


Online Appendix

Appendices


Appendix A Additional Tables and Figures

Figure A1. Kernel Density of $\log(\text{intensity})$ - FULL sample




Sources: AnaCredit data and Authors' calculations.

Figure A2. Kernel Density of $\log(\text{intensity})$ - INC sample




Sources: AnaCredit data and Authors' calculations.

Figure A3. Kernel Density of $\log(\text{intensity})$ - INC sample (IPW corrected)




Sources: AnaCredit data and Authors' calculations.

Figure A4. Kernel Density of $\log(\text{intensity})$ - FULL sample (alternative)



Sources: AnaCredit data and Authors' calculations.

Figure A5. Kernel Density of $\log(\text{intensity})$ - INC sample (alternative)



Sources: AnaCredit data and Authors' calculations.

Table A1. Channel Decomposition

<i>Channels</i>	<i>Estimands</i>
Intensive margin	0.0000 (0.0095)
Reallocation among incumbents	−0.0897*** (0.0273)
Entry/Exit pricing	0.0977*** (0.0365)
Entry/Exit composition	−0.0037 (0.0067)

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

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

Table A2. Quantile Model

Panel A ($q = 0.10$)

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

Panel B ($q = 0.25$)

Model:	(1)	(2)	(3)	(4)
<i>Coefficients</i>				
θ	0.1232	-0.1389	-0.1896*	0.0408
	(0.0915)	(0.0749)	(0.0896)	(0.0649)
Controls	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel C ($q = 0.50$)

Model:	(1)	(2)	(3)	(4)
<i>Coefficients</i>				
θ	0.1334 (0.0775)	-0.1097* (0.0543)	-0.2269** (0.0693)	0.0848 (0.0474)
Controls	Yes	Yes	Yes	Yes
<i>Fixed-effects</i>				
Loan Type	Yes	Yes	Yes	Yes
Country-Size-Time	Yes	Yes	Yes	Yes
Bank-Time	Yes	Yes	Yes	Yes
Bank-Debtor	No	No	No	Yes

Panel D ($q = 0.75$)

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

Panel E ($q = 0.90$)

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

Clustered (Bank) standard errors in parentheses.

*Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$, $p < 0.15$.*

Table A3. Alternative Shock

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

Clustered (Bank) standard errors in parentheses.

*Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

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




Table A4. Staggered by Cohort

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

Point estimates correspond to the estimands described in the Subsection 7.2.

Clustered (Bank) standard errors in parentheses, via delta method.

Fit statistics are taken from the underlying regression (12).

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

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





Figure A8. Monthly Staggered Estimands (95% CI)



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

Table A5. Modules 1&2 vs. Module 3

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

Clustered (Bank) standard errors in parentheses.

*Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

Appendix B Firm-Level GHG Intensity Metric

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

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

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

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

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

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

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

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

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

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

Appendix C Data Description

Appendix C.1 Sample Construction

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

³⁶16,5% for the sample FULL and 20,3% for the sample INC.


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

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

Appendix C.2 Data Description

The kernel densities in Figure C1 confirm that the entire distribution of spreads overlaps strongly across the two samples: the incumbent restriction does not mechanically isolate a segment with systematically higher or lower pricing, but rather trims the extensive margin of short-lived relationships. This is important for the subsequent analyses: differences between FULL and INC estimates can be interpreted as differences in margins of adjustment (transitory vs continuing relationships), rather than artefacts of very different pricing distributions.

Figure C1. Credit Spread (Kernel Density)




Sources: AnaCredit data and Authors’ calculations.

46% depending on EU countries.

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


Figure C2. $\log(\text{intensity})$ (Kernel Density)



Sources: AnaCredit data and Authors' calculations.

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

Figure C3. Firm Size Composition by Carbon Intensity Quartile



Sources: AnaCredit data and Authors' calculations.

Figure C3 further documents how the composition of borrowers varies along the carbon-intensity distribution. In both FULL and INC, moving from the first to the fourth quartile of GHG intensity is associated with a pronounced shift from micro towards larger firms. Low-intensity loans (Q1) are predominantly granted to micro and small enterprises, whereas the share of medium and especially large firms rises steadily with intensity and becomes sizeable in Q4. This monotonic pattern is very similar across the FULL and INC samples, with only a mild tilt towards larger firms in INC, reflecting the fact that long-standing bank–firm relationships are more common among bigger borrowers. Econometrically, this is reassuring: restricting to incumbents does not fundamentally alter the mapping between carbon intensity and firm size, so differences between FULL and INC estimates are unlikely to be driven by a radically different firm-size mix across intensity bins.

