators for operating companies. Utilities are reported separately because they are unusually regulated, emissions-intensive, and central to energy-transition policy. The causal panels are built at four levels: bank-borrower-year for participation, facility-year for contract terms, borrower-year for outcomes, and borrower-lender exposure for substitution. The panel contains each ingredient in public form: the SEC exhibit archive, a validated bank-parent risk set, legally sourced Article 449a treatment assignment, extracted facility terms, and borrower-year lender-composition outcomes. The sample-construction logic follows the relationship-based design of Chodorow-Reich (2014), with the unit shifted from the firm-year cell to the bank-borrower-year cell. The central question is not only whether a borrower receives credit but which banks remain willing to provide it after climate disclosure raises the cost of holding brown exposure.

5. Data and Stylized Facts

The data come from public records rather than a commercial loan vendor, because the empirical question is itself about disclosure. If the paper asks whether a public climate-reporting mandate changes bank behaviour, the underlying contracts should be observable, traceable, and open to audit. The starting point is the EDGAR quarterly master index. Across 2018-2025, we read 1,002,805 filing-form records in the forms where material credit contracts are most likely to appear: 8-K, 10-K, 10-Q, 20-F, 6-K,

and their amendments. These records are merged to the SEC submissions archive, which adds borrower CIKs, legal names, SIC classifications, entity type, tickers, exchanges, legal-entity identifiers where available, and state of incorporation. Financial issuers, identified by SIC codes 6000-6999, are removed from the baseline sample: their balance sheets, regulation, and own role as credit intermediaries make them fundamentally different from operating companies.

We scan 250 operating-company filings in each year from 2018 through 2025, for 2,000 filings in total. The 250 filings are selected deterministically across the full filing year rather than taken from the last few filing days, so the panel is not a year-end archive. The extraction is conservative. For every selected filing, we parse the EDGAR document table and record the exhibit type, filename, description, index size, accession number, and URL. Because EDGAR descriptions are often unhelpfully generic, sometimes saying only “Exhibit 10.1,” we read the generic EX-10 documents and keep an exhibit only when its title or body contains credit-contract language such as credit agreement, loan agreement, term loan, revolving credit, credit facility, or loan and security agreement. Employment agreements, leases, share-purchase contracts, securities-purchase agreements, cooperation agreements, stand-alone pledge agreements, ordinary promissory notes, and other adjacent instruments are dropped unless the exhibit clearly identifies a loan or credit agreement. Each retained contract is preserved as filed, together with its source EDGAR record, so the sample can be checked contract by contract rather than taken on trust.

Table 1. Sample construction from public SEC filings to the audited contract panel

<i>Stage</i>	<i>Count</i>	<i>Role in the empirical design</i>
SEC filing-form records parsed, 2018–2025	1,002,805	Universe used to identify candidate public-company filings
Operating-company filings scanned	2,000	Deterministic full-year scan, 250 filings per year
Retained credit-agreement exhibits	129	Final contract layer used in the tables and figures
Unique SEC filing accessions with retained exhibits	85	Filing-level support for the document archive
Unique borrower CIKs	85	Borrower identifiers for future bank–borrower-year construction
High-emission-sector exhibits	34	Transition-risk contract coverage in the audited sample
Post-2023 exhibits	48	Contract coverage after the Article 449a disclosure period begins

Notes: Counts refer to the final SEC credit-agreement extraction described in the text. Financial issuers are excluded from the operating-company scan.

Table 1 summarizes the sample construction in one place. The table is mechanical because sample construction is where a public-data paper either earns or loses the reader’s trust. The final data set is not a black box: it begins with the SEC filing universe, narrows to a deterministic full-year scan of operating-company filings, and ends with a contract archive that can be inspected row by row. That flow also clarifies the scope of the empirical claims: the paper validates a public archive of filed contracts, not the universe

of all corporate loans. The sample is narrow by construction but fully traceable from the SEC filing universe to the contracts used in the regressions.

Table 2 reports the extraction counts by filing year. The final archive contains 129 credit-agreement exhibits drawn from 85 unique SEC filing accessions and 85 unique borrower CIKs. Two features deserve emphasis before the regressions. First, the sample spans both sides of the Article 449a disclosure period: 81 exhibits are filed before 2023 and 48 from 2023 onward. Second, the archive is not made up only of low-carbon borrowers. Thirty-four exhibits, or 26.4% of the sample, are tied to high-emission SIC sectors: oil and gas, utilities, industrial chemicals, paper, metals, mining, transportation, aviation, equipment, and other industrial activities where transition risk is especially salient for lenders. The classification does not treat pharmaceutical preparations as high-emission merely because they fall inside the broad two-digit chemicals block. The panel therefore contains both pre- and post-Article 449a contracts, and transition-risk borrowers remain present in the public archive after the disclosure regime begins.

Table 2. Full-year SEC credit-agreement extraction by filing year

<i>Year</i>	<i>Filings scanned</i>	<i>Matched filings</i>	<i>Exhibits</i>	<i>High-emission</i>	<i>High-emission share</i>
2018	250	10	28	3	10.7%
2019	250	13	14	3	21.4%
2020	250	12	16	8	50.0%
2021	250	5	5	3	60.0%
2022	250	9	18	9	50.0%
2023	250	13	15	2	13.3%
2024	250	7	11	1	9.1%
2025	250	16	22	5	22.7%
Total	2000	85	129	34	26.4%

Notes: The scan covers 250 operating-company filings in each filing year, selected across the full year rather than from year-end filings. Matched filings count unique accession numbers with at least one retained credit-agreement exhibit. High-emission sectors are identified from the SIC classifications described in the text.

Figure 1 makes the same point more directly. The archive is not concentrated in a single year, and the post-2023 contracts are not simply contracts with low-emission firms. The full-year archive is modest in size, but retained contracts appear in every year from 2018 through 2025, and high-emission borrowers remain visible on both sides of Article 449a implementation. That pattern is not evidence of a treatment effect, but it is a useful fact about the setting: after the EU disclosure regime begins, the public SEC archive still contains brown-credit agreements that support a relationship-level comparison between EU and non-EU banks.

Figure 1: SEC credit-agreement exhibits by filing year and transition-risk sector




Figure 2 isolates the transition-risk composition. The high-emission share is plainly uneven: low in 2018, rising through 2022, falling in 2023 and 2024, and partly recovering in 2025. We do not read this time series as a regulatory effect. Brown-credit agreements enter public filings with a changing sector mix, so the design must absorb borrower-year conditions and industry-year shocks rather than treat aggregate movements in the archive as evidence of bank behaviour.

Figure 2: High-emission-sector share of retained SEC credit-agreement exhibits




Table 3 reports the audit and content checks. All 129 retained exhibits were retrieved in full and verified against their source filings. The median retained exhibit is 234.0 kilobytes and the ninetieth percentile is 1,770.1 kilobytes, which is what one expects from credit agreements that run through definitions, schedules, guarantees, covenants, agency provisions, and signature pages. The median text scan reads 78,522 characters per document. That depth matters: lender and agent information is often buried deep in the contract rather than placed neatly in the exhibit title or opening recital.

The contract-language checks are encouraging, but they also show why the data construction must be careful. Administrative-agent language appears in 105 exhibits (81.4% of the archive); lender language appears in 122 (94.6%). These shares fall short of 100 percent for good reasons: some agreements are abbreviated, some are poorly structured or partly scanned, and some identify parties in attachments that require closer reading. The validated lender panel therefore uses several relationship roles rather than leaning on one field: named lenders, administrative agents, arrangers, and lead banks each reveal a different part of the bank-borrower relationship. Sustainability-linked language appears in only 3 exhibits (2.3%). This does not mean that sustainability-linked loans are nearly absent from the market; a broad SEC credit-agreement extraction, unlike a specialized sustainability-linked loan database, simply recovers few such clauses. The contracts are rich enough to identify bank roles and core loan terms, but not uniform enough to justify a single-field lender measure.

Table 3. Final SEC credit-agreement panel: composition and verification checks

<i>Statistic</i>	<i>Count</i>	<i>Share of exhibits</i>
Matched credit-agreement exhibits	129	100.0%
Unique SEC filing accessions	85	–
Unique borrower CIKs	85	–
Exhibits retrieved in full and verified	129	100.0%
Filed before 2023	81	62.8%
Filed from 2023 onward	48	37.2%
Filed on Form 8-K or 8-K/A	66	51.2%
High-emission SIC sector	34	26.4%
Utility SIC sector	10	7.8%
Administrative-agent language detected	105	81.4%
Lender language detected	122	94.6%
Sustainability-linked language detected	3	2.3%
Filed on an amended SEC form	0	0.0%

Notes: Shares use the 129 retained exhibits as the denominator. The language variables are document-level text signals, not final bank-role classifications.


