



Master in International Finance

RESEARCH PAPER

Academic Year 2025-2026

Selective Repricing of European Syndicated Loans: Borrower Elasticity and Institutional Pricing after the 2022 ECB Tightening

André Correia & Diogo Pereira

Academic Advisor: Professor Johan Hombert

President of the Jury: Professor Augustin Landier

June 1st, 2026

Keywords: Syndicated loans, monetary policy transmission, loan spreads, ECB tightening cycle, private equity sponsorship, demand elasticity.

We thank our academic advisor Johan Hombert for guidance throughout the project.

Author contact:

André Correia, andre.correia@hec.edu, (+351) 924 356 919

Diogo Pereira, diogo.pereira@hec.edu, (+351) 967 716 766

Selective Repricing of European Syndicated Loans: Borrower Elasticity and Institutional Pricing after the 2022 ECB Tightening

André Correia and Diogo Pereira

Abstract

Does a sharp monetary tightening pass through uniformly to syndicated borrowers, or selectively to those least able to defer? Using 7,764 European syndicated tranches priced from 2015 to 2026, we test three structural margins around the 2022 to 2023 ECB cycle: two borrower-elasticity margins (use of proceeds, industry cyclicalities) and one institutional-channel margin (sponsor backing). A three-regime decomposition shows the non-discretionary and non-cyclical differentials widening by 65 and 30 basis points into the plateau, while the sponsor differential spikes 66 basis points during the tightening and then fully reverts, a path a conventional pre-versus-post specification averages into a near-zero coefficient.

Table of Contents

1. Introduction	1
2. Literature Review	2
2.1. The cost of funding and pass-through speed in bank credit	3
2.2. Syndicated-loan spreads as equilibrium prices.....	4
2.3. Sponsor backing, originate-to-distribute and the institutional buyer	5
2.4. Borrower elasticity and the cross-section of outside options	6
3. Data and Institutional Background	7
3.1. The European syndicated loan market.....	7
3.2. The ECB tightening cycle and regime definitions.....	8
3.3. Sample construction and cleaning pipeline.....	9
3.4. Variable construction and classifications.....	9
4. Hypotheses and Empirical Strategy.....	11
4.1. H1. Defensive versus discretionary use of proceeds.....	11
4.2. H2. Non-cyclical versus cyclical industries	12
4.3. H3. Sponsor-backed versus non-sponsored loans.....	13
4.4. Baseline specification.....	14
4.5. Measuring issuance volumes	15
4.6. Standard errors and inference	16
4.7. Identification and endogeneity.....	16
5. Results	17
5.1. Use of proceeds and the inelasticity of defensive borrowers	18
5.1.1. Spread regression	18
5.1.2. Volume evidence: H1 secondary prediction.....	19
5.2. Industry cyclicalities and regulated capital expenditure.....	19
5.2.1. Spread regression with cycle control	20
5.2.2. Volume evidence: H2 secondary prediction.....	21
5.3. Sponsor backing and the institutional channel	21
5.4. Structural premium across regimes	23
5.5. Robustness checks.....	23
5.5.1. Channel orthogonality: pairwise and joint specifications.....	24
5.5.2. Yield-type subsamples for H1	24
5.5.3. Placebo tests on a 2018 to 2021 sample.....	25
5.5.4. Clustered standard errors on the borrower	26
5.5.5. Country fixed effects	27
5.5.6. Joint Wald F-tests across inference regimes.....	28
6. Conclusion.....	28
References.....	31
Tables.....	34
Figures	49
Appendix A. Use of Proceeds Classification	56
Appendix B. Industry Cyclicalities Classification	58

1. Introduction

Between July 2022 and September 2023, the European Central Bank raised the deposit facility rate by 450 basis points across ten consecutive meetings (ECB, 2024), the steepest tightening in its history. The shock was unusual along a second dimension that we exploit throughout the paper: it occurred during an expansion, so the cross-section of spread responses reflects financing-cost variation rather than credit-quality deterioration. The question we address is whether this rate shock was passed through to borrowers uniformly, or whether it concentrated on borrowers whose demand for credit was least elastic at the moment the shock hit.

That question matters because the average conceals the distribution. The market-wide widening of European syndicated spreads in our sample was roughly forty basis points above the pre-shock mean of 285. That average masks a cross-section ranging from near-uniform pass-through to structural premia above one hundred basis points, and it is the shape of this dispersion, not its mean, that determines who actually bears a rate shock. Existing work documents heterogeneity at the bank level (Drechsler, Savov and Schnabl, 2017; Heider, Saidi and Schepens, 2019), and the supervisory record confirms that pass-through speeds varied across borrower segments (ECB, 2023b; IMF, 2023). The borrower-side, tranche-level cross-section, however, has not been mapped.

We map it on 7,764 European syndicated tranches priced between January 2015 and February 2026, drawn from LSEG SDC Platinum and cleaned through the pipeline of Section 3. The dependent variable is the all-in drawn spread in basis points over the reference rate. The variation we exploit is the cross-section of structural features along three margins: the tranche's use of proceeds, the cyclicity of the borrower's industry, and whether the borrower is backed by a private-equity sponsor. The first two map borrower-side features to predicted demand elasticity; the third maps an institutional-channel feature to predicted pricing-technology speed. Each yields a testable cross-sectional prediction for the post-shock spread response, and the two elasticity channels carry a parallel prediction on issuance volume.

Two of the three hypotheses formalize this elasticity logic at different levels of aggregation. H1 operates at the tranche level, contrasting defensive refinancing and working capital against discretionary acquisitions, buyouts and capital expenditure; H2 operates at the industry level, contrasting regulator-mandated capex in utilities, healthcare and transport infrastructure against deferrable consumer-discretionary, technology and project-based spending. For both, inelastic borrowers should pay more in spread and withdraw less from the market during the tightening, so we test the two margins jointly. H3 tests a distinct mechanism: sponsor-backed leveraged loans are distributed to CLOs and credit funds (Bord and Santos, 2012; Ivashina and Sun, 2011) whose required yields reset on a sub-quarterly horizon, while non-sponsored loans reprice with a lag

through bank balance sheets. Because sponsor backing correlates with discretionary use of proceeds, we verify the three channels operate orthogonally rather than as relabelings of a single selection. We test all three with a three-regime decomposition, splitting the window into a pre-shock regime, the tightening phase, and the plateau (boundaries in Section 3.2), interacting each treatment indicator with the two post-shock regimes; H2 adds the euro-area unemployment rate as a cycle control.

The cross-section repriced selectively, along the two predicted signatures. Relative to their pre-shock differential, defensive tranches repriced 38 basis points higher than discretionary during the tightening and 65 higher in the plateau: the differential widened rather than reverted as the environment stabilized, which is the slow demand-elasticity signature predicted by H1. The non-cyclical differential rose by 42 basis points in the tightening and 30 in the plateau, reversing the conventional credit-cyclicality reading that the non-recessionary cycle lets us identify cleanly. The sponsor differential, by contrast, widened 66 basis points during the tightening and then gave the entire spike back, the signature of a fast institutional channel meeting a slow relationship one. Asymmetric contractions in primary issuance volume corroborate the elasticity reading on a second margin: discretionary notional volume fell by 29 percent against 21 for defensive (Table 5), and cyclical by 18 against 1 for non-cyclical (Table 7).

A joint specification confirms the three channels identify independent margins rather than substitute labels for the same selection (Table 10). The H3 result also carries a methodological point: a binary specification on the sponsor margin returns -4 basis points, indistinguishable from zero, averaging the tightening spike and plateau reversion into a misleading null that the three-regime decomposition recovers. The headline results survive five robustness checks (Section 5.5) and underpin the conceptual, empirical and methodological contributions developed in Section 6.

The rest of the paper proceeds as follows. Section 2 reviews the literature. Section 3 describes the data and institutional setting. Section 4 sets out the hypotheses and empirical specification. Section 5 reports the results and robustness checks. Section 6 concludes.

2. Literature Review

Four bodies of work bear on the question, and together they motivate both our hypotheses and our specification choice. The first describes how policy rates propagate through bank balance sheets and at what speed; the second treats syndicated-loan spreads as equilibrium prices that contain bargaining rents on top of credit risk; the third documents the institutional buyer base for sponsor-backed leveraged loans and the originate-to-distribute model; the fourth maps borrower outside options to the price of credit. Two themes run through them: that fast and slow pricing channels coexist, which grounds the regime

decomposition, and that spreads embed bargaining rents sensitive to borrower elasticity, which grounds H1 and H2. We draw throughout on academic work, on official-sector publications by the ECB, IMF and Bank of England, and on industry coverage by PitchBook LCD and the financial press.

2.1. The cost of funding and pass-through speed in bank credit

This subsection traces how policy rates reach loan pricing through bank funding costs and the speed at which they do so, drawing on both the academic literature and the supervisory record from the 2022 to 2023 cycle.

The credit view of monetary policy (Bernanke and Blinder, 1988) distinguishes the rate channel from the credit channel. The latter is operative because banks face balance-sheet and funding constraints that enter the markup over the reference rate. For a loan priced as a spread above its reference rate (EURIBOR for most euro-denominated tranches, with SOFR or other benchmarks for foreign-currency facilities), the spread is the contractual object that absorbs the shock: the reference rate passes through mechanically, while the spread reflects everything else.

The spread is slow to absorb that shock on the bank-funded side for a structural reason. Drechsler, Savov and Schnabl (2017) show that banks benefit from a "deposits channel": when policy rates rise, deposit costs adjust only gradually, so effective funding costs do not increase one-for-one with policy rates. Banks earn higher margins on existing deposit balances and pass the funding-cost shock through to clients only as legacy facilities reset or new facilities are negotiated. The 2022 to 2023 cycle began from a deposit facility rate of -0.50%, an unusual starting point. Heider, Saidi and Schepens (2019) document that in this negative-rate environment, banks reluctant to charge depositors negative rates absorbed compressed margins instead, with consequences for credit supply. Bank funding costs therefore entered the 2022 to 2023 tightening from an unusual baseline and adjusted with a lag. Auer and Conti (2024) reach a parallel conclusion at the macroeconomic level: euro-area lending to non-financial corporates contracted by more than historical regularities would predict during this cycle, and credit-supply indicators from the Bank Lending Survey account for most of the unexplained gap. The bank lending channel was therefore unusually important in this cycle relative to previous tightening episodes, consistent with the deposit-channel lag described above.

The mechanism is summarized in the deposit beta: the share of a policy-rate change that passes through to deposit rates, documented to be materially below one in the early phase of a tightening (Drechsler, Savov and Schnabl, 2017). In this cycle, rates on newly agreed time deposits rose by 372 basis points and overnight-deposit rates by only 50, against the 450-basis-point policy increase (ECB, 2023c). A low deposit beta translates mechanically into slow-adjusting funding costs and, through the bank-funded margin, slow-

adjusting loan spreads. This is the slow side of the institutional channel tested in H3: relationship-bank pricing lags because its funding cost lags, while institutional pricing (Section 2.3) clears against market-paced CLO and credit-fund yields. The asymmetry between the two speeds motivates the three-regime decomposition, which a binary specification would average over while one side is still mid-adjustment.

The supervisory record from the cycle itself corroborates this academic story. The ECB Bank Lending Survey reported successive net tightening balances on credit standards for loans to non-financial corporates from the third quarter of 2022 onwards, a pace of net tightening that the survey described as the highest level since the euro-area sovereign debt crisis in 2011 (ECB, 2023a). The contemporaneous ECB Economic Bulletin documented a forceful transmission of the policy-rate increases to euro-area financing and monetary conditions (ECB, 2023b), and the IMF Global Financial Stability Report documented a sharp differentiation in funding costs between investment-grade and riskier issuers (IMF, 2023), which the ECB documented as consistent with euro-area conditions (ECB, 2023d).

Both the academic literature and the supervisory record therefore point the same way: bank-funded pass-through is slow and lagged, adjusting only as funding costs and legacy facilities reset. This is the slow side of the channel comparison that H3 sets against the fast institutional side documented next in Section 2.3.

2.2. Syndicated-loan spreads as equilibrium prices

If the speed of pass-through is one part of the answer, the object that absorbs it is the other: this subsection establishes that the all-in drawn spread is not a pure default-risk premium but an equilibrium price embedding bargaining rents on top of credit risk.

Sufi (2007) shows that the architecture of the syndicate is itself endogenous to information asymmetry: lead arrangers retain larger shares and concentrate participations more tightly when borrowers are opaque, evidence that monitoring effort and information rents are priced into the contract. The pricing of bank debt also embeds the value of the lending relationship. Schwert (2020) documents that bank loans carry a material price premium over comparable-risk public debt, evidence that liquidity provision, monitoring and the implicit option to renegotiate are priced features of the contract rather than free attributes.

These rents become more valuable in tight states. Chodorow-Reich and Falato (2022) demonstrate that covenant violations transfer effective control rights to lenders and that lenders in weaker health contract credit upon receiving them, evidence that even long-dated facilities behave as short-horizon contracts when the environment moves. The same state-contingency applies before any violation is realized: at origination, lenders price the option to renegotiate, and the value of that option depends on the borrower's outside options at the moment a renegotiation might be needed.

The implication for the present paper is direct. If the all-in drawn spread were a pure default-risk premium, a sharp rate shock would move spreads uniformly, modulated only by changes in default expectations. To the extent that the spread embeds bargaining rents, the same shock should be passed through more aggressively where lenders sit on stronger relative bargaining positions. This is the reduced-form prediction tested in H1 and H2.

2.3. Sponsor backing, originate-to-distribute and the institutional buyer

The bargaining channel of subsection 2.2 operates across the syndicated market, but a second and faster pricing channel runs in parallel through one of its segments; this subsection documents the institutional buyer base for sponsor-backed leveraged loans and the originate-to-distribute model that connects it to bank origination.

Sponsor-backed transactions occupy a distinct corner of the syndicated market: banks originate but distribute to non-bank investors, principally collateralized loan obligations (CLOs) and credit funds. Bord and Santos (2012) document the rise of the originate-to-distribute model in corporate lending and the structural separation it creates between origination and ultimate risk-bearing. That separation is the institutional substrate of H3: sponsor-backed term loans clear against the institutional bid, which marks required yields to market on a sub-quarterly horizon, while non-sponsored loans clear through bank balance sheets that lag.

Before any post-shock dynamic can be detected, that institutional bid operates as a price-compressing factor in calm markets. Demiroglu and James (2010) find that reputable private-equity groups obtain tighter spreads, longer maturities and higher leverage, and Ivashina and Kovner (2011) show that repeated bank-sponsor relationships translate into lower spreads and looser covenants. This calm-markets sponsor discount is what H3 must overcome before any post-shock spike can be identified.

Two mechanisms explain why that discount inverts during tightening. First, Ivashina and Sun (2011) document that institutional demand pressure in leveraged loans compresses primary spreads when issuance clears quickly; by the same logic, the channel widens spreads quickly when institutional buyers retreat or demand higher compensation for floating-rate exposure. Second, Becker and Ivashina (2014) show that bond-eligible firms substitute from bank loans into public bonds when monetary policy tightens or lending standards harden, evidence that the supply curve for bank credit is state-dependent on the access-to-market margin. Sponsor-backed leveraged borrowers typically lack that bond-market margin: they cannot avoid the rate shock by issuing in another market, so the shock must clear in primary loan spreads. This is the rate-shock setting in which the institutional bid dominates pricing on the sponsor side.