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

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

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

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

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

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

Appendix D Multiple Banks vs Single Bank Test

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

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

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

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

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

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

We also estimate the following model using the same sample:

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

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

where λ_{cst} is the country $\times$ size $\times$ time fixed effects.

Based on results from Eqs. (13) and (14), we compare the supply-side shocks $\hat{\mu}_{bt}$. We find that supply-side shocks estimates do not vary substantially. The Pearson correlation coefficient is higher than 0.80 over the entire timespan for the FULL sample, and around 0.52 for the INC sample, implying that the *pseudo* demand-side control has not led to a misspecification (See Figure D1).

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

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

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

We find that the proportion of the variance that is explained by the supply-side shocks from Eq. (14) is around 0.42 and 0.51 (See Within R^2 in Table D1.A). These results suggest that the demand-side controls are similar.

In the second step, we compare the results of Eq. (13), based on multiple-bank firms sample, and Eq. (14), based on the full sample, to the extent that potential differences in the two supply-side shocks estimates can therefore be attributed to variation in supply-side shocks faced by multiple-bank firms and single-bank firms. Figure D2 presents the month-by-month correlation between $\hat{\mu}_{bt}^{FT}$ and $\hat{\mu}_{bt}^{CST}$. The Pearson coefficients are lower on average than in the first step (around 0.80) but they also vary more significantly. These results corroborate Degryse et al. (2019) findings and support the assumption that excluding single-bank firms could lead to biased estimates of the strenght of supply-side shocks. The regression procedure also support this assumption (See Table D1.B).

Table D1. Comparing Supply Shocks


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

Clustered (Bank) standard errors in parentheses.

*Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

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


Figure D1. Monthly correlation between supply shocks (multiple-banks sample)



Source: AnaCredit data and Authors' calculations.

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

Figure D2. Monthly correlation between supply shocks (multiple banks vs single bank)



Source: AnaCredit data and Authors' calculations.

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

Appendix E Proofs of Section 4

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

Proof Eq. (2)

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

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

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

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

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

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

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

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

Define:

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

Therefore, we have:

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

Define:

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

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

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

The zero-profit condition can be rewritten:

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

Solving for M :

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

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

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

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

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

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

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

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

Define

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

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

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

We rewrite the profit function as:

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

Differentiating w.r.t. r gives

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

Because

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

we obtain

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

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

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

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

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

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

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

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

and thus

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

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

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

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

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

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

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

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

Proof Eq. (4)

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

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

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

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

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

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

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

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

Define:

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

Therefore, we have:

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

Define:

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

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

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

The zero-profit condition can be rewritten:

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

Solving for M_τ :

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

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

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

Solving using the quadratic formula yields:

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

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

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

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

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

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

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

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

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

$F(r_\tau^*, \tau)$ for $\frac{dr_\tau^*}{d\tau}\big|_{\tau=0}$ yields

$$\frac{dr_\tau^*}{d\tau}\bigg|_{\tau=0} = - \frac{\frac{\partial F}{\partial \tau}\big|_{(r_0,0)}}{\frac{\partial F}{\partial r}\big|_{(r_0,0)}}$$

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

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

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

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

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

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

$$\frac{\partial F}{\partial \tau}\bigg|_{(r_0,0)} = -(1 + r_0)(1 - \delta r_0) = -M$$

It follows that:

$$\frac{dr_\tau^*}{d\tau}\bigg|_{\tau=0} = \frac{M}{1 - \delta - 2\delta r_0}$$

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

$$r_\tau^* \approx r_0^* + \frac{dr_\tau^*}{d\tau}\bigg|_{\tau=0} \tau = r_0^* + \frac{M}{1 - \delta - 2\delta r_0} \tau = r^* + \alpha \cdot \tau$$

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

Stochastic Dominance

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

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

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

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

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

and conditional payoff distributions.

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

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

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

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

Then, for any $\tau > 0$,

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

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

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

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

The unconditional CDF $F_X(x)$ is therefore

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

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

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

By construction of the main model,

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

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

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

1. For $x < 0$:

$$F_X(x) = 0 = F_{X|\tau}(x).$$

2. For $0 \leq x < X$:

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

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

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

3. For $x \geq X$:

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

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

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

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

Appendix F Disentangling Price and Composition Effects

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

$$\text{Latent price: } y_{bdt}^* = \theta [\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})] + \mu(x) + u_{bdt}, \quad (1)$$

$$\text{Selection: } G_{bdt} = \mathbb{1}\{V_{bdt} > 0\}, \quad (2)$$

$$V_{bdt} = \delta [\text{Post}_t \times \text{TreatBank}_b \times \log(x_{dt})] + \mu(x) + \eta_{bdt}.$$

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

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

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

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

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

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

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

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

Define the conditional grant probability

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

and the joint numerator

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

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

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

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

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

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

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

Differentiating under the integral sign (Leibniz),

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

Combining these pieces,

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

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

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

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

Two thought experiments help fix ideas. In a cross-section at fixed (b, t) across firms with different intensities, a regression for new-loan spreads primarily reflects *between-debtor* differences among those who do obtain a loan at t . Conversely, along the time dimension for a single pair (b, d) observed around the shock, regressing y_{bdt}^* on intensity identifies a *pure price* effect because composition is held fixed. Our unbalanced

panel—with one-off and repeated relationships—mechanically mixes these objects, and the channels decomposition is designed precisely to separate within-price, reallocation-among-incumbents, and entry/exit effects.

