Absorb or Pass Through?
The Incidence of Supervisory Capital Pressure under Easy and Tight Money*

Paul Konietschke† Björn Imbierowicz‡

November 2026

Abstract

Who bears a bank capital requirement depends on the conditions under which it arrives. We follow the EU's supervisory stress tests of 2021 and 2023, the same instrument arriving on opposite sides of the euro area's shift from easy to tight money. After the 2023 exercise, harder-hit banks rebuilt capital and, once its guidance applied, rationed credit: among large listed borrowers, the same firm was 2.2 percentage points less likely to obtain new credit from its harder-hit lenders, renewals spared. Exposed firms did not substitute into bonds—new issuance fell alongside new loans, on the same supervisory calendar. The investment consequence fell on firms with thin cash reserves, significantly negative through the bottom quarter of the cash distribution and fading with liquidity. Around the 2021 exercise, banks absorbed the requirement instead: capacity was in place both times; only its cost moved. The contrast is consistent with state-dependent transmission.

Keywords: bank capital requirements; stress tests; monetary policy; corporate debt structure; investment.
JEL codes: E22; E52; G21; G28; G32.

*First draft: July 2026. Comments welcome.

†Goethe University Frankfurt and Technical University Darmstadt: konietschke@em.uni-frankfurt.de, konietschke@vwl.tu-darmstadt.de

‡Bundesbank: bjoern.imbierowicz@bundesbank.de

1 Introduction

Who bears a bank capital requirement? Like any incidence question, this one is decided by conditions instead of statutes: the bank absorbs the requirement when holding capital is cheap, and passes it down the intermediation chain when it is dear. The past decade made the question unusually concrete. Between 2015 and 2022 the European Central Bank expanded its balance sheet by trillions of euro and held policy rates at or below zero; from 2022 it raised rates by more than four percentage points and began running down its asset portfolio. Throughout both phases, European supervisors kept renewing supervisory capital pressure on banks on a recurring cycle: the European Union-wide stress tests select the largest banking groups, publish scenario-implied capital depletions bank by bank, and feed the results into the capital guidance set in the following supervisory cycle. Monetary transmission is known to depend on the state in which policy operates (Kashyap and Stein, 2000; Tenreyro and Thwaites, 2016) and to run through bank capital itself (Van den Heuvel, 2002), and lending responses to capital regulation are known to depend on bank and market conditions (Jiménez et al., 2017; De Jonghe et al., 2020; Bahaj and Malherbe, 2020). Where the same supervisory capital pressure comes to rest—inside the bank, or with the firms it lends to—and how that incidence varies with the monetary environment is, however, largely unexplored. This paper measures that incidence for listed euro-area firms, following the requirement from bank balance sheets through the credit and bond markets to firm investment.

Two exercises make the measurement possible: the European Union-wide stress tests of 2021 and 2023, the two most similar supervisory capital events European supervision has produced. They share their methodology and disclosure format, nine in ten of the 2021 banks reappear in 2023's slightly broader sample, and they landed two years apart on opposite sides of the monetary transition—the 2021 test amid renewed asset purchases at negative rates, the 2023 test published in the month the corporate-sector portfolio began to run off. Each exercise assigns tested banks a continuous dose, the adverse-scenario depletion of the capital ratio standardized within the exercise. Firms inherit exposure through pre-existing syndicated-loan relationships with tested banks, which carry between half and roughly seventy percent of the sampled firms' syndicated credit. Because scenario depletion varies systematically across national banking systems, all comparisons are made within country (Section 5). The aggregate backdrop poses the question the design answers at the firm level: through the tightening, net bank lending to euro-area firms collapsed from record levels into negative territory while net bond issuance stayed positive—the loan margin closed while the bond market, often described as the spare tire that lets firms ride out a bank-credit contraction (Grosse-Rueschkamp et al., 2019), stayed open.

The paper's central credit result is a within-firm comparison among the large listed borrowers the design measures. After the 2023 exercise, the same firm is 2.2 percentage points less likely to obtain new credit from its harder-hit lenders than from its less affected ones, against a 42 percent baseline origination rate. Because the comparison is across the lenders of the same firm, it absorbs everything the firm does or experiences. Renewals of maturing facilities show no such gradient, a difference the stacked specification tests directly. Banks maintained existing commitments and rationed new money, and the timing is the supervisory calendar's, not the monetary one's: the gradient is absent at the 2022 rate shock and at the publication of results,

and concentrates in the year the guidance informed by the exercise took effect. Behind it stands the bank response. Banks receiving larger 2023 depletions rebuilt capital while the exercise ran, most of the build accruing in the months banks projected their own balance sheets through the scenario, before results were public, and contracted assets once the guidance applied.

The real consequence lands where internal slack is thin. Interacting exposure with the firm's pre-window cash holdings shows the investment effect concentrating among low-cash firms: the marginal effect of exposure is significantly negative through roughly the bottom quarter of the cash distribution, about three percentage points of assets per unit of exposure at the tenth percentile, where cash is three percent of assets. It fades by the median and is indistinguishable from zero from the seventy-fifth percentile upward, where cash exceeds sixteen percent (Figure 4). The moderating interaction is continuous: a median split of the same variable finds nothing, because the moderation lives in the tails. Bond-market access, nearly uncorrelated with cash in this sample, marks a second, coarser dimension of the same capacity: the investment decline is carried almost entirely by exposed firms without access. Access is an ex ante marker of financial flexibility, not an exercised exit: exposed firms did not substitute into the bond market after the credit contraction but issued less there as well, on the same guidance calendar as the loan rationing, while the bond market itself stayed open and the secondary spreads of exposed issuers narrowed.

The 2021 exercise documented the absorption benchmark. Reconstructing each bank's requirement stack from public sources shows that standing capital headroom already covered the guidance the 2021 exercise implied for twenty-nine of the thirty banks the public stack covers: banks could stand still, standing still was free, and the record shows they did: no balance-sheet response on any margin, no capital build in the months banks knew their own results, no within-firm rationing, no financing response, no investment response. The within-firm arm rests on an accommodation-era sample that carries less within-firm dose variation. The benchmark's quiet is not a power artifact: even after inflating the 2021 estimates for that vintage's coarser dose disclosure, the bound on the 2021 annual investment response sits below the 2023 point estimate, and imposing the 2021 disclosure rules on the 2023 dose does not attenuate the 2023 estimate.

Banks absorbed the 2021 exercise; they passed the 2023 exercise through; and the incidence of the pass-through fell on the firms with the least internal slack. That contrast is consistent with state-dependent transmission of capital regulation, and funding conditions are a plausible mechanism: the 2023 lending contraction concentrates in wholesale-funded banks, whose liabilities repriced with the policy rate, and is cushioned by deposit funding, an interaction absent in 2021. The comparison has two important limits. The formal cross-event interaction is significant but estimated relatively imprecisely, so the state-dependence reading rests on the contrast in the transmission chain across the two episodes as well as an indicative pooled interaction. And the tightening period bundles the rate cycle, the balance-sheet runoff, and the era's non-monetary shocks, which no within-country design can separate, so the difference between the two exercises cannot be assigned to the monetary state alone.

The evidence behind the chain is clear: the within-firm rationing gradient and the incidence gradient through roughly the bottom quarter of the cash distribution hold under every inference

procedure we apply. The bond-issuance contraction is weaker, and the average annual investment response is somewhat imprecisely estimated and concentrated geographically. Section 7 sets out the accounting behind each estimate.

The paper joins two literatures that have rarely met. On the supervisory side, stress tests and capital requirement increases reduce credit supply, in the United States (Berrospide and Edge, 2010; Ivashina and Scharfstein, 2010; Acharya et al., 2018a; Flannery et al., 2017; Cortés et al., 2020; Doerr, 2021; Kandrac and Schlusche, 2021) and in Europe (Mésonnier and Monks, 2015; Gropp et al., 2019; Fraisse et al., 2020; Behn et al., 2022; Konietschke et al., 2022; Kok et al., 2023). The effects concentrate between announcement and implementation (Mésonnier and Monks, 2015; Bedayo and Galán, 2024), which is why the treatment here is a process running from the selection of tested banks to the application of guidance. That evidence largely stops at bank balance sheets: firm-level real outcomes in the tradition of Peek and Rosengren (2000); Chodorow-Reich (2014); Acharya et al. (2018b) are scarce for supervisory instruments in the euro area. On the monetary side, the purchase programmes of the expansion phase let eligible firms substitute bonds for loans and freed bank balance-sheet capacity (Grosse-Rueschkamp et al., 2019; Todorov, 2020; De Santis and Zaghini, 2021; Arce et al., 2021; Darmouni and Papoutsi, 2022), and loan-to-bond substitution is generally strongest when bank credit is tight (Becker and Ivashina, 2014; Adrian et al., 2013; Crouzet, 2018); the contraction phase is far less studied, with the available firm-level evidence on quantitative tightening concerning the United States and abstracting from bank capital (Acharya et al., 2023; Eren et al., 2025). Studies that interact capital regulation with the monetary stance do so at the bank-lending level and within a single regime (Aiyar et al., 2016; Imbierowicz et al., 2021; De Jonghe et al., 2020). The closest antecedent is Behn et al. (2025), who show on euro-area credit-register data that macroprudential capital tightening restrains lending more when monetary policy tightens simultaneously. Relative to that result this paper contributes the accommodation-side counterfactual—a documented episode in which the same instrument produced no detectable response anywhere along the chain—the extension from credit volumes to firm investment, and the incidence margin: which firms, sorted by their own internal slack, bear the pass-through. That incidence margin has a home literature of its own, in which cash buffers shield investment from credit-supply contractions (Duchin et al., 2010; Campello et al., 2010) and firms hold liquidity against precisely this contingency (Opler et al., 1999; Almeida et al., 2004; Sufi, 2009; Lins et al., 2010); what this paper adds is not the gradient itself but its attachment to a supervisory quantity instrument, the within-firm credit stage upstream of it, and the absorption counterfactual in which the same instrument produces no gradient at all (Section 6.5).

The paper proceeds as follows. Section 2 describes the supervisory cycle, the monetary environment, and the motivating facts. Section 3 states the mechanism as a simple bank problem dated on the supervisory calendar. Section 4 introduces the data and sample, Section 5 the empirical strategy. Section 6 contains the results, Section 7 the robustness and inference analysis, and Section 8 concludes.

2 Institutional Setting

2.1 The supervisory cycle

The European Union-wide stress tests are run every two to three years by the European Banking Authority together with the European Central Bank's banking supervision arm. Each exercise follows the same calendar. In late January the authorities announce the exercise, name the participating banking groups, and publish the macrofinancial scenarios. Over the following months banks project their balance sheets through a three-year adverse scenario under a common methodology. In mid-summer the results are published bank by bank: for the groups in the European Banking Authority sample, full templates including the capital ratio at the start and its minimum under the adverse scenario; for the additional groups tested directly by the supervisor, coarser ranges. The published depletions are the key quantitative input to the Pillar 2 guidance set in the Supervisory Review and Evaluation Process that concludes in December; the guidance-a supervisory capital expectation sitting above the binding requirement stack rather than a hard requirement—applies from the start of the following year, and the bank-level disclosure itself informs markets and counterparties on the day of publication (Konietschke et al., 2022).

Supervisory capital pressure of this kind is therefore not a one-day event. The sample is known months before the results; the results are public months before the guidance changes; and the new guidance applies months after that. Evidence on comparable instruments shows that lending adjusts mostly inside the announcement-to-implementation window rather than around formal compliance dates (Mésonnier and Monks, 2015; Bedayo and Galán, 2024). The empirical design treats the window boundaries explicitly: the publication of results dates the event, and the analysis probes alternative windows anchored at the exercise launch and at the start of the following supervisory cycle. Throughout the paper, a requirement change or capital-requirement pressure refers to this joint treatment—the published depletion and the guidance it feeds; the hard Pillar 2 requirement is set by the broader supervisory assessment, and the online appendix verifies that stress-test depletions do not move it.

Table 1 lists the two exercises, the anchor pair on which the comparison rests. The adverse-scenario depletion of the fully loaded common equity ratio varies widely across banks within every exercise—the cross-sectional dose the design exploits—and the set of banks with large depletions rotates considerably from one exercise to the next. Because the two exercises share methodology, disclosure, and most of their banks, a contrast across them is not also a contrast across changes in the test itself.

2.2 The monetary environment within the stress testing timeline

The exercises straddle a reversal in the monetary environment. We date the reversal by the corporate-sector purchase programme, its sharpest institutional calendar and the policy governing the bond market that firms losing bank credit would turn to. Net purchases, which had absorbed euro-area corporate issuance since 2016, ended in June 2022; holdings remained stable under reinvestment for the year in which policy rates rose by more than four percentage points; and runoff began in July 2023, the month the 2023 results were published. The 2021 test fell amid renewed pandemic-era purchases at negative rates; the 2023 test landed at the runoff boundary,

Table 1: The anchor exercises: coverage, capital constraint, and credit reach into the sample

	Publication	Banks (N)	Measurable	Median depletion (pp)	Credit share (aggregate)	Top-5 countries by bank count
2021 ST (accommodation)	2021-07-30	100	50	4.84	56.4%	DE (16), FR (11), IT (8), ES (8), BE (6)
2023 ST (tightening)	2023-07-28	110	70	4.34	70.6%	DE (22), IT (12), ES (10), FR (10), NL (6)

Notes: stress-tested banks are all banking groups covered by the relevant EU-wide exercise; the measurable subset reports a point capital-depletion estimate, the remainder only an ordinal capital band (never assigned a fabricated point value). Median depletion is in percentage points of common-equity-tier-1 capital under the adverse scenario, measurable subset only. Credit share (aggregate) is the dollar-weighted share of pre-window sample-firm credit carried by a stress-tested lead bank – a differently weighted statistic from Table 2's firm-average row (see that table's notes).

its entire aftermath inside the contraction. We refer to the two states as accommodation and tightening, with the 2022 rate increases treated as conventional monetary policy—common within country and period, absorbed by the fixed effects—rather than as a regime boundary. One caveat attaches to this classification. The 2023 publication nearly coincides with both the runoff onset and the peak of the rate cycle, so the tightening state bundles a balance-sheet and a rate-level component that these events cannot separate. It bundles the era's non-monetary shocks as well: the energy repricing and the growth slowdown of 2022–23 are common within country and period and are absorbed as such, but they are part of what tightening means here. No within-country comparison can strip the label down to monetary policy alone. The state dependence we document should therefore be read as accommodation versus tightening broadly, in the spirit of the state-dependent transmission literature (Kashyap and Stein, 2000; Tenreyro and Thwaites, 2016), rather than as an isolated effect of portfolio runoff.

What that reversal did to corporate debt finance is the paper's motivating fact. Figure 1 prices *new* bank credit against *new* bond finance at the same moment, with quantities beneath. Through the tightening the two instruments part ways on both margins: at-issue bond spreads through 2023–2025 sit at or below their 2021 levels and net issuance stays positive throughout, while the cost of new bank credit climbs from below 200 toward 350 basis points and net bank lending to non-financial corporations collapses from record levels into negative territory. Both cost series price only the transactions that occur, so the central fact is the divergence in direction, corroborated by the quantities, not the level gap. The same rotation appears in every large euro-area economy (Appendix Figure A1): the loan margin closed while the bond margin stayed open.

The same asymmetry appears at the firm level, keyed to the stress tests themselves: tracked in event time (Appendix Figure A2), the bond share of debt of more and less exposed firms moves in parallel through the 2021 publication, while around the 2023 exercise the two groups part after publication. These are descriptive patterns—composition and selection are not held constant—but they pose the question the design addresses formally: when the bank absorbs the requirement, does anything reach firms at all; and when it does not, who bears the adjustment?

The reversal reaches the banks through their funding. Figure 2 shows the realized funding costs of the tested banking groups sitting at half a percent—and compressed across banks—through the entire 2021 exercise, then repricing to two and a half percent across the 2023 exercise, with wholesale-funded banks pulling away from deposit-funded ones only once the tightening

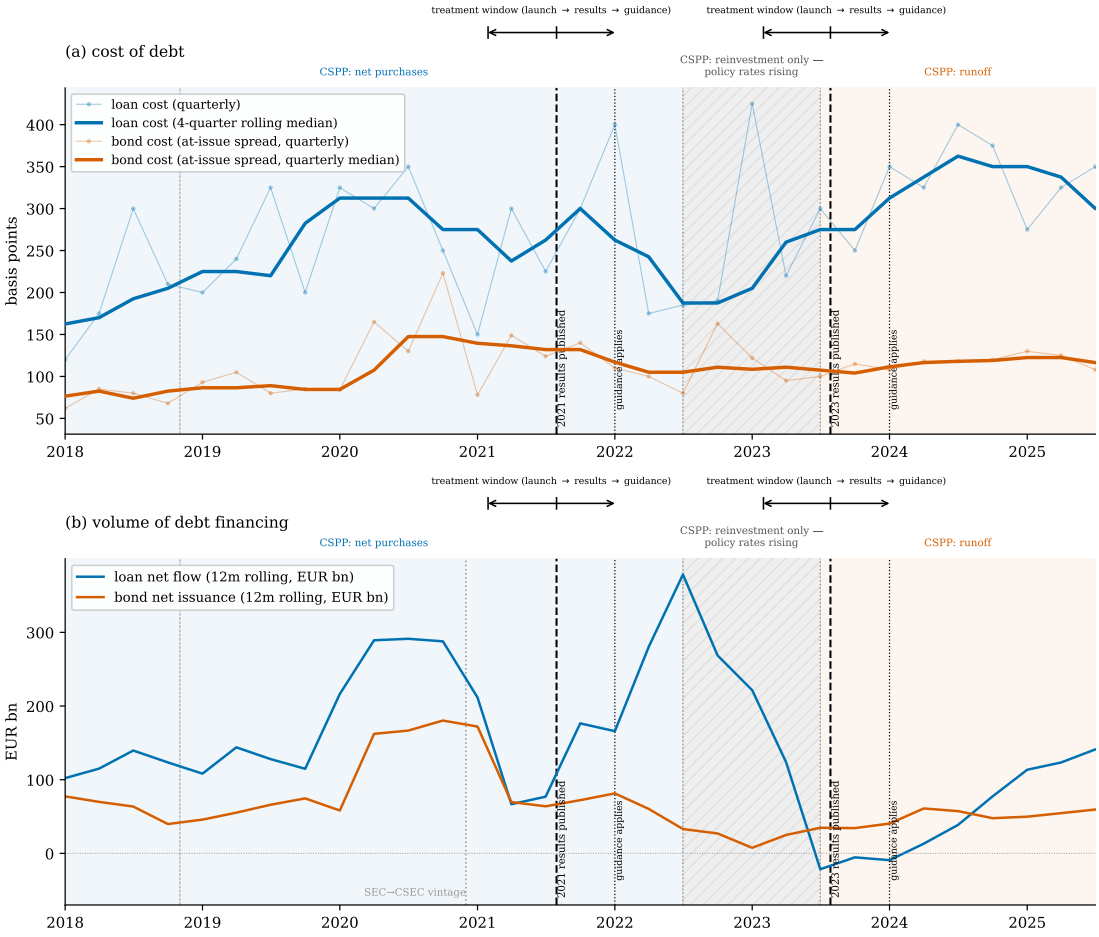


Figure 1: Corporate debt financing conditions across instruments, 2018–2025. The upper panel prices new financing in both markets: the cost of new syndicated loans (all-in spread drawn at origination, volume-weighted quarterly median and its four-quarter rolling median) and the at-issue spread to benchmark on newly issued bonds of non-financial corporate issuers (quarterly median and four-quarter rolling median), in basis points. The lower panel shows twelve-month rolling net flows of bank lending to euro-area non-financial corporations and net issuance of debt securities by the same sector, in billions of euro. Loan spreads are reported for about fifteen percent of tranches (thirty percent by volume); bond at-issue spreads for about two-thirds of issues. Shading marks the phases of the corporate-sector purchase programme: net purchases (blue), reinvestment only (hatched, July 2022 to June 2023, opening in the month of the ECB's first rate increase), and runoff (orange, from July 2023). For the two anchor exercises, the arrow spans the treatment window from the exercise's launch through the publication of results (dashed line) to the date the resulting supervisory guidance applies (dotted line within the arrow's span); the one dotted line outside the treatment arrows, in the lower panel, marks the issuance-series vintage transition. Sources: ECB bank-lending and debt-securities statistics; DealScan origination spreads; SDC and secondary-market bond pricing.

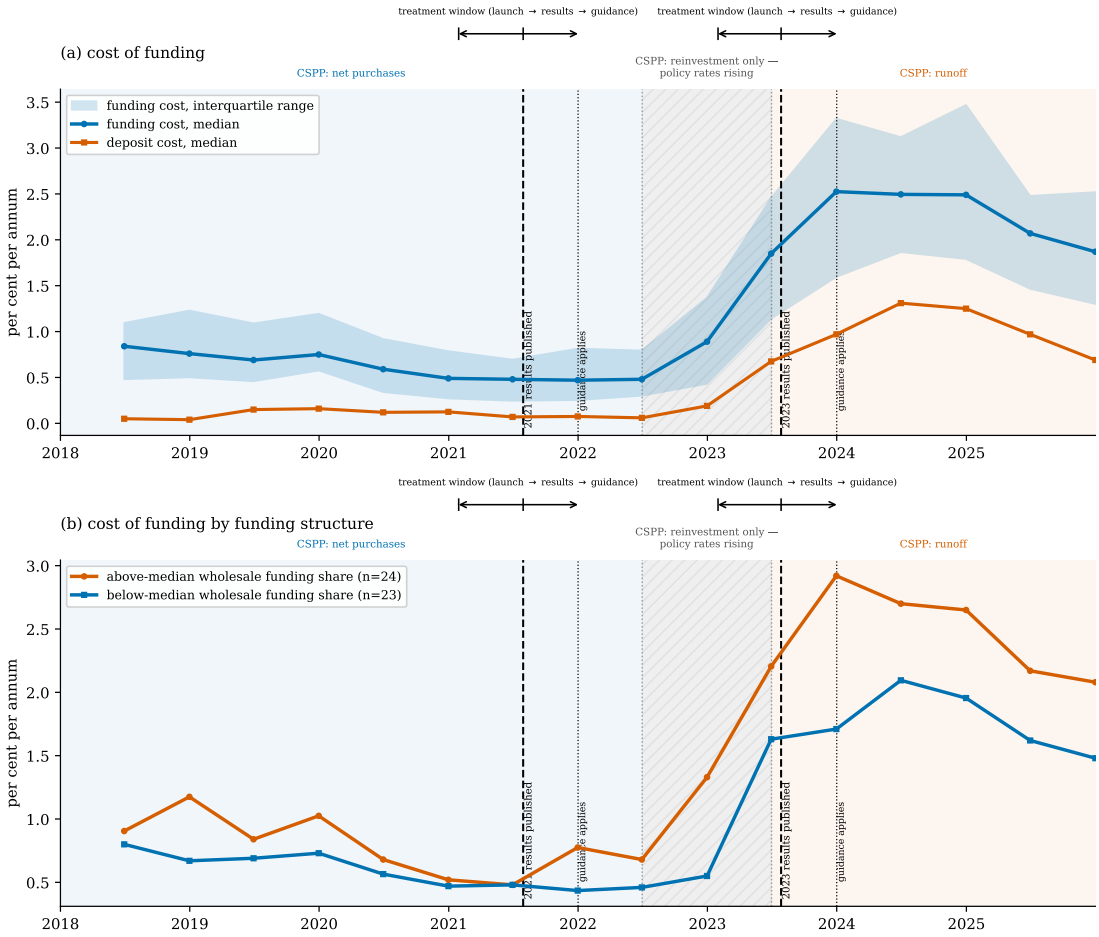


Figure 2: Realized funding costs of the stress-tested banking groups, 2018–2025. Panel (a): the median across tested groups of interest expense over average interest-bearing liabilities, with the interquartile range shaded; the lower line is interest expense on customer deposits over average customer deposits. Panel (b): the same funding-cost median, splitting banks by above- versus below-median wholesale funding share, fixed before 2022. Between 32 and 39 of the tested groups enter per half-year through their consolidated interim statements; a small number of groups whose reported interest expense reflects gross trading positions rather than funding are excluded. Shading marks the phases of the corporate-sector purchase programme as in Figure 1; for the two anchor exercises the arrow spans the treatment window from launch through results publication to the date the resulting guidance applies.

begins. The two exercises were published at opposite ends of the price of balance-sheet capacity: the 2021 test asked for capital when holding it cost nothing; the 2023 test asked for the same thing when every unit of balance sheet had begun to pay the policy rate.

3 A Simple Framework

The framework organizes the bank's choice between absorbing supervisory capital pressure in its existing buffer and adjusting its balance sheet; funding costs set the relative price of the two responses. It is not a structural claim: the funding-cost term that prices standing still is an assumption the data only partly discipline, through the bank-level funding interactions and their 2021 falsification. The bank holds a single asset, with risk weights normalized to one, so the framework collapses the distinction between total and risk-weighted assets and is silent on the risk-weight margin the data track separately. Appendix C states the three assumptions and

contains every derivation.