Industry coverage of the cycle corroborates this mechanism. PitchBook LCD reported a sharp contraction in European CLO new-issue volume and a widening of primary leveraged-loan spreads from late 2022 onwards, with the institutional bid effectively pausing during the most acute phase of the tightening before resuming at materially wider clearing levels in 2023 (PitchBook LCD, 2022, 2023). The financial press tracked the same dynamic, attributing the roughly \$42 billion of hung buyout debt on US and European bank balance sheets in late 2022 to a CLO bid that had retreated faster than bank balance sheets repriced (Bloomberg, 2022). The Bank of England (2023) documented a widening of leveraged-loan spreads and subdued primary issuance, alongside a rise of over 400 basis points in the average stock rate on floating-rate corporate bank loans between January 2022 and May 2023.

The combination yields the H3 prediction directly: the sponsor segment should reprice quickly during the tightening because the institutional buyer base demands it, then compress partially in the plateau as the bank-relationship channel catches up. Whether a binary specification masks this two-regime path is the empirical question taken up in Section 5.3.

2.4. Borrower elasticity and the cross-section of outside options

The bargaining and institutional channels described so far both depend on what the borrower can do instead of accepting the offered terms; this subsection maps the literature on borrower outside options to the use-of-proceeds and industry-cyclicality margins tested in H1 and H2.

Petersen and Rajan (1994) show that the primary benefit of a close lending relationship is greater availability of credit; the attachment it creates is nonetheless a source of lender bargaining power when a borrower's outside options narrow. Schwert (2018) makes the parallel point for corporate borrowers, who sort durably into better-capitalized lenders and pay the resulting relationship rent when their options thin. Diamond (1991) supplies the model: the choice between bank loans and directly placed debt turns on borrower reputation and monitoring needs, so borrowers with thinner market alternatives remain bank-dependent and more exposed to lender pricing power.

For the present paper the relevant outside option is wider than bond-market access. What matters at the tranche level is whether the borrower can defer the transaction, downsize the funding need, draw on internal liquidity or stay away from the market. A refinancing or working-capital tranche has a fixed cash-flow calendar and limited deferral capacity; an acquisition, buyout or capital-expenditure tranche has discretion over timing, size and execution. The same logic carries to industries: regulator-mandated capex in utilities, healthcare and transport infrastructure binds investment regardless of financing cost, while consumer-discretionary, technology and project-based capex can be deferred. This deferability of investment, rather

than the cyclicalities of earnings, is the property our two hand-coded classifications are built to capture (Section 3.4), and the prediction is therefore an elasticity prediction rather than a cyclicalities prediction. The 2022 to 2023 episode, a rate-led tightening during an expansion as Section 3.2 documents, supplies the identification environment in which the two can be separated.

3. Data and Institutional Background

Having established what the literature predicts, we now describe the data used to test those predictions. Subsection 3.1 sets out the structural features of the European syndicated loan market, 3.2 defines the ECB tightening cycle and the three regimes, 3.3 documents the sample-construction pipeline from the raw 28,411-tranche extract to the final 7,764-tranche sample, and 3.4 describes the variables and classifications.

3.1. The European syndicated loan market

The European syndicated loan market operates across two segments that share a common contractual envelope but differ in their pricing technology and in the identity of the marginal investor. The first is the relationship-bank market, which intermediates investment-grade revolvers and term loans for non-financial corporates. Pricing here is bilateral, anchored to the lead arranger's internal funding cost and to the long-run profitability of the relationship; it absorbs short-run swings in market conditions through internal margins and passes them to clients only with a lag, as legacy facilities reset or new facilities are negotiated. The second is the institutional leveraged-loan market, which intermediates broadly syndicated term loan B paper for sponsor-backed and other highly levered borrowers. Pricing here is quasi-public: collateralized loan obligations (CLOs) and credit funds set required yields each day by reference to a mark-to-market on existing positions, and primary spreads adjust to clear new issues against the prevailing secondary market (Bord and Santos, 2012). The two segments overlap on the lender side, since the same banks lead both books, but they are governed by very different marginal-investor dynamics.

A third segment, the European private-credit (direct-lending) market intermediated by non-bank credit funds, has grown rapidly since 2020 (Alternative Credit Council, 2025; European Central Bank Banking Supervision, 2024) and provides a partial substitute for syndicated lending at the leveraged end. Primary-market data on this segment is sparse and inconsistent across providers, and we exclude it from the analysis; the rapid expansion of private credit during our plateau regime may attenuate or amplify our sponsor-channel findings to the extent that marginal sponsor borrowers shifted between syndicated and private-credit financing during 2024 to 2025. We return to this scope limitation in Section 6.

This segmentation is central to our identification strategy: because the two segments adjust at different speeds, the regime in which a spread is measured matters, which is what the three-regime decomposition of Section 4 is built to exploit. Beyond the segmentation, the European market shares the standard syndicated-loan template documented in Sufi (2007): a lead arranger that retains a portion of the loan and exercises monitoring authority on behalf of participants, an investment-grade pool dominated by revolvers, and a leveraged pool dominated by term loan B paper.

3.2. The ECB tightening cycle and regime definitions

Between the announcement of the first hike on 21 July 2022 (effective 27 July 2022) and the announcement of the final hike on 14 September 2023 (effective 20 September 2023), the European Central Bank delivered ten consecutive policy-rate increases, lifting the deposit facility rate from -0.50% to 4.00% (ECB, 2024). The cumulative 450-basis-point increase was the steepest in the institution's history and the most rapid sustained tightening in any major advanced economy in more than two decades. This episode is central to our identification. It was a rate-led tightening rather than a recessionary credit cycle: euro-area output continued to expand and corporate defaults stayed at historically low levels, even as bank lending to non-financial corporates contracted by more than historical regularities predict (Auer and Conti, 2024). The cross-section of spread responses can therefore be read primarily as a response to financing-cost variation rather than to a contemporaneous downturn, the property we exploit in H2, where the prediction inverts the conventional credit-cyclical reading.

We split the sample window into three regimes that mirror this institutional history. The pre-shock regime spans 1 January 2015 to 20 July 2022, the interval during which the deposit facility rate was at zero or negative. The tightening regime spans 21 July 2022 to 25 October 2023, the rapid rate-hike phase. The plateau regime spans 26 October 2023 to 2 February 2026: the Governing Council paused at its meeting of 26 October 2023, held the deposit rate at 4.00% through the spring of 2024, and began a measured descent in mid-2024 that continued through the end of our sample window.

Loans are assigned to a regime by financial close date, on the principle that the close date is when the spread is committed and printed in the primary market; the regime boundaries follow announcement rather than effective dates because loan-market participants reprice on the announcement. The three-regime decomposition is our preferred form throughout, with the binary pre-versus-post specification reported as a comparison column. Figure 1 plots the monthly mean spread of the clean sample against the ECB deposit facility rate with the three regimes shaded: the two series rise together through the tightening phase, and in the plateau the deposit rate descends from its 4 percent peak to about 2 percent while the mean spread drifts down by around 40 bps from its early-plateau high near 350 bps.

3.3. Sample construction and cleaning pipeline

We extract loan-level data from LSEG SDC Platinum on 2 February 2026 (LSEG, 2026), restricting attention to syndicated tranches issued by borrowers domiciled in any Western or Eastern European jurisdiction with announcement dates from 1 January 2015 onwards and all amounts converted to US dollars. The raw extract returns 28,411 tranches. The first cleaning step drops 920 tranches whose financial close status is reported as “not closed” in SDC: these correspond to deals announced but never consummated, and including them would distort both the priced-loan sample and the volume statistics of Tables 5 and 7.

The remaining 27,491 tranches form the full dataset that anchors the volume analyses of Sections 5.1.2 and 5.2.2. For the spread regressions, five further filters are applied (Steps 2 to 6 of Table 1). The first and largest drops tranches with missing or "no data" all-in drawn spread (-19,550 observations); this high attrition is a well-known feature of SDC Platinum, where pricing is reported only for the subset of tranches the lead arranger discloses.

We do not impose an additional support filter on the use-of-proceeds or industry classification at this stage: the H1 sample restricts to defensive or discretionary tranches (3,746), the H2 sample to cyclical or non-cyclical (5,247), while the neutral buckets are excluded from the corresponding hypothesis test but retained in the joint specification of Section 5.5.1 and in the H3 sample (the full 7,764 tranches). Table 1 reports the cleaning pipeline.

3.4. Variable construction and classifications

The dependent variable is the all-in drawn spread reported by SDC in basis points over the relevant reference rate. This variable measures the per-period compensation that the borrower commits to pay above the reference rate for the duration of the loan, including the drawn margin and any recurring fees that are amortised into the spread. It does not include one-off arrangement fees, which are reported separately and reflect syndicate composition more than underlying credit risk.

Three structural classifications enter the right-hand side of the regressions. Two of them, use of proceeds and industry cyclical, are hand-coded against a single organizing test: whether the borrower can defer, downsize, or walk away from the spending the credit funds. The test is one of demand inelasticity, not of credit risk or earnings volatility, and because that distinction is the least conventional choice in the paper we set it out in full here before applying it; Appendices A and B give the complete taxonomies and, in B.1, the defense of the hand-coded approach against a data-driven alternative.

The first classification is the use-of-proceeds classification, derived from the SDC use-of-proceeds field through the atomic taxonomy in Appendix A. It applies the deferability test at the tranche level: because the SDC field describes the funded transaction directly, the classification reads the deferability property straight off the recorded purpose. Defensive uses (maturing principal, working capital, restructuring, trade flows in motion, standby liquidity backing outstanding commercial paper) are tied to a fixed cash-flow calendar that cannot be postponed without triggering distress; discretionary uses (acquisitions, buyouts, project finance, dividend recapitalizations, share repurchases) retain timing flexibility. The defensive bucket is therefore the non-discretionary category in economic substance; we use the term "defensive" throughout to emphasize that the financing protects an existing obligation or operational need rather than funding a new, timing-flexible decision. Because many tranches carry multiple use-of-proceeds tags, the bucket is assigned by an atomic combination rule: a tranche is defensive only if every informative atom is defensive, discretionary only if every informative atom is discretionary, and neutral otherwise, with the informationally empty tags (general corporate purpose, other) treated as silent so that a "general corporate purpose; working capital" tranche reads as defensive rather than being diluted to neutral. Appendix A sets out the atomic taxonomy, the combination rule, and the defense of both against the borderline cases they must resolve.

The second classification is the industry-cyclical classification, derived from the TRBC industry code via Appendix B. It applies the same deferability test one level up, at the industry rather than the tranche: non-cyclical industries (26) face capital-expenditure and working-capital schedules fixed by regulation, contractual replacement cycles or non-discretionary demand; cyclical industries (76) can postpone investment, run down inventory or rephase project finance; neutral (41) collects industries whose investment deferability is genuinely ambiguous. The property we are after is the deferability of investment, not the cyclical of earnings, and the two diverge: a regulated water utility and a tobacco company have similarly low earnings cyclical, but only the utility carries a regulator-mandated capex schedule that binds investment regardless of financing cost, so only the utility belongs in the non-cyclical bucket on our test. This is also why we hand-code the classification rather than sorting industries by their beta to GDP, which would treat the two as equivalent and mismeasure the inelasticity concept H2 is built to isolate (Appendix B.1). As with "defensive," the operative property is inelasticity of investment demand rather than insensitivity to the business cycle; in both labels we follow industry convention while testing the elasticity property underneath.

The third classification is the sponsor flag, constructed from the SDC sponsor field, equal to one whenever a private-equity sponsor is recorded as backing the borrower on the tranche and zero otherwise.

Three controls enter every spread regression. Maturity is measured in years from closing to final maturity. Log tranche size is the natural logarithm of the principal amount in USD millions. Loan yield type is categorical with three levels: Investment Grade (omitted), Leveraged and Highly Leveraged. The H2 preferred specification additionally includes the contemporaneous euro-area unemployment rate, centered on the in-sample mean, as a cycle control (UnempC; source Eurostat series une_rt_m). Two further fields support robustness checks: the Issuer/Borrower PermID for borrower-clustered standard errors, and the Domicile Nation for country fixed effects (both in Section 5.5). The main specifications use HC1 robust standard errors. Table 2 reports the spread moments by regime; Table 3 reports the within-regime composition of the issuance pool.

4. Hypotheses and Empirical Strategy

This section sets out the three hypotheses that organize the empirical analysis and the specifications used to test them. Subsections 4.1 and 4.2 derive cross-sectional predictions for two borrower-elasticity margins (use of proceeds and industry cyclicality); subsection 4.3 introduces a complementary institutional-channel test (sponsor backing) and explains why it belongs in a paper centered on borrower elasticity. Subsection 4.4 presents the spread regression equations and 4.5 sets out how issuance volume is measured for the secondary prediction the two elasticity channels share. Subsections 4.6 and 4.7 then discuss inference and identification.

The hypotheses test two distinct mechanisms in parallel. H1 and H2 run through borrower-side outside options; H3 runs through institutional pricing technology. The two generate opposite regime signatures, persistence versus reversal, that a binary specification collapses into a single coefficient and cannot tell apart, which is why the three-regime decomposition is our preferred form. The elasticity channels additionally carry a secondary prediction on issuance volume (Sections 4.5, 5.1 and 5.2), and Section 5.5 tests whether the three channels operate orthogonally.

4.1. H1. Defensive versus discretionary use of proceeds

The first hypothesis concerns demand elasticity at the tranche level, applying the deferability test of Section 3.4: defensive tranches finance obligations on a fixed cash-flow calendar, while discretionary tranches fund transactions whose timing the borrower can postpone, rescale, or abandon. When the rate environment tightens, lenders extract larger rents from the inelastic side of the market, because the borrower cannot postpone its market visit.

The bargaining literature (Schwert, 2020) and the bank-relationship literature (Petersen and Rajan, 1994; Schwert, 2018) imply that the structural defensive discount visible in the pre-shock baseline should compress, or reverse, in regimes where rollover need binds. We therefore predict the defensive-by-tightening interaction to be positive, and the defensive-by-plateau interaction to be positive and to persist into the plateau because the rollover schedule does not relax once policy stabilizes. The H1 null is no differential repricing across uses of proceeds, which would obtain if the bargaining channel is empirically dominated by a uniform repricing of credit risk premia across all use-of-proceeds categories. The same elasticity logic delivers a parallel secondary prediction on primary issuance volume: during the tightening, discretionary borrowers should withdraw from the market at a higher rate than defensive borrowers; this prediction is tested directly alongside the spread regression in Section 5.1.2.

4.2. H2. Non-cyclical versus cyclical industries

The second hypothesis applies the same deferability test at the industry level, under the classification set out in Section 3.4 and detailed in Appendix B: non-cyclical borrowers face investment schedules pre-committed under regulation, contractual replacement cycles or non-discretionary demand, while cyclical borrowers retain the option to defer capex, run down inventory, or rephase project finance. The bargaining logic invoked in H1 carries over (Petersen and Rajan, 1994; Diamond, 1991): where outside options are thin, lenders extract larger rents in tight states. We therefore predict the non-cyclical-by-tightening and non-cyclical-by-plateau interactions to be positive and to persist into the plateau, and the same elasticity asymmetry to appear on the volume margin (cyclical issuance contracting more than non-cyclical during the tightening).

The prediction runs against the conventional credit-cyclicality reading, under which cyclical borrowers would be expected to pay more in a downturn because their cash flows are more cycle-exposed. The two readings diverge precisely because our classification sorts industries by the deferability of investment rather than by earnings cyclicality (Section 3.4): the borrowers our test labels non-cyclical are those whose investment binds regardless of the rate environment, not those whose earnings are least cycle-exposed. As Section 3.2 documents, the 2022 to 2023 episode was a rate-led tightening rather than a recessionary credit cycle, with euro-area output continuing to grow and corporate defaults remaining at historically low levels. The cross-section of spread responses therefore reflects financing inelasticity more than business-cycle exposure, and the H2 prediction is testable precisely because the rate shock is not contaminated by a contemporaneous downturn. To discipline this reading against any residual labour-market variation, the preferred H2 specification includes the contemporaneous euro-area unemployment rate as a cycle control

(UnempC), centered on the in-sample mean. The non-cyclical-by-regime interactions are essentially unchanged when UnempC is added (Section 5.2.1).