Figure 3 shows how these signals vary by year. Lender language appears more often than administrative-agent language in most years, which fits the structure of many syndicated agreements: the syndicate is listed in one place, while agency authority is defined elsewhere. The figure also makes clear that this is not a simple text-counting exercise. Role visibility changes with filing format, exhibit structure, and document length. The empirical panel therefore treats the extracted text as the first auditable layer. The bank-level layer then maps bank names to consolidated parents using GLEIF Golden Copy files, FFIEC NIC hierarchy data, FDIC institution records, EBA Pillar 3 materials, and manual validation of recurring name variants.

The SEC exhibit archive is the spine of the data set, and it is the layer verified in the analysis reported here. The broader architecture links borrower transition exposure to

EPA Greenhouse Gas Reporting Program data, high-emission SIC and NAICS classifications, EU Taxonomy sector mappings, and climate-risk language from SEC filings. Bank treatment status is linked to the legal scope of CRR Article 449a, EBA Pillar 3 templates and transparency files, and consolidated bank-parent identifiers from GLEIF and supervisory hierarchy files. Firm controls come from SEC XBRL company facts and submissions metadata; macro-financial controls come from public FRED series; physical-risk measures come from the FEMA National Risk Index. The point of this architecture is not only coverage but traceability: the loan contract, the borrower, the bank parent, the emissions proxy, and the climate-risk controls can all be followed back to public records as each layer is added.

The stylized facts impose a clear discipline on the empirical analysis. The public SEC filing system contains a workable credit-agreement archive for the years around the EU Pillar 3 climate-disclosure mandate: post-2023 contracts, high-emission borrowers, and enough lender and agent language to build bank-borrower relationships. At the same time, the archive is not a random draw from all corporate lending, the extraction is limited in size, and the sector mix changes over time. That is why the main specifications are built around within-borrower-year comparisons between EU and non-EU banks, with bank-year and borrower-year shocks absorbed directly. The data section therefore does more than document coverage. It explains why the paper identifies the disclosure effect through relationships rather than aggregate lending volumes.

Figure 3: Contract-language signals in retained SEC credit-agreement exhibits



6. Empirical Analysis

The empirical analysis begins by measuring what the public SEC credit-agreement archive actually contains: the timing, sectoral composition, and contract-language content

of the retained loan exhibits. This is not a substitute for the bank-level estimates; it is the first piece of evidence needed to interpret them. If the public filings contained no brown borrowers after 2023, or if the contracts contained no usable lender and agent language, any claim about cross-border bank-credit transmission would rest on a weak base.

Table 4 shows how the retained agreements are distributed across the pre- and post-disclosure periods and across high-emission and other sectors. The first fact is that the post-2023 archive is not empty: it contains 48 credit-agreement exhibits, of which eight are associated with high-emission SIC sectors. The high-emission share falls from 32.1% before 2023 to 16.7% after. This shift is not a treatment effect: the table describes which contracts appear in SEC filings, not how a given bank changes its lending. It does establish that the public filing environment continues to contain transition-risk borrowers after the EU disclosure regime begins, which the bank-side comparison requires. Post-2023 high-emission contracts remain observable; their lower sample share is sample composition, not a causal lending response.

The second fact is that document structure differs sharply across cells. The median retained high-emission exhibit in the post period is 81.5 kilobytes, against 618.0 kilobytes for other post-2023 sectors and 805.2 kilobytes for high-emission contracts before 2023. This matters because long credit agreements are more likely to contain the agency provisions, lender schedules, covenant definitions, collateral clauses, and amendments needed to construct relationship-level variables. The language signals remain useful: in the post period, lender language appears in all retained high-emission exhibits and in 92.5% of other-sector exhibits, and the administrative-agent signal is present in all eight post-period high-emission exhibits. Those high rates deserve care because the cell is small, but the practical implication is clear. The bank panel cannot rely on one field or one phrase; it must combine named lenders, agents, arrangers, lead banks, document context, and manual validation. Role language survives in the key post-2023 high-emission cell, but contract thickness changes enough to require the validated lender schedule used later.

Figure 4 makes the pattern easier to see. The post-2023 high-emission cell does not lose lender or agent language; if anything, the raw role-language shares are high in that small group of contracts. In simple difference-in-differences terms, the high-emission-minus-other-sector change is 0.2 percentage points for lender language and 10.6 percentage points for administrative-agent language. These are measurement facts about the documents that feed the causal design rather than causal estimates of disclosure

regulation, and they carry a warning: contract observability has several dimensions. A contract may reveal lender roles while still being shorter, thinner, or less informative on other margins. A credible design must therefore distinguish a bank’s decision to lend from the researcher’s ability to observe the lender in a filed contract.

Table 4. Document-level pre/post composition by transition-risk sector

<i>Group</i>	<i>Exhibits</i>	<i>Borrowers</i>	<i>Agent language</i>	<i>Lender language</i>	<i>SLL language</i>	<i>Median KB</i>
Pre-2023, other sectors	55	32	76.4%	92.7%	0.0%	148.4
Pre-2023, high-emission	26	17	80.8%	100.0%	3.8%	805.2
Post-2023, other sectors	40	29	85.0%	92.5%	5.0%	618.0
Post-2023, high-emission	8	7	100.0%	100.0%	0.0%	81.5

Notes: The sample is the final SEC credit-agreement panel. The post period begins in 2023, the baseline implementation year used in the document-level analysis. Borrowers count unique CIKs within each row. Agent, lender, and sustainability-linked loan (SLL) language are document-level text signals. The table describes contract availability and observability; it is not a bank-level treatment-effect estimate.

Figure 4: Pre/post contract-language signals by transition-risk sector




Figure 5 traces this observability issue year by year. The gap between high-emission and other-sector contracts moves substantially over the sample period: in some years high-emission contracts are more likely to contain lender or agent language, in others less. This volatility is exactly why the main empirical strategy does not rest on raw aggregate counts. The treatment estimate must compare Article 449a-treated and non-treated banks within borrower-years, absorbing bank-year and borrower-year shocks. The document-level evidence supports that choice: the contracts are usable, but the availability of relationship language is not constant enough for simpler before-after comparisons.

Table 5 formalizes these patterns with document-level regressions. The first two columns estimate linear-probability models for the presence of administrative-agent and lender language, scaled in percentage points; the last two examine log exhibit size and log scanned text. All specifications include filing-year fixed effects and controls for

utility-sector status and Form 8-K status; the language models also control for log exhibit size. Standard errors are clustered by borrower CIK. The high-emission-by-post coefficient is small and positive for both language outcomes, and neither estimate is statistically precise in this exhibit-level sample. By contrast, the interaction is negative and statistically sharper for document size and scanned text, consistent with the descriptive evidence that post-2023 high-emission contracts are shorter. The interpretation is modest: there is no stable document-level evidence that high-emission post-2023 contracts become more transparent in a way that would mechanically generate the paper's bank-level results. The later relationship estimates are not riding on improved post-2023 visibility of brown-borrower contracts.

The regression table also shows why controlling for document structure matters. Longer exhibits are more likely to reveal administrative-agent language, and the coefficient on log exhibit bytes is positive in the lender-language specification as well, though smaller and not statistically sharp. That pattern fits the simple reality of credit agreements: more complete contracts contain more places where bank roles can appear. Utility-sector contracts display less administrative-agent language, even though utilities are central to transition risk. The message matches the figures. The raw filings are rich but not uniform, so the design must carry document-quality controls and robustness checks alongside the lender classification work.

Figure 6 summarizes the distribution of exhibit size across period and sector cells. Post-2023 high-emission contracts have a smaller median and a thinner upper tail than their pre-2023 counterparts, which is why text depth and role detection cannot be treated as fixed features of the sample. Larger documents may reveal more bank names, but they also contain more amendments, cross-references, schedules, and exhibits that require careful parsing. For the bank-level analysis, this motivates two safeguards. First, every bank-role classification is stored with the passage of contract text on which it is based. Second, the specifications are repeated on narrower samples, agent-only contracts and full-syndicate contracts, to ensure the results are not driven by changing document observability.

Taken together, the results show that the public SEC archive is rich enough for the paper's bank-level evidence but too uneven for aggregate filing counts to carry the causal claim. The archive contains post-2023 high-emission credit agreements, and the contracts contain enough lender and administrative-agent language to construct bank-borrower relationships. Observability, however, varies across sectors and years. That variation is

part of the reason the paper uses the fixed-effect structure in equation (10). The validated bank layer that follows therefore asks the Article 449a question at the relationship level, where treated and non-treated lenders are compared within the same borrower-year cells rather than through aggregate filing patterns.