3.1 The bank's choice

A bank holds equity K against loans L and faces supervisory scrutiny beginning at the ratio $K/L = \rho + g(d)$: ρ is the standing requirement stack and $g(d)$ the Pillar 2 *guidance* fed by the bank's published stress-test depletion d, attaching at its application date τ, the start of the year after publication (Section 2). Debt costs move with the policy rate r at a bank-specific sensitivity β_b; equity carries a premium π that moves far less; and the syndicated market prices loans at an economy-wide spread over r, so the bank chooses loan *quantities*, never borrower-specific prices. Standing headroom covered the implied guidance in both vintages (Table 12, Figure 6)—a fact, not an assumption—so the bank can always stand still and absorb the guidance into its discretionary buffer. What the monetary state prices is what standing still *costs*: a bank near its scrutiny point rolls its marginal funding under the shadow of supervisory engagement, at an expected price $\beta_b r$—nil when money is easy, dear once the marginal liability reprices. The same tightening cheapens the repairs: the net cost of capital $\pi - \beta_b r$ falls as rates rise, and a declined marginal mandate forfeits only a competitive sliver while shrinking the balance sheet the expectation applies to. Appendix C turns this comparison into a bank-specific rate threshold $\bar{r}_b$, decreasing in the funding sensitivity β_b and in the dose d: below it the bank absorbs; above it, it acts on both margins. Deposit-rich banks, whose funding barely reprices (Drechsler et al., 2017), remain absorbers in both states.

3.2 Predictions

Prediction 1 (state-dependent pass-through). Below its threshold $\bar{r}_b$, a bank's published depletion produces no balance-sheet response; above it, lending falls and capital rises with d, and the response steepens in β_b. Because headroom covered the guidance both times, the regime contrast identifies the funding cost crossing the threshold, not feasibility. That interaction is a statement about banks; the framework does not aggregate it to the firm level, and the firm-level version is tested separately.

Prediction 2 (incidence on slack). Firms replace a rationed lead bank's capacity imperfectly, so the investment consequence of exposure falls where internal slack is thin and fades *continuously* as slack thickens: the cross-derivative of investment in exposure and slack is positive (Appendix C), a gradient rather than a threshold, so a median split understates what a continuous measure shows. Pre-window cash holdings are the primary marker of slack, bond-market access a secondary correlate; exposed firms' bond issuance contracted too, a fact the framework does not explain and Section 6 documents as open.

3.3 Two regularities the framework rationalizes

Regularity 1 (quantities, not prices). Because the marginal loan is priced at the market spread, capital pressure raises the bank's reservation spread and shows up as declined deals, not dearer ones: origination probability falls with the lender's d while transacted spreads carry no lender-specific mark-up; and because renewals of maturing facilities carry franchise rents that marginal

new money does not, new money is rationed first and renewals are spared.

Regularity 2 (the guidance clock). The scrutiny that enforces guidance begins at the application date τ, not at publication, and declining a deal early forfeits surplus against a charge that is not yet live. Rationing, where it occurs, is therefore dated by the supervisory calendar: at guidance application, not at the rate shock and not at the publication of results. The capital build, whose convex cost rewards spreading, may run ahead of that clock (Appendix C).

The framework is deliberately partial: it abstracts from general equilibrium, the bond market, and dynamics beyond τ, and the supervisor is a published mapping from depletion to guidance plus a calendar. Its message: supervisory capital pressure is a quantity instrument whose bite is priced by the monetary state—the cost of standing still.

4 Data and Sample

The analysis combines five building blocks. Syndicated-loan relationships come from DealScan, covering euro-area borrowers and the composition of their lending syndicates through early 2026. Firm fundamentals come from Worldscope for listed firms in the twenty euro-area countries at annual frequency and from Compustat Global at quarterly and half-yearly frequency, where reported. Debt structure comes from Capital IQ at the instrument level, which separates bonds from bank debt and dates each instrument; bond-market evidence is complemented by security-level issuance from the SDC new-issues database and weekly secondary-market spreads for the sampled firms' euro-denominated bonds. Bank-level treatment comes from the published stress-test disclosures; bank balance sheets come from the European Banking Authority's transparency exercises at semiannual frequency and from Orbis financial statements at quarterly frequency; the tested groups' realized funding costs come from their consolidated interim statements. Aggregate financing flows come from the European Central Bank's statistics on bank lending to non-financial corporations and on debt-securities issuance. Appendix B documents firm identification, lender consolidation, the match rate of each merge, and the construction of every panel.

From roughly fifteen thousand euro-area borrowers in the syndicated-loan data, the sample retains listed non-financial firms whose equity can be identified and whose lead lenders can be matched to tested banking groups: between roughly 350 and 470 firms per exercise enter the estimation samples, depending on the outcome panel; the descriptive tables report the firms observable in each exercise's anchor year. The restriction to syndicated borrowers selects large firms deliberately—these are the firms with plausible access to both bank and bond finance, for which the reallocation question is well posed.

Table 2 describes the firms. Syndicated borrowing is broad-based across the size distribution of listed firms, but its structure is not: larger firms borrow through more lead relationships and from more banks, and the incidence of bond-market access rises steeply with size. Tested banks are not a niche lender group: across exercises, banks participating in the stress tests account for between half and roughly seventy percent of the sample's aggregate pre-window syndicated credit (Table 1), with the reliance steepest among the sample's smaller borrowers (size split in the online appendix). Supervisory capital pressure on these banks moves the constraint on the marginal supplier of most sample firms' bank credit. Nor is dual-market presence rare: in every year of the sample period, roughly five hundred listed euro-area firms simultaneously hold an

active syndicated facility and have bonds outstanding. What is scarce is loan-price observation, not loan-taking: origination spreads are disclosed for only about fifteen percent of European tranches, which constrains cost-side comparisons throughout.

Two further tables set out the raw material on either side of the relationship. Table 3 describes the tested banks exercise by exercise: how many carry a point dose and how many are disclosed only as ranges, the number of countries they span, the capital they start with, and the depletion the adverse scenario implies. Table 4 does the same on the credit side for the multi-lender sample of the within-firm test: the pre-window tranches, firm–bank pairs, firms, and lead banks, the size and type of the facilities involved, and the rates at which pre-window relationships produced new loans and renewals afterwards. That sample conditions on at least two tested lead banks per firm and on no firm-level outcome, so it is not nested in Table 2's, and its tested-bank credit share is the highest of the three the paper reports—above the firm average for the firms in Table 2 and the aggregate share in Table 1—as conditioning on tested lead banks would suggest.

Table 2: Sample composition: syndicated-loan borrowers around the anchor exercises

	2021	2023
Firms (N)	402	349
Total assets (EUR mn)	3500.9 (1106.3, 11551.4)	4047.7 (1416.5, 14623.7)
Sales (EUR mn)	2154.9 (632.1, 5829.4)	3194.4 (1032.5, 8369.0)
Employment	7358 (2656, 22042)	8375 (3372, 25293)
Leverage (debt / assets)	0.34 (0.23, 0.47)	0.33 (0.21, 0.42)
Investment rate (annual)	0.026 (0.014, 0.048)	0.035 (0.020, 0.056)
Syndicated credit / assets	0.18 (0.09, 0.35)	0.13 (0.07, 0.24)
Lead-bank relationships	6.0 (3.0, 10.0)	5.0 (3.0, 9.0)
Stress-tested lead banks	4.0 (2.0, 7.0)	4.0 (2.0, 8.0)
Credit from tested banks (firm average)	69.0%	78.6%
Bond-market access	76.6%	77.4%

Notes: sample is non-financial, exchange-listed euro-area firms with a measurable stress-test exposure before the relevant publication. The table reports the anchor vintages (2021, 2023) in full; a split of the 2023 roster by asset-size terciles appears in the online appendix. Rows report the median (25th, 75th percentile in parentheses, not a standard error); the last two rows are firm-level means or shares. Credit and bank-relationship counts are pre-window.

Table 3: The tested banks, 2021 and 2023 exercises

	2021	2023
Banks with a point dose (N)	50	70
Banks disclosed as ranges only (N)	51	41
Countries with a point-dosed bank (N)	15	16
Starting fully loaded CET1 ratio, median (interquartile range)	15.71 (13.48–17.88)	15.48 (13.52–17.45)
Adverse-scenario minimum CET1 ratio, median	10.80	10.58
Depletion, mean (SD)	4.86 (2.07)	4.76 (2.47)
Depletion, interquartile range	3.33–6.21	3.17–5.59

Notes: The table describes the banks tested in the 2021 and 2023 EU-wide stress tests. Banks with a point dose (the measurable subset of Table 1) are those whose adverse-scenario depletion of the fully loaded CET1 ratio is disclosed as a figure; banks disclosed as ranges only are the additional SSM banks whose results are published in buckets, for which no continuous dose is imputed and which enter no estimate; the range-only count includes, in each exercise, one entity without a legal-entity identifier that Table 1 omits, which is why that table's rosters are one bank shorter. All remaining rows are computed on the point-dosed banks. The starting ratio is the fully loaded common equity tier 1 (CET1) ratio at the scenario's starting point, the adverse-scenario minimum is the lowest ratio over the scenario horizon, and depletion is the difference between the two; ratios are in percent of risk-weighted assets and depletion in percentage points. Source: stress-test disclosures of the European Banking Authority and the Single Supervisory Mechanism, assembled as described in Section 4.

Table 4: The multi-lender loan sample around the 2021 and 2023 exercises

	2021	2023
Firm–bank pairs (N)	2,433	2,351
Firms (N)	443	367
Lead banks (N)	38	49
Pre-window tranches (N)	5,290	4,597
Median pre-window tranche amount (USD mn)	333.1	416.7
Revolving and credit-line share of pre-window tranches	43.3%	45.9%
Term-loan share of pre-window tranches	31.8%	29.8%
New-origination rate, post window	31.4%	41.9%
Renewal rate among pairs with a maturing tranche	33.2%	41.8%
Credit from tested banks (firm average)	72.0%	83.1%

Notes: The sample is the regression sample of the within-firm test (Table 7): firm–bank pairs with a pre-window lead relationship, a measured dose, and at least two tested lead banks with distinct doses. Tranche counts, amounts, and facility-type shares are computed at the tranche–pair level, so a tranche with several tested lead banks is counted once per pair; amounts are DealScan's converted amounts, in US dollars. Revolving and term-loan shares do not sum to one hundred; the remainder is other facility types (bridge loans, Schuldschein, letters of credit, guarantees). The new-origination rate is the share of pairs with at least one new lead tranche in the post window. The renewal rate is the share of pairs with a pre-window tranche maturing inside the post window (485 pairs in 2021, 777 in 2023) that record an amendment, extension, or refinancing of it. The last row averages, across firms, the share of each firm's pre-window lead credit (split equally among its lead banks) held with tested banks; it differs from the firm average in Table 2 in its sample, which conditions on at least two tested lead banks. Shares and rates are in percent. Source: LSEG DealScan, which reports converted amounts in US dollars and provides no euro-converted field.

5 Empirical Strategy

5.1 Exposure

Each exercise assigns tested banks a continuous treatment intensity: the depletion of the fully loaded common equity ratio between the starting point and its adverse-scenario minimum, standardized within the exercise. Standardizing within the exercise is essential because scenario design and severity change from one test to the next, so raw depletions are not comparable across years; the dose is a bank's position in its own exercise's distribution. Banks whose results are disclosed only as ranges cannot carry a continuous dose and are excluded from the measure, and the affected credit share is reported wherever coverage does identifying work.

Throughout, we call z_b^e the bank's *dose* and reserve *exposure* for the firm-level measure built from it. Firms inherit exposure through their lending relationships. For firm f and exercise e, exposure (equation 1) is the credit-weighted average depletion of the firm's lead banks over a pre-exercise window,

$$E_f^e \;=\; \sum_b w_{fb}^e \, z_b^e, \tag{1}$$

where w_{fb}^e is the share of the firm's pre-window syndicated lead credit held by bank b *among its lead banks with a measured dose*: the weights sum to one over measured lead banks, so exposure is a renormalized average, not an average diluted by lenders without a continuous dose; z_b^e is the bank's standardized depletion. Weights are fixed before the exercise's results are published—for the 2023 exercise before it was even announced—so no bank's dose is known when the weights close and they cannot respond to the treatment. The measure uses lead arrangers because the lead holds the lending relationship: it screens, monitors, and decides renewal, so a lead bank under capital pressure is the relevant margin for the firm's continued access. Syndicate participants are ignored in the baseline; the robustness analysis shows that including them, or excluding the small set of firms whose only tested lenders are participants, leaves the estimates essentially unchanged. The magnitudes support that choice: 6.4 percent of tranches (2.2 percent of volume) have stressed participants but clean leads, roughly three-quarters of syndicates contain two or more stressed lender groups, and participant-level lending shares are observed on about a fifth of rows—too sparse to carry a share-weighted baseline.

Because adverse-scenario depletion varies mostly across national banking systems while firms borrow predominantly from domestic banks, raw exposure partly ranks countries rather than firms. All headline estimates therefore use exposure demeaned within the lead banks' country and include country-by-time fixed effects: identification comes from firms whose banks did unusually badly relative to other banks of the same banking system, compared across firms within the same country and period.

5.2 Estimation

The two exercises are stacked as repeated events. For each exercise, outcomes are aligned in event time around the publication date; each event's window ends before the next exercise would contaminate it, the 2023 window before the publication of the following exercise and the 2021

window at the cap set out below. In the stacked panel we estimate

$$y^e_{ft} \;=\; \beta\,(\text{post}^e_t \times E^e_f) \;+\; \gamma\,(\text{post}^e_t \times E^e_f \times \text{tight}^e) \;+\; \alpha^e_f \;+\; \lambda^e_{ct} \;+\; \varepsilon^e_{ft}, \tag{2}$$

where post^e_t indicates periods after exercise e's publication, tight^e indicates exercises conducted under tightening, and the fixed effects are firm-by-event (α^e_f) and country-by-period-by-event (λ^e_{ct}); β measures the response to exposure under accommodation and γ the additional response under tightening. In estimation, E^e_f is always the within-country demeaned exposure—denoted E_{dm} in the tables—whose in-sample standard deviation is about 0.44 in the 2023 panel and between 0.36 and 0.45 across events; per-standard-deviation magnitudes in the text use the quoted estimate's own panel. Equation (2) is estimated in two forms: event by event, where only β is identified and the anchored panels of Section 6 report it, and pooled across regimes, where γ is identified and appears only in the formal regime-difference tests. Stacking each exercise with its own clean window follows the event-by-event design of Cengiz et al. (2019) and avoids the contamination that arises when already-treated units serve as controls under staggered treatment (Goodman-Bacon, 2021; Callaway and Sant'Anna, 2021)—a live concern here, since every large bank is tested in every exercise with a changing dose. Event-time paths are estimated by local projections with the same clean-window discipline, horizon by horizon (Dube et al., 2025), which accommodates exactly this repeated, non-absorbing treatment.

The monetary chronology is the one Section 2 sets out; the cap on the 2021 window is a supervisory-calendar cap, not a monetary one: the baseline post-window runs to the end of 2022, closing when the next exercise's sample and scenarios became public and anticipation began, and a variant ending with net purchases in June 2022 is reported as the expansion-phase robustness window. Under this chronology fiscal year 2022 closes under reinvestment, before the runoff began, so the 2021 exercise has a usable annual post-year alongside its quarterly and half-yearly outcomes; the 2022 rate increases fall inside it and are absorbed by the fixed effects, with a sector-by-period variant reported as the explicit rate-cycle robustness (Appendix Table A18). Event-time exhibits use a uniform grid of four pre- and four post-publication half-years, estimated through the half preceding the following exercise's publication, with that boundary marked and the half of the following exercise's launch shaded inside the window; the static post-windows keep the caps just described. Annual outcomes use fiscal years aligned to calendar time by period-end dates, financing outcomes use quarters, and sub-annual real outcomes use interim data whose coverage—and its limits—are reported with the results. The remaining exercises of the supervisory cycle stay outside the comparison: those before 2020 ran under a different methodology and the pandemic separates their aftermaths from the 2021 test, while the 2025 exercise's aftermath is still too short to carry annual outcomes and will supply a second tightening-state event once longer observation accrues. Their estimates are collected in the online appendix.

Inference respects the level at which the treatment varies. Standard errors are clustered by firm in the firm-level panels, and by banking group in the bank-level regressions and in the within-firm pair panel, whose dose varies across a firm's banks; Section 7 reports lead-bank and two-way clustering variants for every headline estimate. Because the effective variation sits at the level of roughly fifty to seventy banks per exercise—about twenty-five effective shocks

once exposure weights are accounted for—clustered inference is supplemented three ways. First, randomization tests permute the bank-level dose within country and exercise and rebuild firm exposure from the permuted doses, holding the entire exposure network fixed; for interaction terms, the studentized version of the permutation is the primary inference, since studentization protects the test's level in small effective samples. A within-country regression of the dose on pre-event bank observables probes the exchangeability the permutation assumes, and rejects on exactly one observable: the bank's own prior-vintage dose, which persists across exercises. Because the dose is correlated with its prior-vintage value, we additionally report permutations stratified by prior-dose tercile alongside the unconditional version. Second, exposure-robust standard errors in the sense of Adão et al. (2019) and Borusyak et al. (2022) are computed from the equivalent shock-level regression, whose numerical equivalence to the firm-level estimate we verify exactly; the permutation and the exposure-robust standard errors formalize the same thought experiment—shares held fixed, shocks reassigned—by different routes. Third, we adjust inference for multiple testing within prespecified families of related outcomes (Online Appendix, Section OA.8): the headline coefficients form one family, the mechanism interactions a second, the bond-margin tests of Section 6.4 a third, and the two bank-level funding interactions their own two-cell family, with stepdown family-adjusted p-values (Romano and Wolf, 2005) reported for each.

The design has five boundaries. Its identified object is incidence, not credit-supply elasticities in the sense of matched-register estimation (Khwaja and Mian, 2008; Amiti and Weinstein, 2018): headline exposure aggregates a firm's syndicate, the multi-lender subsample supports a within-firm comparison of a firm's own lenders (Section 6), and the syndicated market the design observes is the marginal source of most sample firms' bank credit—the elasticity of a firm's total borrowing to a single lender's supply shift is a different question. It cannot separate the balance-sheet and rate-level components of the tightening state. Firms do not match to banks at random, and fixing the weights before each exercise rules out reverse causation of the weights but not sorting on unobserved trends; exposure should accordingly be read as an intent-to-treat measure over pre-existing relationships, with the flat pre-event paths—shown for the anchor pair in Figure 5 and for the remaining outcomes in the event-study exhibits—carrying the identifying burden. And the comparison across regimes rests on few clean events, so the formal accommodation-versus-tightening difference is estimated against considerable sampling uncertainty (Section 7). Finally, the multi-lender subsample is a listed-firm frame, and the robustness analysis rebuilds the pair panel on the full borrower universe, listed and unlisted alike, to test how far the design's estimate travels beyond it (Appendix Table A10).

6 Results

The results follow the 2023 transmission chain from the bank down: the bank response, the allocation of credit within firms, the calendar both credit margins obey, the bond margin, and the incidence on investment. The 2021 exercise then supplies the benchmark—the same instrument, the same disclosure format, largely the same banks, landing in the other monetary state and absorbed where it landed.

6.1 The bank response under tightening

Banks receiving larger 2023 depletions rebuilt capital while the exercise ran. The common-equity numerator of the capital ratio rises with the dose—four percent per unit over the exercise year, significant at the one-percent level—while risk-weighted assets barely move (Appendix Table A1). Table 5 dates the build *inside* the exercise: most of it—3.7 of the 4.3 percent—accrues between launch and publication, the months in which banks project their own balance sheets through the scenario and so know their own outcomes before markets do; after publication the capital margin is quiet, and the build is capital-only, with risk-weighted assets flat in every window. A dose-correlated earnings windfall from repricing assets could produce a passive build on a similar calendar, and the public record discriminates part of the way: issuance of additional-tier-one and hybrid instruments *declines* with the dose—the build did not come from fresh hybrid capital, consistent with a retention response—while dividend cuts are too rare in these samples to carry a coefficient. Which earnings funded the build stays open; that banks built it while running their own templates does not.

Table 5: Bank capital around the exercises: the private and the public phase

Window	Δlog CET1 capital	Δlog RWA	N
Panel A: 2021 exercise (accommodation)			
Placebo (year before launch) (2019-12-31 to 2020-12-31)	-0.021 (0.063)	-0.009 (0.023)	43
Pre-publication (2020-12-31 to 2021-06-30)	+0.013 (0.014)	+0.019** (0.010)	44
Publication half-year (2021-06-30 to 2021-12-31)	-0.004 (0.008)	-0.003 (0.007)	43
Following year (2021-06-30 to 2022-12-31)	+0.022 (0.047)	-0.019 (0.020)	44
Panel B: 2023 exercise (tightening)			
Placebo (year before launch) (2021-12-31 to 2022-12-31)	-0.012 (0.020)	+0.003 (0.007)	57
Pre-publication (2022-12-31 to 2023-06-30)	+0.037*** (0.013)	+0.005 (0.010)	61
Publication half-year (2023-06-30 to 2023-12-31)	+0.008 (0.010)	+0.005 (0.007)	61
Following year (2023-06-30 to 2024-12-31)	-0.009 (0.013)	+0.000 (0.014)	58

Notes: outcome is the change in the common-equity (CET1) capital level (log) and in risk-weighted assets (log) per unit of within-country demeaned dose (bank-level cross-section, one observation per bank, country fixed effects, robust HC1 standard errors in parentheses); * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Windows are dated by half-end stocks relative to each exercise's own publication date; the pre-publication window ends four weeks before results are published, when banks' own template results are largely final internally. The placebo window is the year immediately before the exercise's launch.

Once the resulting guidance applied, the same banks contracted. Figure 3 and Table 6 contain the bank-level evidence, for both exercises; the 2021 columns—flat on every margin—belong to the benchmark of Section 6.7. Within country—the discipline all comparisons in this paper obey, because scenario depletions differ systematically across national banking systems—total assets of

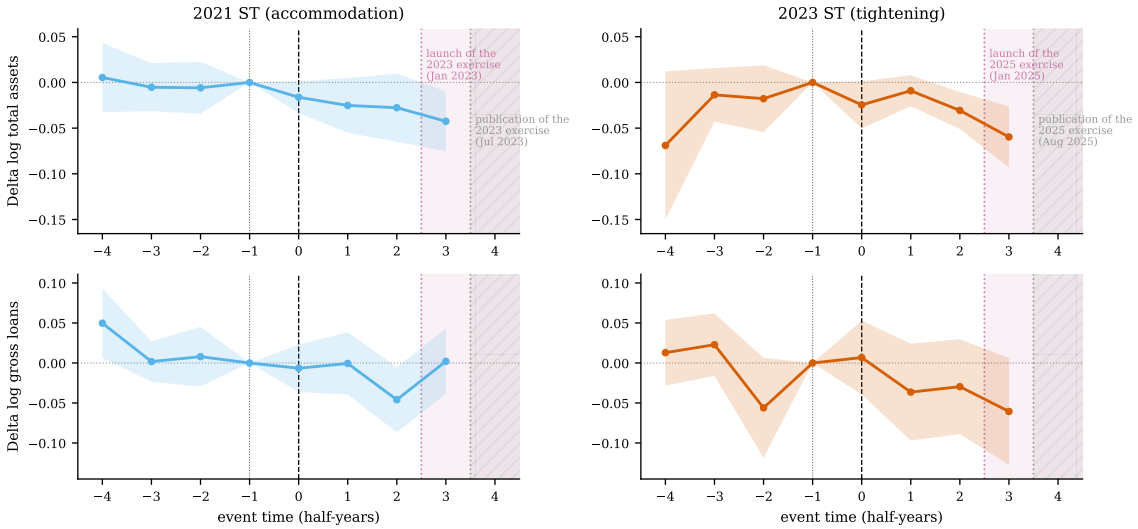


Figure 3: Bank balance sheets around the 2021 and 2023 exercises. Event-time coefficients of cumulative log changes in total assets (top row) and gross customer loans (bottom row) on the standardized adverse-scenario depletion, within country, in half-years relative to the publication half; the reference period is the half preceding publication, and both columns share each row's vertical scale. Shaded bands are 90 percent confidence intervals, clustered by bank; the dashed vertical line marks publication. Each window is estimated through the half preceding the following exercise's publication, with the closing boundary labeled and a lighter interior band marking the half in which the following exercise was launched. Deposits, corporate credit exposures, and risk density appear in the appendix version (Figure A4).

more affected banks fall by three to four and a half percent per unit of dose over the two years after the 2023 exercise, estimated plainly or within country, stable when the sample is extended to every bank with usable annual statements. Gross customer loans decline once country differences are absorbed, at the margin of significance. Deposits are flat: the contraction is not a deposit run but an adjustment on the asset side. A central-bank refinancing repayment cannot produce this joint pattern: it extinguishes a liability against excess reserves, a zero-risk-weight asset, so total assets fall while risk-weighted assets and common-equity capital have no reason to move. One test the bank panel cannot carry is a formal cross-vintage difference. Pooling the two vintages' four-quarter asset and loan changes and interacting the dose with the tightening indicator returns a null on both margins (+0.007 on assets, −0.037 on loans, studentized randomization p of 0.82 and 0.47, on 90 and 89 bank-events), an interaction a panel this size cannot power. The bank-side contrast accordingly rests on the joint pattern of the event studies (Appendix Table A14, notes).

Why banks would absorb in one state and pass through in the other has a natural funding-side reading: standing still is cheap while deposits reprice slowly and central-bank funding is abundant, and costly once liabilities reprice. The lending margin supports it. Among the banks affected by the 2023 exercise, the reduction in customer loans concentrates in banks funded by wholesale markets, whose liabilities repriced with the policy rate, and is significantly cushioned by deposit funding; the same interaction is flat around the 2021 exercise, when the deposit facility rate did not move (Appendix Table A3). This is evidence consistent with the funding mechanism, short of proof: the interactions sit at the six-percent level under the randomization schemes of Section 7, do not survive the stratified variant, and two Portuguese banks move the wholesale cell by a fifth. Its firm-level counterpart is a bounded null (Section 7).