4.3. H3. Sponsor-backed versus non-sponsored loans

The third hypothesis exploits a different mechanism from H1 and H2: the institutional segmentation of the European loan market described in Section 3.1. Sponsor-backed transactions are concentrated in the broadly syndicated leveraged-loan segment, where CLOs and credit funds dominate the buyer base (Bord and Santos, 2012). Institutional required yields are set daily by reference to the secondary market and feed back into primary spreads on a sub-quarterly horizon (Ivashina and Sun, 2011), while non-sponsored corporate loans reprice with a lag through the bilateral bank channel as funding-cost shocks pass through only when existing facilities reset.

These two pricing technologies generate distinct dynamics in response to a sharp rate shock. During the tightening, institutional buyers must reset required yields rapidly to clear new issues, generating a large positive sponsor-by-tightening interaction that reflects the speed differential between the channels rather than a permanent re-rating of sponsor credit. During the plateau, the slow relationship channel completes its pass-through as bank-funded margins on non-sponsored loans catch up with the higher cost of funds, compressing the cross-segment differential from its tightening peak. We therefore predict a positive and large sponsor-by-tightening interaction and a small, possibly negative, sponsor-by-plateau interaction. A negative plateau coefficient would not mean sponsor loans had become structurally cheaper; it would mean only that the relationship channel had not yet caught up when the tightening differential was measured. This is a pricing-technology speed differential, not an underlying inelasticity gap, and it is the signature that distinguishes H3 from H1 and H2.

A natural concern is that the three structural margins are correlated in the cross-section: sponsor-backed deals are typically discretionary in use of proceeds, and the H2 industry classification interacts with both. If the channels were relabelings of a single underlying selection, the channel $\times$ regime interactions should attenuate toward zero once the other channels are stacked into the regression. Section 5.5.1 tests this directly with a pairwise-and-joint orthogonality specification: if the channels survive each other intact, they identify three distinct pricing margins rather than three labels for the same one. The H3 null is uniform repricing across sponsor status; the orthogonality null is that any one channel subsumes another.

4.4. Baseline specification

We estimate up to four specifications by ordinary least squares with HC1 heteroscedasticity-robust standard errors. Equation (3), the three-regime decomposition with controls, is our preferred form for H1 and H3; for H2 it is augmented with the unemployment cycle control as documented in Section 4.2. The binary specification of equation (1) and the no-control three-regime form of equation (2) are reported alongside as comparison columns; equation (4), the full joint orthogonality form, is reported in Section 5.5.1.

Equation (1) is the binary specification. The post-shock indicator *HighRate* equals one for tranches issued from 21 July 2022 onwards and zero before. The structural indicator *S* equals one for the treated category in each hypothesis (defensive in H1, non-cyclical in H2, sponsor in H3) and zero for the comparison category. The vector *X* collects the controls: maturity, log tranche size, and the leveraged and highly leveraged dummies, with investment grade as the omitted yield-type bucket.

$$Spread_i = \alpha + \beta_1 S_i + \beta_2 HighRate_i + \beta_3 (S_i \times HighRate_i) + X_i' \varphi + \varepsilon_i \quad (1)$$

Equation (2) is the three-regime decomposition without controls. The regime indicators *Tight* and *Plat* take the value one for tranches issued during the tightening regime and during the plateau regime respectively, with the pre-shock regime as the omitted category. The interactions (*S* × *Tight*) and (*S* × *Plat*) are the coefficients of central interest: they measure the differential change in the spread for borrowers in the treated category in each regime, relative to the pre-shock baseline differential captured by β_1 . The control-free specification isolates the unconditional differential dynamics across regimes.

$$Spread_i = \alpha + \beta_1 S_i + \gamma_1 Tight_i + \gamma_2 Plat_i + \delta_1 (S_i \times Tight_i) + \delta_2 (S_i \times Plat_i) + \varepsilon_i \quad (2)$$

Equation (3) augments equation (2) with the control vector *X* and is the preferred specification for H1 and H3. The controls absorb the within-regime composition shifts in maturity, tranche size and yield-type bucket that would otherwise contaminate the structural interactions; the magnitude of those shifts, and the reason they matter, is documented in Section 4.7. For H2, the preferred specification additionally includes the unemployment cycle control $\psi \cdot UnempC$, centered on the in-sample mean; the non-cyclical-by-regime interactions are essentially unchanged by its addition (Section 5.2.1), so it disciplines the rate-shock reading without altering the substantive result.

$$Spread_i = \alpha + \beta_1 S_i + \gamma_1 Tight_i + \gamma_2 Plat_i + \delta_1 (S_i \times Tight_i) + \delta_2 (S_i \times Plat_i) + X_i' \varphi + \varepsilon_i \quad (3)$$

Equation (4) is the full joint orthogonality specification reported in Section 5.5.1. It stacks the three channels in a single regression on the full sample of 7,764 tranches, with the unemployment cycle control retained because the H2 channel is present. If the channel × regime interactions are essentially unchanged when the

other channels are introduced, the three margins operate orthogonally rather than as substitute labels for the same underlying selection. Section 5.5.1 reports four variants: three pairwise combinations and one full joint.

$$Spread_i = \alpha + \beta_D Defensive_i + \beta_N NonCyc_i + \beta_S Sponsor_i + \gamma_1 Tight + \gamma_2 Plat + \sum \delta_{C,R} (Channel_C \times Regime_R) + \psi UnempC_i + X_i' \varphi + \varepsilon_i \quad (4)$$

4.5. Measuring issuance volumes

The spread regressions test how the rate shock repriced credit; the volume analysis tests a distinct prediction of the same elasticity logic, namely that inelastic borrowers also withdraw less from the market when the shock hits. This subsection sets out how we measure issuance volume, which differs from the spread measurement in both the sample and the unit of observation.

The volume analysis runs on the full 27,491-tranche dataset that remains after Step 1 of the cleaning pipeline (Table 1), before any spread-availability filter is applied. This is deliberate. The spread regressions condition on a disclosed all-in drawn spread, observed for fewer than thirty percent of European tranches, whereas issuance is recorded whether or not the spread is disclosed. Measuring volume on the full dataset therefore separates two distinct margins: whether a borrower came to market at all, and whether the spread it paid was disclosed. If the elasticity reading is correct, the more flexible borrowers should withdraw from issuance itself during the tightening, not merely from the disclosed subsample, which is the prediction the spread sample alone cannot test.

We measure volume as notional issuance in US dollars rather than as a count of tranches, so that a single large financing weighs more than a small one, consistent with the economic quantity of credit raised. Because the three regimes differ sharply in length, raw regime totals are not comparable: the pre-shock regime spans 90.58 months, the tightening regime 15.15 months, and the plateau regime 27.27 months. We therefore normalize to per-month notional issuance within each regime, dividing total notional by the number of months in the regime. This places the three regimes on a common monthly footing and prevents the long pre-shock window from mechanically dominating the comparison.

To make the cross-regime movement legible, we rebase the per-month figure to an index equal to 100 in the pre-shock regime for each borrower group. The index reports the proportional change in monthly issuance relative to the group's own pre-shock baseline, so the defensive and discretionary series (H1), and the non-cyclical and cyclical series (H2), are each read against their own starting point rather than against each other. The prediction is an asymmetry in the contraction: during the tightening, the elastic group (discretionary in H1, cyclical in H2) should contract more than the inelastic group, and both should recover in the plateau as deferred deals return. Tables 5 and 7 report the per-month notional figures and the rebased

index by regime; Figures 3 and 4 plot the same series at quarterly frequency. We read the volume evidence as corroborating rather than identifying: it is consistent with the elasticity mechanism and addresses the disclosure-selection concern, but it does not by itself rule out every alternative account of the spread differential.

4.6. Standard errors and inference

All standard errors in the main specifications use the HC1 heteroscedasticity-robust correction. Because the dependent variable is bounded after winsorising and dispersion is far larger in the leveraged segment, we cannot rule out heteroscedasticity along both the categorical and continuous dimensions; HC1 adjusts the residual variance estimator with a small-sample factor of $n/(n - k)$ and provides asymptotically valid inference under arbitrary heteroscedasticity. We prefer it to borrower-clustered errors in the headline specifications because the cross-section is highly dispersed: most of the 2,571 borrower clusters contribute only one or two tranches, so within-borrower correlation is limited and the efficiency cost of clustering would outweigh its benefit. We therefore treat clustering as a robustness check rather than the baseline, and re-estimate the headline specifications with one-way borrower-clustered standard errors in Section 5.5.4, where clustering inflates the standard errors but every channel-level conclusion survives. Two-way clustering on borrower and time is not feasible because the tightening regime contains only fifteen months, which leaves too few effective time clusters for valid asymptotic inference along the time axis. The SDC extract provides the stable Issuer/Borrower PermID used for the one-way cluster-robust check.

The sample sizes differ across hypotheses by construction. The H1 sample restricts to tranches whose use of proceeds is unambiguously defensive or discretionary (3,746 tranches). The H2 sample restricts to tranches whose industry is unambiguously cyclical or non-cyclical (5,247 tranches). The H3 sample uses the full clean spread set of 7,764 tranches, since the sponsor flag is defined on every tranche. The orthogonality specifications of Section 5.5.1 are estimated on the full clean spread sample, since neutral classifications enter as the omitted category. Because each restriction is identified by exclusion of an uninformative neutral bucket rather than by an active selection mechanism, we treat the three hypothesis samples as parallel cross-sections and report results in parallel rather than harmonizing them onto a single sample.

4.7. Identification and endogeneity

Three identification concerns warrant explicit discussion. The first is composition. The cross-section shifted markedly across regimes, with the leveraged and sponsor shares rising and the investment-grade share falling (Table 3), reflecting a withdrawal of bank-led non-sponsored lending alongside an expansion of

institutional leveraged paper. Without controls, post-shock interactions would partially capture these shifts. The yield-type controls in equation (3) absorb the within-bucket repricing required for an apples-to-apples comparison, and the regime decomposition isolates the time variation in the spread differential within each bucket.

The second concern is the selection of which tranches enter the sample at all. As Section 3.3 documents, the spread regressions condition on observable spreads, which are reported for fewer than 30% of European tranches in SDC. If the selection mechanism shifted across regimes in a way correlated with the structural classifications, the cross-section of post-shock interactions could partially reflect a shift in disclosure rather than a shift in pricing. Sections 5.1.2 and 5.2.2 show that the more elastic borrowers withdrew from issuance, not just from disclosure, which is the elasticity prediction itself.

The third concern is that the post-shock interactions could reflect pre-existing trends in the cross-section of spreads rather than a response to the rate shock. If so, equivalent interactions should appear when an arbitrary pseudo-shock is imposed on a pre-shock subsample. Section 5.5.3 addresses this with placebo regressions on a 2018 to 2021 sample using a 21 July 2020 pseudo-shock. We read the placebo evidence as cleanly supportive for H2 and H3 and for the H1 plateau coefficient, and as supportive but not conclusive for the H1 tightening coefficient, with the residual ambiguity coming from the unusual macroeconomic environment of the COVID period itself.

A specification choice connected to these concerns deserves explicit mention: we do not include month-year fixed effects in the preferred specifications. The regime indicators are themselves the object of interest, and month-year fixed effects would absorb the common time variation that defines them. The trade-off is that our regime coefficients remain exposed to omitted time-varying macro confounders correlated with the structural margins; the unemployment cycle control of Section 4.2 mitigates this for H2 directly, and the placebo design of Section 5.5.3 mitigates it for all three hypotheses through a counterfactual time-shift.

5. Results

This section presents the regression results hypothesis by hypothesis. Sections 5.1 and 5.2 split each elasticity channel into two sub-subsections, reporting the spread regression first and the corroborating volume evidence second; Section 5.3 reports the sponsor channel; Section 5.4 reads the three channels together through the structural premium across regimes; Section 5.5 collects five robustness checks. The full coefficient blocks for each specification are gathered in Tables 4 to 15 at the back of the paper.

5.1. Use of proceeds and the inelasticity of defensive borrowers

This subsection tests H1 along its two predicted margins: the spread regression in Section 5.1.1 and the corroborating volume evidence in Section 5.1.2.

5.1.1. Spread regression

The H1 spread regression is reported in Table 4. The binary specification in column 1, which uses a single post-shock indicator, returns a defensive-by-post coefficient of +55.9 bps with HC1 standard error 10.1 ($p < 0.001$). Read in isolation, this coefficient already supports the elasticity prediction: defensive borrowers paid materially more than discretionary borrowers after the rate shock, holding controls constant. The pre-shock baseline gap is the structural defensive discount of -50.2 bps, statistically significant at the 1% level, which corresponds to the conventional credit reading that refinancing and working-capital paper signals lower risk than acquisition or buyout paper. The post-shock interaction more than offsets this discount and tips the average post-shock comparison so that defensive borrowers now pay more than discretionary borrowers.

The three-regime specification in column 3 unpacks the binary average into two separate movements. The defensive-by-tightening interaction is +38.2 bps with HC1 standard error 17.9 ($p = 0.032$), and the defensive-by-plateau interaction is +64.6 bps with HC1 standard error 11.3 ($p < 0.001$). Both are positive, and the plateau coefficient is the larger of the two. This is the slow-channel pattern predicted by H1: rather than building sharply during the rate-hike phase and unwinding once the cycle stabilized, the defensive premium widened further into the plateau. The mechanism operates on the discretionary side. Discretionary borrowers need not come to market in 2024 and 2025 unless the deal economics work, and a 4%-plus reference-rate environment shifts those economics against the marginal acquisition or buyout; the defensive side, by contrast, must come to market on its own schedule regardless. A joint Wald F-test of the null that both interactions equal zero rejects at the 1% level ($F(2, 3736) = 17.5, p < 0.001$), so the defensive channel as a whole is statistically distinguishable from a uniform repricing model.

The plateau coefficient is comparable in magnitude to the pre-shock structural defensive discount but with the opposite sign: defensive borrowers paid 50 bps less than discretionary borrowers pre-shock and 14 bps more in the plateau, holding size, maturity and yield type constant. We interpret the reversal as a change in the relative bargaining position of the two groups, with discretionary borrowers more willing to withdraw from the market once the rate environment turned. Panel A of Figure 2 displays the corresponding monthly mean spreads: the defensive line tracks below the discretionary line in the pre-shock period, rises sharply above it during the tightening and remains above it through the plateau, the persistence into the plateau being the slow-channel signature.

5.1.2. Volume evidence: H1 secondary prediction

The elasticity reading also predicts an asymmetric volume response. If discretionary borrowers can defer and defensive borrowers cannot, the tightening regime should see a sharper contraction of discretionary issuance than of defensive issuance; the plateau should then see a rebound concentrated on the discretionary side as deferred deals return. Beyond corroborating the spread reading, this test mitigates the most direct form of the spread-sample selection concern raised in Section 4.6: if discretionary borrowers had merely withdrawn from observable disclosure rather than from issuance, the H1 spread differential would partially reflect a shift in disclosure rather than a shift in pricing. The volume analysis on the full 27,491-tranche dataset, where issuance is observed regardless of spread disclosure, addresses this directly.