Figure 5: High-emission-sector gap in contract-language availability





Table 5. Document-level contract-language and size regressions

	<i>Agent language</i>	<i>Lender language</i>	<i>Exhibit size</i>	<i>Text scanned</i>
	(pp)	(pp)	(log bytes)	(log chars)
High-emission × post-2023	2.84 (14.82)	0.77 (8.60)	-1.44** (0.71)	-1.17** (0.50)
High-emission sector	16.73* (10.13)	8.05* (4.48)	0.89** (0.42)	0.55* (0.29)
Utility sector	-32.88*** (12.38)	0.28 (5.00)	0.08 (0.54)	0.11 (0.50)
Form 8-K or 8-K/A	-12.12 (8.15)	-0.57 (5.57)	0.51 (0.37)	0.56* (0.29)
Log exhibit bytes	5.22** (2.52)	0.90 (0.75)	—	—
Year fixed effects	Yes	Yes	Yes	Yes
Observations	129	129	129	129
R-squared	0.163	0.066	0.123	0.129

Notes: The sample is the final SEC credit-agreement panel. The first two columns are linear-probability models for document-level role-language availability, with outcomes scaled in percentage points. Standard errors, in parentheses, are clustered by borrower CIK. The post period begins in 2023. Year fixed effects absorb common filing-year shocks. The coefficients diagnose contract observability and document structure; they are not estimates of Article 449a effects on bank lending. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Figure 6: Distribution of retained exhibit size by period and transition-risk sector



7. Additional Econometric Diagnostics

The additional diagnostics push the exhibit-level evidence in the directions that matter most before the paper moves to the validated bank-borrower and facility layer. Rather than asking the document archive alone to carry the Article 449a estimate, they ask whether the archive is selected in a way that could distort it, whether small cells carry too much weight, whether one filing year or one borrower drives the language results, whether nonlinear models tell a different story, whether false positives matter, and whether the observed high-emission post-period gaps look unusual under a simple randomization benchmark.

Table 6 starts one step earlier than the retained-contract panel. It uses the full set of 2,000 scanned filings as the denominator and asks which filings produce at least one retained credit-agreement exhibit. Only 85 filings, or 4.2%, enter the retained archive. That low share is unsurprising: most SEC filings contain no new or amended credit agreement. What matters is whether selection into the archive jumps mechanically for high-emission firms after 2023. It does not. The high-emission-by-post coefficient is negative in the filing-level linear-probability model and in the exhibit-count specifications. The estimates are small relative to their standard errors, but they work against a simple visibility story in which post-2023 high-emission firms become mechanically easier to find in EDGAR. Selection into the retained archive does not rise for brown borrowers after Article 449a.

Table 7 treats the small post-period high-emission cell with appropriate caution. In that cell, administrative-agent and lender language are detected in all retained exhibits, while sustainability-linked loan language is not detected. Fisher exact tests do not reject equality between post-period high-emission and other-sector contracts for administrative-agent,

lender, or SLL language. The same holds when the comparison is restricted to high-emission contracts split before and after 2023. The exact tests reinforce the interpretation used throughout: the role-language outcomes are measurement checks, not treatment effects, and the small post-period high-emission cell cannot support a role-language treatment claim on its own.

Table 6. Filing-level selection into the retained credit-agreement archive

	<i>Any retained exhibit</i>	<i>Exhibit count</i>	<i>Exhibit count</i>	<i>Exhibit count</i>
	LPM (pp)	OLS	Poisson	Neg. binomial
High-emission × post-2023	-3.34 (2.50)	-0.06 (0.04)	-0.75 (0.62)	-1.00* (0.53)
High-emission sector	3.46** (1.72)	0.03 (0.03)	0.43 (0.40)	0.64* (0.33)
Utility sector	-1.54 (2.43)	0.03 (0.07)	0.36 (0.62)	0.11 (0.55)
Form 8-K or 8-K/A	-1.84* (1.06)	-0.05* (0.03)	-0.74*** (0.28)	-0.67*** (0.26)
Filing-year fixed effects	Yes	Yes	Yes	Yes
Observations	2000	2000	2000	2000
R-squared	0.009	0.009	—	—

Notes: The sample is the full set of 2,000 scanned SEC filings. The first column is a linear-probability model for whether a filing contributes at least one retained credit-agreement exhibit; the second column uses the retained-exhibit count. The Poisson and negative-binomial columns report log conditional-mean coefficients. Standard errors, in parentheses, are clustered by SEC CIK. The post period begins in 2023. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table 7. Exact tests for small contract-language cells

<i>Outcome</i>	<i>Post high</i>	<i>Post other</i>	<i>Fisher p-value</i>	<i>High pre</i>	<i>High post</i>	<i>Fisher p-value</i>
Administrative-agent language	100.0%	85.0%	0.571	80.8%	100.0%	0.309
Lender language	100.0%	92.5%	1.000	100.0%	100.0%	—
SLL language	0.0%	5.0%	1.000	3.8%	0.0%	1.000

Notes: The sample is the final SEC credit-agreement panel. The first Fisher test compares high-emission and other-sector contracts within the post-2023 period. The second compares pre- and post-2023 contracts within high-emission sectors. Cell entries are document shares in percent.

Table 8 returns to document quality. The post-2023 high-emission interaction is negative for both log exhibit bytes and log scanned characters, echoing the main document-level regressions. These contracts are not becoming longer or richer in the cell where the regulatory comparison matters most. The verification row shows no variation because every retained exhibit was retrieved in full, which locates the empirical uncertainty where it belongs: in what contracts disclose and which contracts are filed, not in missing documents.

Table 8. Document-quality balance around high-emission post-2023 cells

<i>Outcome</i>	<i>Interaction coefficient</i>	<i>Standard error</i>	<i>Observations</i>
Log exhibit bytes	-1.44	0.71	129
Log scanned characters	-1.17	0.50	129
Amendment indicator (pp)	No variation	—	129
Retrieved and verified (pp)	No variation	—	129

Notes: The sample is the final SEC credit-agreement panel. Each coefficient is the high-emission-sector by post-2023 interaction from a document-level balance regression with filing-year fixed effects, utility status, and Form 8-K status. Standard errors are clustered by borrower CIK.

Table 9 reports year-by-year high-emission gaps in administrative-agent and lender language. The pattern is uneven. High-emission contracts show more visible agent language in several years, less in 2022, and small post-period cells in 2023 and 2024. Lender-language gaps stay small because lender language is already common in both groups. The event-year table helps precisely because it is not smooth: it shows why the paper does not lean on a simple aggregate before-after comparison and why the bank-level analysis must absorb filing-year shocks directly. Annual volatility is part of the identification problem, not an afterthought.

Table 9. Year-by-year high-emission gaps in contract-language availability

<i>Filing year</i>	<i>Other exhibits</i>	<i>High-emission exhibits</i>	<i>Agent gap</i>	<i>Lender gap</i>
2018	25	3	12.0	4.0
2019	11	3	36.4	9.1
2020	8	8	12.5	25.0
2021	2	3	50.0	0.0
2022	9	9	-11.1	0.0
2023	13	2	30.8	0.0
2024	10	1	10.0	10.0
2025	17	5	5.9	11.8

Notes: The sample is the final SEC credit-agreement panel. Gaps are high-emission-sector means minus other-sector means, in percentage points. A dash indicates that one side of the yearly comparison is empty.

Table 10 asks whether the main language gaps are pinned to one calendar year. They are not. The administrative-agent coefficient moves from negative to positive as different years are removed, as expected in a small document archive, and the standard errors remain wide. The lender-language coefficient is more stable, ranging from modestly negative to modestly positive across the leave-one-year exercises. The lesson is not the sign of any single row but that dropping one filing year never produces a sharp, stable, economically large high-emission post-period language jump. No single calendar year drives the document-language findings.

Table 10. Leave-one-year-out document-language regressions

<i>Omitted year</i>	<i>Agent gap</i>	<i>Agent SE</i>	<i>Lender gap</i>	<i>Lender SE</i>	<i>Observations</i>
2018	2.84	(16.70)	-0.66	(9.12)	101
2019	4.87	(15.29)	0.79	(9.04)	115
2020	5.95	(14.31)	5.62	(7.35)	113
2021	5.03	(15.22)	-0.11	(8.69)	124
2022	-4.58	(16.20)	-2.98	(10.67)	111
2023	-13.59	(13.01)	4.85	(11.39)	114
2024	5.06	(15.92)	0.77	(9.53)	118
2025	21.62	(20.12)	-3.23	(7.13)	107

Notes: The sample is the final SEC credit-agreement panel. Each row re-estimates the document-level high-emission-sector by post-2023 language regression after excluding one filing year. Outcomes are scaled in percentage points. Standard errors are clustered by borrower CIK.