Table 6: Bank bridge: capital-depletion dose and balance-sheet response, 2021 (accommodation) vs. 2023 (tightening)

Outcome	N	(1) OLS	(2) + country FE	(3) + FE, wtd.
2021 stress test (accommodation)				
Gross loans	28	-0.015	-0.029	-0.031
		(0.017)	(0.035)	(0.037)
Total assets	29	-0.007	0.024	0.049
		(0.015)	(0.042)	(0.079)
Customer deposits	28	0.011	-0.031	-0.027
		(0.018)	(0.050)	(0.064)
NFC exposure	49	-0.031	-0.038	0.000
		(0.020)	(0.029)	(0.030)
REA density (pp)	49	0.006	0.001	0.001
		(0.004)	(0.006)	(0.007)
2023 stress test (tightening)				
Gross loans	35	-0.019	-0.104*	-0.187
		(0.023)	(0.056)	(0.149)
Total assets	36	-0.029**	-0.044***	-0.054***
		(0.012)	(0.014)	(0.017)
Customer deposits	35	-0.036	-0.042	-0.007
		(0.029)	(0.045)	(0.042)
NFC exposure	65	-0.059**	0.004	-0.006
		(0.028)	(0.063)	(0.035)
REA density (pp)	65	0.012*	0.010	0.001
		(0.007)	(0.010)	(0.006)

Notes: sample is stress-tested banks with an Orbis (quarterly) or EBA Transparency Exercise (semiannual) balance-sheet match. Outcome is the log change (or, for REA density, the percentage-point change) in the bank-level item from a pre- to a post-publication snapshot over each vintage's own clean post window (2021: 6 quarters/3 half-years; 2023: 8 quarters/4 half-years); robust (HC1) standard errors in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) is OLS with no controls; column (2) adds bank-country fixed effects; column (3) further weights by pre-window size. Ladder and conventions match the Orbis-based corroboration table in the online appendix (numerically reconciled for the 2023 total-assets row). A blank cell is not estimable; nulls are reported as results, not omitted.

A within-firm comparison of loan and bond costs would sharpen the cost side, but origination pricing is disclosed too rarely to follow individual firms' loan costs over time; the pricing evidence the data do support is assembled with the bond margin below.

6.2 Rationing within the firm

How the contraction reached individual borrowers can be taken one level deeper, and this is the paper's central credit estimate. Firms that borrow from several tested banks at once—about three hundred and seventy in the 2023 window, with a median of four to five lead relationships—allow the allocation to be identified within the firm: comparing the same firm's banks to one another absorbs everything the firm itself does or experiences, leaving only differences across its lenders. This is the matched-lender design of Khwaja and Mian (2008); its identifying assumption is that

the firm's demand does not vary across its own banks in a way correlated with the dose, an assumption pre-window lead relationships are well suited to satisfy. Table 7 reports the result. After the 2023 exercise, the same firm was less likely to obtain new credit from its harder-hit banks: a one-unit difference in dose lowers the probability of a new origination by 2.2 percentage points against a 42-percent base rate (randomization $p = 0.003$), with committed amounts and facility counts confirming, and the gradient unchanged when the firm's relationship sizes are controlled. The estimate holds under every inference procedure of Section 7, and no single bank moves it by more than about a fifth (Appendix Table A15). Renewals of maturing facilities show no such gradient. The origination and renewal responses also differ directly in the stacked specification: the renewal margin's dose response exceeds the origination margin's by +3.5 percentage points (coefficient randomization $p = 0.030$; studentized 0.088), and the renewal-only response is a tight zero whose 80-percent detectable-effect bound, 0.032, is on the order of the origination gradient itself. Banks maintained existing commitments and rationed new money—the ordering relationship lending predicts, existing borrowers protected first (Bolton et al., 2016). The same test's 2021 arm carries an opposite-signed difference (−4.6 percentage points) in a vintage where the origination margin itself is a zero: 0.158 under the studentized permutation, primary for interaction-derived statistics, and 0.039 under the coefficient randomization. The 2021 renewal margin also rests on 485 of the stacked panel's rows against the origination margin's 1,948. We report the difference and do not build on it (Appendix Table A14, Panel B). The gradient is, if anything, stronger outside Spain—for firms and for banks alike—so the supply shift was euro-area-wide.

Table 7: Within-firm pass-through: new credit across the same firm's tested lenders

	2021 (accommodation)	2023 (tightening)
Panel A: New origination (within firm)		
Incidence of new origination	-0.0075	-0.0220***
	(0.0164)	(0.0076)
Asinh(new origination amount)	-0.0924	-0.2150***
	(0.1251)	(0.0565)
Log(1+new tranche count)	-0.0153	-0.0239***
	(0.0228)	(0.0088)
N (pairs / firms / banks)	2,433 / 443 / 38	2,351 / 367 / 49
Panel B: Renewal of maturing facilities		
Renewed within 365 days of a maturing facility	-0.0426	+0.0039
	(0.0285)	(0.0159)
N (pairs / firms / banks)	472 / 102 / 33	768 / 122 / 46

Notes: sample is firm-bank pairs with a pre-window lead relationship at firms with more than one stress-tested lender and within-firm dose variation (Panel B additionally requires a maturing pre-window tranche); firm-by-event fixed effects, bank-cluster robust standard errors (the firm-and-bank two-way cluster reaches the same conclusion at the five-percent level, not shown). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ from the bank-cluster p-value. On the 219 firms present in both event samples, the 2021 and 2023 incidence point estimates are statistically indistinguishable (2023 : −0.016, $p = .068$; 2021 : −0.017, $p = .578$); 2021's null reflects roughly 1.6 times less within-firm dose variance and correspondingly lower power, not a demonstrated zero. The 2023 incidence gradient is robust to controlling for within-firm relationship size (pre-window credit share): −0.021 ($p = .008$) with the control, against −0.022 without it.

Around the 2021 exercise the same specification finds nothing, though on the firms present in both samples the two events' point estimates are statistically indistinguishable, the accommodation-era sample carrying less within-firm dose variation; the within-firm evidence accordingly establishes transmission under tightening and is silent on the comparison across monetary states, which rests on the bank-level aggregates above.[1] Beyond this margin, supervisory corporate credit exposures show no individually significant movement and risk composition barely moves, with one suggestive exception—a tilt of the 2021 origination mix toward riskier loan purposes among more affected banks, reminiscent of risk-taking responses documented under supervisory data (Konietschke et al., 2022).

The estimate above is identified on listed firms, and how far it reaches beyond them is a question the loan data can answer directly. Rebuilding the pair panel on the full euro-area borrower universe gives 7,945 distinct borrowers across the four exercises, most of them unlisted; within the 2023 arm, 1,418 firms carry within-firm dose variation against the 367 listed firms the baseline estimate rests on, and 1,673 against 443 around the 2021 exercise. On the baseline pairs themselves the baseline specification reproduces, at -0.020 per unit dose (randomization $p = 0.028$) against the -0.022 reported above. Extended to the whole universe the response is roughly half that magnitude, -0.012. Splitting the borrowers by listing status—on the wider of the two extended constructions, not on the extension just quoted—puts the response at -0.013 for listed firms and -0.005 for unlisted ones, the latter indistinguishable from zero. Large listed borrowers are where the incidence of the retrenchment lands, and the measured unlisted cell does not show the same strength. How far the extended coefficient falls below the baseline depends on how the extended panel is built, and the two constructions fall on opposite sides of the materiality threshold, so the appendix reports both rather than choosing between them (Appendix Table A10). The extension is a bounded external-validity check, not a second estimate of the same object.

6.3 The guidance calendar

The timing of the rationing identifies the pressure behind it. Traced half-year by half-year, the within-firm gradient is absent around the 2022 rate shock and around the publication of results itself, and concentrates from 2024 onward—the year the guidance informed by the exercise began to apply to the banks. The two margins of the bank's response thus run on different clocks: the capital build answers to what the bank knows, accruing while banks ran their own templates; the credit supply answers to when the supervisor's guidance applies. The rationing does not sit on the central-bank refinancing calendar either, whose repayments were common-timed across banks, not dose-timed.

The bond margin runs on the same clock. Exposed firms' new bond issuance—the subject of the next subsection—contracts most sharply in the same halves in which the within-firm rationing peaks. A joint statistic tests the alignment: comparing the strength of the two margins'

[1] The financing quantity behind the decline lies beyond public data. Firm-level *total* bank borrowing does not track exposure (the diagnostic F statistics sit below one in every specification), but the design could only detect a borrowing shortfall an order of magnitude larger than the investment decline implies, so the totals are uninformative about the financing wedge; the margins that are measurable, new bank credit and new bond issuance, both contracted. Credit-supply elasticities of investment are accordingly not identifiable here; they are a different object from the incidence this paper measures.

half-by-half responses in the guidance-application halves of 2024 against the three preceding post-publication halves, and summing over the two margins, gives $S = +2.94$, with a one-sided randomization p of 0.018 and a family-adjusted p of 0.043 within the family of bond-margin tests in which it was run (Section 6.4). Because the statistic was motivated by the observed timing profiles, we treat it as corroborative rather than predictive evidence; the 2025 exercise provides the out-of-sample check (Section 8). What it establishes is shared timing, not a shared mechanism: the loan rationing and the issuance contraction sit on the guidance calendar together.

6.4 The bond margin

If the requirement never leaves the bank under accommodation but closes the loan margin under tightening, exposed firms' financing mix should be quiet around the 2021 exercise and move around the 2023 one. That is what it does, and under tightening it moves on both external margins at once. The 2021 quiet is what absorption implies: nothing reached the firms, and there was nothing to substitute away from (Section 6.7). Around the 2023 exercise, exposed firms' *new* external finance contracted as a whole. Alongside the rationed bank credit, they issued less in the bond market: about 0.07 fewer new issues per half-year per unit of exposure after publication, against the placebo stillness of the same margins around the 2021 exercise (Table 8). The issuance estimate is significant under exposure-robust standard errors and studentized randomization but not under the family adjustment (Section 7), and is corroborated on amounts and counts. The bond share of exposed firms' debt still drifts up over the same halves, an arithmetic of stocks rather than a rotation (Appendix A.2). The spare tire did not carry the firms losing bank credit.

Table 8: The issuance margin: new bond issues of exposed firms

	2021 (accommodation)	2023 (tightening)
Any new issue (per half-year)	+0.010 (0.029)	-0.017 (0.018)
New issues (count, per half-year)	+0.026 (0.075)	-0.070* (0.041)
Issue volume (log(1 + EUR))	+0.066 (0.199)	-0.138 (0.136)
Firm-halves	3,688	4,992
Firms	466	387
Panel window	2019H1–2022H2	2019H1–2025H1

Notes: firm–half-year panels of new bond issuance (SDC) on each exercise's core roster; the regressor is post-publication × within-country demeaned stress-test exposure (E_{dm}), with firm and country-by-half fixed effects, standard errors clustered by firm; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The 2021 window closes before the 2023 exercise launch. Randomization p-values (bank-dose permutation, 1,000 draws) accompany each cell in the replication package; the 2023 issue-count cell, the estimate quoted in the text, has randomization $p = 0.006$.

The market these firms issued less into was open, and it did not turn them away. Exposed firms' *maturing* bonds were replaced at the same rate as everyone else's over the broad instrument universe; exposed issuers were, if anything, downgraded less than others; and their secondary-market spreads *narrowed* after publication—by roughly twenty-five basis points at the peak, with flat pre-trends and no counterpart around the 2021 exercise (online appendix)—investors

pricing a shrinking stock more richly, not penalizing it. The contraction sits in *new public deals* (Appendix A.2). Whatever closed for these firms, it was not the bond market: investor-side rationing is rejected on every margin the data observe.

Nor was it prices, on either instrument. Loan pricing shows no differential by exposure—across comparable borrowers or within the same bank's own history—with designs that would have detected differentials above roughly a third of the economy-wide loan repricing over the period; at-issue bond pricing for exposed issuers is equally quiet. Prices moved with the monetary state, for everyone; within the detectable range, the supervisory dose sorted who got credit, not what it cost. The bond market did close once in these years, but on the rate calendar, not the supervisory one (Appendix A.2).

What ties the two contractions together is their clock: the same firms that were rationed by their harder-hit banks issued fewer new bonds, and both margins contract on the supervisory guidance calendar (Section 6.3).

Why exposed firms issued less is the one link the record leaves open, and the leading candidate has been probed directly. The institutional reading is complementarity: new public issuance runs through the relationship banks—they arrange the packages new bonds fund, provide the bridges, and dominate every sample firm's bookrunning—so a lead bank rationing new commitments could close the gateway to new deals while leaving routine rollovers, and the market itself, untouched. We test the hypothesis on three margins, and none provides positive evidence for it. One is precise enough to exclude an economically meaningful effect: the within-firm rationing is not concentrated in gateway facilities—revolvers, backstops, credit lines—relative to term loans (-0.002, standard error 0.002), and its detectable-effect bound of 0.007 is a third of the 0.022 gradient the channel would have to route, so an effect of the relevant size would have shown. The facility classification covers 73 percent of the stacked panel, short of the 80 percent standard. The other two are underpowered: the co-occurrence interaction (-0.016, standard error 0.070) carries a detectable-effect benchmark of 0.11, and the new-public-versus-programme split ($+0.011$, standard error 0.046, the wrong sign) a bound of 0.13 against an equivalence bound of 0.10, each wider than the 0.07 decline the channel would have to explain. Appendix Table A11 collects the test cells. Three simpler candidates—demand anticipation, credit quality, and the purchase programme's withdrawal—are examined in Appendix A.2; none accounts for the decline. On the relationship margin, the loan-side dose survives a horse race against a bookrunner-weighted dose built in parallel to it (-0.272 per unit, randomization $p = 0.041$, against the underwriting dose's -0.115), but the underwriting channel is too imprecisely measured to rank against the lending one—its standard error is 0.140 on the 99 firms whose bookrunner history is measurable, so the comparison locates the contraction on the lending relationship without excluding the underwriting one (Appendix A.3). If the decline reflected lost underwriter certification, it should concentrate where lead relationships are strongest. Interacting exposure with pre-window lead-relationship strength finds nothing on any issuance margin, in either vintage (all randomization p-values above 0.2), on a cell underpowered in the same way: a standard error of 0.035 on the count margin against a decline of 0.070. Both margins contract on one clock, prices are quiet, the market is open, and access is selective. The mechanism behind the issuance contraction remains open: the gateway margin was tested and unsupported, and the co-occurrence, deal-type, and

Table 9: Incidence on internal slack: exposure interacted with pre-window cash

Sample	Coef. (SE)	Cluster p	Perm. p (coef.)	Stud. perm. p	N (firms)
Baseline (2023, full sample)	$+0.0174^{***}$ (0.0048)	<0.001	0.002	0.004	1,343 (354)
Ex-Spain (2023)	$+0.0164^{***}$ (0.0050)	0.001	0.007	0.004	1,196 (315)
2021 placebo (re-pinned, not re-estimated)	+0.0031 (0.0021)	0.143	0.223	–	742 (406)
Pre-trend placebo (ex-Spain cut)	+0.0064 (0.0046)	0.161	0.295	0.283	598 (313)

Notes: coefficients on the triple interaction of the post-publication indicator, within-country demeaned exposure (E_{dm}), and the firm's pre-window cash holdings (demeaned; per-standard-deviation statements in the text use the estimation panel's own dispersion), in the annual investment-rate panel around the 2023 exercise, with firm and country-by-year fixed effects. Standard errors clustered by firm in parentheses; the cluster p is the analytic p-value. Permutation p-values reassign the bank-level dose within country and exercise and rebuild exposure (coefficient and studentized variants). The 2021 placebo row restates the 2021-exercise estimate; the pre-trend row estimates a pseudo-post window before the actual event on the ex-Spain cut. The tail-sensitivity diagnostics behind the text are in Appendix Table A6. Sources: annual firm fundamentals matched to stress-test disclosures as described in Section 4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

certification margins were tested at power public data cannot make decisive.

6.5 Incidence on internal slack

Does the financing contraction reach real activity, and for whom? The sharpest statement the data support is distributional. In the annual investment panel of the 2023 exercise, interacting exposure with the firm's pre-window cash holdings—a continuous measure of internal slack—shows the investment consequence concentrating where cash is thin: the triple interaction is $+0.0174$ per standard deviation of cash (clustered $p < 0.001$; randomization $p = 0.002$; studentized randomization $p = 0.004$; Table 9). Figure 4 turns the interaction into the paper's central real-effect exhibit: the marginal effect of exposure on investment, traced across the cash distribution. Through roughly the bottom quarter of the cash distribution the effect is significantly negative under every scheme (Appendix Table A14, Panel A): -3.2 percentage points of assets per unit of exposure at the tenth percentile of cash, where cash is three percent of assets (all four schemes at the five-percent level), and -2.5 at the twenty-fifth (every scheme at ten percent; studentized randomization $p = 0.096$). By the median, where cash is ten percent of assets and the point estimate is -1.8, the delta-method and exposure-robust intervals still exclude zero but the studentized randomization does not ($p = 0.20$), and the median-level estimate does not survive excluding Spain or restricting to the sub-annually covered majority. We therefore restrict the robust incidence result to approximately the bottom quarter of the cash distribution. From the seventy-fifth percentile upward, where cash exceeds sixteen percent of assets, the effect is indistinguishable from zero. The moderation is a continuous gradient, not a threshold: a within-country median split of the same variable finds nothing, because the identification lives in the tails of the cash distribution.

The interaction survives the family adjustment ($p = 0.003$), the stratified randomization ($p = 0.005$), exposure-robust standard errors sharper than the clustered ones, the exclusion of Spain (retaining 89 percent of the sample, studentized $p = 0.004$), and every single-bank exclusion, all of which leave it positive. The 2021 counterpart is quiet, and the interaction's

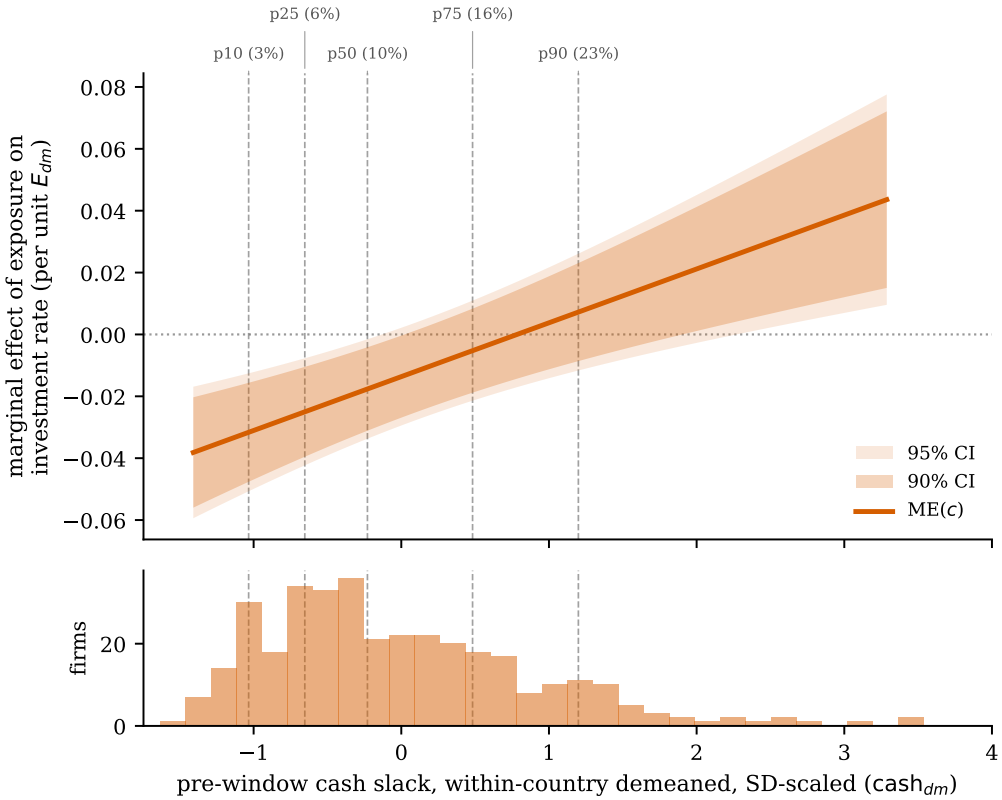


Figure 4: The marginal effect of exposure on investment across the cash distribution. The solid line is the marginal effect of within-country demeaned exposure on the annual investment rate, $\beta_1 + \beta_3 c$, as a function of the firm's pre-window cash slack c (within-country demeaned, scaled by the cross-sectional standard deviation of the cash-to-assets ratio), from the headline specification of Table 9; shaded bands are 90 and 95 percent confidence intervals, clustered by firm. The robustness of the curve's level varies across the distribution: the negative range robust to every inference scheme—delta-method, coefficient and studentized randomization, and exposure-robust—extends through roughly the 25th percentile of cash; between the 25th percentile and the median the bands shown exclude zero but the studentized randomization does not; from the 75th percentile the effect is indistinguishable from zero (Appendix Table A14, Panel A). Dashed vertical lines mark percentiles of the firm-level cash distribution, labeled with the raw cash-to-assets ratio at each percentile; the lower panel counts firms per bin. The horizontal axis is truncated at four for display; the two firms beyond it enter every estimate. Marginal effects are per unit of exposure; one within-country standard deviation of exposure is roughly 0.44 units, so per-standard-deviation magnitudes are a little under half the vertical axis. Sample: the 354 firms of the annual investment panel around the 2023 exercise. Sources as in Table 9.

event-time path sits on zero through the pre-window and turns positive—investment better protected where cash is high—only in the exercise's aftermath (Appendix Figure A3). Because the tails carry the identification, the full tail diagnostics are reported in Appendix Table A6: the interaction survives winsorizing at both conventional cuts, rank and log transformations, and the exclusion of either cash decile alone—including the bottom decile, where the decline concentrates—and weakens only when both tails are removed at once, cuts that remove the identified contrast rather than test it. A pandemic-downturn placebo seals the test-specificity: around the postponed 2020 exercise, with cash fixed at its pre-pandemic level, the same triple is a null with a detectable-effect bound at the size of the 2023 estimate—a downturn alone does not produce the gradient; a stress test under tightening does.

The gradient itself has a home literature. That low-cash firms cut investment hardest when

external finance contracts is the template result of Duchin et al. (2010) for the 2008 crisis, anticipated in the precautionary theory of corporate cash (Opler et al., 1999; Almeida et al., 2004) and echoed in crisis-time survey and liquidity-management evidence (Campello et al., 2010; Lins et al., 2010; Sufi, 2009). What is new here is the gradient's attachment to a supervisory quantity instrument with a bank-level dose, the within-firm credit stage upstream, and the absorption counterfactual in which the same instrument produces no gradient at all. The standing objection in that literature—that low cash proxies for being constrained on unobservables rather than for slack itself—meets three answers in this design. Cash is nearly orthogonal to bond-market access here (correlation -0.04), so it is not access relabeled. In a horse race against firm size, the classic constraint proxy, the cash interaction survives unchanged ($+0.0176$, studentized randomization $p = 0.004$) while the size interaction carries nothing ($+0.002$, studentized $p = 0.74$); age and dividend-payer status are unavailable in the estimation sample (Appendix Table A14 notes). And the 2020 pandemic placebo above is aimed at exactly this reading: a generic constrained-firm gradient should have appeared in the sample's severest downturn, and did not.

A coarser marker points the same way. Splitting the exposed firms of the 2023 exercise by whether they had bond-market access before the test—any outstanding bond in the two pre-window years—the decline is carried almost entirely by firms without access: exposed firms that could issue held investment roughly flat, while exposed firms without access cut it by nearly a percentage point (conditional means in Appendix Table A5); under the coarser tercile construction the no-access-minus-access gap among high-exposure firms is -1.8 percentage points, against a quiet 2021 counterpart (Table 10). Access marks capacity, and it is genuinely distinct from the cash measure, with which it is nearly orthogonal: the access split and the cash gradient are separate facts about the same incidence. Within the tightening vintage the formal access triple is imprecise, so the access margin is a pattern in conditional means and bootstrap contrasts, not a separately identified elasticity. The natural rival reading, that firms without bond access are smaller and more cyclical and cut investment more in any downturn (Gertler and Gilchrist, 1994), meets the same discriminating placebo as above: in 2020, when the scheduled exercise was postponed,[2] the severest downturn in the sample produced no access differential in investment—a null with clean pre-trends, powered to detect a gap of the 2023 magnitude. The split survives conditioning on size, leverage, and the volatility of pre-period sales growth, attenuating only under the profitability control that costs a fifth of the sample, and it holds within the sub-annually covered majority, so it is not the coverage concentration of Section 6.6 relabeled. Around the 2021 exercise no formal gradient exists, consistent with access carrying no insurance value when there is no pass-through to insure against. Operating profitability echoes the pattern at a marginal level: higher for exposed firms under accommodation, lower under tightening, the only survivor of a forty-four-cell sweep of additional real outcomes. Employment shows nothing within either anchor window.

[2]EBA statement of 12 March 2020, postponing the EU-wide stress test to 2021 so that banks could prioritize operations and customer support.

