Bridging Two Markets: Dealer Intermediation across Securities Lending and Repo

May 2026

Georgios Angelis Alexiou*

June 15, 2026

Abstract

Repo and securities-lending markets trade the same sovereign bonds through separate participant universes, bridged by dealers who borrow bonds from beneficial owners and redeploy them as repo collateral. I model the lending fee as the dealer's repo deployment value: the bond's repo specialness, scaled by its redeployment intensity and split across a lending and a repo bargain. Lending fees therefore pass through repo specialness, most strongly for repo-central bonds when lending utilization is high. In matched securities-lending and repo data for euro-area sovereigns (2018–2024), the pass-through concentrates there and is muted elsewhere; bonds physically cross into repo, and fees, specialness, and repo volume comove. Auction reopenings confirm specialness is a supply-driven repo price. A lending fee is, in effect, a repo price.

Keywords: Securities lending, repo markets, search and matching, dealer intermediation, collateral reuse, European sovereign bonds.

JEL Codes: G12, G23, G28, E44.

*Goethe University Frankfurt and European Central Bank. Email: angelisa@econ.uni-frankfurt.de. The views expressed are those of the author and do not necessarily reflect those of the ECB.

1 Introduction

The European repo market intermediates EUR 13.7 trillion in outstanding contracts at any point in time (EUR 13,651 billion on December 10, 2025, the most recent ICMA survey date; ICMA, 2026).[1] The parallel securities-lending market, on the same underlying sovereign bonds, intermediates hundreds of billions of euros on loan against more than EUR 800 billion in lendable inventory (Aggarwal et al., 2021). These two markets are central to collateralized funding, short selling, settlement, safe-asset allocation, and monetary-policy transmission. Yet they are usually studied separately. Securities-lending research studies short-sale demand, lendable supply, and lending fees. Repo research studies funding conditions, collateral scarcity, and specialness. The dealer balance sheet that bridges the two markets is much less visible.

This paper studies that bridge. A dealer that borrows a sovereign bond from a beneficial owner can also deploy the same bond as collateral in repo. If the bond is special in repo, the borrowed bond has an outside deployment value. The central question is whether that deployment value is reflected in the securities-lending fee, and when the pass-through is strongest.

Three facts motivate the mechanism. First, European repo and securities-lending markets are institutionally segmented even when they trade the same collateral. Repo activity is split across CCP-cleared interdealer platforms such as BrokerTec, MTS, and Eurex, triparty arrangements, and bilateral OTC trading. Securities lending is bilateral OTC. Matching efficiency, counterparty access, settlement conventions, collateral schedules, and mandate constraints differ across the two markets. These frictions make the dealer's ability to connect the markets economically meaningful. Second, Angelis Alexiou et al. (2025) show that European repo reuse chains are common but do not look like simple haircut-wedge rent extraction. Haircut and rate differentials move together across reuse legs rather than substituting, which weakens the view that reuse is mainly a liquidity-windfall arbitrage. Third, Aggarwal et al. (2021) document large European government-bond lending markets in which demand for safe bonds rises during collateral-scarcity episodes, but their setting does not match lending to repo at the ISIN-day level. These facts suggest a different interpretation of reuse and intermediation: dealers may borrow bonds because they can deploy the same scarce collateral in repo, not because reuse mechanically produces a haircut rent.

I formalize this idea in a static two-market Nash-bargaining model, motivated by the

[1] The ICMA figure is the sum of outstanding repos and reverse repos on the books of survey participants. Two-sided reporting creates some mechanical double-counting, which ICMA (2026) estimates to be below 30% using SFTR-regulator evidence. The survey is still a lower bound on total market size because it covers a sample of participants. Europe-wide SFTR-reported outstanding on the nearest reporting date was EUR 25.9 trillion for the EU plus UK combined.