Table 5 reports the volume statistics. Defensive notional volume per month was 100 in the pre-shock regime, 79 in the tightening regime, and 96 in the plateau regime. Discretionary notional volume per month was 100 in the pre-shock regime, 71 in the tightening regime, and 112 in the plateau regime. Two patterns deserve attention. Defensive notional issuance held nearly steady across regimes: it contracted by 21 percent during tightening and recovered to within four percent of the pre-shock benchmark in the plateau. Discretionary notional issuance contracted by 29 percent during tightening, and rebounded above the pre-shock benchmark in the plateau as deferred deals returned. The relative magnitudes line up with the elasticity prediction. Defensive borrowers cannot postpone refinancing, so their issuance volume reacts less to the rate environment. Discretionary borrowers can and do postpone, so their issuance volume contracts more sharply during the tightening and rebounds rapidly during the plateau when the rate trajectory clarifies. Figure 3 plots the quarterly notional volume index by use of proceeds and shows the same dynamic at quarterly frequency.

Because issuance on the full dataset is observed whether or not a spread is disclosed, the sharp discretionary contraction during the tightening shows a withdrawal from the market itself, not merely from the disclosed subsample; the H1 spread differential is therefore a within-deal repricing rather than a compositional artefact. A residual asymmetric-disclosure story on the defensive side cannot be fully excluded with volume data alone, but it would have to be quantitatively large and oddly timed to overturn the differential, since defensive notional volume held nearly steady across regimes.

5.2. Industry cyclicalities and regulated capital expenditure

As with H1, we test H2 on both margins in turn: the spread regression with cycle control (Section 5.2.1), then the corroborating volume evidence (Section 5.2.2).

5.2.1. Spread regression with cycle control

The H2 spread regression is reported in Table 6. The binary specification in column 1 returns a non-cyclical-by-post-shock interaction of +32.7 bps with HC1 standard error 6.9 ($p < 0.001$). The structural baseline is a small negative gap of -7.4 bps, significant at the 5% level: in the pre-shock regime, non-cyclical borrowers paid slightly less than cyclical borrowers, consistent with their typically more stable cash flows. The post-shock interaction inverts this ranking: from 2022 onwards, non-cyclical borrowers paid roughly 25 bps more than cyclical borrowers on average, controlling for size, maturity and yield-type bucket.

The preferred three-regime specification in column 3 adds the contemporaneous euro-area twenty-country unemployment rate as a level control (UnempC, Eurostat series `une_rt_m`, centered on the in-sample mean) to discipline the rate-shock interpretation against any residual labour-market variation. The non-cyclical-by-tightening interaction is +41.9 bps with HC1 standard error 14.5 ($p = 0.004$), and the non-cyclical-by-plateau interaction is +30.4 bps with HC1 standard error 7.2 ($p < 0.001$). Both during the tightening and during the plateau, non-cyclical borrowers paid more than cyclical borrowers relative to their pre-shock baseline; both effects are statistically distinguishable from zero at the 1% level, and the persistence across regimes reflects the structural inelasticity of non-cyclical demand for credit. Utilities cannot postpone regulator-mandated grid investments. Healthcare borrowers cannot delay manufacturing capacity tied to product launches. Consumer staples cannot reduce working capital below the level required to service stable retail demand. These borrowers came to market when the rate environment demanded a higher spread, and they paid it. A joint Wald F-test on both interactions rejects the null at the 1% level ($F(2, 5236) = 11.7, p < 0.001$).

The unemployment level enters with a small negative coefficient (-5.9 bps per percentage point of unemployment, HC1 SE 1.4, $p < 0.001$). We do not interpret this control structurally; it enters to absorb residual labour-market variation, and its sign is not central to the H2 reading. What matters for H2 is that the non-cyclical-by-regime interactions are essentially unchanged with or without the cycle control (without UnempC: +41.7 and +30.1 bps; with UnempC: +41.9 and +30.4 bps), so the control disciplines the rate-shock interpretation without altering the substantive result. The H2 channel reflects the rate shock itself, not labour-market dynamics that happen to coincide with it.

That non-cyclical borrowers pay more is the reversal of the conventional cyclical reading, under which cyclical borrowers, with more cycle-exposed cash flows, would pay more in a downturn. The reversal is identifiable here precisely because the 2022 to 2023 episode was a rate-led tightening rather than a recession (Section 3.2). One caveat applies: industry classifications may also capture sector-specific exposures (energy, regulation, supply chains) beyond the deferability channel, so the H2 evidence is best read as

consistent with, rather than uniquely identifying, the elasticity mechanism. Panel B of Figure 2 displays the corresponding monthly mean spreads: cyclical and non-cyclical overlap closely pre-shock, then non-cyclical crosses above during the tightening and stays elevated through the plateau, the moderate-channel pattern anticipated by H2.

5.2.2. Volume evidence: H2 secondary prediction

The H2 elasticity reading carries the same volume implication as H1, at the industry level: cyclical issuance should contract more than non-cyclical during the tightening, with both groups recovering in the plateau. Table 7 reports the volume statistics on the same full 27,491-tranche dataset, restricted to the 27,483 tranches with a non-missing industry classification. Per-month notional volume for non-cyclical borrowers was indexed at 100 in the pre-shock regime, fell only to 98.7 during tightening, and rose to 123.7 in the plateau. Per-month notional volume for cyclical borrowers was 100 in the pre-shock regime, contracted to 82.2 during tightening, and rose to 119.9 in the plateau. Cyclical notional issuance therefore contracted by about 18 percent during the tightening while non-cyclical was essentially flat, the asymmetric pattern predicted by the elasticity reading. Both groups rebound in the plateau, with the rebound notably stronger for the neutral bucket (170.3) that accumulates deferred deals across both classifications.

Figure 4 plots the quarterly notional volume index by industry classification and shows the same dynamic at quarterly frequency, with the cyclical (orange) series visibly dipping through the tightening band while the non-cyclical (navy) series is essentially flat across the same window. The quarterly figure uses a 2019Q1 base for comparability with Figure 3; the regime-average index in Table 7 Panel C, rebased to pre-shock = 100, sharpens the asymmetric-contraction reading. The H2 volume evidence corroborates the spread reading and addresses the same selection concern raised for H1: cyclical borrowers withdrew from issuance during the tightening, which is the elasticity prediction itself, not from observable disclosure.

5.3. Sponsor backing and the institutional channel

The H3 sponsor channel is tested in Table 8. We begin with the binary specification in column 1, which is the methodological diagnostic that motivates the three-regime decomposition. Using a single post-shock indicator, the sponsor-by-post coefficient is -4.0 bps with HC1 standard error 5.3 ($p = 0.45$), statistically indistinguishable from zero. Read literally, this coefficient says that sponsor-backed loans repriced no differently from non-sponsored loans after the 2022 shock, controlling for size, maturity and yield-type bucket. As we now show, this null is the average of two large opposite-signed regime-specific interactions, and taken at face value it misrepresents the underlying transmission.

The three-regime specification in column 3 reveals the underlying dynamics. The sponsor-by-tightening interaction is +66.1 bps with HC1 standard error 10.3 ($p < 0.001$), among the largest interaction coefficients in the paper and narrowly above the H1 plateau coefficient of +64.6 bps. During the tightening phase, sponsor-backed loans repriced by an additional 66 basis points relative to non-sponsored loans, holding controls constant. The institutional buyer base of broadly syndicated leveraged loans, dominated by CLOs and credit funds, demanded a substantial spread widening to clear new issues at the higher policy rate, and that demand was passed through to borrowers in full and quickly. The non-sponsor segment, dominated by relationship lending between corporates and their core banks, repriced more gradually as banks absorbed the funding-cost shock through internal margins before passing it through to clients as legacy facilities reset (Drechsler, Savov and Schnabl, 2017).

The sponsor-by-plateau interaction reverses to -16.2 bps with HC1 standard error 5.7 ($p = 0.005$). The differential built up during the tightening compresses substantially in the plateau, but this compression does not reflect a permanent re-rating of sponsor credit. We read the pattern as the mechanical implication of a fast channel meeting a slow one: by the time the plateau took hold, the institutional channel had completed its initial repricing of required yields, and the bank-funded margins on non-sponsored loans rose with legacy-facility resets and new facilities priced off the higher cost of funds. The cross-segment differential narrows from its tightening peak (+140.7 bps) back toward the pre-shock baseline (+74.6 bps), and slightly past it to +58.4 bps in the plateau. The negative plateau coefficient is therefore consistent with catch-up in the slow channel rather than a re-rating of the fast one: it does not say sponsor loans are cheaper post-shock, they still pay roughly 58 bps more than non-sponsors, only that the institutional channel's lead over the relationship channel, large during the tightening, has substantially eroded. This compression contrasts sharply with the slow demand-elasticity channels of H1 and H2, both of which persist into the plateau because the underlying borrower-side inelasticity does not relax once policy stabilizes. The H3 channel is a pricing-technology speed differential, not a structural inelasticity. A joint Wald F-test of the null that both sponsor interactions equal zero delivers $F(2, 7754) = 29.0$ ($p < 0.001$), the largest channel F-statistic in the paper.

Panel C of Figure 2 displays the corresponding monthly mean spreads. The sponsored line sits above the non-sponsored line throughout, but the gap widens sharply as the sponsored line spikes during the tightening, then narrows through the plateau, the fast-institutional signature. The contrast between the binary null in column 1 and the three-regime path in column 3 is the methodological contribution of the paper. A binary indicator averages the +66 bps tightening spike against the -16 bps plateau differential into a near-zero coefficient that masks the transmission. The decomposition instead recovers two distinct diagnostic patterns, persistence for the slow elasticity channels (H1, H2) and reversal for the fast

institutional one (H3), which carry different economic readings that a single post-shock coefficient cannot tell apart.

5.4. Structural premium across regimes

This subsection reads the three channels together by reporting the implied structural premium (treated minus comparison group) in each regime, which makes the differential dynamics directly comparable across hypotheses. Table 9 reports the three regime-specific premia for each hypothesis and the plateau reversal, computed as the structural coefficient β_1 from the preferred Spec 3 plus the regime-specific interaction; the H2 row uses the unemployment-augmented Spec 3 of Table 6.

Read across the three regimes (Table 9), the channels separate cleanly into the two predicted signatures. The two elasticity channels reverse a pre-shock discount or parity into a positive premium that holds into the plateau: H1 moves from -50 to +14 bps and H2 from -8 to +23 bps. The institutional channel instead amplifies then reverts, the sponsor premium widening from +75 to +141 bps in the tightening before falling back to +58 bps. Figure 6 plots the same premia with 95 percent confidence intervals and overlays the binary implied premium as a step function; the contrast is starkest on H3, where the binary step is nearly flat while the three-regime path peaks above 140 bps and then settles near its baseline.

To translate these coefficients into dollar magnitudes, consider a tranche with \$300 million in principal and a five-year maturity, broadly representative of the median size in our sample. The +65 basis-point defensive plateau premium then represents roughly \$9.75 million in additional interest cost over the life of the loan; the +30 basis-point non-cyclical plateau premium represents about \$4.5 million; and the +66 basis-point sponsor-by-tightening interaction, for a comparable sponsor-backed tranche, represents about \$9.9 million. These magnitudes are economically meaningful at the individual-borrower level and scale linearly with size and maturity; given the per-month volumes in Tables 5 and 7, they aggregate to materially larger flows across the post-shock issuance pool.

Figure 5 consolidates the six headline interactions with HC1 robust 95 percent confidence intervals, placing the three channel signatures side by side: the two elasticity channels widen or hold into the plateau, while the sponsor-by-plateau coefficient is the only one of the six that falls below zero. All six are statistically distinguishable from zero at the 5 percent level.

5.5. Robustness checks

This subsection reports five robustness exercises and a consolidated F-test summary.

5.5.1. Channel orthogonality: pairwise and joint specifications

A natural concern with three channels estimated in parallel on overlapping samples is that the cross-section of the treatment indicators may be heavily correlated. Sponsor-backed deals are typically discretionary in use of proceeds, and the H2 industry classification interacts with both. If the channels were merely relabelings of the same underlying selection, the channel $\times$ regime interactions should attenuate toward zero once the other channels are introduced. Table 10 tests this directly with four specifications estimated on the full clean spread sample of 7,764 tranches: three pairwise stacks (H1 $\times$ H2, H1 $\times$ H3, H2 $\times$ H3) and one full joint specification stacking all three channels. The unemployment cycle control is retained in every specification that includes the H2 channel.

All six channel $\times$ regime interactions retain the sign, the rough magnitude and the statistical significance reported in their respective main tables. The defensive-by-plateau interaction, the largest H1 coefficient, moves from +64.6 bps in Table 4 to +43 to +47 bps across the four orthogonality specifications. The non-cyclical-by-tightening interaction moves from +41.9 bps to +26 to +27 bps. The sponsor-by-tightening interaction is essentially invariant at +66 to +68 bps, and the sponsor-by-plateau interaction stays within the -14 to -19 bps range. Joint Wald F-tests of the null that all interactions vanish reject in every specification ($F(4, 7750) = 12.3$ for the H1 $\times$ H2 pairwise; $F(4, 7751) = 19.5$ for H1 $\times$ H3; $F(4, 7750) = 20.5$ for H2 $\times$ H3; $F(6, 7747) = 17.0$ for the full joint; all $p < 0.001$). The conclusion is that the three channels identify independent margins of selective repricing rather than substitute labels for the same underlying selection.

The H1 attenuation in the joint specifications (defensive-by-plateau from +65 to +43 bps) is larger than that of H2 or H3, indicating that the use-of-proceeds channel shares some cross-sectional variance with the other two; the channel survives intact but loses about a third of its magnitude when the orthogonal margins absorb their share.

5.5.2. Yield-type subsamples for H1

This subsection asks whether the H1 use-of-proceeds dynamic operates uniformly across yield-type buckets or whether it is concentrated in one segment. We re-estimate equation (3) of H1 separately on the Investment Grade subsample (914 observations) and the Leveraged subsample (2,392 observations); the Highly Leveraged bucket is too small to support a separate decomposition. Results are reported in Table 11.

The leveraged subsample produces a defensive-by-plateau interaction of +69.8 bps with HC1 standard error 14.7 ($p < 0.001$), strongly statistically significant and slightly larger than the pooled-sample coefficient of

+64.6 bps. The defensive-by-tightening interaction in the leveraged subsample is +15.9 bps but with a wide standard error of 27.4 and is not statistically significant. The investment-grade subsample produces a defensive-by-plateau interaction of +30.9 bps with HC1 standard error 16.9 ($p = 0.07$) and a defensive-by-tightening interaction of +34.0 bps with a wide standard error of 40.1, not statistically significant. The slow demand-elasticity channel operates across yield types, but it is sharper in the leveraged segment. This pattern is consistent with the bargaining-channel prediction sketched in Section 4.1: the leveraged segment is precisely where lender market power over rollover-needy borrowers is most acute, both because the buyer base is more concentrated and price-sensitive, and because outside funding options are thinner for sub-investment-grade borrowers. The overall direction of the H1 result is preserved across yield-type buckets but with magnitude concentrated in the leveraged segment, consistent with the bargaining-channel mechanism we propose rather than with a generic risk-repricing story that would predict similar effects across yield types.

5.5.3. Placebo tests on a 2018 to 2021 sample

This subsection is a placebo test designed to confirm that the post-shock interactions reflect the rate shock rather than a pre-existing differential trend. We restrict the sample to tranches issued between 1 January 2018 and 31 December 2021, impose a pseudo-shock on 21 July 2020 (parallel to the actual cycle, shifted back by exactly two years), and re-estimate equation (3) for each hypothesis with the pseudo-tightening and pseudo-plateau windows replacing the actual ones. Table 12 reports the placebo coefficients.