Table 10: The access margin: bootstrap contrasts of the investment response by exposure and bond-market access

Contrast	Estimate (SE), pp	95% CI (pp)
2023 vintage		
Low-exposure tercile: no access − access	-1.472* (0.822)	[-3.224, +0.059]
High-exposure tercile: no access − access	-1.756*** (0.530)	[-2.840, -0.754]
No access: high − low exposure tercile	-0.587 (0.926)	[-2.354, +1.346]
High-exposure, no-access cell mean	-1.654 (0.491)	[-2.669, -0.743]
2021 vintage		
Low-exposure tercile: no access − access	-0.200 (0.745)	[-1.749, +1.152]
High-exposure tercile: no access − access	+0.627 (0.423)	[-0.171, +1.519]
No access: high − low exposure tercile	+0.231 (0.692)	[-0.987, +1.729]
High-exposure, no-access cell mean	-0.208 (0.309)	[-0.772, +0.433]

Notes: sample is each vintage's core annual firm panel. Each cell is a bootstrap contrast (3,000 resamples, fixed seed) of the change in the investment rate from pre- to post-publication, in percentage points, across cells defined by exposure terciles (raw, non-demeaned exposure) and a broad bond-market-access definition; the 2021 post window runs through fiscal 2022. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ from the bootstrap p-value, SE in parentheses; the final row in each block is a cell mean, not a contrast. The underlying two-by-two conditional means under the within-country median exposure split and the narrower access definition—a different construction, not one-for-one comparable—are reported in Appendix Table A5, with the committed cells in the replication package.

6.6 The average investment response

Average investment also declines after the 2023 exercise, but this estimate is considerably less robust than the cash gradient. Table 11 assembles the real outcomes for the anchor pair, and Figure 5 shows the central path at half-year frequency: net investment in property, plant, and equipment of exposed firms runs flat into the 2023 test and falls after publication—about one percentage point per unit of exposure in the first post-half (randomization $p = 0.025$) and one and a half by the third—while the same outcome shows nothing anywhere around the 2021 exercise. The annual panel shows a decline of about two-thirds of a percentage point per standard deviation of exposure over the two post-publication years ($p = 0.068$; randomization $p = 0.033$); it is significant under coefficient randomization and the stratified variant, marginal under exposure-robust standard errors, and not significant under studentized randomization or the family adjustment (Section 7). The formal difference between the two legs of the anchor pair—the additional investment response per unit of exposure when the exercise lands under tightening—is −1.5 percentage points, significant at the ten-percent level under clustered and coefficient-randomization inference only; Section 7 reports the full accounting.

The annual decline is also concentrated on every margin along which it can be measured. It

Table 11: The 2021 versus 2023 contrast: firm-level real outcomes

	Accommodation (2021)			Tightening (2023)		
Outcome	*N*	Firms	Coef. (SE)	*N*	Firms	Coef. (SE)
Investment rate (annual)	742	409	-0.003 (0.003)	1,343	354	-0.016* (0.009)
Employment growth (annual)	298	149	0.021 (0.039)	830	266	-0.015 (0.034)
Gross capital expenditure (quarterly)	3,302	338	-0.002 (0.002)	3,232	285	-0.001 (0.001)
Net PP&E investment (quarterly)	3,692	365	-0.003 (0.002)	3,584	302	-0.003* (0.002)
Debt growth (quarterly)	3,785	370	-0.023 (0.035)	3,664	308	-0.016 (0.011)
Robustness: expansion-phase window (2021 only)						
Debt growth (quarterly)	3,260	369	-0.031 (0.035)			
The anchored regime difference (pooled 2021–2023 pair)						
Accommodation-vs-tightening interaction	*N*=4,023		Firms=560			-0.015* (0.008)
Employment growth × tightening	*N*=2,228		Firms=439			-0.066 (0.042)

Notes: sample is each event's own core annual or quarterly firm panel (the 2021 and 2023 column blocks) or the pooled anchored pair (the bottom block). Outcomes are the investment rate, debt growth, and employment growth (annual, Worldscope, no quarterly series); the regressor is post × within-country demeaned stress-test exposure (E_{dm}), with firm-by-event and country-by-event-by-time fixed effects, standard errors clustered by firm. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The bottom block's rows test whether the accommodation- and tightening-era effects differ on the pooled 2021–2023 pair; the expansion-phase row repeats the 2021 quarterly debt-growth cell on a narrower, pre-amendment post-window.

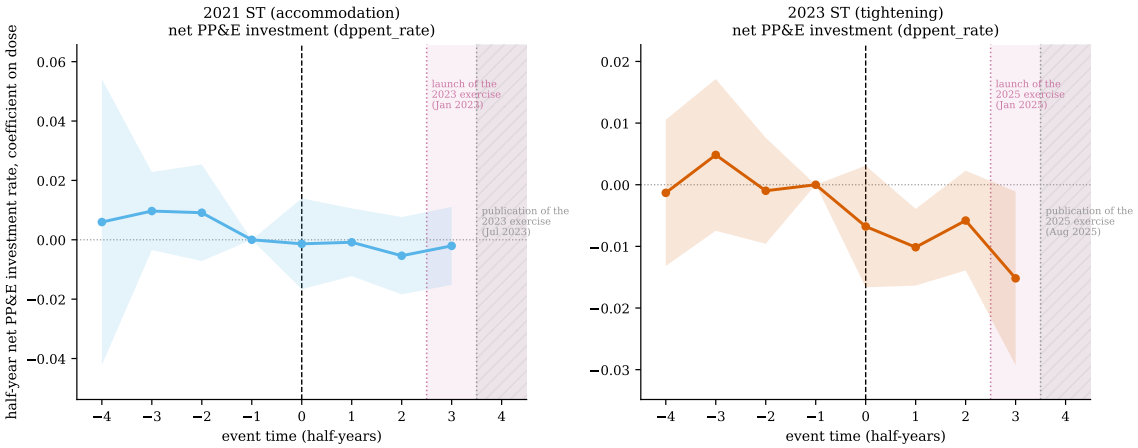


Figure 5: Net investment in property, plant, and equipment around the anchor exercises. Event-time coefficients of the half-year change in net property, plant, and equipment over lagged total assets on within-country demeaned exposure; half-years relative to the publication half, reference at −1; firm-by-event and country-by-half-by-event fixed effects, standard errors clustered by firm, shaded 90 percent confidence intervals. Windows run to the half preceding the following exercise's publication, with the following exercise's launch shaded. Net property additions are the densest sub-annual investment measure available, at the cost of bundling depreciation and disposals with gross investment. Sources: Compustat Global interim statements merged to the exposure spine.

is carried almost entirely by the roughly eighty sample firms without usable sub-annual reporting, which are disproportionately Spanish (Appendix Table A4); excluding Spain, the annual estimate is no longer distinguishable from zero; and excluding a single lender—La Banque Postale—removes four-fifths of the coefficient (Appendix Table A15). The candidate explanation that would have redeemed the concentration—heavier reliance of those firms on tested lead banks—was tested and failed: the reliance interaction carries the wrong sign relative to its mechanical benchmark, and non-Spanish exposed firms in the top reliance tercile show nothing (Appendix Table A13). The real-outcomes claim accordingly rests on the incidence gradient of Section 6.5, which survives excluding Spain and every single-bank cut, not on the average effect. Applying the estimated marginal effect at each firm's own cash position to the below-median-cash exposed firms puts the foregone investment at roughly 270 million euro per year (235 million excluding Spain), about half a billion over the two post-publication years, on the same assets-at-2022 convention as the aggregate that follows. The corresponding arithmetic for the average effect is larger: moving a firm from the 25th to the 75th percentile of exposure lowers annual investment by about 0.6 percentage points of assets, roughly 25 million euro per year at the median sample firm, and on the order of 1.7 billion euro over the two years across all exposed firms (1.3 billion excluding Spain). It applies the flat average coefficient to every exposed firm and is reported only for comparison. The identified object is relative—within-country differences by exposure—so these are reallocation statements, not general-equilibrium ones, computed on roughly 350 large listed firms; the smaller, unlisted, more bank-dependent population lies outside them (Appendix Table A10). A back-of-the-envelope bridge confirms the chain's arithmetic: the extensive-margin rationing of Section 6.2 implies lost new credit of the same order as the investment decline it accompanies.

6.7 The 2021 benchmark: absorption

The 2021 stillness that anchors the comparison is not a limit of measurement; it is what the arithmetic allowed. Reconstructing each bank's requirement stack from public sources—the Pillar 2 requirements set for 2021, combined buffer requirements, and the supervisor's published mapping from stress-test depletion to guidance—shows that the capital headroom banks already held above their distribution-restriction thresholds exceeded the guidance implied by the 2021 exercise for twenty-nine of the thirty banks the public stack covers (Table 12); the thirtieth is Deutsche Bank, whose headroom fell 0.20 percentage points short of the implied guidance under the primary reconstruction—under the alternative Pillar 2 vintage, HSBC Continental Europe joins it at a similar margin. Absorbing the 2021 exercise required no action, and no action is what the record shows: no balance-sheet response on any margin (Figure 3), no capital build in the months banks knew their own results—the one cell that reaches significance under the robust inference of Table 5 is a small pre-publication uptick in risk-weighted assets with no capital counterpart—no movement in distributions, additional-tier-one issuance, portfolio composition, or risk weights. The same reconstruction for 2023 shows headroom of a similar order, which sharpens the contrast: what changed between the exercises is not whether banks *could* stand still but what standing still cost—a reading the funding evidence above supports but does not prove.

Figure 6 traces the intuition over time: headroom net of implied guidance stays positive in

Table 12: Absorption feasibility: headroom versus implied guidance, by dose tercile

	N	Depletion (pp CET1)	Implied P2G (pp)	Headroom (pp)	Feasible (headroom > P2G)
Panel A: 2021 exercise (accommodation)					
Tercile 1	10	3.02	0.88	6.20	10/10 (100%)
Tercile 2	12	4.84	1.25	3.40	12/12 (100%)
Tercile 3	8	6.33	1.88	5.56	7/8 (88%)
All terciles (pooled)	30				29/30 (97%)
Panel B: 2023 exercise (tightening)					
Tercile 1	16	2.94	0.50	4.51	16/16 (100%)
Tercile 2	17	4.22	1.25	4.57	17/17 (100%)
Tercile 3	13	5.87	1.25	5.00	13/13 (100%)
All terciles (pooled)	46				46/46 (100%)

Notes: depletion, implied P2G, and headroom columns report within-tercile medians, in percentage points of CET1. Sample is the subset of dosed banks with a complete public requirement stack (Pillar 1 + disclosed P2R + combined buffer requirement [CCoB + CCyB + O-SII/G-SII] + published SREP P2G bucket-map assignment) at the half immediately preceding the exercise's publication – roughly two-thirds of each vintage's dosed sample (2021: 33/50; 2023: 49/70), Bucket-4 "no cap" banks excluded (implied P2G not point-identified). Pillar 2 requirements are the values set for each exercise's own year, taken from the supervisor's disclosures as published at the time. Dose terciles are within-vintage terciles of adverse-scenario CET1 depletion (the project's standard dose), z-standardized within vintage. Median headroom is the pre-publication CET1 ratio net of every disclosed hard MDA-relevant requirement (Pillar 1, P2R, combined buffer), i.e. headroom *before* Pillar 2 guidance is subtracted; median implied P2G is the bucket-map range midpoint at the bank's own depletion. "Feasible" banks have pre-existing headroom exceeding implied P2G, i.e. can in principle absorb the stress-test-implied capital add-on without new capital action; descriptive only, no permutation inference (no test statistic is defined for this cross-sectional accounting comparison).

every tercile of both vintages, and its management splits along the regime line—the hardest-hit banks ride the buffer down after 2021 and defend it after 2023, echoing the pandemic buffer-usability evidence (Couaillier et al., 2022). The absorption can also be watched bank by bank at the guidance threshold: the nine banks pushed by their 2021 depletion into a higher guidance bucket met the step out of their discretionary buffers—which fell by about one percentage point, roughly the step itself—while banks that stayed inside their bucket held flat; neither group raised capital, and no component of the regulatory capital position responds to the continuous 2021 dose (Appendix Table A2).

The firm side of the benchmark is equally quiet. Around the 2021 exercise the bond share of exposed firms' debt shows no response however the event window is drawn, the issuance margins are still (Table 8), the within-firm specification finds nothing (Section 6.2), and investment does not respond: the 2021 annual estimate is an imprecise zero, and the quarterly and half-yearly estimates are indistinguishable from zero under clustered inference, with one exception: the gross-capex cell reaches significance under randomization alone ($p = 0.045$). The zeros are informative, not merely underpowered: inflating the 2021 detectable-effect limits for the vintage's coarser dose disclosure still bounds the largest annual investment response the 2021 data could have hidden below the 2023 point estimate (Section 7).

Several qualifications travel with the benchmark. The public requirement stack covers two-thirds of the dosed banks, whose doses are indistinguishable from the uncovered third's (Welch p-values 0.20 and 0.89 in the 2021 and 2023 vintages). The accommodation-era nulls also

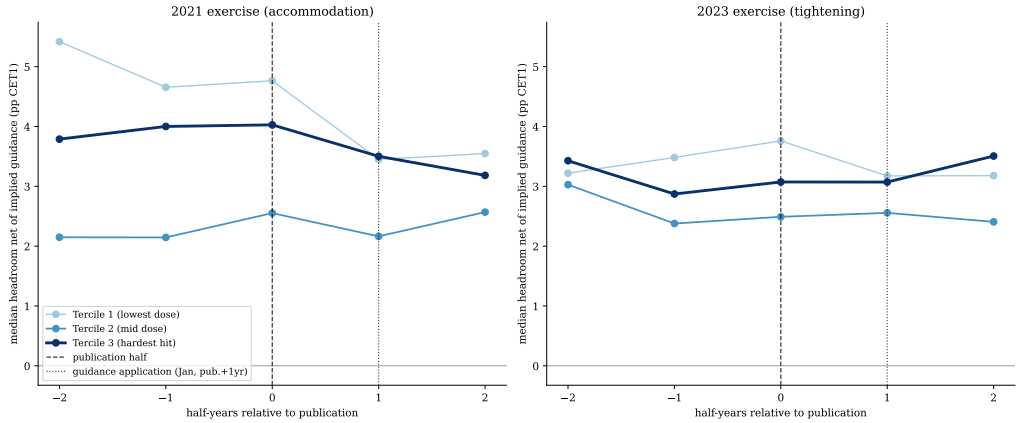


Figure 6: Capital headroom net of implied guidance around the anchor exercises. Median CET1 headroom above the full requirement stack (Pillar 1, Pillar 2 requirement, combined buffers), net of the Pillar 2 guidance implied by the exercise's published depletion, in percentage points, by within-vintage dose tercile; half-years relative to the publication half. The dashed line marks publication, the dotted line the January the resulting guidance takes effect; values above zero mean standing headroom already covered the implied guidance. Sample is the stack-complete subset of Table 12, roughly two-thirds of each vintage's dosed banks. Sources: supervisory transparency-exercise capital ratios, disclosed Pillar 2 requirements, macroprudential buffer registers, and the published guidance bucket map.

differ in sharpness: the risk-weight zero is tight, the total-assets zero looser. In one bank-level source the raw cross-section shows more affected banks shrinking after 2021. That pattern sits across countries—high-depletion banking systems repaying central-bank funding on their own schedules—and the within-country comparison absorbs it; the transparency-based 2021 zeros are tight in the raw column already. Among the smaller banks that file only annual statements, a within-country negative appears at longer 2021 horizons, marginal and measured over a window that contains the rate shock, so for the smallest lenders the accommodation-era zero is less sharply determined. Pandemic-era policy attaches as well: distribution restrictions in force through September 2021 elevated measured headroom and closed the payout margin by policy, so the accommodation-state quiet bundles cheap funding with pandemic-era support measures the design does not separate; the pandemic control specifications are collected in the robustness compendium. Finally, one regularity cuts against the quiet: sales growth of firms exposed to the 2021 exercise runs below that of unexposed firms through mid-2022, in two independent accounting sources—the one signal that exposure correlated with firm outcomes under accommodation. It propagates nowhere: 2021 investment, employment, financing mix, and the exposed banks' balance sheets are all quiet, every 2021 pre-publication path is flat, and conditioning the 2021 investment estimates on pre-window sales growth leaves the zeros intact. The regularity remains unexplained; it does not contaminate the benchmark.

7 Inference and Robustness

The schemes sort the estimates into three groups: robust central results, supporting evidence, and exploratory results.[3] The sorting is worth setting out in full, because the dose varies at the

[3]Every check in this section and in Appendix A is backed by an estimate in the replication package; Appendix Table A17 tabulates the main ones, with coefficients, standard errors, and samples, and the online appendix

bank level, the number of dose-carrying banks per exercise is modest, and the exposure that transmits it is a shift-share object.

Appendix Table A16 therefore subjects each headline coefficient to the full apparatus of Section 5: firm-, lead-bank-, and two-way-clustered standard errors; exposure-robust standard errors from the equivalent shock-level regression (Adão et al., 2019; Borusyak et al., 2022), whose numerical equivalence to the firm-level estimate is verified exactly; coefficient and studentized randomization; the family adjustment; the prior-dose-stratified randomization; and the effective number of bank-level shocks, near twenty-five in every panel.

The within-firm gradient and the cash-slack interaction hold at conventional significance under every scheme—studentized, family-adjusted, and stratified p-values between 0.003 and 0.022—with exposure-robust standard errors matching or sharper than the clustered ones. These are the paper's robust central results. Between the twenty-fifth percentile and the median the incidence gradient's level no longer survives the sharpest schemes or the exclusion of Spain, which is why the incidence claim is scoped to the bottom quarter of the cash distribution. The issuance count sits between: significant under exposure-robust standard errors and studentized randomization ($p = 0.037$), it does not survive the family stepdown. It joins the bank-side capital build and asset contraction, and the 2021 absorption benchmark, as supporting evidence; the benchmark rests on the conjunction of feasibility with quiet on every measured margin, not on any single null (Section 6.7). Two estimates are demoted—the cross-event interaction, which fails four schemes at once, and the annual investment static, stopped by the studentized randomization (0.226)—and both are reported as exploratory throughout the paper. No central claim depends on either.

The stratified column exists because the randomization's exchangeability assumption was tested and failed on exactly one observable, the bank's own prior-vintage dose; doses persist across exercises, so the prespecified fallback (Online Appendix, Section OA.8) re-runs every randomization within country-by-prior-dose-tercile cells, and every central estimate survives that discipline while the cross-event interaction does not. For the comparison across monetary states, the minimum detectable difference at conventional power remains about 2.3 percentage points of investment against a point estimate of 1.5 (Appendix Table A7, Panel B; Appendix A.5 reports the contrast's rescalings).

The same comparison can also be estimated directly rather than read off two separately estimated arms, by pooling the two events in a single within-firm model whose interaction term is the regime difference itself. That interaction is right-signed and exploratory: -0.004 per half-year per unit dose, with a studentized family-adjusted randomization p of 0.34, and the pooled frame's overlap of pairs present in both vintages is 34.4 percent, below the prespecified threshold (Online Appendix, Section OA.8), so the composition caveat attaches to it as well. What the design does deliver on this margin is a bound on the accommodation arm rather than a difference: the 2021 within-firm response is excluded from being as large as the baseline 2023 estimate, its 90 percent randomization interval lying entirely inside ± 0.022, though smaller responses are not excluded. The reading across monetary states accordingly rests on the mechanism chain, not on this interaction (Appendix Table A10, Panel B).

collects the remainder, together with a record of which test families, decision rules, and thresholds were fixed before estimation.

Two construction-side disciplines close the remaining sorting and measurement concerns (Appendix Table A14, Panels C and D). Residualizing the 2023 dose on the 2021 dose within country and re-estimating the two central cells on the orthogonalized remainder leaves both intact—the within-firm gradient at -0.024 (analytic $p = 0.012$) and the cash interaction at $+0.016$ ($p = 0.004$)—so the persistent dose component that motivated the stratification is not what carries them; the stratification scheme's own support is reported alongside, with 24 of its 35 country-by-prior-tercile cells singletons carrying 21.7 percent of the dose variance. Measurement filters point the same way: restricting to firms with at least three-quarters of lead credit measured leaves the weld and cash cells intact (-0.026 and $+0.017$) while the annual static, an exploratory estimate, attenuates further (-0.006), and a bound that scores unmeasured lead banks at dose zero preserves the pattern.

Single banks and single countries are the design's natural pressure points, and both are mapped exhaustively. Across all seventy single-bank exclusions the within-firm gradient stays between -0.019 and -0.027, the cash-slack interaction stays positive in every cut, and the issuance count moves within a range that never crosses zero (Appendix Table A15, Figure A5). The annual investment static is the exception, and the influence has a name: excluding La Banque Postale removes four-fifths of the coefficient—while the same bank's exclusion strengthens the issuance and cash-slack estimates, and of the eighty-one firms without usable sub-annual reporting only three borrow from it, so the bank concentration and the coverage concentration barely overlap. The cross-event interaction's most influential exclusion, Groupe Crédit Agricole, moves that coefficient by a third toward zero. Geography tells one story: in the leave-one-country-out jackknife, eight of ten exclusions leave 97 to 109 percent of the 2023 annual coefficient in place, excluding France *strengthens* it by two-fifths, and Spain alone cuts it to a third and takes its significance with it—the same concentration Section 6.6 reports from the coverage side—while the within-firm rationing is, if anything, stronger outside Spain. Nor is Spain's temporary bank levy behind the concentration: every dosed Spanish lender was liable, so the levy is common to the within-country comparison, and controlling for firms' reliance on levy-liable lenders leaves the estimate unchanged (Estrada and Pérez Montes, 2024).

The two vintages do not disclose equally—the 2021 exercise left more banks with range-only results, so its dose covers a smaller share of firms' measured lead credit—and differential measurement could in principle manufacture part of the cross-vintage contrast. Appendix Table A12 tests this in both directions. Imposing the 2021 disclosure regime on the 2023 dose (masking the ten banks whose 2023 results would have been range-only under the 2021 rules and rebuilding exposure) moves the 2023 annual coefficient from -0.0156 to -0.0196: it *strengthens*, so coarser disclosure cannot be manufacturing the contrast by attenuation. The shift is 0.326 of the raw 2023–2021 contrast, just under the materiality threshold of one-third. Two hundred random masks of the same size put the degraded coefficient at -0.0140 on average with no sign flips, and the identity mask lies outside the random-mask support—the actually-affected banks are not a random set, a selection on measurability the design tolerates because exposure is demeaned within country on the measured set and every comparison is made within it. In the other direction, inflating each 2021 null's detectable-effect limit by the vintage's mean measured credit share (0.73) puts the largest annual investment response the 2021 data could have hidden

at 0.0120 per unit of exposure, below the 2023 point estimate of 0.0156—a yardstick that is itself an exploratory estimate, and one that rests on the annual margin alone, since the quarterly bounds are absolutely tighter but the corresponding 2023 quarterly points sit inside them. Even under the least favorable attenuation correction, a 2023-sized effect in 2021 would have been detected. The comparison is arithmetic in the share it divides by, so its sensitivity is worth stating plainly. The bound scales the raw detectable-effect limit by $1/\lambda$, the vintage's mean measured credit share, and clears the 2023 point estimate so long as the true share exceeds about 0.56; beneath that flip point the ordering reverses. The flip point sits well below the mean the arithmetic uses, 0.73, and below the vintage's median of 0.75, though above its twenty-fifth percentile of 0.55—the ordering is comfortable at the central tendency the bound is built on, while a quarter of the estimation sample's firms are measured less completely than the flip point itself.

The sector reading of the tightening-era results—the 2023 scenario was itself built around the energy shock—is arbitrated in Appendix A.1: additive sector-by-period effects leave the 2023 static intact at both digit rungs, sector composition explains 12.5 percent of the within-country exposure variance, and a dose built only from banks' non-corporate stress losses reproduces the response (Appendix Table A18); the one fragility—the pooled cross-event interaction annihilated under the replacement convention—stands. The added mechanism interactions are reported in Appendix Table A13 and discussed in Appendix A.3: the funding-repricing triple does not carry to the firm level (bounded and uninformative rather than contradicting), the reliance interaction fails its mechanical benchmark, and the bookrunner horse race is one-sided—the issuance contraction sits on the lending relationship, with the underwriting dose bounded rather than excluded. Four facts separating the test from the tightening—the rate-shock placebo ($p = 0.88$), the weak-bank horse race, the guidance timing, and the rate-risk decomposition, under which measured rate-risk exposures explain 1.4 percent of the within-country dose variation (Appendix Table A9)—are collected in Appendix A.4. Window variants, exposure redefinitions, the no-variation margins, and the placebo catalog are in Appendix A.5 and the robustness compendium (Appendix Table A17). The wider specification surface—roughly 1,080 variants across outcomes, windows, and exposure definitions—remains exploratory and uncorrected, and is reported in the replication package; the paper draws no conclusion from an isolated cell of it.

8 Conclusion

Whether bank regulation restrains the real economy depends on the conditions under which it arrives. We follow supervisory capital pressure from bank balance sheets to corporate financing and investment across the 2021 and 2023 EU stress tests. After the 2023 exercise, harder-hit banks rebuilt capital while the exercise ran, most of it before their results were public, and rationed new credit once the resulting guidance applied, sparing renewals. Exposed firms did not replace the lost financing in bond markets: new issuance fell alongside new loans, on the same guidance calendar, while the market itself stayed open. The mechanism behind that contraction remains open. Investment fell most among firms with little cash, significantly negative through roughly the bottom quarter of the cash distribution under every inference procedure we apply, fading by the median and absent among cash-rich firms. None of these margins moved comparably

after the 2021 exercise, when standing headroom covered the implied guidance and absorbing it required no action; the 2021 investment zeros carry detectable-effect bounds that exclude an effect of the 2023 size even after correcting for that vintage's coarser disclosure. In 2021, the buffer was the bank's own headroom; in 2023, it was the firm's own slack.