Appendix G Heterogeneous Effects of the Climate Stress Test

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

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

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

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

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

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

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

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

Table G1. Asymmetric Model

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

Clustered (Bank) standard errors in parentheses.

*Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

Table G1 reports the results on the INC sample⁴³ and the within-relationship specification. In INC, we find economically meaningful and statistically significant post-treatment slope shifts in both regimes, with a clear asymmetry. Thus, the post-shock change in the spread-intensity semi-elasticity for treated banks is more negative when the debtor's intensity is *decreasing* than when it is increasing. Interpreted through the lens of transition-risk integration, this pattern is consistent with treated banks conditioning their slope response not only on the level of emissions intensity, but also on its *direction of change*, reacting more strongly when intensity improvements materialize.

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

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

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

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

Table G2. Channel Decomposition (up vs. down)

<i>Channels</i>	<i>Estimands</i>
Intensive margin (up)	0.0000 (0.0100)
Intensive margin (down)	-0.0005 (0.0121)
Difference intensive (down - up)	-0.0004 (0.0104)
Reallocation among incumbents (up)	-0.0815*** (0.0286)
Reallocation among incumbents (down)	-0.1196*** (0.0220)
Difference reallocation (down - up)	-0.0381** (0.0182)

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

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

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

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

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

In this setup, $\theta^{(q)}$ measures how the post-CST slope differential between treated and control banks changes for loans to borrowers whose carbon intensity lies above the q -th quantile, relative to otherwise similar low-intensity loans.⁴⁴ Table G3 extends the static channel decomposition from Eq. (8) to the quantile-based Eq. (15). The results suggest that the CST-induced adjustment in carbon-risk pricing operates primarily through extensive margins, with marked heterogeneity across the emissions distribution.

Table G3. Channel Decomposition by Quantiles

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

Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

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

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

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

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

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

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

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

relationship lending (see, for example, [Fuchs et al., 2024](#)).

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

Appendix H Channels Decomposition

Appendix H.1 Algebraic Decomposition within the Incumbent Sample

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

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

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

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

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

We then work with the transformed variables

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

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

$$\hat{\theta}^R = \frac{\tilde{Z}'\tilde{y}}{\tilde{Z}'\tilde{Z}}. \quad (16)$$

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

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

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

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

We therefore define the *within-relationship* components:

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

so that

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

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

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

and set

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

By construction,

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

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

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

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

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

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

hence

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

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

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

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

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

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

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

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

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

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

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

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

where

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

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

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

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

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

Expand the numerator and denominator:

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

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

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

Define the within-relationship slope

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

and the between-relationship slope

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

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

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

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

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

Substituting into the numerator:

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

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

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

For the denominator, Lemma H.1 yields

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

Combining these expressions, we obtain

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

Define

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

Then we can rewrite

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

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

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

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

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

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

used in Eq. (7) is entirely driven by between-relationship variation among incumbents: it vanishes if and only if the pure between-relationship slope θ^B coincides with the within-relationship slope $\hat{\theta}^I$, and its magnitude is scaled by the share of total variation in Z that comes from between-relationship differences. In this precise sense, $\hat{\theta}^R - \hat{\theta}^I$ captures reallocation of credit across incumbent relationships within bank×time cells, over and above within-relationship repricing.

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

Appendix H.2 Decomposition of the Entry/Exit Margin

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

Setup. As in Appendix [Appendix H.1](#), let Z_{bdt} denote the triple interaction

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\hat{\theta}^F = \frac{Z'y}{Z'Z} = \frac{Z_I'y_I + Z_E'y_E}{Z_I'Z_I + Z_E'Z_E}. \quad (18)$$

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

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

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

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

where

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

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

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

Proof. We start from the definition (18):

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

By construction of the group-specific slopes,

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

so that we can write

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

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

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

Define

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

Then we can rewrite

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

which establishes part (i).