Table 11 gives the corresponding borrower and functional-form checks. Dropping one borrower CIK at a time leaves the administrative-agent coefficient centered close to the baseline and keeps the lender-language coefficient near zero; no single borrower turns the document-level language result into a precise positive finding. The logit and probit checks are informative in a different way: because role language is very common, the high-emission post-period cell generates separation or collinearity in binary-response models. That is why the paper keeps the linear-probability model as the transparent baseline for document-observability diagnostics. The separation reflects the structure of the data, not a suppressed result. Neither a single borrower nor a nonlinear functional form overturns the exhibit-level measurement evidence.

Table 11. Borrower influence and alternative functional forms

<i>Panel A. Leave-one-borrower-out influence</i>					
Outcome	Baseline	Minimum	Median	Maximum	Max. shift
Administrative-agent language	2.84	-7.45	2.96	7.40	10.29
Lender language	0.77	-6.37	0.74	3.32	7.14
<i>Panel B. Alternative functional forms</i>					
Outcome and model	Coefficient/AME	Standard error	Observations		
Administrative-agent, LPM	2.84	14.82	129		
Administrative-agent, Logit AME	—	—	—		
Administrative-agent, Probit AME	—	—	—		
Lender, LPM	0.77	8.60	129		
Lender, Logit AME	—	—	—		
Lender, Probit AME	—	—	—		

Notes: The sample is the final SEC credit-agreement panel. Panel A reports the distribution of the high-emission-sector by post-2023 coefficient after dropping one borrower CIK at a time. Panel B reports the linear-probability coefficient and average marginal effects from logit and probit models when those nonlinear specifications produce an estimable marginal effect; dashes indicate separation or collinearity. Outcomes are scaled in percentage points.

Table 12 shows why the forensic screening of credit-agreement exhibits matters. Reconstructing the pre-refinement archive by appending the 18 manually removed false positives would change the language coefficients materially, especially for lender language. The corrected final panel, and the conservative panel that drops retained broad adjacent-label matches, tell a calmer story: the lender gap is near zero and the agent gap remains imprecise. This is the strongest econometric reason to keep the audit trail in the paper. False positives are not cosmetic: they can move point estimates in a small document panel, so the conservative contract screen matters for the estimates, not just for housekeeping.

Finally, Table 13 reports randomization inference. The exercise permutes high-emission status within filing year 500 times and re-estimates the document-level language regressions after each permutation. The observed administrative-agent gap has a two-sided permutation p-value of 0.896; the lender-language gap, 0.966. These values are

unexciting, and they should be. The document-level evidence is not manufacturing a rare post-2023 high-emission language pattern. It shows a contract archive that is usable, audited, and sober about its own limits: the observed language gaps look ordinary under a within-year randomization benchmark.

The diagnostics sharpen the boundary of the paper’s empirical claim. The SEC contract layer is strong enough to support the validated bank and facility layer that follows: lender-name extraction, bank-parent harmonization, legal Article 449a treatment assignment, relationship-presence estimation, and facility-term extraction. At the same time, the extra tests make it harder to overclaim from the document layer alone. The public filings reveal real contracts, real borrower-sector variation, and real bank-role language, but also small cells, year-to-year volatility, and sensitivity to false-positive screens. Far from being weaknesses to hide, they are the reason the next section keeps the lender validation and treatment crosswalk fully visible.

Table 12. Sensitivity to false-positive credit-agreement screens

<i>Sample</i>	<i>Exhibits</i>	<i>Audit removals included</i>	<i>Agent gap</i>	<i>Agent SE</i>	<i>Lender gap</i>	<i>Lender SE</i>
Corrected final panel	129	0	2.84	(14.82)	0.77	(8.60)
Pre-refinement reconstruction	147	18	-8.70	(18.30)	-12.42	(12.78)
Conservative final panel	124	0	5.41	(13.98)	0.01	(9.05)

Notes: The sample combines the final panel with the audit record of removed exhibits. The pre-refinement reconstruction appends the 18 manually removed false positives to the corrected 129-exhibit panel. The conservative final panel drops retained broad adjacent-label matches. Coefficients are high-emission-sector by post-2023 language gaps in percentage points.

Table 13. Randomization inference for document-language gaps

<i>Outcome</i>	<i>Observed gap</i>	<i>Permutation p-value</i>	<i>Usable draws</i>
Administrative-agent language	2.84	0.896	500
Lender language	0.77	0.966	500

Notes: The sample is the final SEC credit-agreement panel. The null distribution is built by permuting high-emission-sector status within filing year and re-estimating the document-level language regression. The reported p-value is two-sided and compares the absolute permuted coefficient with the absolute observed coefficient.

8. Validated Bank-Relationship and Facility-Term Layer

This section carries the evidence from observable documents to bank relationships and facility terms. The SEC agreements are used to build a contract-by-contract lender schedule, a legally sourced Article 449a treatment crosswalk, a facility-term panel, and a borrower-year credit-supply panel. The approach is conservative. Bank names are first detected in the contract text, but a row is retained only when the surrounding evidence window places the bank in a credit role: exact administrative agent, collateral agent, agent for the lenders, named lender, arranger, syndication agent, documentation agent, issuing bank, or a closely related lender role. Mentions without that role information are removed.

The resulting layer still rests on public contracts, not private DealScan schedules, but it is no longer a raw dictionary exercise.

Table 14 reports the validation audit. The raw dictionary pass produces 431 document-bank rows. Validation keeps 364 and excludes 67 mention-only rows. Among the validated observations, 69 are exact administrative-agent roles and 295 are named lender-or-arranger roles. This is the contract-by-contract check the paper needed: the analysis does not treat every bank name in a filed agreement as a lending relationship. It keeps the rows where the contract language says the bank is actually doing bank work. The bank panel is built from validated credit roles, not incidental bank mentions.

Table 15 moves from document-bank rows to contracts. The retained archive contains 129 credit-agreement exhibits; 92 have at least one validated bank parent, and 15 contain at least one legally treated Article 449a parent. The same extraction yields the fields needed for a first facility-term analysis: disclosed amounts, pricing-spread language, maturities, financial covenants, collateral or guarantees, and sustainability-linked language. These fields do not replace lawyer-verified term sheets. They do something more modest but valuable: they make the public contracts estimable as a structured lender-and-term panel. Most retained contracts now contribute a validated bank parent, an extractable facility term, or both.

Table 14. Validation audit for the validated bank-role layer

<i>Validation statistic</i>	<i>Value</i>
Raw dictionary document-bank rows	431
Validated document-bank rows	364
Excluded mention-only rows	67
Validated administrative-agent rows	69
Validated named lender/arranger rows	295
Validated exhibits with at least one mapped parent	92
Validated observed bank parents	39

Notes: The table reports the classification of contract evidence windows. A row is retained when the bank parent appears as an exact administrative-agent role or as a named lender, arranger, syndication, documentation, issuing-bank, or related bank role in the contract opening evidence. Mention-only rows are excluded from the validated bank layer.

Treatment status is assigned from legal and supervisory sources rather than geography alone. Table 16 matches the observed EU/EEA banking parents to the EBA 2025 Transparency Exercise and uses CRR Article 449a together with Commission Implementing Regulation (EU) 2022/2453 as the legal basis for the treatment definition, as set out by European Parliament and Council (2013) and European Commission (2022). All 12 treated parents observed in the validated schedule are matched to the EBA transparency list (European Banking Authority, 2025). This matters because treatment

does not attach to “a European-sounding bank name” but to a parent banking group inside the Article 449a Pillar 3 ESG disclosure architecture. Treatment status is legally sourced and auditable at the consolidated bank-parent level.

Table 15. Contract-by-contract validated lender schedule and term coverage

<i>Validation or term statistic</i>	<i>Contracts</i>
Retained credit-agreement exhibits	129
With at least one validated bank parent	92
With exact validated administrative agent	64
With validated lender/arranger parent	91
With Article 449a-treated parent	15
With extractable facility amount	93
With extractable pricing-spread language	62
With maturity language	81
With financial-covenant language	61
Secured or guaranteed	102
Sustainability-linked language	3

Notes: The sample is the contract-by-contract lender schedule. A bank parent is retained only after evidence-window validation. Facility amounts, pricing spreads, maturities, covenants, collateral, guarantees, and sustainability-linked language are extracted from the SEC contract text and should be read as structured contract-text fields rather than hand-coded legal opinions.