The contrast between the two exercises is consistent with state-dependent transmission of supervisory capital pressure, and funding conditions are a plausible mechanism: the 2023 contraction concentrates in wholesale-funded banks and is cushioned by deposits, an interaction absent in 2021. The design does not isolate monetary policy from the broader tightening environment, and the formal cross-event interaction is significant but estimated relatively imprecisely; the state-dependence reading rests on the contrast in the transmission chain across the two episodes as well as an indicative pooled interaction. The average annual investment effect is concentrated along dimensions that Sections 6.6 and 7 set out and that the tested candidate explanation does not account for; the real-outcomes claim accordingly rests on the incidence gradient, which survives excluding Spain and every single-bank cut.

The real-incidence result is local to the large listed syndicated borrowers the design observes: on the syndicated margin itself the measured response for unlisted borrowers is smaller and imprecise, -0.005 against -0.013 for listed ones (Appendix Table A10). Two readings of that gap point in different directions—coarser measurement off the listed frame attenuates the unlisted coefficient toward zero, while bilateral bank credit, invisible to this design, may be where an unlisted firm's pass-through appears at all—so the evidence neither establishes comparable pass-through to the broader unlisted population nor signs the direction of the difference.

Two policy implications follow from the headroom arithmetic. First, the incidence of supervisory capital pressure is state-contingent, so its calibration should be. The welfare cost of asking banks to hold more capital is not a constant of the instrument: theory prices it off banks' funding conditions (Van den Heuvel, 2008; Begenau, 2020), and the case for cyclically varying requirements rests on exactly the wedge these episodes display (Repullo and Suárez, 2013; Malherbe, 2020). The two exercises provide an empirical counterpart at the level of guidance: the same instrument, feasible to absorb both times, was absorbed when standing still was free and passed through once it was dear, onto the large listed borrowers this design measures and, among them, onto those with the least internal slack. Guidance calibrated in capital units alone is therefore calibrated in real units that vary with the monetary state; a supervisor who wants constant real pressure would have to lean against the funding cycle, and one who wants countercyclical pressure gets part of it for free under accommodation and more than intended under tightening.

Second, the buffer evidence bears on releasable-buffer design. Banks rode their discretionary headroom down after 2021 and defended it after 2023—buffers were usable exactly when usability was cheap, echoing the pandemic record (Couaillier et al., 2022) and the logic by which capital regulation shapes banks' own appetite for buffers (Bahaj and Malherbe, 2020). If buffers are to absorb pressure when it bites, the releasable component has to be released against the funding state, not only against measured losses; and the incidence result adds a distributional reason to care, since within the listed sample the firms that bear the pass-through are those with the least cash to self-insure with. That is an incidence statement, not optimal-policy arithmetic.

The 2025 exercise provides a natural out-of-sample test: published in August 2025, inside the contraction, it should reproduce the 2023 pattern first on the financing margins—within-firm rationing on its guidance calendar and the issuance response—with the annual investment margin observable once the 2026 fiscal years close, by mid-2027. The first result is in. Before the new guidance begins to apply in January 2026 the guidance-calendar mechanism predicts nothing, and the publication-half financing cell is null: -0.0004 per unit dose, with a randomization p of 0.93. The arm's pre-period diagnostics do not yet run clean, so this is a first installment, not a confirmation (Appendix Table A10, Panel B). What the record supports meanwhile is a statement about incidence: supervisory capital pressure can operate as a quantity instrument whose incidence depends on what absorbing it costs—in these two episodes the capacity to absorb was in place both times, and it was the cost of standing still that moved with the monetary environment.

References

Viral V. Acharya, Allen N. Berger, and Raluca A. Roman. Lending implications of U.S. bank stress tests: Costs or benefits? *Journal of Financial Intermediation*, 34:58–90, 2018a. doi: 10.1016/j.jfi.2018.01.004.

Viral V. Acharya, Tim Eisert, Christian Eufinger, and Christian Hirsch. Real effects of the sovereign debt crisis in europe: Evidence from syndicated loans. *The Review of Financial Studies*, 31(8):2855–2896, 2018b. doi: 10.1093/rfs/hhy045.

Viral V. Acharya, Rahul Singh Chauhan, Raghuram G. Rajan, and Sascha Steffen. Liquidity dependence and the waxing and waning of central bank balance sheets. NBER Working Paper 31050, National Bureau of Economic Research, 2023.

Rodrigo Adão, Michal Kolesár, and Eduardo Morales. Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010, 2019. doi: 10.1093/qje/qjz025.

Tobias Adrian, Paolo Colla, and Hyun Song Shin. Which financial frictions? parsing the evidence from the financial crisis of 2007 to 2009. *NBER Macroeconomics Annual*, 27:159–214, 2013.

Shekhar Aiyar, Charles W. Calomiris, and Tomasz Wieladek. How does credit supply respond to monetary policy and bank minimum capital requirements? *European Economic Review*, 82: 142–165, 2016.

Heitor Almeida, Murillo Campello, and Michael S. Weisbach. The cash flow sensitivity of cash. *Journal of Finance*, 59(4):1777–1804, 2004.

Mary Amiti and David E. Weinstein. How much do idiosyncratic bank shocks affect investment? evidence from matched bank-firm loan data. *Journal of Political Economy*, 126(2):525–587, 2018. doi: 10.1086/696272.

Óscar Arce, Ricardo Gimeno, and Sergio Mayordomo. Making room for the needy: The credit-reallocation effects of the ECB's corporate QE. *Review of Finance*, 25(1):43–84, 2021. doi: 10.1093/rof/rfaa020.

Saleem Bahaj and Frederic Malherbe. The forced safety effect: How higher capital requirements can increase bank lending. *Journal of Finance*, 75(6):3013–3053, 2020.

Bo Becker and Victoria Ivashina. Cyclicality of credit supply: Firm level evidence. *Journal of Monetary Economics*, 62:76–93, 2014. doi: 10.1016/j.jmoneco.2013.10.002.

Mikel Bedayo and Jorge E. Galán. The impact of the countercyclical capital buffer on credit: Evidence from its accumulation and release before and during COVID-19. Banco de España Working Paper 2411, Banco de España, 2024.

Juliane Begenau. Capital requirements, risk choice, and liquidity provision in a business-cycle model. *Journal of Financial Economics*, 136(2):355–378, 2020.

Markus Behn, Rainer Haselmann, and Vikrant Vig. The limits of model-based regulation. *Journal of Finance*, 77(3):1635–1684, 2022. doi: 10.1111/jofi.13124.

Markus Behn, Stijn Claessens, Leonardo Gambacorta, and Alessio Reghezza. Macroprudential and monetary policy tightening: More than a double whammy? ECB Working Paper Series 3043, European Central Bank, 2025. Also BIS Working Paper 1257; CEPR Discussion Paper 20181.

Jose M. Berrospide and Rochelle M. Edge. The effects of bank capital on lending: What do we know, and what does it mean? *International Journal of Central Banking*, 6(4):5–54, 2010.

Patrick Bolton, Xavier Freixas, Leonardo Gambacorta, and Paolo Emilio Mistrulli. Relationship and transaction lending in a crisis. *Review of Financial Studies*, 29(10):2643–2676, 2016.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213, 2022. doi: 10.1093/restud/rdab030.

Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021. doi: 10.1016/j.jeconom.2020.12.001.

Murillo Campello, John R. Graham, and Campbell R. Harvey. The real effects of financial constraints: Evidence from a financial crisis. *Journal of Financial Economics*, 97(3):470–487, 2010.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454, 2019. doi: 10.1093/qje/qjz014.

Gabriel Chodorow-Reich. The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *The Quarterly Journal of Economics*, 129(1):1–59, 2014. doi: 10.1093/qje/qjt031.

Kristle Romero Cortés, Yuliya Demyanyk, Lei Li, Elena Loutskina, and Philip E. Strahan. Stress tests and small business lending. *Journal of Financial Economics*, 136(1):260–279, 2020. doi: 10.1016/j.jfineco.2019.08.008.

Cyril Couaillier, Marco Lo Duca, Alessio Reghezza, and Costanza Rodriguez d'Acri. Caution: Do not cross! capital buffers and lending in covid-19 times. Working Paper 2644, European Central Bank, 2022.

Nicolas Crouzet. Aggregate implications of corporate debt choices. *Review of Economic Studies*, 85(3):1635–1682, 2018.

Olivier Darmouni and Melina Papoutsi. The rise of bond financing in europe. ECB Working Paper Series 2663, European Central Bank, 2022.

Olivier De Jonghe, Hans Dewachter, and Steven Ongena. Bank capital (requirements) and credit supply: Evidence from pillar 2 decisions. *Journal of Corporate Finance*, 60:101518, 2020. doi: 10.1016/j.jcorpfin.2019.101518.

Roberto A. De Santis and Andrea Zaghini. Unconventional monetary policy and corporate bond issuance. *European Economic Review*, 135:103727, 2021. doi: 10.1016/j.euroecorev.2021.103727.

Sebastian Doerr. Stress tests, entrepreneurship, and innovation. *Review of Finance*, 25(5): 1609–1637, 2021. doi: 10.1093/rof/rfab007.

Itamar Drechsler, Alexi Savov, and Philipp Schnabl. The deposits channel of monetary policy. *Quarterly Journal of Economics*, 132(4):1819–1876, 2017.

Arindrajit Dube, Daniele Girardi, Òscar Jordà, and Alan M. Taylor. A local projections approach to difference-in-differences. *Journal of Applied Econometrics*, 40(7):741–758, 2025. doi: 10.1002/jae.70000.

Ran Duchin, Oguzhan Ozbas, and Berk A. Sensoy. Costly external finance, corporate investment, and the subprime mortgage credit crisis. *Journal of Financial Economics*, 97(3):418–435, 2010.

Egemen Eren, Denis Gorea, and Daojing Zhai. How do quantitative easing and tightening affect firms? BIS Working Paper 1286, Bank for International Settlements, 2025.

Ángel Estrada and Carlos Pérez Montes. Un análisis de la evolución de la actividad bancaria en España tras el establecimiento del gravamen temporal de la Ley 38/2022. Documentos Ocasionales 2408, Banco de España, 2024.

Mark Flannery, Beverly Hirtle, and Anna Kovner. Evaluating the information in the Federal Reserve stress tests. *Journal of Financial Intermediation*, 29:1–18, 2017.

Henri Fraisse, Mathias Lé, and David Thesmar. The real effects of bank capital requirements. *Management Science*, 66(1):5–23, 2020. doi: 10.1287/mnsc.2018.3222.

Mark Gertler and Simon Gilchrist. Monetary policy, business cycles, and the behavior of small manufacturing firms. *The Quarterly Journal of Economics*, 109(2):309–340, 1994. doi: 10.2307/2118465.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021. doi: 10.1016/j.jeconom.2021.03.014.

Reint Gropp, Thomas Mosk, Steven Ongena, and Carlo Wix. Banks response to higher capital requirements: Evidence from a quasi-natural experiment. *Review of Financial Studies*, 32(1): 266–299, 2019. doi: 10.1093/rfs/hhy052.

Benjamin Grosse-Rueschkamp, Sascha Steffen, and Daniel Streitz. A capital structure channel of monetary policy. *Journal of Financial Economics*, 133(2):357–378, 2019. doi: 10.1016/j.jfineco.2019.03.006.

Björn Imbierowicz, Axel Löffler, and Ursula Vogel. The transmission of bank capital requirements and monetary policy to bank lending in Germany. *Review of International Economics*, 29(1): 144–164, 2021. doi: 10.1111/roie.12500.

Victoria Ivashina and David Scharfstein. Bank lending during the financial crisis of 2008. *Journal of Financial Economics*, 97(3):319–338, 2010.

Gabriel Jiménez, Steven Ongena, José-Luis Peydró, and Jesús Saurina. Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments. *Journal of Political Economy*, 125(6):2126–2177, 2017. doi: 10.1086/694289.

John Kandrac and Bernd Schlusche. The effect of bank supervision and examination on risk taking: Evidence from a natural experiment. *Review of Financial Studies*, 34(6):3181–3212, 2021. doi: 10.1093/rfs/hhaa090.

Anil K. Kashyap and Jeremy C. Stein. What do a million observations on banks say about the transmission of monetary policy? *American Economic Review*, 90(3):407–428, 2000. doi: 10.1257/aer.90.3.407.

Asim Ijaz Khwaja and Atif Mian. Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–1442, 2008. doi: 10.1257/aer.98.4.1413.

Christoffer Kok, Carola Müller, Steven Ongena, and Cosimo Pancaro. The disciplining effect of supervisory scrutiny in the EU-wide stress test. *Journal of Financial Intermediation*, 53:101015, 2023. doi: 10.1016/j.jfi.2022.101015.

Paul Konietschke, Steven Ongena, and Aurea Ponte Marques. Stress tests and capital requirement disclosures: Do they impact banks' lending and risk-taking decisions? ECB Working Paper Series 2679, European Central Bank, 2022.

Karl V. Lins, Henri Servaes, and Peter Tufano. What drives corporate liquidity? an international survey of cash holdings and lines of credit. *Journal of Financial Economics*, 98(1):160–176, 2010.

Frédéric Malherbe. Optimal capital requirements over the business and financial cycles. *American Economic Journal: Macroeconomics*, 12(3):139–174, 2020.

Jean-Stéphane Mésonnier and Allen Monks. Did the EBA capital exercise cause a credit crunch in the euro area? *International Journal of Central Banking*, 11(3):75–117, 2015.

Tim Opler, Lee Pinkowitz, René M. Stulz, and Rohan Williamson. The determinants and implications of corporate cash holdings. *Journal of Financial Economics*, 52(1):3–46, 1999.

Joe Peek and Eric S. Rosengren. Collateral damage: Effects of the Japanese bank crisis on real activity in the United States. *American Economic Review*, 90(1):30–45, 2000.

Rafael Repullo and Javier Suárez. The procyclical effects of bank capital regulation. *Review of Financial Studies*, 26(2):452–490, 2013.

Joseph P. Romano and Michael Wolf. Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282, 2005.

Amir Sufi. Bank lines of credit in corporate finance: An empirical analysis. *Review of Financial Studies*, 22(3):1057–1088, 2009.

Silvana Tenreyro and Gregory Thwaites. Pushing on a string: U.S. monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74, 2016. doi: 10.1257/mac.20150016.

Karamfil Todorov. Quantify the quantitative easing: Impact on bonds and corporate debt issuance. *Journal of Financial Economics*, 135(2):340–358, 2020. doi: 10.1016/j.jfineco.2019.08.003.

Skander J. Van den Heuvel. The bank capital channel of monetary policy. Working paper, University of Pennsylvania, The Wharton School, 2002.

Skander J. Van den Heuvel. The welfare cost of bank capital requirements. *Journal of Monetary Economics*, 55(2):298–320, 2008.

A Additional results

This appendix collects the descriptive corroboration of the motivating facts and the additional exhibits referenced in the text: the aggregate financing rotation country by country (Figure A1), the descriptive bond-share paths around the exercises (Figure A2), the capital response of the tested banks and the regulatory-component detail (Tables A1 and A2), the bank-level funding-mechanism interactions (Table A3), the coverage decomposition of the 2023 investment effect (Table A4), the access-margin conditional means (Table A5), the cash-slack interaction in event time and its full tail diagnostics (Figure A3, Table A6), the full bank-margin event studies (Figure A4), the annual regime table for the anchor pair (Table A7), the dose decomposition referenced in Section 7 (Table A8), the disclosure-degradation test and attenuation bounds (Table A12), the bond-margin complementarity test cells behind Section 6.4 (Table A11), the mechanism-interaction disclosure (Table A13), the additional inference and construction checks (Table A14), the leave-one-bank-out summary and bank-by-bank distributions (Table A15, Figure A5), the inference-completion table behind Section 7 (Table A16), and the robustness compendium (Table A17). Appendices A.1 through A.5 carry the robustness discussions moved from the main text: the sector margin, the mechanism interactions, the facts separating the test from the tightening, and the window, composition, and placebo checks. Further material—the exercises outside the anchored pair and their pooled estimates, the hard-requirement placebo, the gross-capital-expenditure event studies, the window-robustness grid, and the Orbis-based corroboration of the bank bridge—is collected in the online appendix.

Table A1: The capital side of the bank bridge: change in CET1 capital and risk-weighted assets by capital-depletion dose

	2021 (accommodation)	2023 (tightening)
Δlog CET1 capital (numerator)	+0.034	+0.043***
	(0.052)	(0.012)
Δlog RWA (denominator)	-0.000	+0.010
	(0.022)	(0.013)
Banks	44	61

Notes: bank-level cross-sectional regressions (one observation per bank, country fixed effects, robust HC1 standard errors in parentheses) of the log change in CET1 capital and in risk-weighted assets on continuous within-vintage capital-depletion dose, among tested banks with a continuous (non-range) dose; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The tightening exercise shows an adjustment on both sides of the capital ratio; the accommodation exercise shows no such response in either component.

A.1 The sector margin

Because the 2023 scenario was itself built around the energy and inflation shock, sector-by-period variation is the natural rival reading of the tightening-era results, and it produced this paper's most consequential disclosed fragility: under the replacement convention—sector-by-period fixed effects replacing country-by-period—the pooled cross-event interaction is annihilated, falling to -0.001 (0.011). Table A18 disciplines that fact from three directions. First, adding sector-by-period effects *on top of* the headline country-by-period specification leaves the 2023 static's

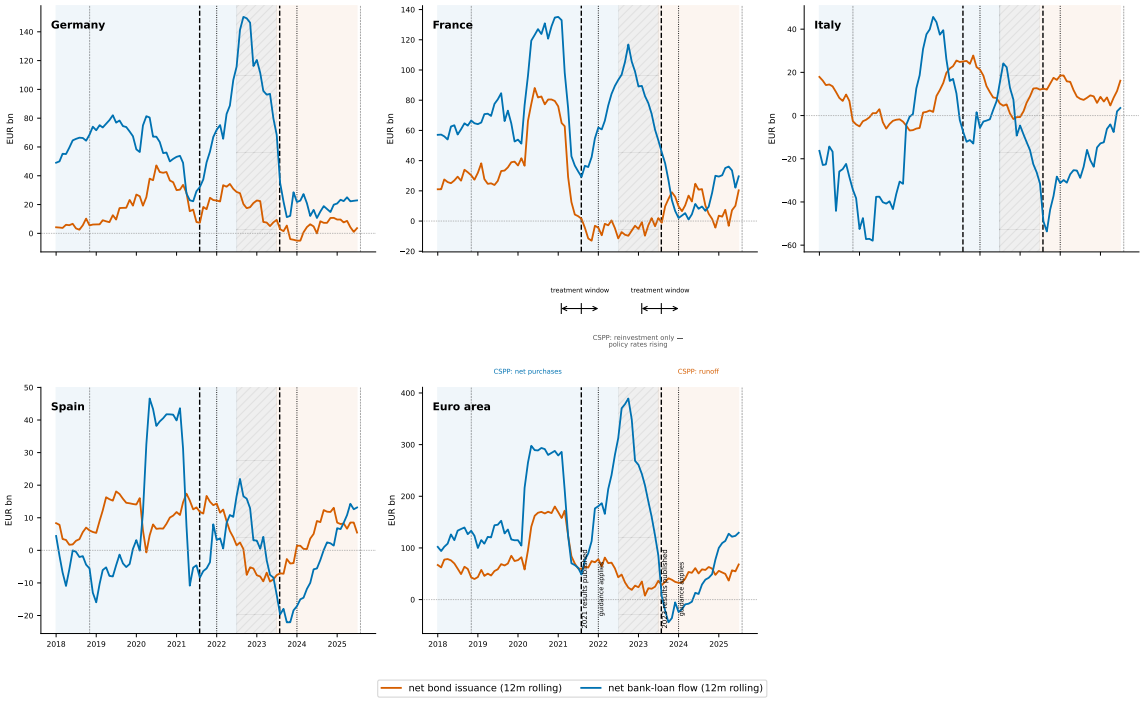


Figure A1: Corporate financing rotation by country. Twelve-month rolling net flows of bank lending to non-financial corporations and net debt-securities issuance by the same sector, in billions of euro, for the four largest euro-area economies and the euro area. Shading marks the phases of the corporate-sector purchase programme: net purchases (blue, through June 2022), reinvestment only (hatched, opening in the month of the ECB's first policy-rate increase), and runoff (orange, from July 2023). Each anchor exercise's treatment window—launch through results publication to guidance application—is drawn in the euro-area panel. Sources: ECB bank-lending and debt-securities-issuance statistics by country.

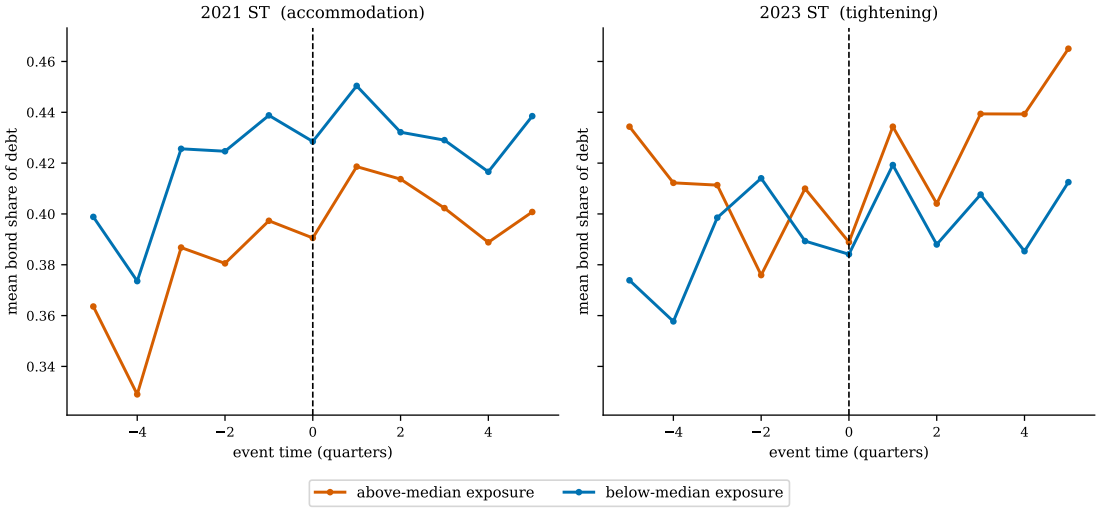


Figure A2: Bond share of debt around the anchor exercises. Each panel shows the mean bond share of total debt, by quarter relative to publication, for firms above versus below the within-country median of stress-test exposure in that exercise's pre-window roster. The vertical line marks the publication quarter; panels display the five quarters on either side of publication. Group sizes vary with reporting coverage across quarters: for the anchor exercises, between 96 and 210 firms per group and quarter (2021) and between 82 and 177 (2023); composition therefore drifts across quarters, which is one reason the figure is motivation rather than evidence.

Table A2: The regulatory capital position around the exercises

	Dose (country FE)	Above median (country FE)	Crosser vs. stayer (country FE)	MDE80 (dose)	N
Panel A: 2021 exercise (accommodation)					
CET1 ratio	+0.732 (0.723)	+1.697* (0.883)	-0.438 (0.507)	2.026	45
Tier 1 ratio	+0.625 (0.620)	+1.668** (0.734)	-0.223 (0.530)	1.738	33
Total capital ratio	+0.545 (0.649)	+1.610** (0.806)	-0.566 (0.663)	1.818	33
log RWA	-0.026 (0.023)	-0.058 (0.044)	+0.061** (0.024)	0.065	45
MDA distance	+0.842 (0.731)	+1.277 (0.913)	-0.455 (0.510)	2.049	32
Discretionary buffer	+0.201 (0.771)	+0.868 (0.548)	-1.124** (0.513)	2.160	22
Panel B: 2023 exercise (tightening)					
CET1 ratio	-0.298 (0.365)	+0.210 (0.697)	—	1.023	59
Tier 1 ratio	+0.084 (0.164)	+0.757 (0.623)	—	0.459	45
Total capital ratio	+0.291 (0.317)	+0.624 (0.709)	—	0.887	45
log RWA	+0.001 (0.014)	-0.013 (0.027)	—	0.039	59
MDA distance	-0.262 (0.445)	+0.909 (0.707)	—	1.247	44
Discretionary buffer	+0.282 (0.356)	+0.905*** (0.294)	—	0.997	26
Dose x tightening (pooled)	-1.030 (0.801)	-1.488 (1.118)	–	2.243	104

Notes: bank-level cross-sectional regressions (one observation per bank, country fixed effects, robust HC1 standard errors in parentheses) of the change (pre-publication half to one year after) in each regulatory capital component on within-country demeaned dose (continuous), an above-within-vintage-median dose indicator, and a P2G-bucket-crosser-vs-stayer indicator; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Basel III leverage ratio is unavailable in this sample at these dates and is omitted. MDE80 is the smallest effect (dose column, same units as the coefficient) detectable with 80 percent power at this sample size. The crosser column compares banks pushed by the exercise into a higher SREP Pillar 2 guidance bucket against banks that were not; this comparison is defined for 22 banks in 2021 (9 crossers) but only 3 banks crossed in 2023 (†), so the 2023 crosser column is a materially underpowered cell. The bottom row is the pooled 2+2 dose×tightening interaction (country×vintage fixed effects), with its above-median-binary companion; both use the CET1 ratio outcome.