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

subsample). Therefore,

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

This completes the proof. $\square$

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

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

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

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

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

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

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

$$\hat{\theta}^k = \frac{\sum_{i \in \mathcal{I}_k} w_i^k Z_i y_i}{\sum_{i \in \mathcal{I}_k} w_i^k Z_i^2}, \quad (19)$$

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

We consider three samples:

- **FULL sample** ($k = F$): all granted loans; weights $w_i^F \equiv 1$.
- **INC sample** ($k = R$): only incumbents; weights $w_i^R \equiv 1$ for $i \in \mathcal{I}_I$ and no observations outside $\mathcal{I}_I$.

- **INC-IPW sample** ($k = R, w$): same incumbents as in INC, but each observation $i \in \mathcal{I}_I$ is assigned a positive weight $w_i^{R,w}$ constructed as

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

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

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

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

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

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

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

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

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

Hence each θ^k depends on two components:

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

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

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

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

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

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

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

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

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

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

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

By (21), we can write

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

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

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

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

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

Proof. By definition,

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

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

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

Thus

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

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

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

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

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

Proof. From (21),

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

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

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

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

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

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

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

Thus we can write

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

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

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

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

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

share the same pricing function m_I , Lemma H.6 implies that

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

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

Combining these two pieces yields the formal decomposition.

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

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

where:

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

Proof. The identity

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

is algebraic. Under Assumption H.4, we have

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

By Lemma H.5, the difference

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

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

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

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

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

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

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

Appendix I Placebo Shock

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

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

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

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

Table I1. Placebo shock: November 2019

	FULL	INC	INC-IPW	INT
<i>Panel A. Triple-difference estimates</i>				
log(GHG intensity)	0.0538** (0.0162)	0.0624*** (0.0149)	0.0684*** (0.0151)	−0.0098 (0.0204)
log(GHG intensity) × TreatBank	0.0452* (0.0216)	0.0280 (0.0217)	0.0234 (0.0212)	0.0138 (0.0209)
log(GHG intensity) × Post	−0.0032 (0.0203)	0.0046 (0.0165)	0.0053 (0.0173)	0.0070 (0.0084)
log(GHG intensity) × TreatBank × Post	−0.0246 (0.0225)	−0.0047 (0.0198)	−0.0053 (0.0203)	−0.0018 (0.0094)
Country–size–month FE	Yes	Yes	Yes	Yes
Bank–month FE	Yes	Yes	Yes	Yes
Instrument FE	Yes	Yes	Yes	Yes
Bank–debtor FE	No	No	No	Yes
Observations	966,863	539,384	539,384	539,383
R^2	0.3246	0.3324	0.3341	0.8326
Within R^2	0.0071	0.0082	0.0084	0.0001
<i>Panel B. Channel decomposition and wild-cluster bootstrap inference</i>				
Channel	Estimate	Share	Bootstrap s.e.	p -value
Intensive margin	−0.0018	0.072	0.0092	0.846
Reallocation among incumbents	−0.0029	0.119	0.0110	0.768
Entry/exit pricing at adjusted composition	−0.0193**	0.785	0.0084	0.016
Entry/exit composition in observables	−0.0006	0.024	0.0012	0.645
Total effect, FULL	−0.0246	1.000	0.0220	0.285

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

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

We consider a competitive banking market in which each bank decides whether to incorporate transition risk into its lending decisions. We use a simple binary-action coordination game with strategic complementarities, which is a special case of the noncooperative supermodular games analyzed by [Milgrom and Roberts \(1990\)](#). We model transition-risk integration as a coordination problem with strategic complementarities⁴⁶: when only a few banks integrate transition risks into their lending behavior, (i) discipline on carbon-intensive borrowers is weak because credit demand can reallocate to non-adopting banks, and (ii) the private effectiveness of screening and pricing is limited because firms' incentives to adjust remain weak. When many banks integrate transition risks into their lending behavior, leakage shrinks and the price signal becomes stringent economy-wide, increasing both private and social returns. This structure can generate multiple Nash equilibria (see, [Cooper and John, 1988](#); [Milgrom and Roberts, 1990](#)), which in turn motivates regulatory intervention to select the socially desirable equilibrium. The framework is deliberately stylized and should be interpreted as a proof-of-concept model: it abstracts from richer dynamics and institutional details, but it provides formal micro-founded substance for the argument that uncoordinated market outcomes may justify public intervention.

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

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

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

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

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

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

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

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

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

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

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

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

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

$$\Pi_i(C; x) = q_C(x)[m - \ell_C(x)] - \sigma, \quad (3)$$

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

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

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

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

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

From Eqs. (1)–(2),

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

so that

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

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

Differentiating (5) with respect to x yields

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

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

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

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

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

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

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

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

is continuous and strictly increasing in x on $[0, 1]$. In particular, the game exhibits in-

creasing differences in the sense of [Milgrom and Roberts \(1990\)](#): the private incentive to adopt C is higher when a larger fraction of other banks adopt C .