Table 16. Legally sourced Article 449a treatment crosswalk for observed bank parents

<i>Contract bank parent</i>	<i>Country</i>	<i>EBA transparency institution / LEI</i>	<i>Validated documents</i>
Banco Santander	ES	Banco Santander, S.A. (5493006QMFDDMYWIAM13)	6
BNP Paribas	FR	BNP Paribas (R0MUWSFPU8MPRO8K5P83)	4
BBVA	ES	Banco Bilbao Vizcaya Argentaria, S.A. (K8MS7FD7N5Z2WQ51AZ71)	3
Deutsche Bank	DE	DEUTSCHE BANK AKTIENGESELLSCHAFT (7LTWFZYICNSX8D621K86)	2
ING	NL	ING Groep N.V. (549300NYKK9MWM7GGW15)	2
Rabobank	NL	Cooperatieve Rabobank U.A. (DG3RU1DBUFHT4ZF9WN62)	2
UniCredit	IT	UNICREDIT, SOCIETA' PER AZIONI (549300TRUWO2CD2G5692)	2
ABN AMRO	NL	ABN AMRO Bank N.V. (BFXS5XCH7N0Y05NIXW11)	1
Commerzbank	DE	COMMERZBANK Aktiengesellschaft (851WYGNLUQLFZBSYGB56)	1
Credit Agricole	FR	Groupe Credit Agricole (FR969500TJ5KRTCJQWXH)	1
DNB	NO	DNB BANK ASA (549300GKFG0RYRRQ1414)	1
Natixis	FR	Groupe BPCE (FR9695005MSX1OYEMGDF)	1

Notes: Treatment is assigned to observed EU/EEA banking parents matched to the EBA 2025 Transparency Exercise list under CRR Article 449a and Commission Implementing Regulation (EU) 2022/2453. Non-EU/EEA parents are coded outside the treatment group.

Table 17 shows the coverage of the validated relationship layer. The validated layer contains 364 document–bank observations, where each observation is a single bank parent identified in a single filing. These observations span 39 distinct bank parents and 92 retained exhibits, and involve 61 unique borrower firms (identified by CIK). Twelve observed parents are treated Article 449a parents, but they account for only 7.1% of document-bank rows, and 15 retained exhibits contain at least one treated parent. The sample therefore has enough treated-bank presence to estimate the first relationship tests, but the treated cell is thin, a fact that should shape the interpretation rather than be hidden from it. The main estimates need borrower-year fixed effects precisely because treated banks are a small share of observed relationships.

Table 18 makes the parent-by-parent validation visible. For each main parent, it reports the raw dictionary rows, the rows kept after evidence-window validation, the mention-only rows excluded, the exact administrative-agent rows, and the number of validated documents. The most frequent parents are large U.S. and Canadian banks: Bank of America, JPMorgan Chase, Wells Fargo, PNC, Citigroup, Citizens, U.S. Bancorp, and Fifth Third appear often. The treated Article 449a group is thinner but observable: BNP Paribas, Banco Santander, BBVA, Deutsche Bank, ING, UniCredit, Commerzbank, Rabobank, ABN AMRO, Credit Agricole, DNB, and Natixis all appear at least once after validation. The table lets the reader inspect both sides of the lender market: the treated European groups and the non-treated lenders that can absorb or replace their exposures. That visibility is what makes substitution testable.

Table 17. Validated bank-relationship extraction coverage

<i>Coverage statistic</i>	<i>Value</i>
Retained credit-agreement exhibits	129
Exhibits with at least one validated mapped bank parent	92
Validated document–bank relationship rows	364
Unique borrower CIKs with mapped bank parents	61
Observed bank parents	39
Observed Article 449a-treated parents	12
Relationship rows linked to Article 449a-treated parents	7.1%
Exhibits with at least one Article 449a-treated parent	15
Mean mapped bank parents per exhibit	2.82

Notes: The sample is the validated bank-role layer built from the SEC contract archive. Article 449a treatment is reported for observed EU/EEA parent banking groups; UK, Swiss, Canadian, Japanese, Australian, and U.S. parents are outside the treated group in this version.

Table 18. Parent-by-parent validation audit for the validated bank layer

Bank parent	Country	Treat.	Raw rows	Kept rows	Excluded	Agent rows	Validated documents
Bank of America	US	No	55	49	6	25	49
JPMorgan Chase	US	No	40	34	6	19	34
PNC Financial Services	US	No	30	27	3	3	27
Wells Fargo	US	No	32	26	6	2	26
Citigroup	US	No	21	20	1	6	20
Citizens Financial	US	No	19	16	3	3	16
Bank of Montreal	CA	No	15	15	0	0	15
Fifth Third Bancorp	US	No	17	15	2	1	15
U.S. Bancorp	US	No	18	14	4	1	14
KeyCorp	US	No	12	12	0	0	12
Truist Financial	US	No	13	12	1	1	12
Royal Bank of Canada	CA	No	11	11	0	3	11
Banco Santander	ES	Yes	7	6	1	0	6
BNP Paribas	FR	Yes	7	4	3	0	4
BBVA	ES	Yes	4	3	1	0	3
Deutsche Bank	DE	Yes	4	2	2	0	2
ING	NL	Yes	3	2	1	0	2
Rabobank	NL	Yes	2	2	0	0	2
UniCredit	IT	Yes	3	2	1	0	2
ABN AMRO	NL	Yes	1	1	0	0	1
Commerzbank	DE	Yes	2	1	1	0	1
Credit Agricole	FR	Yes	1	1	0	0	1
DNB	NO	Yes	1	1	0	0	1
Natixis	FR	Yes	1	1	0	0	1
Other observed non-treated parents	–	No	111	87	24	5	87

Notes: The table reports the parent-level validation audit built from the contract text. Raw rows are dictionary hits in bank-role evidence windows. Kept rows survive the validation rule; excluded rows are

mentioned-only observations; agent rows are exact administrative-agent or equivalent agent-role matches. The treatment flag identifies observed Article 449a EU/EEA parent banking groups.

Table 19 reports the facility-term layer. The contracts yield 197 detected facility rows across 129 underlying exhibits, with 160 rows containing an extractable amount. Revolving-credit and term-loan language are common, and the extraction also identifies asset-based and delayed-draw structures when contracts name them. The empirical unit is therefore no longer a document that happens to contain lender language but a validated lender schedule connected to loan size, maturity, pricing, covenant, collateral, and sustainability-linked text fields. The paper now measures the contractual margins through which banks can adjust credit without simply entering or exiting a relationship.

Table 19. Facility-level term-extraction coverage

Facility statistic	Value
Detected facility rows	197
Underlying retained exhibits	129
Rows with extractable amount	160
Rows with pricing-spread language	122
Rows with maturity language	139
Rows with financial-covenant language	118
Rows secured or guaranteed	167
Revolving-credit rows	72
Term-loan rows	58
ABL rows	12
Delayed-draw rows	7

Notes: The sample is the facility-term panel extracted from the SEC credit agreements. A single contract can contribute more than one facility row when the text names more than one facility type.

Table 20 estimates the relationship-level risk-set test using the legal Article 449a crosswalk. The unit is a bank-parent-exhibit cell: every observed bank parent is paired with every retained exhibit, and the dependent variable equals 100 if the bank parent is validated in that exhibit. The coefficient of interest is the interaction of treated Article 449a parent, high-emission borrower sector, and the post-2023 period. The first column uses bank and filing-year fixed effects. The next columns absorb borrower-year and bank-year variation, which is closer to the baseline causal specification in equation (10).

The estimates tell a useful causal story. With only bank and year fixed effects, the treated-bank interaction is positive, at 5.6 percentage points. Once borrower-year fixed effects are introduced, the estimate becomes -14.6 percentage points, and it remains negative at -16.1 percentage points with bank-year and borrower-year fixed effects. The active-bank restriction is also negative, though less precise. Aggregate comparisons and within-borrower-year comparisons therefore do not say the same thing, which is the point of the identification strategy. If the question is whether treated banks change their lending

to the same borrowers, the comparison must be made inside borrower-years, not across different borrowers who happen to file different contracts. Evidence for Article 449a transmission appears only once borrower demand is held fixed, and it runs through relative treated-bank participation rather than aggregate contract counts.