The H2 placebo is broadly clean. The non-cyclical-by-pseudo-tightening interaction is +22.6 bps with HC1 standard error 12.3 ($p = 0.07$, marginally significant at the 10% level), and the non-cyclical-by-pseudo-plateau interaction is -35.0 bps with HC1 standard error 24.8 ($p = 0.16$). Neither is statistically distinguishable from zero at the conventional 5% level, although the pseudo-tightening coefficient warrants the caveat that the placebo is not perfectly clean. The H2 result in the actual sample (+41.9 bps, $p = 0.004$) is nearly twice the magnitude of the pseudo-tightening coefficient and is significant at the 1% level, supporting the reading that the H2 effect is a genuine response to the rate shock rather than the continuation of a pre-existing trend.

The H3 placebo is a clean negative null and is, if anything, more informative than a flat null. The sponsor-by-pseudo-tightening interaction is -30.5 bps with HC1 standard error 10.1 ($p = 0.003$), strongly negative and the opposite sign from the actual sponsor-by-tightening interaction of +66.1 bps. The interpretation is that during the early COVID disruption, sponsor-backed loans repriced less than non-sponsored loans relative to the pre-COVID baseline, a pattern consistent with the relationship- and reputation-based pricing advantages that sponsor-backed borrowers carry into stressed markets (Demiroglu and James, 2010;

Ivashina and Kovner, 2011) rather than with the speed differential that drives the actual-sample result. The institutional channel of H3 is rate-shock-specific, not crisis-specific: a pure liquidity event of the COVID type generates the opposite of the H3 pattern, whereas the rate shock of 2022 generates it. The placebo evidence is consistent with the H3 finding reflecting the specific institutional dynamic of a sustained rate-shock cycle rather than a generic crisis episode.

The H1 placebo is partial and warrants a careful reading. The defensive-by-pseudo-tightening interaction is +22.1 bps with HC1 standard error 14.8 ($p = 0.14$, not statistically significant), and the defensive-by-pseudo-plateau interaction is -85.1 bps with HC1 standard error 23.5 ($p < 0.001$, opposite sign from the actual H1 plateau coefficient). The opposite-sign pseudo-plateau coefficient indicates that the defensive-discretionary relationship was rebounding through 2021 as COVID stress unwound, which is the opposite of what occurred in the actual 2022 to 2023 sample, where the defensive premium widened further into the plateau. The defensive-by-tightening interaction in the actual sample (+38.2 bps) overlaps in magnitude with the not-significant COVID pseudo-tightening interaction (+22.1 bps), so we cannot fully exclude a residual liquidity-event component from the H1 tightening coefficient. The defensive-by-plateau interaction in the actual sample (+64.6 bps), by contrast, is sharply opposite in sign from the COVID pseudo-plateau coefficient (-85.1 bps) and is therefore unambiguously rate-shock-driven. The plateau coefficient is the cleanly identified part of H1; the tightening coefficient should be read more cautiously, with the placebo delivering supportive but not conclusive backing. Because the COVID period is itself an environment in which a defensive elasticity premium might plausibly appear, the placebo's failure to reproduce the actual-sample plateau result under those conditions strengthens rather than weakens the H1 plateau identification.

5.5.4. Clustered standard errors on the borrower

This subsection relaxes the HC1 assumption that two tranches from the same borrower have independent residuals. Using the stable Issuer/Borrower PermID, we re-estimate the binary and three-regime specifications of each hypothesis with one-way cluster-robust standard errors on the borrower, across 2,571 clusters (2,110 multi-tranche borrower groups and 461 PermID-unassigned singletons). As set out in Section 4.6, two-way clustering on borrower and time is not feasible given the fifteen-month tightening regime. Table 13 reports the HC1-versus-cluster comparison on the headline interactions, and Figure 7 overlays the HC1, cluster-robust and country-fixed-effect estimates for the same six coefficients.

Three findings emerge. First, the cluster correction inflates standard errors relative to HC1 by between roughly 27% and 56% on the headline interactions (median inflation about 35% across the full set). The magnitudes suggest mild within-borrower correlation rather than a dominant clustering structure. Second,

every preferred-specification headline result of Sections 5.1 to 5.3 retains its significance level under clustering, with one exception: the defensive-by-tightening interaction of H1, which is significant at the 5% level under HC1 ($p = 0.032$), drops to marginal significance at the 10% level under clustering (cluster SE 22.6, $p = 0.091$). The plateau coefficient of the same channel remains highly significant under both inference regimes (cluster SE 15.1, $p < 0.001$), and a joint Wald F-test on both H1 interactions still rejects at the 1% level under clustering ($F(2, 1699) = 9.5$, $p < 0.001$), so the channel as a whole survives the inference change. The H2 and H3 plateau and tightening interactions all retain at least 5% significance under clustering. Third, the H3 diagnostic null on the binary specification is preserved: the sponsor-by-post coefficient of -4.0 bps with HC1 standard error 5.3 ($p = 0.45$) becomes -4.0 bps with cluster standard error 7.8 ($p = 0.61$) under clustering. The three-regime decomposition is required to recover the underlying institutional dynamics under both inference regimes; clustering does not rescue the binary specification.

5.5.5. Country fixed effects

This subsection addresses cross-country compositional heterogeneity. The cross-section of European syndicated borrowers is unevenly distributed across 38 domicile nations, with the United Kingdom accounting for 29.4% of the clean sample, France 11.5%, the Netherlands 8.8%, and the next seven countries (Germany, Spain, Turkey, Ireland, Switzerland, Luxembourg, Italy) accounting for a further 30%. If the regime composition of any single country differs from the pooled composition, the regime-by-structural-margin interactions of Sections 5.1 to 5.3 could partially reflect a within-country compositional shift rather than within-borrower repricing. The Domicile Nation field in the SDC extract allows us to address this concern directly.

We re-estimate the preferred three-regime specification of each hypothesis with up to 37 country dummies (35 for H2; the United Kingdom omitted as reference category); for H2 we additionally retain the unemployment cycle control. The country block is jointly highly significant in every hypothesis: a partial Wald F-test on the joint nullity of all country dummies delivers $F = 41.7$ in H1, $F = 23.8$ in H2, and $F = 97.6$ in H3, all with $p < 0.001$. The larger country coefficients sit on Ireland (-30.8 bps relative to the United Kingdom in the H3 specification), Spain (-19.5 bps), Switzerland (-27.7 bps), Turkey (-75.2 bps) and the Netherlands (-15.7 bps). Where the underlying cell is reasonably populated, these are consistent with composition tilts (an offshore-leveraged-loan tilt for Ireland, private-banking-led financings for Switzerland, investment-grade trade financings for Turkey), but we read them as descriptive rather than as estimates we would interpret structurally. The country dummies on micro-jurisdictions such as Gibraltar, Monaco, Serbia and Ukraine carry very large coefficients but rest on at most five observations each. Table 14 reports the country fixed effects.

The relevant question is whether the regime channels survive the introduction of country fixed effects. The answer is affirmative for all three hypotheses. The sponsor-by-tightening interaction is +65.3 bps under country fixed effects, compared with +66.1 bps under the pooled HC1 specification, and the sponsor-by-plateau interaction is -17.0 bps compared with -16.2 bps. The non-cyclical-by-tightening interaction is +36.6 bps, compared with +41.9 bps, and the non-cyclical-by-plateau interaction is +32.1 bps, compared with +30.4 bps; both H2 interactions remain significant at the 1% level and the latter slightly tightens to its country-FE estimate. The only material attenuation is on the defensive-by-plateau interaction of H1, which moves from +64.6 bps to +48.1 bps once country fixed effects are absorbed; the channel remains strongly significant ($p < 0.001$) and the qualitative reading is unchanged. The joint Wald F-tests on the channel interactions under country fixed effects deliver $F = 10.0$ in H1, $F = 11.8$ in H2 and $F = 28.9$ in H3, all rejecting at the 1% level. The geographic invariance of the channel coefficients indicates that the channels operate at the borrower-and-tranche level rather than at the country level: the cross-country composition of the issuance pool does not drive the results.

5.5.6. Joint Wald F-tests across inference regimes

This subsection consolidates the joint F-tests into a single comparison table. Where the individual t-tests in Tables 4, 6 and 8 evaluate each coefficient in isolation, a joint Wald F-test asks whether a channel as a whole is distinguishable from uniform repricing across the two post-shock regimes. Table 15 reports these tests for each hypothesis under three inference regimes (HC1 robust, cluster-robust on PermID, and HC1 with country fixed effects), alongside the orthogonality F-tests of Section 5.5.1. Every test rejects its joint null at the 1 percent level, confirming that each channel is robust to the inference regime and that the three jointly identify independent margins of selective repricing.

6. Conclusion

We asked whether a sharp monetary tightening falls uniformly across syndicated-loan borrowers, or whether it concentrates on those with the least scope to step away. On 7,764 European tranches priced between January 2015 and February 2026, the answer is that it concentrates, and that it does so along two economically distinct patterns. The borrower-elasticity margins repriced persistently: the defensive-discretionary and non-cyclical-cyclical differentials both reversed their pre-shock signs and held into the plateau, by 65 and 30 basis points respectively. The institutional margin repriced transiently: sponsor-backed loans widened 66 basis points during the tightening and then surrendered the entire spike, the signature of a fast buyer base outrunning a slow relationship channel that later caught up. Asymmetric issuance contractions corroborate the elasticity reading on a second margin, and an orthogonality test

confirms that the three margins are independent rather than relabelings of one selection. The persistence-versus-reversal split, invisible to a binary specification, is the organizing result.

Three contributions follow from these results. The conceptual one is the organizing framework itself: two borrower-elasticity margins and one institutional-channel margin, unified by the prediction that elasticity produces persistence and pricing technology produces reversal, a prediction the regime signatures bear out. The empirical one is the setting: testing spread and volume jointly on the steepest euro-area tightening on record, in a non-recessionary environment that lets the H2 elasticity reading be separated from the conventional cyclical reading rather than confounded with it. The methodological one is the cautionary case the H3 channel provides: a binary coefficient of -4 basis points, indistinguishable from zero, unpacks into a +66 spike and a -16 reversion, so a single post-shock window can return a confident null precisely where two pricing channels of different speed coexist.

The findings carry recommendations for three audiences. For central banks and supervisors, a single average pass-through coefficient is an insufficient statistic for the transmission of a rate shock: two cycles with identical average widening can distribute it very differently across borrowers, so transmission monitoring should track the cross-sectional dispersion of pass-through, not just its mean. The volume evidence sharpens the point, since the most flexible borrowers withdraw from issuance rather than absorb the higher spread, so quantity-based measures can understate the effective cost faced by borrowers who cannot defer. For market practitioners, a defensive refinancing in a non-cyclical industry now prices at a structural premium that did not exist before the shock, so pre-shock comparables understate the current spread by margins that run into millions of dollars on a single mid-size financing (Section 5.4). For the literature, the three-regime decomposition is worth adopting wherever a fast pricing channel coexists with a slow one, and joint specifications discipline against relabeling artefacts.

Three limitations warrant acknowledgement. First, the dependent variable is restricted to the all-in drawn spread; non-price terms such as commitment fees, covenants, maturities and structural protections are not consistently reported in SDC and would require manual extraction from loan-by-loan filings. Second, the spread sample is conditional on observable spreads, reported for roughly twenty-nine percent of European tranches in our window: the volume analyses of Sections 5.1.2 and 5.2.2 mitigate the most direct selection threat, but a residual asymmetric-disclosure story cannot be excluded with volume data alone. Third, the borrower-level dimension is restricted to the variables in the SDC extract, which limits our ability to test the more granular within-channel predictions of the bargaining mechanism and to control for alternative financing channels such as bond-market access and equity issuance. The most material instance is European private credit, excluded for lack of consistent primary-market data (Section 3.1): if marginal sponsor borrowers migrated from syndicated to private-credit financing during 2024 and 2025, the plateau

compression we attribute to the relationship channel catching up could partly reflect a change in the composition of who remained in the syndicated sponsor market. We cannot rule this out with the present data, and flag it as the leading qualification to the H3 plateau reading.

These limitations map onto natural extensions. Extending the dependent variable to non-price terms would test predictions the spread alone cannot: defensive borrowers should also receive less favorable covenants and shorter maturities post-shock. Linking SDC to a regulatory register such as ECB AnaCredit would isolate the disclosure-selection mechanism from the issuance-volume one, and borrower-level financial-statement controls would test the within-channel predictions of the bargaining mechanism. Finally, the fast-versus-slow framework could be applied to the early-easing path that began in mid-2024, a second three-regime cycle that would discipline the speed differential from the opposite direction.

The broader takeaway concerns the unit of analysis. Borrowers come to market with different degrees of flexibility, and lenders price that flexibility into the spreads they extract, so a rate shock redistributes bargaining rents across the cross-section in ways aggregate statistics do not capture. Whether monetary policy operates evenly or selectively is therefore a cross-sectional question, and the three-regime decomposition is the simplest specification that answers it without averaging fast and slow channels into a misleading uniform reading.

References

- Alternative Credit Council, 2025. Financing the Economy 2025. AIMA and Houlihan Lokey.
<https://www.aima.org/article/press-release-strong-growth-sees-private-credit-market-reach-us-3-5-trillion.html>
- Auer, S., Conti, A. M., 2024. Bank lending in an unprecedented monetary tightening cycle: Evidence from the euro area. Banca d'Italia, Questioni di Economia e Finanza (Occasional Papers) No. 856.
<https://www.bancaditalia.it/pubblicazioni/qef/2024-0856/>
- Bank of England, 2023. Financial Stability Report, July 2023. Bank of England, London.
<https://www.bankofengland.co.uk/financial-stability-report/2023/july-2023>
- Becker, B., Ivashina, V., 2014. Cyclicalities of credit supply: Firm level evidence. Journal of Monetary Economics 62, 76-93. <https://www.nber.org/papers/w17392>
- Bernanke, B. S., Blinder, A. S., 1988. Credit, money, and aggregate demand. American Economic Review 78(2), 435-439. <https://www.nber.org/papers/w2534>
- Bloomberg, 2022, November 29. Banks stuck with \$42 billion debt seize chance to offload it. Bloomberg News, New York. <https://www.bloomberg.com/news/articles/2022-11-29/banks-stuck-with-42-billion-debt-seize-chance-to-offload-it>
- Bord, V. M., Santos, J. A. C., 2012. The rise of the originate-to-distribute model and the role of banks in financial intermediation. Federal Reserve Bank of New York Economic Policy Review 18(2), 21-34.
<https://www.newyorkfed.org/medialibrary/media/research/epr/2012/EPRvol18n2.pdf>
- Chodorow-Reich, G., Falato, A., 2022. The loan covenant channel: How bank health transmits to the real economy. Journal of Finance 77(1), 85-128. <https://doi.org/10.1111/jofi.13074>
- Demiroglu, C., James, C. M., 2010. The role of private equity group reputation in LBO financing. Journal of Financial Economics 96(2), 306-330. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1032781
- Diamond, D. W., 1991. Monitoring and reputation: The choice between bank loans and directly placed debt. Journal of Political Economy 99(4), 689-721. <https://doi.org/10.1086/261775>
- Drechsler, I., Savov, A., Schnabl, P., 2017. The deposits channel of monetary policy. Quarterly Journal of Economics 132(4), 1819-1876. <https://www.nber.org/papers/w22152>
- European Central Bank, 2023a. The euro area bank lending survey -First quarter of 2023. European Central Bank, Frankfurt am Main, May 2023.
https://www.ecb.europa.eu/stats/ecb_surveys/bank_lending_survey/html/ecb.blssurvey2023q1~22c176b442.en.html
- European Central Bank, 2023b. Economic Bulletin, Issue 3/2023. European Central Bank, Frankfurt am Main.
<https://www.ecb.europa.eu/press/economic-bulletin/html/eb202303.en.html>