Proof. Continuity follows from the continuity of $q_C(\cdot)$, $q_{NC}(\cdot)$, $\ell_{NC}(\cdot)$ and $\eta(\cdot)$. Differentiating (5) with respect to x yields

$$\frac{\partial}{\partial x} \Delta \Pi(x) = \frac{D}{N} \left[\lambda \ell'_{NC}(x) + \lambda \eta(x) + (1 - \lambda + \lambda x) \eta'(x) \right],$$

see Eq. (6). Under Assumption J.1, we have $\ell'_{NC}(x) < 0$, $\eta(x) \geq 0$, $\eta'(x) > 0$, and $1 - \lambda + \lambda x > 0$ for all $x \in [0, 1]$. Hence the term in brackets is strictly positive, implying $\partial \Delta \Pi(x) / \partial x > 0$ for all $x \in [0, 1]$. This establishes that $\Delta \Pi(x)$ is strictly increasing. $\square$

Assumption J.3 (Stag-hunt boundary conditions). The payoff differential satisfies

$$\Delta \Pi(0) < 0 \quad \text{and} \quad \Delta \Pi(1) > 0.$$

Proposition J.4 (Existence and uniqueness of the adoption threshold). Suppose Assumptions J.1 and J.3 hold. Then there exists a unique $\bar{x} \in (0, 1)$ such that $\Delta \Pi(\bar{x}) = 0$. Moreover,

$$\Delta \Pi(x) < 0 \quad \text{for } x < \bar{x}, \quad \Delta \Pi(x) > 0 \quad \text{for } x > \bar{x}.$$

In the large- N approximation where each bank treats x as given, the individual best response is NC whenever $x < \bar{x}$, C whenever $x > \bar{x}$, and both actions are optimal when $x = \bar{x}$.

Proof. By Proposition J.2, $\Delta \Pi(x)$ is continuous and strictly increasing on $[0, 1]$. Assumption J.3 implies $\Delta \Pi(0) < 0$ and $\Delta \Pi(1) > 0$. By the Intermediate Value Theorem, there exists at least one $\bar{x} \in (0, 1)$ such that $\Delta \Pi(\bar{x}) = 0$. Strict monotonicity rules out multiple roots, so $\bar{x}$ is unique. Monotonicity further implies $\Delta \Pi(x) < 0$ for all $x < \bar{x}$ and $\Delta \Pi(x) > 0$ for all $x > \bar{x}$. The characterization of best responses follows immediately from the sign of $\Delta \Pi(x)$. $\square$

Given Proposition J.4, the symmetric best-response correspondence can be written as

$$\text{BR}(x) = \begin{cases} \{NC\} & \text{if } x < \bar{x}, \\ \{C, NC\} & \text{if } x = \bar{x}, \\ \{C\} & \text{if } x > \bar{x}. \end{cases}$$

Proposition J.5 (Nash equilibria and coordination structure). Under Assumptions J.1 and J.3, the following statements hold:

- (i) The profile in which all banks choose NC (i.e. $x = 0$) is a symmetric pure-strategy Nash equilibrium.
- (ii) The profile in which all banks choose C (i.e. $x = 1$) is also a symmetric pure-strategy Nash equilibrium.
- (iii) In a finite N -player game, there exists a symmetric mixed-strategy equilibrium in which each bank adopts C with probability $p \in (0, 1)$ such that $\Delta \Pi(p) = 0$. In the

large- N limit, this probability converges to the unique threshold $\bar{x}$ characterized in Proposition J.4. Under standard adjustment dynamics, $\bar{x}$ acts as an unstable tipping point separating the basins of attraction of $x = 0$ and $x = 1$.

Proof. (i) If all banks choose NC , then $x = 0$ and any unilateral deviation to C faces an adoption rate $x_{-i} \approx 0$. By Assumption J.3, $\Delta\Pi(0) < 0$, so deviating to C strictly lowers the deviator's payoff. Hence $x = 0$ is a pure-strategy Nash equilibrium.

(ii) If all banks choose C , then $x = 1$ and any unilateral deviation to NC faces $x_{-i} \approx 1$. Assumption J.3 implies $\Delta\Pi(1) > 0$, so adopting C yields a strictly higher payoff than NC . Deviating to NC is therefore unprofitable, and $x = 1$ is a pure-strategy Nash equilibrium.