Table 20. Legal Article 449a bank-risk-set estimates

	<i>Bank FE + year FE</i>	<i>Bank FE + borrower-year FE</i>	<i>Bank-year FE + borrower-year FE</i>	<i>Active banks only</i>
Article 449a bank × high-emission × post-2023	5.601 (4.295)	-14.643*** (4.087)	-16.077*** (4.341)	-7.712 (11.584)
Bank fixed effects	Yes	Yes	No	No
Bank-year fixed effects	No	No	Yes	Yes
Borrower-year fixed effects	No	Yes	Yes	Yes
Observed bank universe	All	All	All	≥3 documents
Observations	5031	5031	5031	3483
R-squared	0.097	0.216	0.269	0.288

Notes: The sample is the validated bank-parent risk set under the legally sourced Article 449a treatment crosswalk. Each row is a bank-parent-exhibit cell. The dependent variable equals 100 when the bank parent is validated in the exhibit. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table 21 brings the facility terms into the same logic. The coefficient is the interaction of Article 449a-treated bank participation, high-emission borrower status, and the post-2023 period. The point estimates are imprecise, unsurprising given the small treated cell, but the magnitudes are economically meaningful enough to report. The interaction is associated with a 1.38 log-point larger disclosed facility amount, a 172 basis-point wider pricing-spread midpoint, a 2.7-year longer extracted maturity, a 41 percentage-point lower financial-covenant-language indicator, little movement in security, and a 13.5 percentage-point higher sustainability-linked-language indicator. These estimates do not settle the facility-level treatment effect; their value is that the margins on which banks can respond are now observable: quantity, price, maturity, control rights, collateral, and transition-linked contracting. The facility layer turns the paper from a lender-presence exercise into a contracting exercise, even where precision remains limited.

Table 21. Facility-term outcomes around Article 449a-treated bank participation

<i>Outcome</i>	<i>Article 449a bank × high-emission × post-2023</i>	<i>SE</i>	<i>Observations</i>
Log facility amount	1.375	(2.130)	160
Pricing-spread midpoint (bp)	171.672	(221.296)	116
Maturity (years)	2.705*	(1.548)	99
Financial-covenant language (pp)	-41.246	(50.520)	197
Secured or guaranteed (pp)	0.972	(39.267)	197
Sustainability-linked language (pp)	13.515	(12.125)	197

Notes: The sample is the facility-term panel. The coefficient is the interaction of Article 449a-treated bank participation, borrower high-emission status, and the post-2023 period. Specifications include all lower-order terms, filing-year fixed effects, utility status, and Form 8-K status. Standard errors are clustered by borrower CIK. These estimates describe contract terms around treated-bank participation and should not be read as a standalone causal treatment effect. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table 22 aggregates the validated schedule to the borrower-year level and provides the clearest evidence on credit supply and substitution. The post-2023 high-emission cell shows little movement in the probability of having any Article 449a-treated parent, a 0.40 increase in the treated-parent count, and a treated-parent share estimate below one percentage point. The non-treated parent count, by contrast, rises by 3.76. The substitution margin is therefore much larger than the treated-bank expansion margin, a pattern consistent with the leakage mechanism in the theoretical framework. If treated banks do not expand relative presence in brown borrower-years, non-treated lenders become more important in the observed contracts. The strongest credit-supply result is not treated-bank retreat in isolation but a reweighting of the lender set toward non-treated parents in high-emission borrower-years after Article 449a.

Table 22. Borrower-year credit-supply and lender-substitution outcomes

<i>Outcome</i>	<i>High-emission × post-2023</i>	<i>SE</i>	<i>Borrower-years</i>
Any Article 449a-treated parent (pp)	2.693	(21.544)	85
Article 449a-treated parent count	0.402	(0.608)	85
Non-treated parent count	3.757**	(1.854)	85
Article 449a-treated parent share (pp)	0.823	(5.995)	85
Log disclosed facility amount	0.377	(1.155)	73
Secured/guaranteed contract share (pp)	-36.708*	(20.643)	85

Notes: The sample is the borrower-year panel aggregated from the validated contract lender schedule. The coefficient is the high-emission-sector by post-2023 interaction. Specifications include filing-year fixed effects and utility status, with standard errors clustered by borrower CIK. The non-treated parent count is the substitution margin: it rises when validated non-Article 449a parents appear more often in the same borrower-year cell. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

9. Robustness

The robustness analysis asks a narrower but essential question: are the document-level patterns reported above artifacts of one timing choice, sample screen, or transition-risk proxy? The answer matters because the document-level results are the bridge between raw SEC contracts and the validated bank-borrower-year layer. If the basic language patterns vanished when we changed the post-period cutoff or dropped one filing-year shock, the contract archive would be too fragile to support the main identification strategy.

Table 23 first checks whether the main restrictions leave enough data to be informative. The full sample contains 129 retained exhibits. The Form 8-K-only restriction leaves 66, so the main results do not come from a small number of annual or quarterly reports. Excluding utilities leaves 119 exhibits and keeps a high-emission share of 20.2%, which matters because utilities are both central to transition risk and institutionally unusual. The verification restriction leaves the sample unchanged,

confirming that the exercise is not driven by missing or incomplete source documents. The robustness samples remain large enough to test the main document-language patterns.

Figure 7 shows the same coverage check visually. The most demanding restrictions reduce the sample, as they should, but none collapses the archive. This matters for interpretation: a robustness exercise is useful only if it changes the sample enough to be a real test while leaving enough observations for the comparison to remain meaningful. The nonutility and Form 8-K-only samples do exactly that. They move the data away from possible special cases without emptying the post-2023 high-emission cell.

Table 23. Robustness restrictions: retained sample coverage and composition

<i>Restriction</i>	<i>Exhibits</i>	<i>Borrowers</i>	<i>High-emission</i>	<i>Post-2023</i>	<i>Agent language</i>	<i>Lender language</i>
Full sample	129	85	26.4%	37.2%	81.4%	94.6%
Exclude 2020	113	73	23.0%	42.5%	83.2%	95.6%
Exclude 2021–2022	106	71	20.8%	45.3%	83.0%	93.4%
Exclude utilities	119	81	20.2%	39.5%	83.2%	94.1%
Form 8-K only	66	51	22.7%	42.4%	78.8%	93.9%
Retrieved and verified	129	85	26.4%	37.2%	81.4%	94.6%

Notes: The sample is the final SEC credit-agreement panel. Borrowers are unique CIKs within each restricted sample. Shares use retained exhibits as the denominator. The table is a coverage audit: it shows whether the document archive remains usable after changing sample screens.

Figure 7: Retained exhibit coverage under robustness restrictions




Table 24. Robustness of document-level language gaps to timing and sample restrictions

<i>Specification</i>	<i>Agent gap</i>	<i>Agent SE</i>	<i>Lender gap</i>	<i>Lender SE</i>	<i>Observations</i>
Baseline: post 2023	2.84	(14.82)	0.77	(8.60)	129
Anticipation: post 2022	-10.45	(14.94)	-7.07	(9.43)	129
Delayed: post 2024	-12.60	(12.38)	5.19	(11.01)	129
Exclude 2020	5.95	(14.31)	5.62	(7.35)	113
Exclude utilities	-6.76	(13.66)	2.64	(10.33)	119
Form 8-K only	-10.38	(17.33)	-6.44	(10.44)	66

Notes: The sample is the final SEC credit-agreement panel. Coefficients are high-emission-sector by post-period interactions from document-level linear-probability models for administrative-agent and lender language, with outcomes scaled in percentage points. All specifications include filing-year fixed effects, utility-sector status, Form 8-K status, log exhibit size, and amended-form status when it varies. Standard errors, in parentheses, are clustered by borrower CIK. The table tests whether document observability is sensitive to timing choices and sample restrictions.

Table 24 reports the main document-level robustness tests. The dependent variables are the two contract-observability outcomes used in the empirical analysis, administrative-agent language and lender language, both in percentage points; the coefficient in each column is the interaction between high-emission status and the relevant post-period definition. In the baseline, the interaction is 2.84 percentage points for administrative-agent language and 0.77 for lender language. These estimates are small and not statistically sharp. The signs carry no stable economic message, and they need not: what matters is that the exhibit archive produces no large mechanical jump in lender or agent visibility for high-emission borrowers after 2023. Timing and sample restrictions do not turn document observability into a spurious Article 449a result.

The timing checks tell a similar story. Moving the post-period cutoff to 2022, which allows for anticipation after the publication of the EBA implementing standards, gives a lender-language interaction of -7.07 percentage points. Moving it to 2024, which allows for slower implementation, gives 5.19. The coefficients differ, as they should in a small document-level archive, but the story is disciplined: the sample shows no stable, spurious post-regime increase in lender visibility for brown borrowers. Excluding 2020 gives a positive but imprecise lender-language estimate, easing the concern that the baseline is a COVID-era filing artifact. Restricting to Form 8-K exhibits pushes the lender-language gap negative, to -6.44 percentage points.