Table A3: The funding mechanism at the bank level: dose interactions with funding structure

Panel A: 2023 headline interactions and 2021 placebo Interaction	Coef. (SE)	Cluster p	N (banks)	Countries	Note
Dose × wholesale-funding share (2023 headline)	-0.0769*** (0.0269)	0.004	50	16	gateA_ws_pre; country FE; HC1 bank cross-section cluster
Dose × deposit-cushion share (2023 headline)	+0.0905*** (0.0291)	0.002	50	16	gateA_info_dep_pre; companion cell; country FE
Dose × wholesale share, 2021 placebo (companion measure)	-0.0493 (0.3335)	0.882	13	8	c_dm:b4_ws_eba, 2_fe; F7 common sample; perm. p=0.787
Dose × deposit cushion, 2021 placebo	–	–	–	–	no committed companion cell in this window (not estimated; disclosed, not fabricated)

Panel B: permutation battery (r3 family 3, src/177_r3_conditions.py) Interaction	Coef.-perm p	Stud.-perm p	Stratified p	RW-fam.-3 p	LOBO summary
Dose × wholesale-funding share	0.156	0.063	0.284	0.063	range [-0.0938,-0.0680], 2/50 flagged ≥20%: Banco Comercial Português, SA (PT) -21.9%; Caixa Geral de Depósitos, SA (PT) -21.9%
Dose × deposit-cushion share	0.043	0.042	0.172	0.059	0/50 flagged ≥20%

Notes: bank-level regressions of the change in log customer loans over the four post-publication quarters on the standardized depletion dose interacted with the bank's pre-event funding structure from the EBA transparency exercise: the wholesale funding share and the customer-deposit share, each measured before the exercise. Country fixed effects; heteroskedasticity-robust standard errors in parentheses (bank-level cross-sections; 50 banks in 16 lender countries in the 2023 columns). The 2021 placebo re-estimates the wholesale interaction on the accommodation vintage (13 banks with the required funding data; its randomization p is 0.787); no 2021 companion exists for the deposit interaction and the cell is left blank rather than approximated. Panel B subjects the 2023 interactions to the alternative inference schemes of Section 7: coefficient and studentized randomization p-values from direct bank-level relabeling within country, the prior-dose-stratified variant, the two-cell family-adjusted p-values, and the leave-one-bank-out summary with flagged banks named. The interactions sit at the six-percent studentized and family-adjusted level and do not survive the stratified variant; they are evidence consistent with the funding mechanism, short of proof. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ (clustered/robust inference as stated). Source: the bank funding and balance-sheet panels of Section 4.

Table A4: Selection decomposition: the 2023 annual investment effect by sub-annual (Compustat quarterly) reporting coverage

Subsample	Firms	Coef. (SE)
Full 2023 sample	354	-0.016* (0.009)
Covered (≥6 clean quarters)	273	-0.005 (0.004)
Non-covered (<6 clean quarters)	81	-0.040*** (0.007)

Composition	Covered	Non-covered
Bond-market access	77.7%	77.8%
Median total assets (EUR bn)	3.99	3.89
Share in Spain	8.8%	18.5%
Mean exposure (E_{dm})	-0.293	-0.250

Notes: sample is the 2023 annual estimation sample underlying Table 11, split by Compustat quarterly reporting coverage ("covered" firms report ≥6 clean quarters in the 2023 quarterly window; "non-covered" is the complement). Both rows use the same annual specification as the headline 2023 result (post × E_{dm}, firm-by-event and country-by-event-by-year fixed effects, standard errors clustered by firm); * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The composition panel below reports descriptive conditional means, not a regression. Non-covered firms resemble covered firms in size and bond-market access but are disproportionately Spanish, consistent with semiannual reporting practice concentrating the tightening-era decline in a subsample invisible to quarterly data.

Table A5: The access margin: conditional means by exposure and bond-market access

Panel A: 2021 stress test (accommodation)

	Bond-market access	No bond-market access
Above-median exposure	-0.109 (N=149)	-0.721 (N=48)
Below-median exposure	-0.566 (N=140)	-0.570 (N=46)

Panel B: 2023 stress test (tightening)

	Bond-market access	No bond-market access
Above-median exposure	+0.131 (N=140)	-0.857 (N=37)
Below-median exposure	+0.388 (N=129)	+0.136 (N=37)

Notes: descriptive conditional means (not starred) of the change in the investment rate from pre- to post-publication, in percentage points, by within-country exposure (above/below median) and bond-market access (any outstanding bond in the two pre-window years), with firm counts in parentheses; the 2021 post window runs through fiscal 2022. These are the cells formerly shown alongside the bootstrap contrasts of Table 10, whose tercile construction and broader access definition are not one-for-one comparable. Sample is each vintage's core annual firm panel.

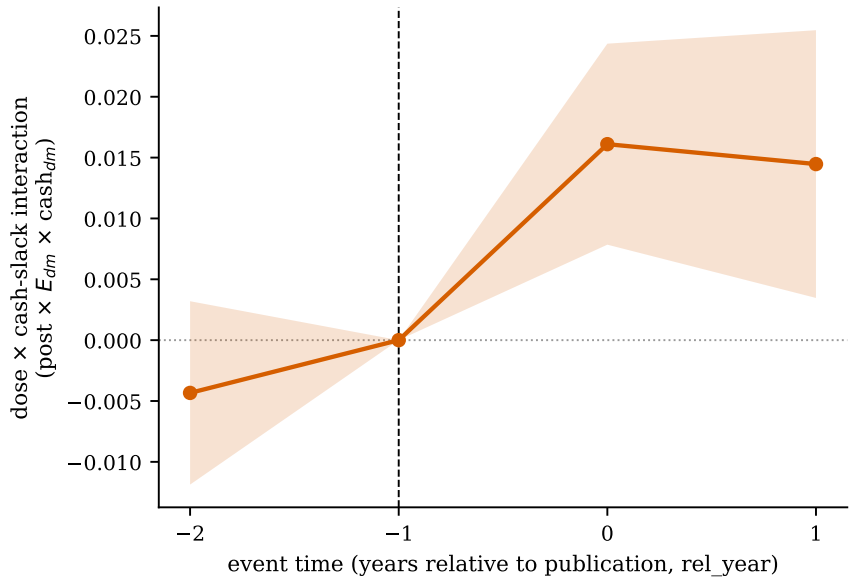


Figure A3: The cash-slack interaction in event time. Coefficients on the interaction of within-country demeaned exposure with the firm's pre-window cash holdings, estimated year by year in the annual investment panel around the 2023 exercise; years relative to the publication year, reference at −1; firm and country-by-year fixed effects; shaded 90 percent confidence band, clustered by firm. A positive coefficient means investment is better protected, per unit of exposure, where pre-window cash is higher. Sources as in Table 9.

Table A6: Incidence on internal slack: baseline cuts and tail-sensitivity diagnostics

Panel A: baseline, ex-Spain, and placebo cuts Sample	Coef. (SE)	Cluster p	Perm. p (coef.)	Stud. perm. p	N (firms)
Baseline (2023, full sample)	$+0.0174^{***}$ (0.0048)	<0.001	0.002	0.004	1,343 (354)
Ex-Spain (2023)	$+0.0164^{***}$ (0.0050)	0.001	0.007	0.004	1,196 (315)
2021 placebo (re-pinned, not re-estimated)	+0.0031 (0.0021)	0.143	0.223	–	742 (406)
Pre-trend placebo (ex-Spain cut)	+0.0064 (0.0046)	0.161	0.295	0.283	598 (313)
Panel B: tail-sensitivity diagnostics (r3 battery) Variant	Coef. (SE)	Cluster p	Perm. p (coef.)	Stud. perm. p	N (firms)
Drop top 1% (influence, committed re-pin)	$+0.0170^{***}$ (0.0047)	<0.001	–	–	353
Drop top 5% (influence, committed re-pin)	$+0.0177^{***}$ (0.0065)	0.007	–	–	349
Winsorize [1,99] (influence, committed re-pin)	$+0.0170^{***}$ (0.0047)	<0.001	–	–	354
Winsorize [5,95] (r3 battery)	$+0.0171^{***}$ (0.0053)	0.001	0.003	0.008	354
Trim [1,99] (obs. dropped, r3 battery)	$+0.0142^{**}$ (0.0056)	0.011	0.012	0.033	346
Trim [5,95] (obs. dropped, r3 battery)	$+0.0103^{*}$ (0.0054)	0.058	0.088	0.143	318
Rank-transformed (within-country, r3 battery)	$+0.0115^{**}$ (0.0051)	0.024	0.025	0.051	354
log(1+cash) (r3 battery)	$+0.0173^{***}$ (0.0048)	<0.001	<0.001	0.004	354
Exclude top decile (r3 battery)	$+0.0132^{**}$ (0.0057)	0.022	0.007	0.031	318
Exclude bottom decile (r3 battery)	$+0.0146^{***}$ (0.0050)	0.003	0.012	0.008	318
Exclude both deciles (r3 battery)	+0.0093 (0.0058)	0.112	0.057	0.191	282
Pandemic-downturn placebo (2020, MDE 2.8se=0.0171)	+0.0026 (0.0061)	0.671	0.643	0.727	415

Notes: coefficients on the triple interaction of the post-publication indicator, within-country demeaned exposure (E_{dm}), and the firm's pre-window cash holdings (demeaned), in the annual investment-rate panel around the 2023 exercise, with firm and country-by-year fixed effects. Panel A repeats the four core rows of the main text's Table 9: baseline, ex-Spain, the 2021 placebo (carried over from the 2021 estimate, not re-estimated), and the pre-trend placebo on the ex-Spain cut. Panel B reports the tail-sensitivity diagnostics: influence diagnostics taken from the baseline surface (no randomization p-value is available for those cells, marked by dashes), winsorized, trimmed, rank- and log-transformed variants, single- and both-decile exclusions, and the 2020 pandemic-downturn placebo with its 80-percent detectable-effect bound stated in the row label. Standard errors clustered by firm in parentheses; the cluster p is the analytic p-value; permutation p-values reassign the bank-level dose within country and exercise and rebuild exposure (coefficient and studentized variants). Sources: annual firm fundamentals matched to stress-test disclosures as described in Section 4. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

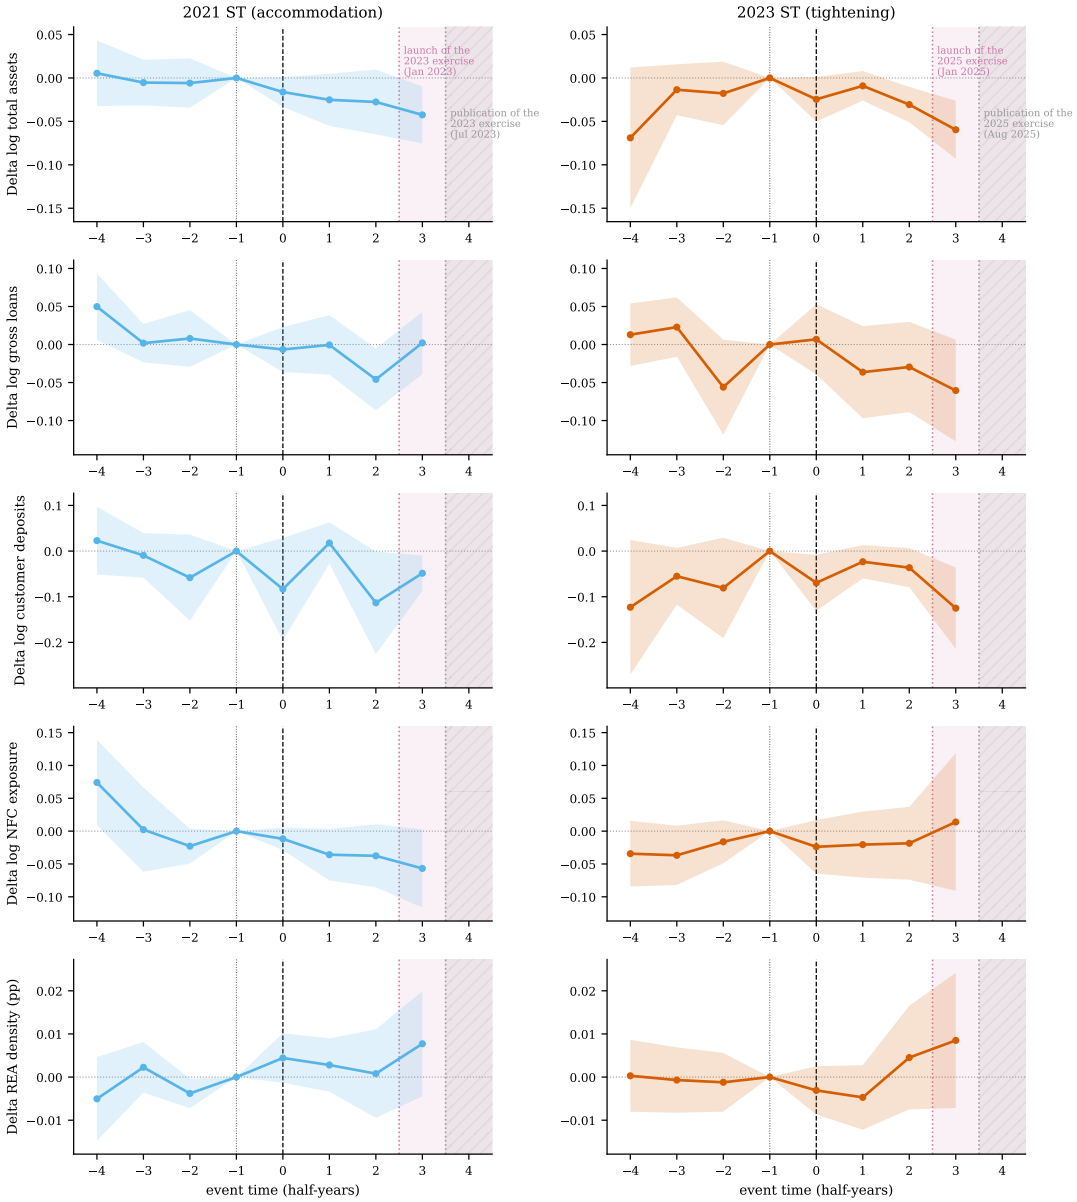


Figure A4: Bank balance sheets around the anchor exercises, all margins. As Figure 3, adding customer deposits (quarterly financial statements) and corporate credit exposures and risk density (semiannual supervisory disclosures).

magnitude intact—−0.0154 at the one-digit rung and −0.0191, larger than the baseline, at the two-digit rung ($p = 0.070$)—so the annihilation is a property of removing the country axis the design identifies on, applies to the pooled interaction, and does not carry over to the 2023 response itself. Second, sector composition accounts for 12.5 percent of the within-country exposure variance the design uses: the dose is not a relabeled sector shock. Third, a dose built only from banks' *non-corporate* stress losses reproduces the full-dose response (−0.0131 against −0.0156), and in the joint specification the corporate-loss component is dead while the non-corporate component carries the effect—the opposite of what a sectoral confound, which must travel through banks' corporate books, would predict; the individual cells are imprecise under studentized randomization ($p = 0.17$–0.21), so this is a point-pattern argument and is stated as one. The loss-rate decomposition points the same way (Table A8): firms attached to banks with

Table A7: Investment response to the capital-constraint dose: the anchored pair

Panel A: Investment rate response, by stress-test vintage

	2021	2023
Exposure × post	-0.003 (0.003)	-0.016* (0.009)
Event-time window	[-2, 0]	[-2, -1, 0, 1]
Firm-years	742	1,343
Firms	409	354
Post-window	Complete	Complete

Panel B: Accommodation versus tightening (anchored 2021–2023 pair)

	Exposure × post × tightening
Coefficient	-0.015* (0.008)
Firm-years	4,023
Firms	560
Minimum detectable effect	0.023
Equivalence bound	0.029

Notes: sample is the core annual firm panel for each stress-test vintage; outcome is the investment rate (capital expenditure over lagged total assets). The regressor is within-vintage, lender-country-demeaned stress-test exposure interacted with a post-publication indicator (Panel A) or additionally with a tightening indicator (Panel B); both panels include firm-by-event and country-by-year-by-event fixed effects. Standard errors, clustered by firm, in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Panel B's minimum detectable effect and equivalence bound are the smallest difference detectable at 80% power and the largest difference the estimate can rule out at 90% confidence.

riskier corporate books invest less in both anchor episodes—a borrower-risk gradient that is not state-dependent—while the depletion dose is. What finally rules out a purely sectoral reading is evidence the sector axis cannot touch: the within-firm comparison absorbs the firm—and with it the sector—entirely and still finds the gradient on the guidance calendar, and exposure predicts nothing at the rate shock itself.

A.2 The bond margin: further candidates and the share drift

Three simpler readings of the issuance decline were examined alongside the complementarity family of Table A11. The simpler candidates fare no better: the decline is no steeper among the firms that cut investment, and exposure predicts nothing at the rate shock itself, so simple demand does not account for it; exposed issuers were downgraded less, so credit quality does not; and exposure to the purchase programme's withdrawal adds a correctly signed but unconfirmed increment.

The share margin behaves as the stock arithmetic implies. The bond share of exposed firms'

Table A8: Decomposing the dose: corporate-book risk versus the capital requirement

Panel A: Excess NFC loss share (loss share minus exposure share)

	2021	2023	2021+2023 diff. (anchored)
Exposure × post	-0.001	-0.000	+0.001
	(0.003)	(0.003)	(0.004)

Panel B: NFC loss rate, with a book-share control

	2021	2023	2021+2023 diff. (anchored)
Exposure × post	-0.011**	-0.006	+0.003
	(0.006)	(0.006)	(0.006)
Book-share control	-0.073	+0.059	+0.010
	(0.047)	(0.056)	(0.037)

Notes: sample and specification mirror the hero table (Table A7): outcome is the investment rate, firm-by-event and country-by-year-by-event fixed effects, standard errors clustered by firm in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The exposure measure decomposes projected stress-test losses into corporate (NFC) versus household portfolios; Panel A nets out each bank's NFC book-share exposure from its NFC loss share, Panel B controls for it directly. The 2021+2023 column repeats the anchored regime-difference specification of Table A7, Panel B (2021 accommodation vs. 2023 tightening, bridged/v2 exposure), with each NFC variant in place of baseline exposure; the two single-event columns to its left remain the original (v1) NFC-heterogeneity static cells and are not re-derived here.

debt still drifts *up* through the post-publication halves, by about 1.7 percentage points per standard deviation of exposure three halves out, imprecisely estimated; but that tilt is the arithmetic of stocks when short loans run off faster than long bonds, not a rotation.

Two further characterizations behind Section 6.4's summary clauses are recorded here. On the composition of what exposed firms did issue: the issuance exposed firms retained skews heavily toward drawdowns of existing medium-term-note programmes, away from newly underwritten offerings, and routine rollovers ran on throughout. On the calendar of the one closure the period does contain: high-yield issuance troughed in the half-year of the fastest hikes, a full year before runoff began, and normalized by 2024, while investment-grade issuance never closed—which is what makes the contraction of exposed firms within an open market informative.

A.3 Mechanism interactions

Three further interactions test mechanisms the framework and the review of the evidence raise, and their results are reported whatever they say (Table A13). First, the funding-repricing mechanism does not detectably carry to the firm level: the firm-level triple and the within-firm pair triple are both insignificant under studentized randomization and the second family's adjustment (Panel A), and the detectable-effect bounds do not reach twice the 2023 point estimates, so the cells are uninformative rather than contradicting—the funding story stands on the bank-level interaction and its 2021 falsification, not on a firm-level cross-derivative. One anomalous cell: the 2021 falsification arm of the firm triple is positive at studentized $p = 0.051$, in an underpowered arm outside the prespecified families; we flag it and do not build on it. Second, the reliance interaction

Table A9: Is the dose interest-rate risk? Measured rate-risk exposure and the 2023 estimate

Panel A: correlation of the dose with bank rate-risk exposure (bank level)		
	Raw dose	Within-country dose
Wholesale funding share (pre-2022)	0.158	−0.060
FVOCI debt-securities share of assets	−0.096	−0.066
Debt-securities share incl. amortised cost	−0.286	0.054
Sovereign-book weighted maturity	0.081	−0.028

Panel B: the 2023 investment estimate conditional on rate-risk exposure (firm level)		
	Dose, conditional	Dose alone, same sample
+ Wholesale funding share	−0.0162*	−0.0165*
	(0.0084)	(0.0088)
+ FVOCI securities share	−0.0165*	−0.0166*
	(0.0087)	(0.0087)
+ Sovereign maturity	−0.0147*	−0.0157*
	(0.0083)	(0.0085)
+ All three	−0.0153*	−0.0165*
	(0.0081)	(0.0088)
+ All three and bank health	−0.0186**	−0.0192**
	(0.0084)	(0.0084)

Notes: Panel A reports Pearson correlations across the 63 measurable banks of the 2023 continuous-dose roster between each rate-risk measure and the bank's standardized adverse-scenario depletion, raw and within-country demeaned—the latter being the only variation the design uses. All measures are taken from pre-publication data: the wholesale funding share from the last pre-2022 transparency reference date, the securities shares and the sovereign-book maturity (trading book excluded; bucket-midpoint weighted) from the 2022 year-end templates the exercise itself was run on. Jointly, the three measures of Panel B explain 1.4 percent of the within-country dose variation. Panel B re-estimates the annual 2023 investment specification of Table 11 adding the firm's exposure-weighted lead-bank average of each measure (interacted with the post indicator); the right column re-estimates the dose alone on the identical subsample, so the comparison isolates the control from the coverage restriction. Standard errors clustered by firm; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Randomization p-values (bank-dose permutation within country, 1,000 draws) accompany every cell in the replication package and remain between 0.012 and 0.042 across all conditional cells; within-firm and bank-level counterparts show the same stability. Given the near-orthogonality in Panel A, the stability in Panel B is confirmation rather than a demanding test, and is reported as such.

that Section 6.6 uses carries the wrong sign relative to its mechanical benchmark—so it cannot even mechanically reproduce the coverage concentration—and non-Spanish exposed firms in the top reliance tercile are a statistical zero (Panel B): reliance does not explain the Spanish concentration, which remains a disclosed fragility. Third, the bond-underwriting margin: a bookrunner-weighted dose, built from the same disclosures over lead bookrunning roles, correlates with the loan-side exposure at 0.31 in the 2023 vintage and 0.38 in 2021—related, far from collinear—and the horse race is one-sided: the loan-side dose predicts the issuance contraction on the bookrunner-measurable sample alone (randomization $p = 0.028$) and jointly ($p = 0.041$), while the bookrunner dose carries nothing ($p = 0.275$; randomization $p = 0.402$), although three-quarters of the sample's bond deals involve a tested bookrunner. The asymmetry locates the contraction on the lending relationship, but it does not rank the two channels: the bookrunner dose's standard error is 0.140 against a loan-side point estimate of −0.272 on the 99 firms whose bookrunner history is measurable, so the underwriting channel is bounded rather than excluded.

Table A10: The within-firm estimate on the full borrower universe, the pooled regime parameter, and the registered 2025 prediction

	Estimate (std. error)	Randomization p	N
Panel A: the 2023 within-firm origination response, by borrower frame			
Committed pairs (reproduction)	−0.0202 (0.0074)	0.028	2,413
Full universe, ITT-commensurable	−0.0117 (0.0045)	0.110	6,668
Full universe, wide construction	−0.0083 (0.0040)	0.233	7,421
Listed borrowers	−0.0131 (0.0063)	0.106	2,790
Unlisted borrowers	−0.0049 (0.0048)	0.537	4,631
Panel B: the pooled regime parameter and the bounded arms			
Regime parameter γ, per half-year	−0.0042 (0.0023)	0.225 [0.341]	109,552
2021 accommodation arm	+0.0065 (0.0049)	—	7,507
2025 publication half	−0.0004 (0.0035)	0.928	8,588

Notes: The outcome throughout is the incidence of a new origination in the firm–bank pair. Units are window-level incidence per unit of dose, except the regime parameter γ of Panel B, which is per half-year per unit of dose and so is not commensurable with the window-level rows. Panel A rebuilds the pair panel on the full euro-area borrower universe the syndicated loan data covers: 7,945 distinct borrowers across the four exercises, most of them unlisted. Within the 2023 arm, 1,418 firms carry within-firm dose variation against the 367 of the committed listed frame; the corresponding 2021 counts are 1,673 against 443. Row 1 re-estimates the committed specification on the committed pairs alone: the 2,351 committed pairs map to 2,413 platform rows via shared listed identifiers, and coverage is 100 percent. Row 2 is the intent-to-treat-commensurable construction, defined exactly as the committed frame minus only the listed filter; row 3 is a wider construction that additionally admits pairs the committed frame's own definition excludes. The extended coefficient is 41 to 58 percent of row 1's committed-restricted coefficient, −0.0202, depending on which construction is used; that retention is measured against row 1 and not against the frozen reference −0.0220 (the baseline within-firm estimate), from which the two constructions diverge by 46.8 percent (row 2) and 62.3 percent (row 3), landing on opposite sides of the prespecified 50 percent threshold. The construction dependence is disclosed here rather than resolved, and no clean full-universe replication is claimed. Rows 4 and 5 split the wide construction of row 3, not row 2—their samples sum to its 7,421 rows—and both sides of the split are reported. In Panel B, the bracketed figure is the family-adjusted randomization p from the round's five-member stepdown, on which γ is exploratory; the pooled two-event model's both-vintage pair overlap is 34.4 percent, below its prespecified threshold, so a composition caveat attaches to γ. The 2021 arm carries no permutation p because it is adjudicated by confidence-interval inversion: its 90 percent studentized-permutation interval, $[-0.0042, +0.0173]$, lies entirely inside ± 0.0220, which excludes an accommodation-arm response as large as the committed 2023 estimate but nothing smaller. The 2025 row tests a registered prediction of no response before the guidance applies in January 2026; that arm's pre-period diagnostics do not pass the flatness screen, and the cell is reported with that flag attached. Standard errors: Panel A rows are firm-fixed-effect window-level cross-sections—the committed within-firm specification—with CRV1 errors clustered on the bank group. Panel B's regime parameter comes from the pooled two-event model carrying firm-by-event-by-half and firm–bank pair fixed effects, clustered on the bank group (CRV1, fast-demeaned and permutation-consistent); the 2021 row is that same committed within-firm specification estimated on the 2021 arm's window-level cross-section with the same clustering, which is what makes it commensurable with Panel A; the 2025 row is a firm-fixed-effect cross-section clustered on the bank group. No significance stars: the randomization p column carries the inference. Sources: src/183–189, studentized permutation with seed 42 and 1,000 draws throughout; committed cells in `leg_a_regime.csv`, `d_bridges.csv`, and `round3_family5_verdicts.csv`.