(iii) In a finite N -player game, the usual arguments for two-action coordination games imply the existence of a symmetric mixed-strategy equilibrium in which each bank adopts C with probability $p \in (0, 1)$ that solves the indifference condition $\Delta\Pi(p) = 0$, interpreted as the adoption rate among other banks. As $N \rightarrow \infty$, this adoption probability converges to the unique threshold $\bar{x}$ in Proposition J.4, which separates regions where C or NC are strictly preferred. Under standard myopic best-response or replicator dynamics, $\bar{x}$ is unstable, while $x = 0$ and $x = 1$ are locally stable, yielding the familiar stag-hunt coordination structure. $\square$

Consequently, the game features two robust pure-strategy Nash equilibria:

$$x = 0 \quad (\text{all } NC), \quad x = 1 \quad (\text{all } C),$$

which correspond to a standard stag-hunt coordination structure. If others do not integrate transition risk (x close to 0), it is privately unattractive to be an early adopter; if others do integrate (x close to 1), deviating to non-adoption becomes unattractive because leakage disappears and the private benefits of risk integration are high.

Social welfare and systemic-risk externality. Let $E(x)$ denote the social benefit from lower systemic risk when the fraction of adopters is x , with $E'(x) > 0$ and $E''(x) \geq 0$, allowing for convexity when systemic resilience improves sharply once climate risk pricing becomes sufficiently widespread. We allow $E(x)$ to be an aggregate externality term that can scale with market size N (e.g. $E(x) = N \tilde{E}(x)$), so that $E'(x)/N$ remains $O(1)$ under the large- N approximation. We interpret $\ell_{NC}(x)$ and $\ell_C(x)$ as privately internalized expected credit losses. Thus, the welfare term $E(x)$ captures residual systemic externalities not priced by individual banks, such as fire-sale spillovers, network default cascades, and other real-economy costs of crises (Allen and Gale, 2000; Eisenberg and Noe, 2001).

Define social welfare as the sum of banks' profits and a residual systemic externality term:

$$W(x) = \sum_{i=1}^N \Pi_i(a_i; x) + E(x). \quad (7)$$

Under symmetry, let $\Pi_C(x) \equiv \Pi_i(C; x)$ and $\Pi_{NC}(x) \equiv \Pi_i(NC; x)$ for any bank i . Then welfare can be written as

$$W(x) = xN \Pi_C(x) + (1 - x)N \Pi_{NC}(x) + E(x).$$

Differentiating $W(x)$ w.r.t. x yields

$$\begin{aligned} W'(x) &= \frac{d}{dx} \left(xN\Pi_C(x) + (1-x)N\Pi_{NC}(x) + E(x) \right) \\ &= N\Pi_C(x) + xN\Pi'_C(x) - N\Pi_{NC}(x) + (1-x)N\Pi'_{NC}(x) + E'(x) \\ &= N[\Pi_C(x) - \Pi_{NC}(x)] + N[x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)] + E'(x). \end{aligned} \quad (8)$$

Rewrite Eq. (8) more compactly as

$$W'(x) = N \Delta\Pi(x) + N[x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)] + E'(x). \quad (9)$$

Dividing by N yields an expression for the marginal social gain per bank of an increase in the adoption rate:

$$\frac{1}{N}W'(x) = \Delta\Pi(x) + \underbrace{x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)}_{\text{endogenous payoff effects as } x \text{ changes}} + \underbrace{E'(x)/N}_{\text{systemic-risk externality}}. \quad (10)$$

Eq. (10) highlights the wedge between private and social incentives (Pigou, 2017). An individual bank, treated as atomistic in a large- N environment, compares $\Pi_C(x)$ and $\Pi_{NC}(x)$ taking x as given. Its decision rule is therefore characterized by the sign of $\Delta\Pi(x)$ alone. In contrast, a social planner internalizes not only the direct payoff differential $\Delta\Pi(x)$, but also (i) the impact of a higher x on all banks' payoffs, captured by $x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)$, and (ii) the marginal reduction in systemic risk, $E'(x)/N$.

To make (10) more explicit, recall that under symmetry

$$\Pi_C(x) = q_C(x)[m - \ell_C(x)] - \sigma, \quad \Pi_{NC}(x) = q_{NC}(x)[m - \ell_{NC}(x)],$$

with

$$\ell_C(x) = \ell_{NC}(x) - \eta(x), \quad \ell'_{NC}(x) < 0, \quad \eta(x) \geq 0, \quad \eta'(x) > 0,$$

and

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

Differentiating $\Pi_C(x)$ and $\Pi_{NC}(x)$ w.r.t. x gives

$$\begin{aligned} \Pi'_C(x) &= q'_C(x)[m - \ell_C(x)] + q_C(x)[- \ell'_C(x)] \\ &= q'_C(x)[m - \ell_{NC}(x) + \eta(x)] - q_C(x)[\ell'_{NC}(x) - \eta'(x)], \end{aligned} \quad (11)$$

$$\Pi'_{NC}(x) = q'_{NC}(x)[m - \ell_{NC}(x)] + q_{NC}(x)[- \ell'_{NC}(x)]. \quad (12)$$

Given $q'_C(x) = q'_{NC}(x) = \frac{D}{N}\lambda > 0$, $\ell'_{NC}(x) < 0$, $\eta(x) \geq 0$ and $\eta'(x) > 0$, Eqs. (11)–(12) highlight two effects of an increase in x on private payoffs: (i) changes in loan volumes $q_C(x)$ and $q_{NC}(x)$ through demand reallocation, and (ii) changes in expected climate transition monetary losses through $\ell_{NC}(x)$ and $\eta(x)$.