The most revealing restriction is the exclusion of utilities. Without utilities, the lender-language gap stays small, at 2.64 percentage points, while the agent-language gap turns negative. Rather than weakening the data foundation, this restriction clarifies it. Utilities are a distinctive part of the transition economy: heavily regulated, capital intensive, emissions intensive, and often central to public transition policy. Their contracts therefore carry different observability properties from the rest of the high-emission sector. For the bank-level analysis, that is a reason to report results with and without utilities, not to discard the public-contract approach.

Figure 8 plots the interaction coefficients and confidence intervals across the timing and sample restrictions, keeping the robustness exercise from becoming a list of point estimates. The estimates move, especially when the post period is delayed to 2024 or utilities are removed, but the confidence intervals are wide and the estimates mostly oscillate around zero. That is the message the research design needs: the document-level evidence is not producing an artificial improvement in lender observability precisely for

high-emission borrowers after 2023. The main bank-level estimates are not an artifact of better post-2023 text extraction.

Table 25 asks whether the language gaps depend on how transition risk is defined. The broad high-emission SIC proxy produces the baseline gaps. The utility proxy produces much larger negative interactions, especially for lender language, while the high-emission nonutility proxy produces small positive interactions. This split is substantively plausible. Utility borrowers are not simply another emissions-intensive industry; they face distinctive regulatory oversight, capital-investment cycles, and contract structures. The table therefore argues against treating “brownness” as a single mechanical label. In the bank-level analysis, emissions intensity, high-emission sector status, EU Taxonomy exposure, and text-based transition language should be reported separately rather than collapsed into one opaque score. The economic meaning of “brown” depends on the borrower type, especially in the utility sector.

Taken together, the robustness evidence strengthens the empirical foundation while keeping the interpretation appropriately modest. The checks show that the SEC contract archive remains usable under meaningful restrictions, that the language patterns are not pinned to a single post-period cutoff, and that utilities deserve separate treatment. They also show why the causal test belongs at the bank-borrower relationship level. Document observability varies across sector, period, and filing format, so a credible Article 449a estimate must compare EU and non-EU banks within the same borrower-year and must report sensitivity to lender-role definitions, transition-risk proxies, timing, and utility-sector inclusion. The robustness section does not merely defend the exhibit-level analysis; it specifies the safeguards the bank-level panel must satisfy before any claim about credit reallocation can persuade.

The scope of the evidence is also clear. The paper covers SEC-filed contracts of public operating companies, not the full private loan market. It verifies document-level lender and agent language, and the bank layer harmonizes the observed validated parents to a legal Article 449a treatment crosswalk. Detailed commitment schedules and subsidiary allocations sometimes sit outside the opening pages of a credit agreement, so facility allocation is measured from the filed text available in the public archive. The baseline transition-risk measure is SIC-based, with utilities reported separately because their contracts and regulation differ from other high-emission borrowers. Rather than undoing the main evidence, these boundaries define its external validity: the paper identifies cross-border credit transmission in observable public contracts and shows where borrower

emissions, EU Taxonomy exposure, and richer facility allocations can extend the analysis.

Figure 8: Robustness of high-emission-by-post contract-language coefficients




Table 25. Alternative transition-risk proxies in document-level robustness tests

<i>Transition-risk proxy</i>	<i>Agent gap</i>	<i>Agent SE</i>	<i>Lender gap</i>	<i>Lender SE</i>	<i>Observations</i>
High-emission SIC	2.84	(14.82)	0.77	(8.60)	129
Utility SIC	60.07***	(18.89)	-7.98	(8.73)	129
High-emission, nonutility	-8.77	(13.34)	3.05	(10.26)	129

Notes: The sample is the final SEC credit-agreement panel. The post period begins in 2023. Coefficients are risk-proxy by post-period interactions from document-level language regressions with filing-year fixed effects and the same document controls used in Table 24. Standard errors, in parentheses, are clustered by borrower CIK. The table is reported to show that the transition-risk proxy changes the observability pattern and should not be collapsed into one opaque brownness score. *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

10. Conclusions and Policy Implications

This paper asks a practical question with large consequences: when banks must disclose climate-related exposures, do they change the credit they provide, or do those exposures simply move elsewhere in the financial system? The question is especially sharp for the European Union’s Article 449a regime because EU banks operate far beyond Europe. If disclosure changes their appetite for brown credit, the effects reach borrowers in markets that did not write the rule. If it does not, the regulation improves bank reports without changing the allocation of real capital. The paper’s central contribution is to bring that question to the level where it can be answered: the bank-borrower relationship inside actual loan contracts.

The empirical evidence gives a clear answer in the sample observable through public contracts. Article 449a exposure is associated with lower treated-bank relative participation in the post-2023 high-emission cell once borrower-year and bank-year variation are absorbed. The sign reversal between aggregate and within-borrower-year

specifications is central. Without borrower-year fixed effects, treated banks appear more present in the relevant contracts; with borrower demand held fixed, the treated-bank interaction turns negative. That pattern is exactly why international bank-credit questions need relationship-level identification: aggregate lender counts confuse who files which contract with how banks adjust to regulation.

The borrower-year evidence points to substitution as the most visible credit-supply margin. Treated-parent counts rise only modestly and treated-parent shares barely move, while non-treated parent counts rise by 3.76 in post-2023 high-emission borrower-years. This does not mean that every EU bank withdraws from every brown borrower; the finding is more precise and, for policy, more important: in the observed contracts, non-treated lenders become more prominent exactly where a unilateral European disclosure rule would be expected to pressure treated banks. That is the margin on which a disclosure mandate can become international leakage rather than real transition.

The facility evidence shows how the same regulatory pressure can surface inside the loan contract. The paper extracts 197 facility rows from the SEC agreements and measures loan amount, pricing-spread language, maturity, covenants, security, and sustainability-linked terms. The facility estimates are not precise enough for a standalone causal claim, but the magnitudes are informative: treated-bank participation in the post-2023 high-emission cell is associated with larger disclosed amounts, wider spread midpoints, longer maturities, and more sustainability-linked language. These margins matter because banks can respond to climate disclosure without fully exiting a relationship. They can reprice, adjust maturity and control rights, require security, or write transition-linked clauses.

The policy implication is straightforward. Climate-disclosure mandates should not be judged only by whether regulated banks report greener portfolios. A portfolio can look cleaner because borrowers decarbonize, because regulated banks leave, or because another lender takes the exposure, and those are different economic outcomes. Real transition requires changes in borrower behaviour or in the terms on which brown credit is supplied. Divestment reduces treated-bank exposure but may also reduce finance for transition-heavy firms. Substitution leaves the real allocation of brown credit close to unchanged while moving the exposure outside the regulated balance sheet. The evidence here is most consistent with the substitution margin being economically important.

The international-money-and-finance lesson is that regional financial regulation travels through global bank relationships, but the destination is not mechanically greener credit. European disclosure rules reach U.S. borrowers because global banks participate

in the same syndicated and bilateral credit markets. Once the rule crosses that border, the response depends on the full lender set. If non-treated lenders can absorb the exposure, the rule changes the geography of balance-sheet ownership more clearly than the underlying financing of brown activity. That is why comparable disclosure templates, supervision of cross-border exposures, and attention to lender substitution are not administrative details: they determine whether climate-financial regulation affects the real economy.

The final lesson is methodological. Public contracts are messy, but their messiness is visible. Every retained SEC exhibit can be traced to its source filing; the false-positive exclusions are documented; the bank-parent mapping is shown; the Article 449a treatment crosswalk is legally sourced; and the facility terms are extracted from filed agreements rather than inferred from proprietary fields. That transparency is part of the evidence. It lets readers follow the path from raw filings to bank relationships, facility terms, and substitution outcomes. In a policy area where measurement choices often shape the conclusion, auditability is not a side benefit but what makes the paper's central claim credible: mandatory bank climate disclosure is transmitted through cross-border credit relationships, and the strongest observed response is lender-composition adjustment rather than an immediate disappearance of brown credit.