- European Central Bank, 2023c. Economic Bulletin, Issue 8/2023, Box: Monetary dynamics during the tightening cycle. European Central Bank, Frankfurt am Main. https://www.ecb.europa.eu/press/economic-bulletin/focus/2024/html/ecb.ebbox202308_08~09682c131a.en.html
- European Central Bank, 2023d. Financial Stability Review, November 2023. European Central Bank, Frankfurt am Main. <https://www.ecb.europa.eu/press/financial-stability-publications/fsr/html/ecb.fsr202311~bfe9d7c565.en.html>
- European Central Bank, 2024. Key ECB interest rates and monetary policy decisions, July 2022 to October 2023; corroborated in ECB Annual Report 2024. https://www.ecb.europa.eu/stats/policy_and_exchange_rates/key_ecb_interest_rates/html/index.en.html
- European Central Bank Banking Supervision, 2024. Interview with Elizabeth McCaul, conducted by Martin Arnold and Costas Mourselas, Financial Times, 10 July 2024. ECB Banking Supervision. <https://www.bankingsupervision.europa.eu/press/interviews/date/2024/html/ssm.in240710~8a289e1673.en.html>
- Eurostat, 2026. Unemployment by sex and age (monthly data). Series une_rt_m, euro area 20 aggregate. Eurostat, Luxembourg. https://ec.europa.eu/eurostat/databrowser/view/une_rt_m/default/table?lang=en
- Heider, F., Saidi, F., Schepens, G., 2019. Life below zero: Bank lending under negative policy rates. Review of Financial Studies 32(10), 3728-3761. <https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2173.en.pdf>
- International Monetary Fund, 2023. Global Financial Stability Report: Safeguarding Financial Stability amid High Inflation and Geopolitical Risks. International Monetary Fund, Washington, D.C., April 2023. <https://www.imf.org/en/Publications/GFSR/Issues/2023/04/11/global-financial-stability-report-april-2023>
- Ivashina, V., Kovner, A., 2011. The private equity advantage: Leveraged buyout firms and relationship banking. Review of Financial Studies 24(7), 2462-2498. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1017857
- Ivashina, V., Sun, Z., 2011. Institutional demand pressure and the cost of corporate loans. Journal of Financial Economics 99(3), 500-522. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1299300
- LSEG, 2026. SDC Platinum Financial Securities (Loans), delivered via LSEG Workspace. London Stock Exchange Group, London (data extracted 2 February 2026). <https://www.lseg.com/en/data-analytics/products/sdc-platinum-financial-securities>
- Petersen, M. A., Rajan, R. G., 1994. The benefits of lending relationships: Evidence from small business data. Journal of Finance 49(1), 3-37. <https://doi.org/10.1111/j.1540-6261.1994.tb04418.x>
- PitchBook LCD, 2022. 2023 European Leveraged Loan Outlook: Tough road to recovery. PitchBook Data, Inc. (LCD), 16 December 2022. <https://pitchbook.com/news/articles/2023-european-leveraged-loan-outlook-tough-road-to-recovery>
- PitchBook LCD, 2023. European Credit Markets Quarterly Wrap, Third Quarter 2023. PitchBook Data, Inc., Seattle. <https://pitchbook.com/news/reports/q3-2023-european-credit-markets-quarterly-wrap>

- Schwert, M., 2018. Bank capital and lending relationships. *Journal of Finance* 73(2), 787-830.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2690490
- Schwert, M., 2020. Does borrowing from banks cost more than borrowing from the market? *Journal of Finance* 75(2), 905-947. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3059607
- Sufi, A., 2007. Information asymmetry and financing arrangements: Evidence from syndicated loans. *Journal of Finance* 62(2), 629-668. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=596202

Tables

Table 1. Sample construction pipeline.

Step	Filter	Observations remaining	Change vs previous	Comment
Step 0	SDC Platinum raw output (Western and Eastern Europe, USD-converted, announced 1 Jan 2015 onwards)	28,411		<i>Per SDC search filters</i>
Step 1	Drop loans that never closed (Financial Close Date is “not closed”)	27,491	-920	<i>Full dataset for volume analysis (Tables 5 and 7)</i>
Step 2	Drop tranches with missing or “no data” all-in drawn spread	7,941	-19,550	<i>Spread is unobserved</i>
Step 3	Drop Undisclosed Yield Type and Near Investment Grade	7,800	-141	<i>Sparse or ambiguous buckets</i>
Step 4	Drop missing or zero stated maturity	7,765	-35	<i>Maturity enters every regression</i>
Step 5	Winsorise spread at 1st and 99th percentile (30 bps and 775 bps)	7,765	0	<i>Caps outliers; no observations dropped</i>
Step 6	Drop tranches with missing industry classification	7,764	-1	
Final	Clean spread sample for H1, H2 and H3 spread regressions	7,764		<i>Sample used in Tables 2 to 9 and most robustness checks</i>

Step 0 reflects the raw SDC Platinum extract dated 2 February 2026 with Western and Eastern European borrower domiciles. The full 27,491-tranche dataset after Step 1 anchors the volume analysis reported in Tables 5 and 7. Steps 2 to 6 are the spread-regression filters and are applied in the order shown. The final clean spread sample of 7,764 tranches is the sample for H1, H2 and H3 regressions.

Table 2. Spread moments by regime.

Regime	N	Mean	Median	SD	P10	P25	P75	P90
Pre-shock	5,362	284.8	275.0	152.2	95.0	165.0	375.0	475.0
Tightening	681	316.5	300.0	169.9	110.0	175.0	450.0	550.0
Plateau	1,721	329.0	350.0	125.3	150.0	250.0	400.0	475.0
All	7,764	297.4	300.0	149.5	100.0	175.0	400.0	475.0

All-in drawn spreads in basis points, winsorised at the 1st and 99th percentiles. The pre-shock regime spans 1 January 2015 to 20 July 2022, the tightening regime spans 21 July 2022 to 25 October 2023, and the plateau regime spans 26 October 2023 to 2 February 2026. N denotes the number of tranches in each regime in the final clean spread sample.

Table 3. Composition of issuance pool by regime.

Category	Pre-shock		Tightening		Plateau		All	
	<i>N</i>	<i>Share</i>	<i>N</i>	<i>Share</i>	<i>N</i>	<i>Share</i>	<i>N</i>	<i>Share</i>
<i>Panel A. Use of proceeds</i>								
Defensive	753	14.0%	73	10.7%	155	9.0%	981	12.6%
Discretionary	2,136	39.8%	211	31.0%	418	24.3%	2,765	35.6%
Neutral	2,473	46.1%	397	58.3%	1,148	66.7%	4,018	51.8%
<i>Panel B. Industry classification</i>								
Cyclical	2,493	46.5%	285	41.9%	702	40.8%	3,480	44.8%
Non-Cyclical	1,248	23.3%	126	18.5%	393	22.8%	1,767	22.8%
Neutral	1,621	30.2%	270	39.6%	626	36.4%	2,517	32.4%
<i>Panel C. Yield type</i>								
Investment Grade	1,702	31.7%	221	32.5%	262	15.2%	2,185	28.1%
Leveraged	3,180	59.3%	355	52.1%	1,194	69.4%	4,729	60.9%
Highly Leveraged	480	9.0%	105	15.4%	265	15.4%	850	10.9%
<i>Panel D. Sponsor flag</i>								
Non-sponsored	3,718	69.3%	496	72.8%	745	43.3%	4,959	63.9%
Sponsored	1,644	30.7%	185	27.2%	976	56.7%	2,805	36.1%

Counts and within-regime shares for the four categorical variables in the final clean spread sample ($N = 7,764$). Each panel sums to 100 percent within each regime. The leveraged share rises from 59.3 to 69.4 percent and the sponsor share rises from 30.7 to 56.7 percent between the pre-shock and plateau regimes; both shifts motivate the inclusion of yield-type controls and the three-regime decomposition.

Table 4. Effect of defensive use of proceeds on the all-in drawn spread by regime.

Variable	Spec 1 <i>Binary, with controls</i>	Spec 2 <i>Three-regime, no controls</i>	Spec 3 <i>Three-regime, with controls</i>
Defensive	-49.98*** (5.04)	-83.82*** (6.34)	-50.18*** (5.04)
HighRate (post-shock)	-4.12 (5.23)		
Tight		-17.92 (12.85)	9.28 (9.68)
Plat		2.74 (7.07)	-11.06* (5.70)
Defensive × HighRate	55.88*** (10.10)		
Defensive × Tight		117.14*** (22.34)	38.24** (17.87)
Defensive × Plat		94.88*** (14.12)	64.56*** (11.27)
<i>Maturity, log tranche size, yield-type</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Observations	3,746	3,746	3,746
R-squared	0.458	0.049	0.458

Sample of 3,746 tranches whose use of proceeds is classified as defensive or discretionary; the neutral category is excluded. The dependent variable is the all-in drawn spread in basis points, winsorised at the 1st and 99th percentiles. HCl robust standard errors in parentheses. Specification 1 uses a single binary post-shock indicator; Specifications 2 and 3 decompose the post-shock window into a tightening regime (21 July 2022 to 25 October 2023) and a plateau regime (26 October 2023 to 2 February 2026). Specification 3 is the preferred specification. Yield-type controls comprise the leveraged dummy and the highly leveraged dummy, with investment grade as the omitted base category. A joint Wald F-test on the two interactions in Specification 3 delivers $F(2, 3736) = 17.5$ ($p < 0.001$). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5. Volume analysis by regime and use of proceeds (H1 secondary prediction).

Regime	Defensive	Discretionary	Neutral	All categories
<i>Panel A. Aggregate count of tranches</i>				
Pre-shock (90.58 months)	2,052	6,834	8,882	17,768
Tightening (15.15 months)	278	1,018	1,662	2,958
Plateau (27.27 months)	592	2,418	3,755	6,765
<i>Panel B. Per-month notional issuance (USD millions)</i>				
Pre-shock	12,215	33,335	51,964	97,514
Tightening	9,654	23,838	54,251	87,743
Plateau	11,771	37,326	81,149	130,246
<i>Panel C. Notional volume index (pre-shock = 100)</i>				
Pre-shock	100.0	100.0	100.0	100.0
Tightening	79.0	71.5	104.4	90.0
Plateau	96.4	112.0	156.2	133.6

Volume analysis on the full dataset of 27,491 tranches (after Step 1 of the cleaning pipeline, before any spread filter) so as to remove the spread-availability selection that affects the spread sample. Counts in Panel A are within-regime totals; per-month figures in Panel B normalize for unequal regime length (pre-shock 90.58 months, tightening 15.15 months, plateau 27.27 months). Panel C rebases per-month notional values to pre-shock = 100. Discretionary notional issuance contracts by 29 percent during the tightening and rebounds 12 percent above the pre-shock benchmark in the plateau; defensive notional issuance contracts less (21 percent) and recovers to within 4 percent of the benchmark. The asymmetric contraction during the tightening identifies the elasticity asymmetry that underlies the H1 spread differential in Table 4.

Table 6. Effect of non-cyclical industry classification on the all-in drawn spread by regime, with euro-area unemployment cycle control.

Variable	Spec 1 <i>Binary, with controls</i>	Spec 2 <i>Three-regime, no controls</i>	Spec 3 (preferred) <i>Three-regime, with controls and UnempC</i>
NonCyc	-7.39** (3.73)	-17.56*** (5.29)	-7.64** (3.73)
HighRate (post-shock)	7.66* (4.20)		
Tight		16.05 (10.53)	-2.88 (8.90)
Plat		30.60*** (5.63)	-8.23 (5.69)
NonCyc × HighRate	32.73*** (6.90)		
NonCyc × Tight		13.45 (19.80)	41.86*** (14.45)
NonCyc × Plat		42.94*** (9.33)	30.40*** (7.18)
<i>UnempC (centred)</i>			-5.86*** (1.39)
<i>Maturity, log tranche size, yield-type</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Observations	5,247	5,247	5,247
R-squared	0.470	0.019	0.472

Sample of 5,247 tranches whose borrower industry is classified as cyclical or non-cyclical; the neutral category is excluded. Specifications and inference mirror those of Table 4. The preferred Specification 3 adds the contemporaneous euro-area 20 monthly unemployment rate (Eurostat series *une_rt_m*), centered on the in-sample mean, as a cycle control to absorb residual labour-market variation. The non-cyclical-by-regime interactions are essentially unchanged with or without UnempC ($NC \times Tight$ 41.66 $\rightarrow$ 41.86; $NC \times Plat$ 30.12 $\rightarrow$ 30.40), so the cycle control disciplines the rate-shock interpretation without altering the substantive result. A joint Wald F-test on the two interactions in Specification 3 delivers $F(2, 5236) = 11.7$ ($p < 0.001$). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7. Volume analysis by regime and industry classification (H2 secondary prediction).

Regime	Non-Cyclical	Cyclical	Neutral	All categories
<i>Panel A. Aggregate count of tranches</i>				
Pre-shock (90.58 months)	4,531	8,453	4,780	17,764
Tightening (15.15 months)	843	1,222	892	2,957
Plateau (27.27 months)	1,800	2,650	2,312	6,762
<i>Panel B. Per-month notional issuance (USD millions)</i>				
Pre-shock	28,970	44,181	24,321	97,472
Tightening	28,589	36,311	22,816	87,716
Plateau	35,845	52,963	41,421	130,229
<i>Panel C. Notional volume index (pre-shock = 100)</i>				
Pre-shock	100.0	100.0	100.0	100.0
Tightening	98.7	82.2	93.8	90.0
Plateau	123.7	119.9	170.3	133.6

Volume analysis on the full dataset of 27,491 tranches, restricted to the 27,483 tranches with a non-missing industry classification. Per-month figures normalize for unequal regime length as in Table 5. Cyclical notional issuance contracts by about 18 percent during the tightening, while non-cyclical contracts by only about 1 percent. The asymmetric contraction during the tightening identifies the elasticity asymmetry that underlies the H2 spread differential in Table 6: cyclical borrowers exercise the option to defer, while non-cyclical borrowers tied to regulator-mandated capex come to market regardless of the rate environment. Both groups rebound in the plateau.

Table 8. Effect of sponsor backing on the all-in drawn spread by regime.

Variable	Spec 1 <i>Binary, with controls</i>	Spec 2 <i>Three-regime, no controls</i>	Spec 3 (preferred) <i>Three-regime, with controls</i>
Sponsor	74.61*** (3.81)	161.74*** (3.73)	74.60*** (3.82)
HighRate (post-shock)	12.02*** (3.47)		
Tight		23.17*** (7.09)	14.91*** (5.46)
Plat		25.09*** (5.39)	9.94** (4.05)
Sponsor × HighRate	-4.00 (5.32)		
Sponsor × Tight		52.02*** (11.35)	66.12*** (10.28)
Sponsor × Plat		-40.68*** (6.77)	-16.19*** (5.70)
<i>Maturity, log tranche size, yield-type</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Observations	7,764	7,764	7,764
R-squared	0.517	0.259	0.523

Full clean spread sample ($N = 7,764$). Specification 1 uses the binary post-shock indicator and returns a sponsor-by-post-shock interaction that is statistically indistinguishable from zero, the average of two large opposite-signed regime-specific interactions recovered by Specification 3. A joint Wald F-test on the two sponsor interactions in Specification 3 delivers $F(2, 7754) = 29.0$ ($p < 0.001$). The orthogonality of the sponsor channel against H1 and H2 is examined in Table 10. HC1 robust standard errors in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 9. Structural premium reversal across the three hypotheses.