Substituting (11)–(12) into (9) yields an explicit, micro-founded expression for $W'(x)$ in terms of $\ell_{NC}(\cdot)$, $\eta(\cdot)$, and the demand system $q_C(\cdot)$, $q_{NC}(\cdot)$. Although the exact closed

form is not required for the results that follow, equations (11)–(12) make it immediate that $W'(x)$ includes terms proportional to $\ell'_{NC}(x)$ and $\eta'(x)$ that are ignored by individual banks when they take x as exogenous, in addition to the purely systemic-risk component $E'(x)$.

Proposition J.6 (Private versus social incentives). Under symmetry, social welfare is given by

$$W(x) = xN \Pi_C(x) + (1-x)N \Pi_{NC}(x) + E(x),$$

and the marginal social gain per bank of a marginal increase in the adoption rate is

$$\frac{1}{N}W'(x) = \Delta\Pi(x) + [x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)] + \frac{E'(x)}{N}.$$

Any interior solution $x^* \in (0, 1)$ to the planner's problem

$$x^* \in \arg \max_{x \in [0,1]} W(x)$$

must satisfy

$$\Delta\Pi(x^*) = -\left[x^*\Pi'_C(x^*) + (1-x^*)\Pi'_{NC}(x^*)\right] - \frac{E'(x^*)}{N}. \quad (13)$$

In particular, at any interior social optimum x^* , the private payoff differential $\Delta\Pi(x^*)$ is strictly negative.

Proof. The expression for $W(x)$ follows directly from the definition of social welfare and symmetry. Differentiating with respect to x yields

$$W'(x) = N \Delta\Pi(x) + N[x\Pi'_C(x) + (1-x)\Pi'_{NC}(x)] + E'(x),$$

which implies the stated formula for $W'(x)/N$. An interior optimum $x^* \in (0, 1)$ for the planner's problem must satisfy the first-order condition $W'(x^*) = 0$, which can be rewritten as (13). Under Assumption J.1, the term in brackets is strictly positive, and $E'(x^*)/N > 0$ by construction. Hence the right-hand side of (13) is strictly negative, implying $\Delta\Pi(x^*) < 0$. $\square$

Private vs. social optima. A symmetric decentralized equilibrium adoption rate is characterized, in the large- N approximation, by the fixed-point condition that each bank best responds to x based solely on the sign of $\Delta\Pi(x)$, which generates the low-adoption equilibrium $x = 0$ and the high-adoption equilibrium $x = 1$ described above. By contrast, a social planner chooses x to maximize $W(x)$ and internalizes both the impact of higher adoption on all banks' payoffs and the marginal reduction in systemic risk. As a result, there exist parameter constellations under which the decentralized economy is trapped in the low-adoption equilibrium $x = 0$ even though the planner's problem is maximized at a strictly higher x^* , potentially at $x^* = 1$ if $E(\cdot)$ is sufficiently large and/or convex.

Appendix K Wild-Cluster Bootstrap Inference

This appendix describes the wild-cluster bootstrap procedure used throughout the paper for inference on all channel decompositions, in the static, event-study, and staggered specifications.

General Setup

All empirical specifications in the paper are linear fixed-effects models of the form

$$Y = X\beta + D\gamma + F\lambda + u,$$

where Y is the outcome (credit spread), X collects the moderator(s) of interest (typically log-emissions intensity and its interactions with treatment and time indicators), D denotes additional controls, F are high-dimensional fixed effects, and u is the disturbance term. Estimation is conducted by OLS or WLS (when using IPW). Let $\hat{\beta}$ denote the estimated coefficient vector from a given specification, and let $\hat{\Theta} = f(\hat{\beta})$ be a linear functional of $\hat{\beta}$:

$$\hat{\Theta} = A\hat{\beta},$$

for some known matrix A . In all applications, $f(\cdot)$ maps regression coefficients into a set of four channels based on post-shock (or event-time-specific) slope coefficients measured in four samples. Let $\hat{\theta}_F$, $\hat{\theta}_R$, $\hat{\theta}_I$, and $\hat{\theta}_R^{\text{IPW}}$ denote the relevant slope coefficients in these four specifications (for a given post-shock indicator, event time, or staggered treatment-slope term). The four channels are always defined by the same linear mapping as in Eq. (8)