Table A11: Bond-margin complementarity tests

	Margin	Coefficient	SE	Randomization p	Family-adjusted p
Co-occurrence propensity	new issues	−0.016	(0.070)	0.74	0.95
Gateway vs. term facilities	new originations	−0.002	(0.002)	0.32	0.61
New-public vs. programme deals	new issues	+0.011	(0.046)	0.86	0.95
Shared guidance clock (S)	both margins	+2.94	—	0.018	0.043

Notes: the four members of the prespecified bond-margin test family (Section 5). Row 1: triple interaction of the post indicator, within-country demeaned exposure, and the firm's pre-window propensity of bond issues to co-occur with new or amended bank facilities, on the issuance-count panel of Table 8 (183 firms with a measurable co-occurrence history). Row 2: difference in the within-firm rationing gradient between gateway facilities (revolvers, backstops, credit lines, bridges) and term loans, in the stacked pair panel of Table 7; the facility-type classification covers 73 percent of the stacked panel against the 80 percent standard, so the cell is reported as coverage-limited. Row 3: triple interaction of the post indicator, exposure, and an indicator for newly underwritten public deals as against medium-term-note programme drawdowns. Row 4: the joint clock statistic of Section 6.3—the difference in the mean strength of the half-by-half response between the guidance-application halves of 2024 and the three preceding post-publication halves, summed over the loan and issuance margins; its randomization p is one-sided by construction, and the dash marks an entry not defined for the statistic. Randomization p-values are studentized, permuting the bank-level dose within country and vintage; family-adjusted p-values are Romano–Wolf stepdown values over the four members. The clock statistic is consistent with complementarity but does not identify it—shared timing is implied by any mechanism tying both margins to the guidance calendar. The complementarity hypothesis is unsupported: none of rows 1–3 reaches the supporting standard, and the 2021 falsification arms are quiet throughout. Sources: the syndicated-loan, facility, and bond-issuance panels of Section 4.

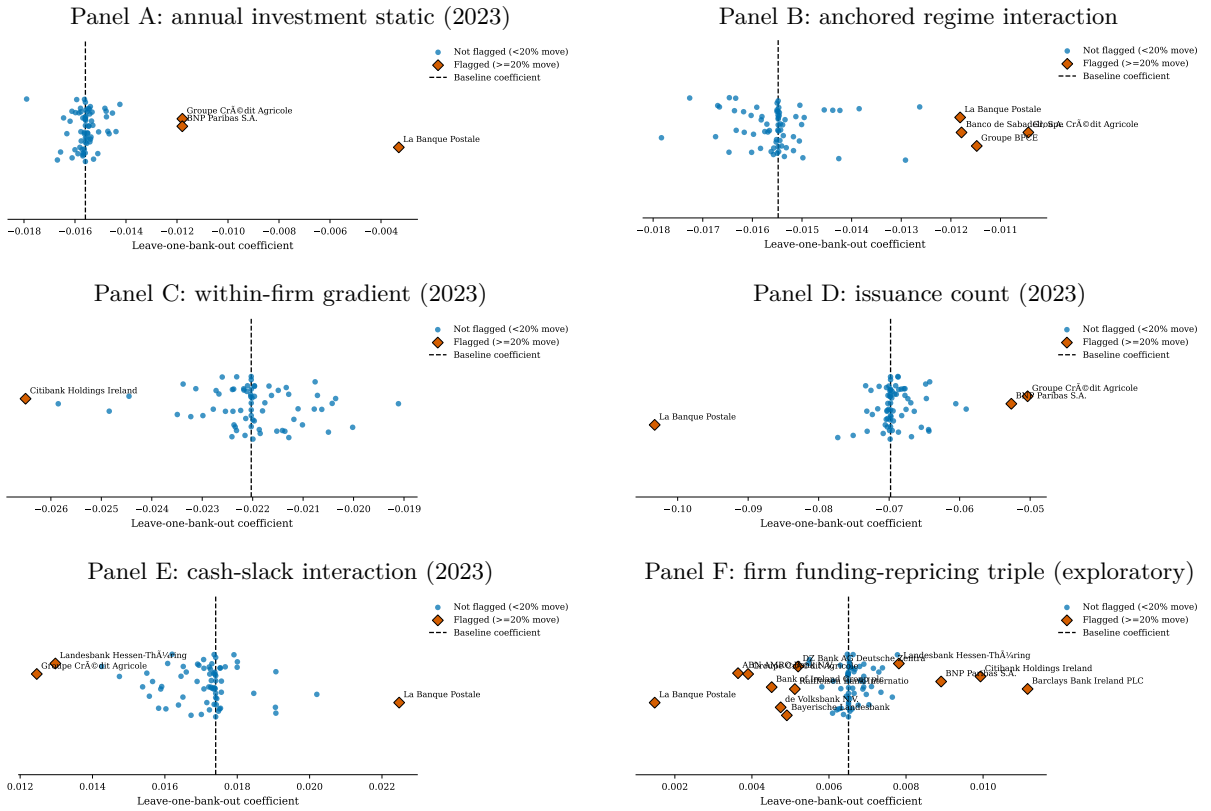


Figure A5: Leave-one-bank-out distributions for the headline coefficients. Each point is the coefficient re-estimated with one dosed banking group excluded (firms' exposure re-weighted over their remaining measured lead banks); the dashed vertical line marks the baseline coefficient, and diamonds flag exclusions that move the coefficient by at least 20 percent of its baseline value, labeled with the bank's name. Summary statistics for each panel are in Table A15. Sources: the estimation panels of Section 6.

Table A12: Dose measurability across vintages: degradation test and attenuation bounds

Panel A: 2023 measurability-degradation falsification			
Metric	Coef. (SE) / value	N / draws	Note
Baseline (unmasked rebuild)	-0.0156 (0.0085)	–	–
Identity-mask degraded (10 banks masked)	-0.0196** (0.0081)	–	shift vs. pin -0.0040
Materiality ratio (|shift|/|raw 2023–2021 contrast|)	0.3264	–	BORDERLINE (thr. 0.33)
Sign-flip falsification (#4)	False	–	–
Attenuation falsification (#4)	–	–	AMBER (away from zero)
Random-mask distribution (200 draws)	mean -0.0140 (sd 0.0042)	200	p5/p25/med/p75/p95 below
percentiles [p5, p25, median, p75, p95]	[-0.0185, -0.0172, -0.0144, -0.0125, -0.0050]		
Sign-flip share (200 draws)	0.0%	200	–
Materiality share (≥0.33, 200 draws)	19.5%	200	–
Outside-support check (identity vs. random)	0.0%	200	OUTSIDE_SUPPORT
Panel B: 2021 attenuation/equivalence bounds (λ=0.7273, mean share_measurable)			
Outcome	Coef. (SE), 2021	MDE80/λ	TOST90/λ
Investment rate (annual)	-0.0035 (0.0031)	0.0120	0.0118
Gross capex rate (quarterly)	-0.0022 (0.0015)	0.0058	0.0064
Net PP&E investment (quarterly)	-0.0028 (0.0022)	0.0083	0.0087
Sales growth (quarterly)	-0.0603 (0.0246)	0.0946	0.1385
Debt growth (quarterly)	-0.0227 (0.0347)	0.1336	0.1096
Sales growth (expansion-only robustness)	-0.0745 (0.0327)	0.1261	0.1765
Debt growth (expansion-only robustness)	-0.0305 (0.0354)	0.1362	0.1220

Notes: Panel A re-estimates the 2023 annual investment static after imposing the 2021 disclosure regime on the 2023 dose: the identity mask removes the ten banks whose 2023 results would have been range-only under the 2021 rules and rebuilds exposure over the remaining measured lead banks; the random-mask distribution repeats the exercise for 200 random masks of the same size. The materiality ratio compares the identity-mask shift to the raw 2023–2021 contrast (threshold one-third). Panel B reports the 2021 anchor-pair nulls with detectable-effect bounds inflated by the vintage's mean measured credit share ($\lambda = 0.7273$): MDE80 is the minimum detectable effect at 80 percent power, TOST90 the 90-percent equivalence bound; both are upper bounds on the true 2021 response per unit of exposure under full attenuation correction. Standard errors clustered by firm in parentheses; stars where analytic p-values are defined: ** $p < 0.05$. Sources: the annual and quarterly estimation panels of Section 6.

A.4 Separating the test from the tightening

The 2023 results record supervisory pressure interacting with the monetary state; they are not the tightening itself wearing supervisory clothes. Four facts separate the two. First, exposure to the 2023 test does not predict a firm's response to the rate shock: dated at the first policy-rate increase—a year before publication—the firm-level placebo is indistinguishable from zero ($p = 0.88$). Second, the dose is not a relabeled weak-bank indicator: entered jointly with the observable dimensions of bank health, the depletion dose retains its size and sharpens (randomization $p = 0.008$). Third, the timing within the tightening is the supervisory calendar's, not the monetary one's (Section 6.3). Fourth, the dose is not disguised interest-rate risk: the 2023 adverse scenario embedded a rate shock, but the overlap is a between-country phenomenon—measured rate-risk exposures jointly explain 1.4 percent of the within-country dose variation the design uses, and entering them directly leaves every headline estimate intact (Table A9). Rates

Table A13: Mechanism interactions: funding-repricing and reliance

Panel A: A2 within-vintage state-variable triples (funding-repricing moderator W)					
Cell	Coef. (SE)	Stud. perm. p	RW-adj. p (Fam. 2)	N (firms)	Note
Firm triple (β_3, 2023 headline)	+0.0065 (0.0059)	0.261	0.598	338	–
Firm triple, 2021 falsification	+0.0104** (0.0045)	0.051	n/a	388	MDE80/λ=0.018 vs. 2$\lvert\hat{\beta}_3^{23}\rvert$=0.013: UNINFORMATIVE – BOUNDED
Pair triple (δ_3, 2023 headline)	-0.0210 (0.0252)	0.492	0.598	365	–
Pair triple, 2021 falsification	-0.0238 (0.0291)	0.589	n/a	443	MDE80/λ=0.115 vs. 2$\lvert\hat{\delta}_3^{23}\rvert$=0.042: UNINFORMATIVE – BOUNDED
Panel B: A4 reliance triple (tested-sample-membership moderator R)					
	Coef. (SE)	Stud. perm. p	RW-adj. p (Fam. 2)	N (firms)	Note
Reliance triple (β_3, headline)	+0.0090* (0.0054)	0.149	0.415	354	–
Mechanical benchmark $\hat{\beta}_3^{mech}$	-0.0023	–	–	–	sign cond. NOT met; $\lvert\hat{\beta}_3\rvert - \lvert\hat{\beta}_3^{mech}\rvert$=+0.0067
Non-Spanish top-tercile decision cell	-0.0062 (0.0058)	0.470	n/a	136	–

Notes: Panel A interacts the 2023 (and, as falsification, 2021) exposure terms with the lead banks' pre-event funding-repricing sensitivity W (wholesale-funding share, demeaned), at the firm level (annual investment panel) and within the pair panel of Section 6. Panel B interacts the 2023 exposure terms with the firm's pre-window reliance on tested lead banks R; the mechanical benchmark is the coefficient implied by coverage arithmetic alone under the null of no reliance channel, and the final test restricts to non-Spanish firms in the top reliance tercile. Standard errors clustered by firm in parentheses; studentized permutation reassigns the bank-level dose within country and exercise; family-2 columns report the stepdown-adjusted p-value within the prespecified family of added interactions. MDE notes state when a 2021 falsification arm cannot detect twice the corresponding 2023 point estimate (uninformative, bounded). * $p < 0.10$, ** $p < 0.05$. Sources: the estimation panels of Section 6 matched to bank funding structures and stress-test rosters.

set the state; the test selected the banks; the guidance set the clock.

A.5 Windows, composition, and further checks

Because the treatment stretches from the announcement of an exercise to the implementation of the resulting guidance, no single event window is canonical; every headline specification is estimated over four windows per exercise—tight, baseline, launch-anchored, and wide (Mésonnier and Monks, 2015; Bedayo and Galán, 2024). The annual contrast is insensitive to the window choice by construction, and the quarterly statics are stable however the window is drawn. The results likewise survive the natural redefinitions of exposure: measuring it over all syndicate members rather than lead arrangers, dropping firms whose tested lenders are all non-lead participants, or weighting banks by observed lender shares moves the 2023 headline coefficient by a few percent, while replacing scenario depletion with the distance of the adverse-scenario minimum to a fixed threshold yields a flat null everywhere—the transmission runs off how much capital the scenario destroys, not off the level the bank retains. Stress-test depletions do not move the following year's published Pillar 2 requirement (online appendix): the results feed the guidance and the public signal, not the binding requirement. Three design margins carry no usable variation and each returns the expected nothing: essentially every significant euro-area lender is stress-tested, so no tested-versus-untested comparison exists; template versus range-only disclosure does not differentiate bank responses; and a coarse dose built from the disclosed ranges adds no precision. The identification margin is, and remains, the continuous dose within tested banks.

The remaining checks concern sample and placebos. The sub-annual real outcomes are reported for five candidate investment measures, and no measure reverses a sign reported in the text. A placebo publication dated inside the pre-window—2019 against 2020—produces no

Table A14: Additional inference and construction checks

	Coef.	SE	Analytic p	Coef. rand. p	Stud. rand. p	Exp.-rob. SE
Panel A: marginal effect of exposure on investment, by cash percentile (2023, full roster)						
p10 (cash 3% of assets)	−0.032	(0.010)	0.001	0.002	0.035	(0.014)
p25 (6%)	−0.025	(0.009)	0.005	0.003	0.096	(0.009)
p50 (10%)	−0.018	(0.008)	0.032	0.022	0.196	(0.007)
p75 (16%)	−0.005	(0.008)	0.526	0.421	0.626	(0.010)
p90 (23%)	+0.007	(0.010)	0.453	0.322	0.468	(0.024)
ex-Spain: p25	−0.014	(0.005)	0.009	0.026	0.080	(0.010)
ex-Spain: p50	−0.007	(0.005)	0.125	0.213	0.315	(0.008)
covered subsample: p25	−0.013	(0.005)	0.008	0.010	0.053	(0.008)
covered subsample: p50	−0.006	(0.004)	0.197	0.295	0.402	(0.006)
Panel B: renewal versus origination, stacked pair panel						
2023: origination margin	−0.032	(0.010)	0.003	—	—	—
2023: difference, renewal − origination	+0.035	(0.016)	0.038	0.030	0.088	—
2023: renewal margin only	+0.002	(0.011)	0.829	—	—	—
2021: origination margin	+0.001	(0.018)	0.974	—	—	—
2021: difference, renewal − origination	−0.046	(0.026)	0.088	0.039	0.158	—
Panel C: orthogonalized dose (2023 dose residualized on 2021 within country)						
Within-firm gradient	−0.024	(0.009)	0.012	—	—	—
Cash-slack triple	+0.016	(0.005)	0.004	—	—	—
Panel D: measurement filters (2023)						
Weld, share measured ≥ 0.75	−0.026	(0.010)	0.016	—	—	—
Cash triple, share measured ≥ 0.75	+0.017	(0.006)	0.003	—	—	—
Annual static, share measured ≥ 0.75	−0.006	(0.005)	0.171	—	—	—
Annual static, dilution bound	−0.015	(0.010)	0.127	—	—	—
Cash triple, dilution bound	+0.020	(0.006)	< 0.001	—	—	—

Notes: additional inference and construction checks on the estimation samples of the main text. Panel A evaluates the level of the marginal-effect curve of Figure 4, $\beta_1 + \beta_3 c$, at percentiles of the pre-window cash distribution: analytic p is the delta-method p-value from the firm-clustered covariance; randomization p-values reassign the bank-level dose within country and vintage (coefficient and studentized variants, the studentized primary for interaction-derived statistics); the exposure-robust column reports the shock-level standard error of the same quantity via exact reparametrization. The ex-Spain and covered-subsample rows re-evaluate the curve excluding Spain and excluding annual-only reporters (≥ 6 clean quarters), respectively. Panel B stacks the origination and renewal margins of the pair panel of Table 7 and tests their dose responses' difference; the renewal-only 2023 response carries an 80-percent detectable-effect bound of 0.032. Because the stacked frame pools the origination and renewal margins in a single specification, its margin coefficients are not the headline within-firm estimate: Panel B's 2023 origination coefficient (−0.032) is the stacked frame's own margin coefficient, while the headline gradient reported in Section 6.2 and Table 7 is −0.022 from the origination panel estimated on its own. Panel C rebuilds the dose as the residual of the within-country demeaned 2023 depletion on its 2021 counterpart over the 47 banks measured in both vintages (retained lead weights renormalized); the stratified-permutation support disclosure accompanying it: 24 of 35 country-by-prior-tercile cells are singletons, carrying 21.7 percent of the weight-mass-weighted dose variance. Panel D restricts to firms whose measured share of lead credit is at least three-quarters, and bounds the exposure aggregation by scoring unmeasured lead banks at dose zero; in the pair panel, 127 of 2,351 pairs (24 firms) lack share data entirely and form an explicit stratum, disclosed rather than conflated with sub-threshold pairs; the pair-level design has no dilution analogue, since it regresses on the bank's own dose. Two cells not tabulated here: the pooled bank-level cross-vintage difference tests referenced in Section 6 (+0.007 on assets, studentized randomization $p = 0.82$, $n = 90$; −0.037 on loans, $p = 0.47$, $n = 89$), and the ex-Spain marginal-effect curve's top decile, where the point estimate turns positive and significant (+0.016 at the ninetieth cash percentile) — the "absent among cash-rich firms" statement holds on the full roster (+0.007, not significant) and is scoped to it. Standard errors clustered as in the tables reporting them. Sources: the estimation panels of Section 6; every cell is reproduced in the replication package.

Table A15: Leave-one-bank-out influence analysis

Target	Baseline coef. (SE)	Min	Max	Banks flagged ≥20%	Max-influence bank
2023 annual static	-0.0156 (0.0085)	-0.0179	-0.0033	3 / 70	La Banque Postale (+78.8%)
Anchored 2021+2023 interaction	-0.0155 (0.0083)	-0.0178	-0.0104	4 / 73	Groupe Crédit Agricole (+32.6%)
Weld incidence (2023)	-0.0220 (0.0076)	-0.0265	-0.0191	1 / 70	Citibank Holdings Ireland Limited (-20.3%)
Issuance count (2023)	-0.0698 (0.0413)	-0.1033	-0.0504	3 / 70	La Banque Postale (-47.9%)
Cash-slack triple (2023)	+0.0174 (0.0048)	0.0125	0.0225	3 / 70	La Banque Postale (+29.1%)
A2 firm triple (2023, exploratory)	+0.0065 (0.0059)	0.0015	0.0112	12 / 70	La Banque Postale (-77.3%)

Notes: each row re-estimates one headline specification of Section 6 excluding one dosed banking group at a time (the firms whose exposure the excluded group carries are re-weighted over their remaining measured lead banks). Min and Max are the extreme coefficients across exclusions; banks are flagged when a single exclusion moves the coefficient by at least 20 percent of its baseline value, with the named bank the largest such move (sign indicates the direction of the coefficient change on exclusion). Baseline standard errors in parentheses as in the underlying estimates. The final row is the firm-level funding-repricing triple of Table A13, included for completeness. Sources: the estimation panels of Section 6.

effects. Controlling for each firm's pre-window saturation under the corporate-sector purchase programme leaves the dose estimate intact on the subsample where saturation is measured, and saturation itself predicts nothing there. Winsorizing outcomes and excluding firms flagged in the identifier reconciliation both leave the pattern of Table A7 in place.

The cross-state contrast survives its natural rescalings, for what an exploratory estimate's stability is worth: the two vintages' raw depletion dispersions are comparable, re-estimating per percentage point of depletion moves the coefficient by one percent, restricting the 2023 arm to the banks dosed in both exercises sharpens the contrast (randomization $p = 0.028$), and reweighting the 2023 roster to the 2021 coverage distribution retains nine-tenths of it. None of this changes the estimate.

B Data construction

This appendix documents the construction of the dataset in the order the pipeline builds it: firm identification, lender consolidation, treatment, exposure, outcome panels, and the bank-level panels. Every merge in the pipeline logs its match rate, and the replication package reproduces each step from raw inputs.

B.1 Firm identification

The unit of observation is the listed firm, keyed by its equity ISIN. Syndicated-loan borrowers (15,062 distinct euro-area borrowers) are resolved to equity ISINs through four routes applied in strict priority order. The primary route uses the permanent organization identifiers that the loan data carry for about three quarters of borrowers, resolved to equity ISINs through the securities reference data of the same vendor; it identifies 1,957 borrowers. A second route maps borrowers' legal entity identifiers to equity ISINs through the public Global LEI Foundation records (28 additional borrowers, with a further 1,670 borrowers matched only to debt securities, retained as evidence of bond-market activity rather than as firm keys). A third route bridges the

Table A16: Completing the inference: headline coefficients under every scheme

		Coef. (SE)	Perm. p (coef.)	Stud. perm. p (studentized)	RW-adj. p (Family 1)	Stratified p (prior-dose tercile)	Eff. shocks (banks)
2023 annual static (investment rate) (cell_id root: annual_static_2023)							
Firm cluster (baseline sample)	(N=1,343)	-0.0156* (0.0085)	0.033	0.226	0.289	0.022	25.16
Firm cluster (lead-bank-matched sample)	(N=1,263)	-0.0180** (0.0083)	–	–	–	–	–
Lead-bank cluster (lead-bank-matched sample)	(N=1,263)	-0.0180** (0.0077)	–	–	–	–	–
BHJ exposure-robust, HC1 (shock-level)	(N=1,343)	-0.0156 (0.0073)	–	–	–	–	–
BHJ exposure-robust, CR1 (lender-country)	(N=1,343)	-0.0156 (0.0090)	–	–	–	–	–
Anchored 2021+2023 interaction (cell_id root: anchored_interaction)							
Firm cluster (baseline sample)	(N=4,023)	-0.0155* (0.0083)	0.055	0.133	0.289	0.425	25.16
Firm cluster (lead-bank-matched sample)	(N=3,722)	-0.0162* (0.0086)	–	–	–	–	–
Lead-bank cluster (lead-bank-matched sample)	(N=3,722)	-0.0162*** (0.0055)	–	–	–	–	–
BHJ exposure-robust, HC1 (shock-level)	(N=4,023)	-0.0155 (0.0164)	–	–	–	–	–
BHJ exposure-robust, CR1 (lender-country)	(N=4,023)	-0.0155 (0.0189)	–	–	–	–	–
Weld incidence (2023) (cell_id root: weld_incidence_2023)							
Bank-group cluster (native pair-level design)	(N=2,351)	-0.0220*** (0.0076)	0.004	0.018	0.020	0.022	25.29
Firm cluster (same sample)	(N=2,351)	-0.0220** (0.0103)	–	–	–	–	–
Two-way (firm + bank-group) cluster	(N=2,351)	-0.0220** (0.0091)	–	–	–	–	–
BHJ exposure-robust, HC1 (shock-level)	(N=2,351)	-0.0220 (0.0076)	–	–	–	–	–
BHJ exposure-robust, CR1 (lender-country)	(N=2,351)	-0.0220 (0.0073)	–	–	–	–	–
Issuance count (2023) (cell_id root: issuance_count_2023)							
Firm cluster (baseline sample)	(N=4,992)	-0.0698* (0.0413)	0.006	0.037	0.289	0.280	25.16
Firm cluster (lead-bank-matched sample)	(N=4,680)	-0.0767* (0.0439)	–	–	–	–	–
Lead-bank cluster (lead-bank-matched sample)	(N=4,680)	-0.0767** (0.0345)	–	–	–	–	–
BHJ exposure-robust, HC1 (shock-level)	(N=4,992)	-0.0698 (0.0215)	–	–	–	–	–
BHJ exposure-robust, CR1 (lender-country)	(N=4,992)	-0.0698 (0.0204)	–	–	–	–	–
Cash-slack triple (2023) (cell_id root: cash_slack_triple)							
Firm cluster (baseline sample)	(N=1,343)	+0.0174*** (0.0048)	0.002	0.004	0.003	0.005	25.00
Firm cluster (lead-bank-matched sample)	(N=1,263)	+0.0163*** (0.0047)	–	–	–	–	–
Lead-bank cluster (lead-bank-matched sample)	(N=1,263)	+0.0163*** (0.0053)	–	–	–	–	–
BHJ exposure-robust, HC1 (shock-level)	(N=1,343)	+0.0174 (0.0039)	–	–	–	–	–
BHJ exposure-robust, CR1 (lender-country)	(N=1,343)	+0.0174 (0.0038)	–	–	–	–	–

Notes: each block reports one headline coefficient under alternative standard errors and randomization schemes; outcomes, samples, and fixed effects are those of the corresponding exhibits in Section 6. Rows within a block: firm-clustered standard errors on the baseline sample; firm- and lead-bank-clustered standard errors on the subsample matched to a lead banking group; exposure-robust standard errors from the equivalent shock-level regression, heteroskedasticity-robust (HC1) and clustered by lender country (CR1). The within-firm block's native design clusters by banking group, with firm and two-way variants beneath. Columns: randomization p-values reassigning the bank-level dose within country and exercise (coefficient and studentized variants); the stepdown family-adjusted p-value within the prespecified headline family; the randomization p stratified on terciles of the banks' prior-vintage dose; and the effective number of bank-level shocks. The coefficient- and studentized-permutation columns use this table's own fresh randomization sequence; last digits can differ from the underlying estimates' sequences (Appendix B.7 reconciles the pairs). Stars accompany analytic p-values only: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; dashes mark combinations not defined for a row. Sources: the estimation panels of Section 6.