Hypothesis (treated vs comparison)	Pre-shock gap (β_1)	Tightening gap ($\beta_1 + \delta_t$)	Plateau gap ($\beta_1 + \delta_p$)	Plateau reversal (δ_p)
H1. Defensive vs Discretionary	-50.2	-11.9	+14.4	+64.6***
H2. Non-Cyclical vs Cyclical	-7.6	+34.2	+22.8	+30.4***
H3. Sponsor vs Non-Sponsored	+74.6	+140.7	+58.4	-16.2***

*The net structural premium in a given regime equals the structural coefficient β_1 from the Spec 3 preferred specification (Tables 4, 6 and 8) plus the regime-specific interaction; the H2 row uses the unemployment-augmented Spec 3 of Table 6. Positive values mean the treated group pays more than the comparison group. The reversal column reports the change in the structural premium between pre-shock and plateau, which is identical to the plateau interaction coefficient δ_p . H1 exhibits a sharp reversal: the structural defensive discount flips to a positive plateau premium. H2 shows reversal and persistence: pre-shock parity flips to a positive premium in tightening and stabilizes in plateau, a moderate-pace rate channel. H3 shows amplification followed by reversion: the structural sponsor premium widens sharply in tightening before reverting in plateau, consistent with the fast institutional channel where CLO and credit-fund yields re-rate within the tightening band and the slow relationship channel catches up in the plateau. Significance on the reversal coefficient: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (HC1 standard errors from the preferred specification).*

Table 10. Channel orthogonality: pairwise and joint specifications.

Variable	Spec 1 <i>Pairwise H1 × H2</i>	Spec 2 <i>Pairwise H1 × H3</i>	Spec 3 <i>Pairwise H2 × H3</i>	Spec 4 <i>Full joint</i>
Panel A. Channel level effects				
Defensive	-27.80*** (4.47)	-17.01*** (4.32)		-17.68*** (4.34)
NonCyc	-3.22 (3.47)		-6.46** (3.28)	-6.96** (3.30)
Sponsor		73.09*** (3.87)	74.58*** (3.81)	73.03*** (3.86)
Panel B. Channel × Regime interactions				
Defensive × Tight	25.57 (16.10)	32.91** (14.98)		34.75** (15.08)
Defensive × Plat	46.89*** (10.14)	42.91*** (10.07)		43.45*** (10.09)
NonCyc × Tight	25.99* (13.36)		25.72** (11.91)	27.03** (11.95)
NonCyc × Plat	34.09*** (6.61)		30.24*** (6.45)	30.71*** (6.44)
Sponsor × Tight		68.23*** (10.27)	66.24*** (10.24)	68.44*** (10.23)
Sponsor × Plat		-13.54** (5.69)	-18.73*** (5.74)	-16.06*** (5.73)
UnempC (centred)	-4.91*** (1.11)		-4.45*** (1.08)	-4.53*** (1.08)
Maturity, log size, yield-type	Yes	Yes	Yes	Yes
Observations	7,764	7,764	7,764	7,764
R-squared	0.482	0.525	0.526	0.528

Each pairwise specification stacks two channels in a single regression to verify that each channel identifies an independent margin of selective repricing. *UnempC* is included in every spec that contains the *H2* channel, consistent with the main *H2* specification (Table 6). All six channel × regime interactions retain the sign, the rough magnitude and the statistical significance reported in their respective main tables: *Defensive × Plat* moves from 64.6 to 43-47, *NonCyc × Tight* from 41.9 to 26-27, *Sponsor × Tight* from 66.1 to 66-68, *Sponsor × Plat* from -16.2 to -14 → -19. Joint Wald *F*-tests of the null that all interactions vanish reject in every specification ($F(4, 7750) = 12.3$ for Spec 1, $F(4, 7751) = 19.5$ for Spec 2, $F(4, 7750) = 20.5$ for Spec 3, $F(6, 7747) = 17.0$ for Spec 4; all $p < 0.001$). The three channels operate orthogonally rather than as substitute labels for the same underlying selection. HCl robust standard errors in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 11. Yield-type subsamples for H1: Investment Grade and Leveraged.

Variable	Investment Grade subsample	Leveraged subsample
Defensive	-25.79*** (4.69)	-51.49*** (7.75)
Tight	-5.03 (6.60)	33.96** (15.71)
Plat	11.04 (10.29)	-9.49 (7.39)
Defensive × Tight	33.98 (40.06)	15.87 (27.35)
Defensive × Plat	30.92* (16.93)	69.82*** (14.71)
<i>Maturity and log tranche size</i>	<i>Yes</i>	<i>Yes</i>
<i>Yield-type controls</i>	<i>Omitted (single yield type)</i>	<i>Omitted (single yield type)</i>
Observations	914	2,392
R-squared	0.146	0.113

*H1 Specification 3 estimated separately on the Investment Grade and Leveraged subsamples of the H1 sample. The Highly Leveraged bucket is too sparse on the defensive cell to estimate H1 separately. Yield-type dummies are dropped because each subsample is single-yield-type by construction. The defensive-by-plateau interaction is large and significant in the Leveraged subsample (about 70 bps) and marginally significant in the Investment Grade subsample (about 31 bps, $p = 0.07$). The within-tightening interaction is positive in both subsamples but imprecisely estimated; the slow demand-elasticity channel is identified primarily on the persistence of the differential into the plateau. The channel operates across yield types but is sharper in the leveraged segment, consistent with the bargaining-channel prediction. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

Table 12. Placebo tests on a synthetic 2018 to 2021 sample with a pseudo-shock at 21 July 2020.

Variable	Placebo H1 <i>Defensive vs Discretionary</i>	Placebo H2 <i>Non-cyclical vs Cyclical</i>	Placebo H3 <i>Sponsor vs Non-sponsored</i>
Structural (Def / NonCyc / Sponsor)	-44.27*** (9.28)	-16.85** (6.79)	80.25*** (7.00)
PsTight	9.61 (6.92)	14.10* (8.28)	28.58*** (6.09)
PsPlat	15.80 (12.83)	24.02 (16.12)	9.92 (13.02)
Structural × PsTight	22.11 (14.81)	22.62* (12.32)	-30.52*** (10.12)
Structural × PsPlat	-85.09*** (23.52)	-35.02 (24.84)	-7.65 (18.34)
<i>Maturity, log tranche size, yield-type</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Observations	1,423	1,714	2,534
R-squared	0.455	0.469	0.508

*The placebo sample restricts observations to financial closings between 1 January 2018 and 31 December 2021. A pseudo-shock is inserted at 21 July 2020 (parallel to the actual cycle, shifted back by two years); the pseudo-tightening regime spans 21 July 2020 to 25 October 2021 and the pseudo-plateau regime spans 26 October 2021 to the end of the placebo window. Specifications mirror the preferred three-regime regression for each hypothesis. The H2 placebo is broadly clean at the 5 percent level. The H3 placebo is a sharp negative null: sponsor-backed transactions repriced less during the pseudo-tightening, opposite in sign to the actual sponsor-by-tightening interaction in Table 8. The H1 placebo is partial and is discussed in the main text. HC1 robust standard errors in parentheses. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

Table 13. HC1 versus one-way borrower-clustered standard errors on the headline interactions.

Coefficient	Estimate	HC1 SE [p]	Cluster SE [p]	SE inflation
<i>H1 (Defensive vs Discretionary)</i>				
Def × HighRate (Spec 1, binary)	+55.88***	10.10 [< 0.001]	13.66 [< 0.001]	+35.3%
Def × Tight (Spec 3)	+38.24**	17.87 [0.032]	22.60 [0.091]	+26.5%
Def × Plat (Spec 3)	+64.56***	11.27 [< 0.001]	15.08 [< 0.001]	+33.7%
<i>H2 (Non-cyclical vs Cyclical)</i>				
NC × HighRate (Spec 1, binary)	+32.73***	6.90 [< 0.001]	10.53 [0.002]	+52.6%
NC × Tight (Spec 3)	+41.86***	14.45 [0.004]	19.49 [0.032]	+34.8%
NC × Plat (Spec 3)	+30.40***	7.18 [< 0.001]	11.17 [0.007]	+55.6%
<i>H3 (Sponsor vs Non-sponsor)</i>				
Sp × HighRate (Spec 1, diagnostic null)	-4.00 (n.s.)	5.32 [0.453]	7.75 [0.606]	+45.8%
Sp × Tight (Spec 3)	+66.12***	10.28 [< 0.001]	13.64 [< 0.001]	+32.6%
Sp × Plat (Spec 3)	-16.19***	5.70 [0.005]	8.15 [0.047]	+42.9%

*Specification 1 (binary) and Specification 3 (three-regime with controls; with UnempC for H2) re-estimated with one-way cluster-robust standard errors on the Issuer/Borrower PermID. The clean sample contains 2,571 clusters: 2,110 PermID groups for tranches issued by identifiable borrowers, plus 461 singleton clusters for tranches with unassigned PermID. Point estimates are identical to the HC1 specifications by construction; only standard errors and test statistics change. SE inflation is the percentage increase in the cluster standard error relative to the HC1 standard error. Two-way clustering on borrower and calendar quarter is not estimated because the tightening regime contains only fifteen months. The cross-channel and full-joint specifications are reported in Table 10 (Orthogonality). Significance on the Estimate column reflects the HC1 inference regime: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.*

Table 14. Country fixed effects on the preferred three-regime specification.

Variable	H1: Defensive vs Discretionary	H2: NonCyc vs Cyclical (UnempC)	H3: Sponsor vs Non-sponsored
Structural (Def / NonCyc / Sponsor)	-28.18*** (5.31)	-6.54* (3.83)	66.47*** (3.94)
Tight	9.45 (9.60)	-2.03 (8.74)	13.05** (5.44)
Plat	-11.80** (5.71)	-8.20 (5.73)	9.78** (3.96)
Structural × Tight	31.45* (18.04)	36.57** (14.34)	65.27*** (10.35)
Structural × Plat	48.08*** (11.26)	32.10*** (7.22)	-16.95*** (5.64)
<i>UnempC (centred)</i>		-5.64*** (1.40)	
<i>Maturity, log tranche size, yield-type</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
<i>Country fixed effects (UK omitted)</i>	<i>Yes (37)</i>	<i>Yes (35)</i>	<i>Yes (37)</i>
Observations	3,746	5,247	7,764
R-squared	0.492	0.489	0.544
<i>Partial F-test on country dummies</i>	<i>F(37, 3699) = 41.7</i>	<i>F(35, 5201) = 23.8</i>	<i>F(37, 7717) = 97.6</i>
<i>p-value</i>	<i>< 0.001</i>	<i>< 0.001</i>	<i>< 0.001</i>

The preferred Specification 3 of each hypothesis re-estimated with country fixed effects. The H2 column additionally retains the UnempC cycle control consistent with the main H2 specification (Table 6). The country variable is the Domicile Nation of the borrower as reported in the SDC extract. The clean sample contains tranches issued by borrowers domiciled in 38 nations; the United Kingdom (29.4 percent of the full sample) is the omitted reference category. H2 includes fewer country dummies because some nations have no cyclical or non-cyclical industry tranches in the H2 sample. HC1 robust standard errors in parentheses. The partial F-test in the bottom rows tests the joint nullity of all country dummies and rejects at the 1 percent level in every hypothesis. The regime channels survive the introduction of country fixed effects in every hypothesis; the only material attenuation is on the defensive-by-plateau interaction of H1 (64.6 to 48.1). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Table 15. Joint Wald F-tests across three inference regimes and channel orthogonality.

Hypothesis	Joint null	HC1 robust	Cluster (PermID)	Country FE
H1 (3,746 obs)	$Def \times Tight = Def \times Plat = 0$	$F(2, 3736) = 17.5^{***}$	$F(2, 1699) = 9.5^{***}$	$F(2, 3699) = 10.0^{***}$
H2 (5,247 obs)	$NC \times Tight = NC \times Plat = 0$	$F(2, 5236) = 11.7^{***}$	$F(2, 1811) = 5.1^{***}$	$F(2, 5201) = 11.8^{***}$
H3 (7,764 obs)	$Sp \times Tight = Sp \times Plat = 0$	$F(2, 7754) = 29.0^{***}$	$F(2, 2570) = 16.6^{***}$	$F(2, 7717) = 28.9^{***}$
Channel orthogonality (Table 10), full clean sample $N = 7,764$				
Pairwise H1×H2 (Spec 1)	All four H1+H2 interactions = 0	$F(4, 7750) = 12.3^{***}$	$F(4, 2570) = 5.2^{***}$	$F(4, 7713) = 8.9^{***}$
Pairwise H1×H3 (Spec 2)	All four H1+H3 interactions = 0	$F(4, 7751) = 19.5^{***}$	$F(4, 2570) = 12.0^{***}$	$F(4, 7714) = 17.4^{***}$
Pairwise H2×H3 (Spec 3)	All four H2+H3 interactions = 0	$F(4, 7750) = 20.5^{***}$	$F(4, 2570) = 12.2^{***}$	$F(4, 7713) = 20.3^{***}$
Full joint (Spec 4)	All six channel × regime interactions = 0	$F(6, 7747) = 17.0^{***}$	$F(6, 2570) = 10.4^{***}$	$F(6, 7710) = 15.5^{***}$

Joint Wald F-tests of the null that all channel × regime interaction coefficients are simultaneously zero, evaluated under three inference regimes: HC1 heteroscedasticity-robust standard errors, one-way cluster-robust standard errors on Issuer/Borrower PermID, and HC1 standard errors with country fixed effects added to the preferred specification. The H2 row uses the unemployment-augmented Spec 3 (Table 6). The top block tests each hypothesis in isolation; the bottom block (Channel orthogonality) tests the pairwise and full joint specifications of Table 10 to verify that the three channels identify independent margins. Every test rejects the joint null at the 1 percent level in every inference regime. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.


Figures

Figure 1. Monthly mean all-in drawn spread and the ECB deposit facility rate, 2015 to 2026.




Cross-tranche monthly mean all-in drawn spread of the clean spread sample ($N = 7,764$), shown as raw monthly values (light navy) and as a 6-month moving average (dark navy). The orange dashed line is the ECB deposit facility rate on the right axis. Background shading marks the three regimes: pre-shock (neutral), tightening (warm peach) and plateau (cool blue). Mean spreads track the policy stance closely from 2022 onwards: the spread rises from the pre-shock band around 280 bps to a tightening-band average above 316 bps, and remains elevated near 329 bps in the plateau even as the deposit facility rate is cut. Co-movement is visible but the spread does not retrace fully when the policy rate is cut, indicating persistence in the post-shock pricing structure.

Figure 2. Monthly mean all-in drawn spread by structural margin (use of proceeds, industry, sponsor backing), 2015 to 2026.



Three panels showing monthly mean spreads split by the three structural margins used in the H1, H2 and H3 contrasts. Each panel plots the raw monthly mean (light) and its 6-month moving average (bold) for each group. Panel A (use of proceeds) shows defensive spreads sitting far below discretionary spreads pre-shock and spiking upward during the tightening and remaining above them in the plateau. Panel B (industry) shows non-cyclical spreads crossing above cyclical spreads during the tightening and remaining elevated. Panel C (sponsor) shows sponsor spreads spiking sharply during the tightening before reverting in the plateau.

Figure 3. Quarterly notional volume index by use of proceeds, 2015Q1 to 2025Q4, base 2019Q1 = 100.