In the static specification, each $\hat{\theta}$ is a single post-shock slope coefficient. In the event-study specification, we obtain collections $\{\hat{\theta}_{F,k}\}$, $\{\hat{\theta}_{R,k}\}$, $\{\hat{\theta}_{I,k}\}$, and $\{\hat{\theta}_{R,k}^{\text{IPW}}\}$ indexed by event time k , and the same mapping is applied for each k . In the staggered design, the objects aggregated into FCST or ECST estimands are again linear combinations of underlying quantities, so that the channel decomposition remains linear in $\hat{\beta}$.

Wild-Cluster Bootstrap Algorithm

Inference relies on the wild-cluster bootstrap at the bank level, following [Cameron, Gelbach, and Miller \(2008\)](#) and [MacKinnon and Webb \(2017\)](#). The same algorithm is used for all specifications; only the mapping $f(\cdot)$ changes depending on whether we work with a single post-shock coefficient, an event-time profile, or a staggered aggregate. In all settings—static, dynamic, and staggered—the wild-cluster bootstrap therefore provides a unified, cluster-robust inference procedure for the four-channel decomposition reported in the main text.

Algorithm 1 Wild-Cluster Bootstrap for Linear Channel Functionals

- 1: **Inputs:** outcome Y , moderator(s) X , treatment and timing indicators, bank-level clusters $c(i)$, fixed effects F , weight matrix W (identity or IPW), number of replications B , wild multiplier law $\mathcal{G}$ (Rademacher in the baseline), confidence level $1 - \alpha$.
- 2: **Baseline estimation:**
 1. Estimate the relevant linear model (e.g., Eq. (6), (9), or (12)) on each sample, using the same set of regressors and fixed effects as in the main text.
 2. Recover the slope coefficients $\hat{\theta}_F$, $\hat{\theta}_R$, $\hat{\theta}_I$, and $\hat{\theta}_R^{\text{IPW}}$, as well as fitted values $\hat{Y}$ and residuals $\hat{u}$ on the effective estimation support of each regression.
 3. Map $(\hat{\theta}_F, \hat{\theta}_R, \hat{\theta}_I, \hat{\theta}_R^{\text{IPW}})$ into channel estimates $\hat{\Theta} = f(\hat{\beta})$ using the four linear combinations defined above.
- 3: **Cluster structure:** construct the union of all bank identifiers appearing in the four estimation samples and index them by $c = 1, \dots, C$. This ensures that the same multiplier is drawn for a given bank across all four regressions in each replication.
- 4: **for** $b = 1, \dots, B$ **do**
- 5: Draw i.i.d. cluster multipliers $\{g_c^{(b)}\}_{c=1}^C$ from $\mathcal{G}$ and set $g_{it}^{(b)} = g_{c(i)}^{(b)}$ for every observation (i, t) in each of the four samples.
- 6: For each sample $s \in \{F, R, I, R^{\text{IPW}}\}$, generate bootstrap outcomes on the original estimation support:

$$Y_{it}^{s,*(b)} = \hat{Y}_{it}^s + \hat{u}_{it}^s g_{it}^{(b)}.$$

- 7: Re-estimate the four regressions on $(Y^{s,*(b)}, X, F, W)$ with exactly the same fixed effects and weighting scheme, obtaining bootstrap slopes $\hat{\theta}_F^{*(b)}$, $\hat{\theta}_R^{*(b)}$, $\hat{\theta}_I^{*(b)}$, and $\hat{\theta}_R^{\text{IPW},*(b)}$ (or their event-time/staggered analogues).
- 8: Map these into a bootstrap draw of the channels:

$$\Theta^{*(b)} = f\left(\hat{\theta}_F^{*(b)}, \hat{\theta}_R^{*(b)}, \hat{\theta}_I^{*(b)}, \hat{\theta}_R^{\text{IPW},*(b)}\right),$$

applied either to the post-shock coefficients, to the whole set $\{\hat{\theta}_{\cdot,k}^{*(b)}\}_{k \in \mathcal{K}}$, or to the staggered aggregates.

9: **end for**

- 10: **Inference:** for each channel (and, in the dynamic case, for each k), compute:
 - the bootstrap standard error as the sample standard deviation of $\{\Theta^{*(b)}\}_{b=1}^B$;
 - percentile (or basic/reversed-percentile) confidence intervals as the empirical $\alpha/2$ and $1 - \alpha/2$ quantiles of the bootstrap distribution.
-