Table A17: Robustness compendium: appendix backing for the checks reported in Sections 6 and 7

Check	Coef. (SE)	N
Panel A: Inference with few banks		
Anchored 2021–2023 interaction, sector × period fixed effects	-0.0012 (0.0110)	3,404
Anchored interaction, dose in raw percentage points of depletion	-0.0062* (0.0037)	4,023
Anchored interaction, 2023 arm restricted to banks also dosed in 2021	-0.0183** (0.0085)	3,935
2023 static on the same overlap roster	-0.0199** (0.0084)	1,287
Anchored interaction, 2023 roster reweighted to the 2021 coverage distribution	-0.0143* (0.0084)	4,023
Specification surface, 1,080 interaction cells: share with randomization $p < 0.10$ / < 0.05	4.35% / 1.67%	1,080
Panel B: Composition of the dose and of the syndicate		
2023 static, exposure measured over all syndicate members (not lead arrangers only)	-0.0162* (0.0086)	1,380
2023 static, firms whose tested lenders are all non-lead participants dropped	-0.0157** (0.0079)	1,294
Correlation of lender-share-weighted with equal-split exposure (2023, demeaned)	0.9985	590
2023 static with distance-to-threshold replacing scenario depletion	0.0023 (0.0048)	1,471
Panel C: Separating the test from the tightening		
Firm-level placebo dated at the first policy-rate increase (2022H2)	0.0011 (0.0071)	1,023
2023 static, dose entered jointly with pre-window CET1, ROA, deposit ratio and size	-0.0226** (0.0104)	1,297
Panel D: Sample and placebo checks		
Gross capital expenditure, 2021, half-yearly	-0.0087 (0.0074)	2,646
Gross capital expenditure, 2023, half-yearly	-0.0007 (0.0013)	2,522
Net investment in productive assets, 2021, half-yearly	-0.0082 (0.0079)	2,912
Net investment in productive assets, 2023, half-yearly	-0.0098 (0.0059)	2,759
Capex/net-investment hybrid, 2021, half-yearly	-0.0082 (0.0071)	2,995
Capex/net-investment hybrid, 2023, half-yearly	-0.0080 (0.0053)	2,801
Asset growth (Δ log assets), 2021, half-yearly	0.0034 (0.0121)	3,016
Asset growth (Δ log assets), 2023, half-yearly	-0.0330 (0.0273)	2,824
PP&E growth (Δ log PP&E), 2021, half-yearly	-0.0139 (0.0242)	2,912
PP&E growth (Δ log PP&E), 2023, half-yearly	-0.0003 (0.0132)	2,759
sign agreement across all 20 measure × event × frequency cells	19 of 20 negative	20
2023 annual investment, Spain excluded	-0.0045 (0.0048)	1,196
2023 annual investment, France excluded	-0.0218** (0.0107)	1,051
2023 half-yearly net PP&E investment, Spain excluded	-0.0061 (0.0053)	1,985
Placebo event dated in the pre-window (2019 vs. 2020), annual investment	0.0029 (0.0028)	297
same placebo, stacked difference-in-differences (pre-window 2019–2020)	-0.0038 (0.0037)	909
2023 annual investment controlling for pre-window CSPP saturation	-0.0298*** (0.0113)	589
the saturation term itself, same specification	0.0050 (0.0067)	589
Annual investment, outcomes winsorized (estimation baseline)	-0.0076* (0.0041)	2,014
same cell on raw, unwinsorized outcomes	-0.0116* (0.0061)	2,014
same cell excluding firms flagged in the identifier name-matching step	-0.0084** (0.0041)	1,908
Panel E: Results-section checks		
Pandemic downturn placebo (2020, no stress-test vintage): access differential in annual investment	0.0021 (0.0054)	1,191
minimum detectable gap at 80% power, same cell	0.0152	1,191
Access gradient, 2023: exposure × no-access triple interaction, no added controls	-0.0119 (0.0172)	2,022
the same triple interaction conditioning on firm size	-0.0137 (0.0179)	2,022
the same triple interaction conditioning on leverage	-0.0127 (0.0168)	2,022
the same triple interaction conditioning on the volatility of pre-period sales growth	-0.0115 (0.0162)	2,013
the same triple interaction conditioning on profitability	-0.0054 (0.0176)	1,638
Investment magnitude, 2023: moving a firm from the 25th to the 75th percentile of exposure	-0.0062	–
the same move in euro at the median sample firm's assets (EUR millions)	-24.87	–

Notes: each row reproduces a committed result cell rather than a new estimate; the source file and the filter isolating the cell travel with the replication package (`robustness_compendium.csv`), which also carries the randomization p-values, seeds and draw counts left out here. Coefficients are per within-vintage standard deviation of adverse-scenario depletion unless the row states another scale, standard errors are clustered as in the source specification, and stars follow that clustered inference: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A's last row counts the disclosed interaction surface: 47 of the 1,080 cells clear the ten-percent randomization threshold and 18 clear the five-percent one. Those cells sit in the primary design block: the share falls from 11.94 percent among continuous-dose cells to 0.56 percent under binary treatments, 0.56 percent under sector-by-period effects and 1.11 percent on the quarterly-covered subsample. Panel D prints the half-yearly cell of each sub-annual investment measure; its quarterly counterparts, the twenty-cell sign tally, the remaining surface axes, the pre-bridge provenance flags and the rows reported in the online appendix are in the CSV.

Table A18: The sector margin: additive fixed-effect ladder and loss-composition arbitration

Panel A: sector ladder (2023 annual static, term post:E_dm_eq)	Coef. (SE)	N (firms)
Rung 0: no sector control (G1 pin)	-0.0156* (0.0085)	354
Rung 1: SIC 1-digit division, REPLACEMENT	-0.0097 (0.0093)	310
Rung 1: SIC 1-digit division, ADDITIVE	-0.0154 (0.0098)	310
Rung 2: SIC 2-digit sector, REPLACEMENT	-0.0122 (0.0104)	310
Rung 2: SIC 2-digit sector, ADDITIVE	-0.0191* (0.0105)	310
Panel B: A5 arbitration cells (2023, decomposable-bank subsample)	Coef. (SE)	Stud. perm. p / Perm. p
(i) Non-NFC dose, full sample	-0.0131** (0.0066)	0.208 / 0.025
(ii) Standard dose, non-NFC-driven subsample	-0.0163* (0.0083)	0.187 / 0.024
(iii) Joint spec: NFC component	-0.0021 (0.0044)	0.766 / 0.657
(iii) Joint spec: non-NFC component	-0.0126** (0.0060)	0.169 / 0.035

Notes: Panel A re-estimates the 2023 annual investment static under sector fixed-effect rungs: replacement rungs substitute sector-by-period for country-by-period effects; additive rungs keep country-by-period and add sector-by-period. Panel B decomposes each bank's adverse-scenario losses by portfolio: (i) a dose built from non-corporate losses only, on the full sample; (ii) the standard dose on the subsample of firms whose exposure is driven by non-corporate-loss banks; (iii) both components entered jointly. Sample is the subsample of banks whose loss decomposition is disclosed. A regression of exposure on two-digit sector dummies given country fixed effects yields an incremental R^2 of 0.125. Standard errors clustered by firm in parentheses; * $p < 0.10$, ** $p < 0.05$. Sources: stress-test loss disclosures matched to the annual estimation panel of Section 6.

current loan data to its legacy predecessor and that predecessor's published link to accounting identifiers (1,196 borrowers). A residual route accepts exact normalized-name matches against the listed-firm universe only when the name is unique on both sides, the borrower and firm country agree, and the firm's reporting history overlaps the borrower's facility activity (89 borrowers, all flagged for review). In total 3,270 borrower records resolve to 2,681 distinct listed firms, of which 1,558 appear in the accounting panel; conflicts between routes (256 cases) are resolved in favor of the higher-priority route and flagged. Where a firm carries an accounting identifier from the legacy bridge it is retained as an attribute, but no merge is keyed on it.

B.2 Lender consolidation

Lenders are consolidated to the tested banking group. A crosswalk of 10,782 lender records maps lead arrangers to group legal entity identifiers, hand-reviewed for the roughly one hundred groups that matter, with unmatched lenders retained and flagged rather than dropped (they carry about five percent of lead volume). Ownership changes are represented as validity spans—for example,

a lender acquired in 2021 maps to its acquirer only for facilities originated after the acquisition date—and every join against the crosswalk is date-aware. Lead arrangers are identified from a fixed list of twenty arranger roles; syndicate participants are excluded from the baseline exposure and used only in the robustness variant and the contamination flag described in the text.

B.3 Treatment

Bank-level treatment is read directly from the published stress-test disclosures. For banks in the European Banking Authority samples, the fully loaded common equity ratio at the starting point and at its adverse-scenario minimum are taken from the bank-by-bank templates; because template item codes change with every exercise, each vintage is decoded against its own data dictionary. Depletion is the difference between the two ratios, standardized within the exercise; banks disclosed only in ranges by the supervisor carry no continuous treatment and their credit is tracked as unmeasurable. The decomposition of projected losses by counterparty sector uses the credit-risk templates, restricted to aggregate country rows to avoid double counting the per-country breakdowns, summing internal-ratings-based and standardized portfolios; losses are measured as the change in the stock of provisions between the starting point and the final adverse-scenario period. The corporate-loss measure used in the text is normalized against the corporate share of the bank's book, since the raw corporate loss share largely reflects business-model composition.

B.4 Exposure

For each exercise, firm exposure is built from lead-arranger credit observed in a pre-exercise window that closes before the exercise's results are published: 2019–June 2021 for the 2021 exercise (the sample was named in January 2021, but no bank's depletion was known until the end-July publication) and 2021–2022 for 2023, before that exercise's January launch. Credit is attributed in equal shares across a facility's lead arrangers, because within-syndicate allocations are reported for only about a fifth of lender records; a share-weighted variant using reported allocations where complete correlates above 0.99 with the equal-split baseline. Exposure is the credit-weighted average of lead banks' standardized depletions, demeaned within the lead banks' country, with weights renormalized over the lead banks that carry a measured dose and the measured share of each firm's lead credit stored alongside; firm–bank weight matrices are stored so that randomization inference can rebuild exposure from permuted bank-level treatments exactly. Firms whose measured lead credit is zero in the window carry no exposure value rather than a zero.

B.5 Firm outcome panels

Accounting outcomes come from annual and interim reports of listed firms in the twenty euro-area countries. Because the same security occasionally carries more than one vendor record, records are deduplicated by preferring the canonical December fiscal close and the more complete record; growth rates and lagged ratios are then re-derived at the ISIN level across consecutive periods only, which prevents lag chains from crossing record switches (the vendor's own precomputed

lags occasionally do). Fiscal periods are assigned to calendar time by their period-end dates, never by fiscal-year labels. Investment is capital expenditure over lagged total assets and is annual; interim filings of European firms rarely report capital expenditure and never report employment, so sub-annual real outcomes are not available from this source. Debt structure comes from instrument-level records that classify bonds and bank debt and date each instrument; the bond share of debt and net issuance indicators are built by aggregating instruments to the firm-quarter.

B.6 Bond-market activity and spreads

A firm's bond-market access is measured by the union of three sources of evidence, kept separate for robustness: instrument-level records showing bonds outstanding in the relevant window; security-level records of euro-denominated bonds alive in the window, drawn from a discovery pass over issuer identifiers that includes matured bonds (58,925 securities, essential because active-only searches miss the majority of historical issues); and debt securities registered to the firm's legal entity identifier. Weekly secondary-market spreads cover 321 firms; the panel deliberately over-represents larger issuers, a selection stated wherever spreads are used.

B.7 Estimation and exhibit conventions

Conventions that apply across the paper's estimates are stated here once. Randomization p-values use one thousand draws with a fixed seed, permuting the bank-level depletion within country and vintage and rebuilding firm exposure through the fixed firm–bank weight matrices; they are reported in the text, while tables carry clustered standard errors and stars only. Where the same statistic is computed both in a source exhibit and in the inference-completion table, the two use different fresh draw sequences from the same fixed seed protocol, so last digits can differ—the within-firm gradient's coefficient-randomization p of 0.003 (source) versus 0.004 (table), the annual static's 0.033 versus the table's 0.029—Monte Carlo resampling differences, not changed estimates. Annual event time indexes fiscal years relative to the exercise's first post-publication fiscal year, which sits at index 0; fiscal years are assigned to calendar time by their period-end dates. For the 2023 exercise, index 0 is FY2023: the fiscal year containing the July publication counts as post because it lies entirely after the exercise's January launch—no part of it is clean pre—and carries five post-publication months; the base year FY2022 sits at -1 and the window is $[-2, +1]$, ending with FY2024 before the following exercise's aftermath. For the 2021 exercise the straddling fiscal year is excluded from both sides rather than assigned: FY2021 splits roughly evenly around the end-July publication, so it is dropped, the base year FY2020 sits at -2, and index 0 is FY2022, the first full post year and the last before the January 2023 launch of the next exercise; the window $[-2, 0]$ therefore contains no observation at -1. Blank table cells denote horizons that a truncation rule or data availability removes, never dropped observations. The access exhibit's sensitivity panel uses terciles of the raw exposure measure and the union access definition, a deliberately different construction from the main panels' within-country median split, so its cells are not comparable one for one. The corporate-loss decomposition never uses the raw corporate loss share, which tracks business models nearly tautologically; both normalized variants appear. Firm-level levels are not currency-harmonized for the roughly five percent of

non-euro reporters, and syndicated credit volumes are converted at deal dates, so credit-to-asset ratios are approximate for those firms. The aggregate credit share and the average firm's credit share differ by eight to thirteen percentage points because larger firms hold more relationships, including more untested lenders, diluting their individual shares relative to the credit-weighted pool.

B.8 Bank panels

Semiannual bank balance sheets are assembled from the transparency exercises, decoded by label rather than by item code for the same reason as the treatment templates. One template idiosyncrasy matters: the item labeled as combining both regulatory approaches in fact carries only standardized-approach rows in every round, so total exposures are reconstructed from the by-class items; without this correction, corporate exposures would exceed totals for a quarter of bank-half-years. Quarterly bank financial statements come from vendor records keyed to registered identifiers, matched to group legal entity identifiers for 141 of 146 groups; interim flows are reported year-to-date and are de-cumulated within the fiscal year, with a flow treated as missing whenever the preceding filing is absent, and stocks taken as reported. Consolidated filings are preferred over unconsolidated ones where both exist. Vendor mnemonics are validated by reproducing the vendor's own published ratios and by reconciling year-end interim values against the annual statements before any field is used.

C Model Derivations

Notation as in Section 3. The bank enters with equity K_0 and loans L_0 (risk weights normalized, so L doubles as risk-weighted assets), and its ratio carries headroom h over the standing stack: $K_0/L_0 = \rho + h$. The test publishes the depletion d; guidance $g(d)$, with $g' > 0$, attaches from the application date τ. The marginal debt cost is $c_b(r) = \beta_b r$; the required return on equity is π; new lending earns the market spread s over the policy rate. Planned originations at the unconstrained optimum are n^*; the bank chooses originations $n \le n^*$ and retention or issuance $k \ge 0$, so loans end at $L = L_0 - (n^* - n)$ and equity at $K = K_0 + k$. Three assumptions carry the derivations, and they are stated as assumptions rather than buried in the algebra.

Assumption 1 (the equity premium is rate-insensitive). π is constant in r, with $\pi \ge \beta_b r$ over the observed range. The net carrying cost of a unit of equity—the premium paid minus the debt cost it displaces—is then $\phi_b(r) = \pi - \beta_b r \ge 0$, *decreasing* in r. Two remarks. First, this is the honest sign: swapping a debt liability for equity saves $\beta_b r$ and costs the premium, so tightening makes capital relatively cheaper to hold, fastest for wholesale-funded banks; the derivation does not, and must not, rest on a carrying cost of capital that rises with r. Second, if the premium instead moved one-for-one with debt costs, ϕ_b would be flat in r: the cost side of the retention margin would lose its state dependence, and activation would rest on Assumption 2 alone; the assumption that it does not is used, and flagged, here.

Assumption 2 (the cost of standing thin is the rate on the marginal liability). The bank values a discretionary buffer b^* above the scrutiny point $\rho + g(d)$. From τ onward, operating with buffer shortfall $u > 0$ costs $\kappa\, c_b(r)\, \psi(u)$ in expectation, with $\kappa > 0$ and ψ increasing, strictly convex,

and differentiable on $u > 0$, extended by $\psi(u) = 0$ for $u \leq 0$: a bank near its trigger point rolls its marginal funding under the shadow of supervisory engagement—intensified dialogue, capital-plan demands, and the market's reading of both—and the expected price of that fragility is proportional to the rate on the liability being rolled. At $r \approx 0$, with official funding abundant and deposits repricing not at all, the price is nil for every bank whatever its buffer; under tightening it rises with r in proportion to β_b. This is the framework's one substantive ingredient, and it is an assumption about supervision-cum-funding, not an accounting identity.

Assumption 3 (convex adjustment costs on both margins). Retention or issuance k carries a friction $\Gamma(k)$, strictly convex and differentiable with $\Gamma'(0) = 0$: the first euro is the marginal planned distribution; deeper cuts and fresh issuance signal. Cutting originations forfeits surplus $S(n^* - n)$, strictly convex and differentiable with $S'(0) = \varsigma > 0$ small: deals are ordered by the relationship rents they carry, so the first mandate declined is marginal new money earning only the competitive sliver ς, while renewals of maturing facilities carry franchise quasi-rents and sit deeper in the queue.

The guidance dents the desired buffer whenever $b^* > h - g(d)$; write $M(d) = \big(g(d) + b^* - h\big)L_0$ for the shortfall at inaction and maintain $M(d) > 0$ as the case of interest—a bank with $M(d) \leq 0$ never faces the trade-off and is an unconditional absorber, which only strengthens the absorption leg. Feasibility is the documented fact $h \geq g(d)$: the hard requirement is never at risk, and $M(d)$ is a discretionary shortfall, not a breach. Because the desired capital is $(\rho + g(d) + b^*)L$, a euro of loans shed releases only $\rho + g(d) + b^*$ euros of it—the unit conversion the comparison requires—so the shortfall after acting is $u = M(d) - k - (\rho + g(d) + b^*)(n^* - n)$. At τ the bank solves

$$\min_{k \geq 0,\; n \leq n^*} \quad \phi_b(r)\, k \;+\; \Gamma(k) \;+\; S(n^* - n) \;+\; \kappa\, c_b(r)\, \psi(u),$$

a convex program. Let $\eta \equiv \kappa\, c_b(r)\, \psi'(u)$ be the marginal value of buffer repair at the optimum; the first-order conditions are

$$\phi_b(r) + \Gamma'(k) \;\geq\; \eta, \qquad S'(n^* - n) \;\geq\; \eta\,(\rho + g(d) + b^*),$$

each with equality where its margin is active. Per euro of buffer repair, shedding loans costs $S'/(\rho + g(d) + b^*)$—the forfeited surplus scaled up by an order of magnitude, since $\rho + g(d) + b^*$ is on the order of one-eighth—while a euro retained repairs a euro at cost $\phi_b(r) + \Gamma'(k)$.

Inaction under accommodation. At $r \approx 0$, $c_b(r) \approx 0$ and hence $\eta \approx 0$, while both marginal costs stay bounded away from zero ($\phi_b(0) = \pi$; $S'(0) = \varsigma$): the optimum is $k = 0$, $n = n^*$ for every bank and every dose. The guidance is absorbed into the buffer, $\partial n/\partial d = 0$ and $\partial K/\partial d = 0$—Prediction 1's accommodation leg.

Activation under tightening. Evaluated at inaction, the contraction margin switches on when $\kappa\beta_b r\, \psi'(M(d))\,(\rho + g(d) + b^*) > \varsigma$ and the retention margin when $\kappa\beta_b r\, \psi'(M(d)) > \pi - \beta_b r$; the corresponding rates are

$$\bar{r}_b^{\,n} \;=\; \frac{\varsigma}{\kappa\,\beta_b\,\psi'\big(M(d)\big)\,(\rho + g(d) + b^*)}, \qquad \bar{r}_b^{\,K} \;=\; \frac{\pi}{\beta_b\big(1 + \kappa\,\psi'(M(d))\big)},$$

and the threshold of the text is $\bar{r}_b = \min\{\bar{r}_b^{\,n}, \bar{r}_b^{\,K}\}$. Both rates are decreasing in the funding

sensitivity β_b—the wholesale gradient—and in the dose d: the dose deepens the dent ($M'(d) = g'(d)L_0 > 0$, ψ' increasing), which lowers both rates, and raises the capital released per unit of lending forgone, which lowers $\bar{r}_b^n$ alone. For deposit-rich banks β_b is small and $\bar{r}_b$ lies beyond any observed rate: they remain absorbers in both states. Which margin activates first depends on parameters; the configuration the evidence selects—capital building and lending contracting together in 2023, neither margin moving in 2021—has the 2023 rate above both activation rates for the responding banks and the 2021 rate, effectively zero, below both.

The mix above the threshold. With both margins active the comparative statics are signed by convexity; because the guidance moves the conversion factor $\rho + g(d) + b^*$ by an amount second order relative to the dent $M(d)$, hold it fixed at R for the statics. Suppose d rises and the shadow value η did not: then k does not rise ($\Gamma'(k) = \eta - \phi_b(r)$) and neither does the cut ($S'(n^* - n) = \eta R$), while $M(d)$ rises strictly, so u rises, $\psi'(u)$ rises, and $\eta = \kappa\, c_b(r)\, \psi'(u)$ rises after all—a contradiction. Hence η increases in d, and with it both actions: above the threshold, capital *rises* with the dose and lending *falls* with it, a mix rather than a corner—the two-clocks capital build and the origination decline are one first-order condition read on its two margins. For a given shortfall, moreover, η scales with $\kappa\beta_b r$: the same dose produces a stronger response the higher the bank's funding sensitivity and the further the rate has risen (exactly so when the frictions and the scrutiny cost are quadratic)—the bank-level interaction of Prediction 1. The framework aggregates none of this to the firm level; the firm-level version of the interaction is an empirical question the paper reports separately.

Pricing. Write the shadow charge on the capital a new loan consumes as $\mu_b(d, r) \equiv \eta\,(\rho + g(d) + b^*)$, zero at $r \approx 0$, positive once the contraction margin activates, and increasing in d. A deal carrying relationship rent $\nu \geq 0$ is originated if its surplus covers the charge, $\varsigma + \nu \geq \mu_b(d, r)$; because the spread s is set market-wide, a pressured bank cannot recover $\mu_b(d, r)$ in price without losing the mandate, and declines instead. Origination probability therefore falls in d once $\mu_b(d, r) > \varsigma$, transacted spreads carry no lender-specific component, and marginal new money ($\nu \approx 0$) is rationed before renewals, whose franchise rents keep their reservation spread below the market's—Regularity 1 and the renewal margin of the within-firm test.

Timing. The scrutiny term exists only from the application date τ, and the two margins face it differently. Declining an origination forfeits surplus immediately while repairing the buffer against a charge that is not yet live, and originations are a flow the bank can cut at τ itself: rationing therefore concentrates once guidance applies, not at publication and not at the rate shock—Regularity 2. Retention is the opposite case: with a convex per-period friction, a bank that already knows its depletion spreads the build forward from the exercise itself rather than paying a steep one-off adjustment at τ—the anticipatory capital clock of the two-clocks evidence. The static program dates the penalty; this smoothing reading of the build's timing is an interpretation of the convexity in Γ, not a derived proposition.

Incidence. Firms invest $I = q(\sigma + x)$ out of internal slack σ and external finance x, with $q' > 0$, $q'' < 0$, and $x = x_0 - \theta\, E_f$, where E_f is the firm's exposure to its lenders' depletions—the model counterpart of the empirical E_f^e of Section 5—and $\theta \in (0, 1]$ reflects imperfect replacement of a rationed lead bank. Then $\mathrm{d}I/\mathrm{d}E_f = -\theta\, q' < 0$, largest where $\sigma + x$ is smallest, so the decline concentrates among low-slack firms; and $\partial^2 I/\partial E_f\, \partial\sigma = -\theta\, q''(\sigma + x) > 0$: the exposure effect

fades continuously as slack rises—a gradient in the level of σ, not a split at its median, so a binary split understates what the continuous measure carries. Pre-window cash holdings are the primary observable marker of σ in the empirical mapping, bond-market access a secondary correlate—Prediction 2.