References

1. Battiston, S., Mandel, A., Monasterolo, I., Schuetze, F., & Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change*, 7, 283-288. <https://doi.org/10.1038/nclimate3255>
2. Berg, F., Kölbel, J. F., & Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315-1344. <https://doi.org/10.1093/rof/rfac033>
3. Bernstein, A., Gustafson, M. T., & Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2), 253-272. <https://doi.org/10.1016/j.jfineco.2019.03.013>
4. Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2), 517-549. <https://doi.org/10.1016/j.jfineco.2021.05.008>
5. Bolton, P., & Kacperczyk, M. (2023). Global pricing of carbon-transition risk. *The Journal of Finance*, 78(6), 3677-3754. <https://doi.org/10.1111/jofi.13272>

6. Bruno, B., & Lombini, S. (2023). Climate transition risk and bank lending. *Journal of Financial Research*, 46(S1), 59-106. <https://doi.org/10.1111/jfir.12360>
7. Campiglio, E., Dafermos, Y., Monnin, P., Ryan-Collins, J., Schotten, G., & Tanaka, M. (2018). Climate change challenges for central banks and financial regulators. *Nature Climate Change*, 8, 462-468. <https://doi.org/10.1038/s41558-018-0175-0>
8. Chava, S. (2014). Environmental externalities and cost of capital. *Management Science*, 60(9), 2223-2247. <https://doi.org/10.1287/mnsc.2013.1863>
9. Chava, S., & Roberts, M. R. (2008). How does financing impact investment? The role of debt covenants. *The Journal of Finance*, 63(5), 2085-2121. <https://doi.org/10.1111/j.1540-6261.2008.01391.x>
10. Cheung, Y.-L., Tan, W., & Wang, W. (2020). Where do banks value corporate social responsibility more? Evidence on the role of national culture. *Journal of Banking & Finance*, 118, 105810. <https://doi.org/10.1016/j.jbankfin.2020.105810>
11. Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *The Quarterly Journal of Economics*, 129(1), 1-59. <https://doi.org/10.1093/qje/qjt031>
12. Christensen, H. B., Hail, L., & Leuz, C. (2021). Mandatory CSR and sustainability reporting: Economic analysis and literature review. *Review of Accounting Studies*, 26(3), 1176-1248. <https://doi.org/10.1007/s11142-021-09609-5>
13. Dafermos, Y., Nikolaidi, M., & Galanis, G. (2018). Climate change, financial stability and monetary policy. *Ecological Economics*, 152, 219-234. <https://doi.org/10.1016/j.ecolecon.2018.05.011>
14. Delis, M. D., de Greiff, K., & Ongena, S. (2018). Being stranded on the carbon bubble? Climate policy risk and the pricing of bank loans. Swiss Finance Institute Research Paper, 18–10. <https://doi.org/10.2139/ssrn.3125017>
15. Ehlers, T., Packer, F., & de Greiff, K. (2022). The pricing of carbon risk in syndicated loans: Which risks are priced and why? *Journal of Banking & Finance*, 136, 106180. <https://doi.org/10.1016/j.jbankfin.2021.106180>
16. Engle, R. F., Giglio, S., Kelly, B. T., Lee, H., & Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies*, 33(3), 1184-1216. <https://doi.org/10.1093/rfs/hhz072>
17. European Banking Authority. (2025). 2025 EU-wide Transparency Exercise: List of institutions and mapped templates. <https://www.eba.europa.eu/risk-analysis-and-data/eu-wide-transparency-exercise>

18. European Commission. (2022). Commission Implementing Regulation (EU) 2022/2453 laying down implementing technical standards on prudential disclosures of ESG risks. https://eur-lex.europa.eu/eli/reg_impl/2022/2453/oj
19. European Parliament and Council. (2013). Regulation (EU) No 575/2013 on prudential requirements for credit institutions and investment firms. <https://eur-lex.europa.eu/eli/reg/2013/575/oj>
20. Flammer, C. (2021). Corporate green bonds. *Journal of Financial Economics*, 142(2), 499-516. <https://doi.org/10.1016/j.jfineco.2021.01.010>
21. Giglio, S., Maggiori, M., Rao, K., Stroebel, J., & Weber, A. (2021). Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34(8), 3527-3571. <https://doi.org/10.1093/rfs/hhab032>
22. Goss, A., & Roberts, G. S. (2011). The impact of corporate social responsibility on the cost of bank loans. *Journal of Banking & Finance*, 35(7), 1794-1810. <https://doi.org/10.1016/j.jbankfin.2010.12.002>
23. Grewal, J., Riedl, E. J., & Serafeim, G. (2019). Market reaction to mandatory nonfinancial disclosure. *Management Science*, 65(7), 3061-3084. <https://doi.org/10.1287/mnsc.2018.3099>
24. Griffin, P. A., Lont, D. H., & Sun, E. Y. (2017). The relevance to investors of greenhouse gas emission disclosures. *Contemporary Accounting Research*, 34(2), 1265-1297. <https://doi.org/10.1111/1911-3846.12298>
25. Hartzmark, S. M., & Sussman, A. B. (2019). Do investors value sustainability? A natural experiment examining ranking and fund flows. *The Journal of Finance*, 74(6), 2789-2837. <https://doi.org/10.1111/jofi.12841>
26. Hong, H., Li, F. W., & Xu, J. (2019). Climate risks and market efficiency. *Journal of Econometrics*, 208(1), 265-281. <https://doi.org/10.1016/j.jeconom.2018.09.015>
27. Houston, J. F., & Shan, H. (2022). Corporate ESG profiles and banking relationships. *The Review of Financial Studies*, 35(7), 3373-3417. <https://doi.org/10.1093/rfs/hhab125>
28. Ilhan, E., Sautner, Z., & Vilkov, G. (2021). Carbon tail risk. *The Review of Financial Studies*, 34(3), 1540-1571. <https://doi.org/10.1093/rfs/hhaa071>
29. Ivashina, V., & Scharfstein, D. (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, 97(3), 319-338. <https://doi.org/10.1016/j.jfineco.2009.12.001>

30. Jiménez, G., Ongena, S., Peydró, J.-L., & Saurina, J. (2014). Hazardous times for monetary policy: What do twenty-three million bank loans say about the effects of monetary policy on credit risk-taking? *Econometrica*, 82(2), 463–505. <https://doi.org/10.3982/ECTA10104>
31. Khwaja, A. I., & Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4), 1413-1442. <https://doi.org/10.1257/aer.98.4.1413>
32. Krueger, P., Sautner, Z., & Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3), 1067-1111. <https://doi.org/10.1093/rfs/hhz137>
33. Matsumura, E. M., Prakash, R., & Vera-Muñoz, S. C. (2014). Firm-value effects of carbon emissions and carbon disclosures. *The Accounting Review*, 89(2), 695-724. <https://doi.org/10.2308/accr-50629>
34. Murfin, J. (2012). The supply-side determinants of loan contract strictness. *The Journal of Finance*, 67(5), 1565-1601. <https://doi.org/10.1111/j.1540-6261.2012.01767.x>
35. Murfin, J., & Spiegel, M. (2020). Is the risk of sea level rise capitalized in residential real estate? *The Review of Financial Studies*, 33(3), 1217-1255. <https://doi.org/10.1093/rfs/hhz134>
36. Nini, G., Smith, D. C., & Sufi, A. (2009). Creditor control rights and firm investment policy. *Journal of Financial Economics*, 92(3), 400-420. <https://doi.org/10.1016/j.jfineco.2008.04.008>
37. Painter, M. (2020). An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2), 468-482. <https://doi.org/10.1016/j.jfineco.2019.06.006>
38. Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2), 550-571. <https://doi.org/10.1016/j.jfineco.2020.12.011>
39. Reghezza, A., Altunbas, Y., Marques-Ibanez, D., Rodriguez d’Acri, C., & Spaggiari, M. (2022). Do banks fuel climate change? *Journal of Financial Stability*, 62, 101049. <https://doi.org/10.1016/j.jfs.2022.101049>
40. Sautner, Z., van Lent, L., Vilkov, G., & Zhang, R. (2023). Firm-level climate change exposure. *The Journal of Finance*, 78(3), 1449-1498. <https://doi.org/10.1111/jofi.13219>

41. Scholtens, B. (2009). Corporate social responsibility in the international banking industry. *Journal of Business Ethics*, 86(2), 159-175. <https://doi.org/10.1007/s10551-008-9841-x>
42. Seltzer, L. H., Starks, L. T., & Zhu, Q. (2022). Climate regulatory risk and corporate bonds. NBER Working Paper, 29994. <https://doi.org/10.3386/w29994>
43. Sufi, A. (2007). Information asymmetry and financing arrangements: Evidence from syndicated loans. *The Journal of Finance*, 62(2), 629-668. <https://doi.org/10.1111/j.1540-6261.2007.01219.x>
44. Tang, D. Y., & Zhang, Y. (2020). Do shareholders benefit from green bonds? *Journal of Corporate Finance*, 61, 101427. <https://doi.org/10.1016/j.jcorpfin.2018.12.001>
45. Zerbib, O. D. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking & Finance*, 98, 39-60. <https://doi.org/10.1016/j.jbankfin.2018.10.012>