Quarterly USD-denominated notional issuance volume on the full dataset ($N = 27,491$ tranches after the financial-close filter, before any spread filter), indexed to 2019Q1 = 100. The raw quarterly index is shown faint and the 4-quarter moving average bold. Defensive (refinancing and working-capital) volume contracts noticeably during the tightening window then recovers, while discretionary (M&A and capex) volume contracts more sharply and rebounds higher in the plateau as deferred deals return. Note that the quarterly base in this figure (2019Q1 = 100) differs from the regime-average base in Table 5 Panel C (pre-shock average = 100); both convey the same qualitative dynamic. Per-month notional values in Table 5 quantify the pattern: defensive volume recovers to within 4 percent of the pre-shock benchmark while discretionary volume rebounds 12 percent above benchmark, consistent with deferred-discretionary issuance returning in the easing window.

Figure 4. Quarterly notional volume index by industry classification, 2015Q1 to 2025Q4, base 2019Q1 = 100.




Quarterly USD-denominated notional issuance volume by industry classification on the full dataset ($N = 27,483$ tranches with non-missing industry), indexed to 2019Q1 = 100. The raw quarterly index is shown faint and the 4-quarter moving average bold. Cyclical volume contracts visibly during the tightening band, while non-cyclical volume is essentially flat across the same window. Both groups rebound in the plateau. Note that the quarterly base in this figure (2019Q1 = 100) differs from the regime-average base in Table 7 Panel C (pre-shock average = 100), and that the 2019Q1 base values differ across categories (non-cyclical sits well above 100 throughout because its 2019Q1 quarter was a relatively low-volume quarter); within-series changes across regimes carry the message. Per-month notional values in Table 7 quantify the asymmetry: cyclical contracts by about 18 percent during the tightening while non-cyclical contracts by only about 1 percent.

Figure 5. Headline interaction coefficients with 95% confidence intervals (preferred Spec. 3).




Forest plot of the six headline interaction coefficients from the preferred three-regime specification of H1 (Table 4, dark navy), H2 (Table 6, medium blue, with the UnempC cycle control) and H3 (Table 8, orange). Horizontal bars are HC1 robust 95 percent confidence intervals. The H1 channel widens monotonically from tightening (+38 bps) to plateau (+65 bps), consistent with a slow demand-elasticity reset. The H2 channel jumps to a tightening peak (+42 bps) that moderates into a sustained plateau premium (+30 bps). The H3 channel reverses sign across regimes: sponsor-backed transactions reprice sharply during the tightening (+66 bps) and then partially reverse during the plateau (-16 bps). All six coefficients are statistically distinguishable from zero at the 5 percent level.

Figure 6. Structural premium across the three regimes by hypothesis (Specification 3 with 95% CI).



Structural premium for each hypothesis (treated minus comparison group) across the three regimes, with shaded 95 percent confidence intervals. The premium equals the structural coefficient β_1 from Specification 3 plus the regime-specific interaction; positive values mean the treated group pays more. The H2 panel uses the unemployment-augmented Spec 3 (Table 6). The grey dashed line is the binary pre-vs-post implied premium from Specification 1. Three channel signatures emerge: H1 shows a monotonic positive transition (slow demand elasticity), H2 shows reversal and partial persistence (moderate rate channel), and H3 shows amplification followed by reversion (fast institutional channel). The binary implied line masks the regime-specific dynamics in every channel.

Figure 7. Headline interaction coefficients across three inference regimes.



Same six headline interactions as Figure 5 re-estimated under three inference regimes: HC1 robust standard errors (dark navy circles), one-way cluster-robust standard errors on Issuer/Borrower PermID (medium blue squares) and HC1 standard errors with country fixed effects added to the preferred specification (orange triangles). Cluster-robust standard errors inflate the HC1 intervals by 27 to 56 percent (see Table 13) but every channel retains its qualitative reading. Country fixed effects modestly attenuate the H1 defensive-by-plateau interaction (+64.6 bps to +48.1 bps) while the H2 non-cyclical-by-plateau interaction tightens to +32.1 bps; the four remaining interactions are essentially unchanged. The H3 sponsor-by-plateau reversal of -16 bps is identified equally sharply in all three inference regimes.

Appendix A. Use of Proceeds Classification

Each tranche's SDC use-of-proceeds field is decomposed into atomic uses, and each atom is mapped to a defensive, discretionary or neutral bucket by the deferability test of Section 3.4. This appendix sets out the methodology and limitations (A.1), the atomic taxonomy (A.2), and the rule for combining atoms in multi-purpose tranches (A.3).

A.1. Methodology choice and limitations

The classification captures the deferability of the funded transaction: whether the borrower can postpone, downsize or abandon the spending without incurring distress. Because the SDC use-of-proceeds field describes the funded transaction directly, the deferability property can be read straight off the recorded purpose: refinancing maturing principal is non-deferrable on its face, an acquisition is timing-flexible on its face, and most atoms sort without ambiguity. An atom is coded defensive if it services an existing obligation or operating need on a fixed calendar, and discretionary if it funds a new decision whose timing the borrower controls. The few borderline atoms are resolved by the same test: recapitalization is defensive (capital-structure repair) but dividend recapitalization is discretionary (a deferrable shareholder distribution); project finance and construction are discretionary because their initiation is postponable; trade-finance atoms are defensive because they fund shipments already in motion. Because the buckets are populated overwhelmingly by their unambiguous core, the H1 contrast does not hinge on these marginal calls.

Two limitations remain. First, the silent-tag treatment of general corporate purpose (A.3) assumes the informative atom carries a mixed-tag tranche's economic content, which is reasonable but not guaranteed. Second, the SDC field is self-reported by the lead arranger and is occasionally coarse, so a recorded purpose may understate the true mix of uses; to the extent this measurement error is classical, it biases the H1 interactions toward zero rather than inflating them. The conservative combination rule mitigates both concerns by excluding genuinely mixed tranches as neutral rather than forcing them to one side, at the cost of sample size.

A.2. Atomic taxonomy

Defensive atoms (D). *Refinancing; Refinance / Retiring Bank Debt; Working Capital; Restructuring; Recapitalization* (excluding dividend recapitalization, which is coded discretionary); *Standby and Commercial Paper Support; Debtor in Possession; Exit Financing; Export and Import Financing; Finance Linked-Trade*. These atoms share the property that the borrower must come to market on a fixed schedule

that cannot be postponed without materially raising operating or financial risk. Refinancing and refinance / retiring bank debt are tied to maturing facilities. Working capital is tied to the ongoing operational cycle. Restructuring and recapitalization are tied to capital-structure repair. Standby and commercial paper support is backup liquidity for outstanding commercial paper. Debtor-in-possession and exit financing are bankruptcy-driven. Trade finance atoms are tied to specific shipments already in motion.

Discretionary atoms (X). *Acquisition Finance; Leveraged Buyout; Sponsored Buyout; Management Buyout; Infrastructure LBO; Real Estate and Property Acquisition; Project Finance; Capital Expenditures; Construction; Spinoff; Dividend Recapitalization; Stock Repurchase; Proceed to Shareholders; Ship Financing; Aircraft Financing; Leases; Mortgage Financing.* These atoms share the property that the borrower retains the option to delay, scale down, or walk away from the underlying transaction. Acquisition financings, buyouts and dividend recapitalizations are timing-flexible by nature. Capital expenditure, project finance and construction are postponable. Asset-acquisition financings (ships, aircraft, real estate, equipment leases, mortgages) involve discretionary timing of the underlying purchase. Stock repurchases and dividend recapitalizations are shareholder distributions that can be deferred.

Neutral atoms (N). *General Corporate Purpose; Other.* These atoms are informationally empty: they are silent tags in the combination rule and do not change the bucket assignment of an otherwise informative tranche.

A.3. Combination rule for multi-purpose tranches

A tranche is coded defensive if all its informative atoms are defensive, discretionary if all are discretionary, and neutral otherwise, with general corporate purpose and other treated as silent tags that do not change an informative tranche's assignment. Tranches mixing defensive and discretionary atoms are coded neutral as genuinely mixed-purpose, and tranches carrying only neutral atoms are uninformative; both are excluded from the H1 contrast. The silent-tag rule is what lets the combination recover operational content: without it, the general corporate purpose tag would reclassify a large share of otherwise-informative tranches as neutral and weaken the contrast. Applied to the clean spread sample, the rule yields 981 defensive and 2,765 discretionary tranches, with the remaining 4,018 neutral and excluded (Table 3, Panel A).

Appendix B. Industry Cyclical Classification

This appendix reproduces the bucket lists applied in the construction of the non-cyclical, cyclical and neutral industry classifications used in H2. Industries are at the TRBC Activity level (the level reported in the SDC Platinum extract). The non-cyclical bucket comprises 26 industries, the cyclical bucket 76, and the neutral bucket 41. The classification is based on the demand-elasticity test outlined in Section 4.2: non-cyclical industries face capital-expenditure and working-capital schedules dictated by regulation, contractual replacement cycles or non-discretionary demand, while cyclical industries can postpone investment, run down inventory or rephase project finance.

B.1. Methodology choice and limitations

The classification of industries into non-cyclical, cyclical and neutral buckets reflects a substantive methodological choice, summarized in Section 3.4 and set out in full here. An ideal classification would be data-driven. One natural candidate would be industry betas to GDP, computed on Compustat or a European equivalent, on the logic that low-GDP-beta industries are also more likely to face non-deferrable borrowing needs. We do not pursue this route for two reasons.

First, the data alternative is materially weaker for European borrowers than for U.S. borrowers. Compustat coverage of private European firms is limited, and Worldscope coverage is inconsistent across the 38 borrower domicile countries in our sample. A data-driven classification would therefore introduce its own coverage selection on top of the SDC spread-disclosure selection already discussed in Section 4.6.

Second, and more substantively, the object we want to capture is not the cyclicity of firm earnings but the deferability of investment. The two concepts overlap but are not identical, and where they diverge it is deferability that our test tracks. The water-utility-versus-tobacco contrast drawn in Section 3.4 is the clean case, but the divergence runs in both directions: a homebuilder and a semiconductor fabricator both have highly cyclical earnings, yet a fab mid-construction faces switching and obsolescence costs that make its capex far harder to defer than a homebuilder's land bank, so earnings cyclicity alone would misrank them on the property H2 cares about. A GDP-beta measure keys on the earnings dimension and would therefore mismeasure the deferability concept H2 is designed to test, in addition to the European coverage problem noted above.

The hand-coded taxonomy in this appendix applies the deferability test directly at the industry level. Borderline cases (financial-services subsectors whose cyclicity depends on the credit they extend, real-estate subsectors whose cyclicity varies by asset class and geography, conglomerates and diversified

holdings whose business mix straddles cyclical and non-cyclical activities, and a small number of activity codes too aggregated to assign cleanly) are excluded as neutral rather than forced into one bucket. We acknowledge two limitations of this choice. First, the bucket assignments embed judgment at the margins, and a different researcher applying the same deferability test would likely produce a slightly different bucket list. Second, the taxonomy is not robust to alternative interpretations of the H2 effect: sector-specific exposures during 2022 to 2023 (energy prices, regulation, supply-chain stress) may correlate with our classifications in ways the hand-coded approach cannot disentangle, a caveat we flag in Section 5.2. Future work could test the H2 reading against an external data-driven classification, conditional on European industry-beta data becoming more comprehensively available.

B.2. Non-cyclical industries (26)

Advanced Medical Equipment and Technology; Airport Operators and Services; Biotechnology and Medical Research; Brewers; Distillers and Wineries; Drug Retailers; Electric Utilities; Fishing and Farming; Food Processing; Food Retail and Distribution; Healthcare Facilities and Services; Highways and Rail Tracks; Household Products; Independent Power Producers; Integrated Telecommunications Services; Marine Port Services; Medical Equipment, Supplies and Distribution; Multiline Utilities; Natural Gas Utilities; Non-Alcoholic Beverages; Oil and Gas Transportation Services; Personal Products; Pharmaceuticals; Tobacco; Water and Related Utilities; Wireless Telecommunications Services.

B.3. Cyclical industries (76)

Advertising and Marketing; Agricultural Chemicals; Airlines; Aluminium; Apparel and Accessories; Apparel and Accessories Retailers; Appliances, Tools and Housewares; Auto and Truck Manufacturers; Auto Vehicles, Parts and Service Retailers; Auto, Truck and Motorcycle Parts; Broadcasting; Casinos and Gaming; Coal; Commercial Printing Services; Commodity Chemicals; Communications and Networking; Computer and Electronics Retailers; Computer Hardware; Construction and Engineering; Construction Materials; Construction Supplies and Fixtures; Consumer Publishing; Courier, Postal, Air Freight and Land-Based Logistics; Department Stores; Discount Stores; Diversified Chemicals; Diversified Mining; Electrical Components and Equipment; Electronic Equipment and Parts; Entertainment Production; Footwear; Forest and Wood Products; Gold; Ground Freight and Logistics; Heavy Electrical Equipment; Heavy Machinery and Vehicles; Home Furnishings; Home Furnishings Retailers; Home Improvement Products and Services Retailers; Homebuilding; Hotels, Motels and Cruise Lines; Household Electronics; IT Services and Consulting; Industrial Machinery and Equipment; Integrated Hardware and Software; Iron and Steel; Leisure and Recreation; Marine Freight and Logistics; Mining Support Services and

Equipment; Miscellaneous Specialty Retailers; Non-Gold Precious Metals and Minerals; Non-Paper Containers and Packaging; Office Equipment; Oil and Gas Drilling; Oil and Gas Exploration and Production; Oil and Gas Refining and Marketing; Oil Related Services and Equipment; Online Services; Paper Packaging; Paper Products; Passenger Transportation, Ground and Sea; Phones and Handheld Devices; Real Estate Services; Recreational Products; Renewable Energy Equipment and Services; Renewable Fuels; Restaurants and Bars; Semiconductor Equipment and Testing; Semiconductors; Shipbuilding; Software; Specialty Chemicals; Specialty Mining and Metals; Textiles and Leather Goods; Tires and Rubber Products; Toys and Children Products.

B.4. Neutral industries (41)

Aerospace and Defense; Banks; Business Support Services; Business Support Supplies; Charity Organisations; Civic and Social Organisations; Closed End Funds; Commercial REITs; Consumer Goods Conglomerates; Consumer Lending; Corporate Financial Services; Diversified Industrial Goods Wholesale; Diversified Investment Services; Diversified REITs; Employment Services; Environmental Services and Equipment; Financial and Commodity Market Operators and Service Providers; Financial Technology (Fintech); Government and Government Finance; Government Administration Activities; Integrated Oil and Gas; Investment Banking and Brokerage Services; Investment Holding Companies; Investment Management and Fund Operators; Life and Health Insurance; Miscellaneous Educational Service Providers; Miscellaneous Fintech Infrastructure; Multiline Insurance and Brokers; Mutual Funds; National Security and International Affairs; Pension Funds; Personal Services; Professional and Business Education; Professional Information Services; Property and Casualty Insurance; Real Estate Rental, Development and Operations; Reinsurance; Residential REITs; Schools, Colleges and Universities; Specialised REITs; UK Investment Trusts.

The neutral bucket comprises industries whose cyclicality is genuinely ambiguous: financial-services subsectors whose cyclicality depends on the type of credit they extend; real-estate operating and development subsectors whose cyclicality varies sharply with the asset class and the geography; conglomerates and diversified holding structures whose underlying business mix straddles cyclical and non-cyclical activities; and a small number of activity codes that are too aggregated to receive a clean assignment. The H2 contrast in Section 5.2.1 excludes the neutral bucket; the cyclical-versus-non-cyclical contrast is recovered on the 5,247-tranche subsample.