OTC search-and-bargaining environments of Duffie et al. (2002, 2005). Conditional on a securities-lending match, the lending fee splits surplus between the beneficial owner and the dealer. The dealer's surplus is a repo deployment value for the same bond—its repo specialness, scaled by how intensively the dealer can redeploy it as collateral and shared with the repo counterparty in a separate repo bargain. When repo specialness is positive, this deployment value enters the lending bargain. The model predicts that positive specialness passes into lending fees most strongly for bonds with high predetermined repo depth when lending utilization is high. The empirical tests rest on this bargaining prediction.

The empirical analysis uses matched ISIN-day securities-lending and repo data for euro-area sovereign bonds from 2018 to 2024, together with the German, French, and Italian auction calendars. The headline mechanism test is the repo-to-lending pass-through, graded by the dealer's repo deployment. Splitting bonds at the median of a predetermined repo-depth centrality measure—built from each bond's first quarter of observed life—the state-dependent specialness pass-through (positive specialness interacted with the high-utilization state) is large and significant for repo-central bonds and statistically indistinguishable from zero for the rest, with a between-group contrast of +13.9 ($t = +2.5$) on the German and Italian post-issuance panel; it concentrates in the positive-specialness region where the model says the deployment channel operates. This contrast is the structural cross-partial of the model. The demand-side counterpart—lending fees respond to demand shocks (lnSLQ) only in the high-utilization state (+2.15 bps, $t = +5.90$, $N = 728{,}502$; a one-standard-deviation shock moves the fee by 20–45% of the same bond's bid-ask spread)—is the reduced form of the same channel and, placed head-to-head against the specialness state interaction, is dominated by it. A cross-market simultaneity test then shows that the same high-utilization lending-demand state moves lending fees, repo specialness, and repo volume for the same bond-day, and a bridge regression confirms that securities-lending quantity predicts repo volume—the physical crossing of bonds from lending into repo.

Two quasi-experiments identify the two halves of the channel. On the repo side, Treasury auction reopenings—exogenous, calendar-set increases in a bond's collateral supply—relieve its repo specialness in Germany, France, and Italy alike (each with $t < -6$), with a dose-response in pre-auction scarcity and null placebos at pseudo-event dates, confirming that specialness is a genuine supply-driven repo price rather than a passive correlate of the fee. On the fee side, the ECB Securities Lending Facility (SLF) and the July 2023 quantitative-tightening onset identify the lending-fee response: for Italian BTPs, SLF-listed bonds show a state-dependent fee response of +6.45 bps in the high-utilization state after QT begins ($t = 3.56$, wild-cluster-bootstrap $p = 0.001$), with flat pre-trends. This is the paper's identified causal estimate of a state-dependent lending-fee response.

The same response is detectable in Germany and France with opposite signs; because the SLF/QT shock bundles a lending-supply and a repo-deployment margin, I read the cross-country sign variation as ruling out a common Europe-wide confound—opposite signs cannot be produced by a single pan-European shock—rather than as a structural prediction, and the German cell, which fails parallel-trends, is treated as corroborative.

The scope is deliberately precise. The centrality evidence identifies which bonds load the deployment-value channel; because lead and lag specialness behave similarly, it is not a timing-causal estimate of daily specialness. The channel is documented most strongly in dealer-active German and Italian collateral markets, and broad euro-area replication is not established in the present data. The Italian DiD also identifies a composite policy treatment, combining lending-supply and repo-deployment margins. A state-dependent single-market supply-elasticity model can explain part of the same pattern, so the case for the deployment-value channel comes from the joint evidence across repo specialness, repo depth, utilization, cross-market outcomes, and the Italian policy shock.

The paper contributes along four dimensions. It links repo and securities lending through dealer deployment incentives for the same sovereign bond. It brings the Duffie, Gârleanu, and Pedersen search-and-bargaining logic to a two-market same-asset setting in which the securities-lending bargain inherits the dealer's repo deployment value. It reframes collateral reuse in this market as the trace of cross-market bridging rather than as a primitive haircut-wedge windfall, complementing and sharpening the liquidity-windfall benchmark of Infante (2019). It also provides matched ISIN-day evidence that the repo-to-lending fee pass-through is state-dependent, concentrated in positive specialness and high utilization, and strongest for bonds with high predetermined repo depth.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 documents the institutional setting in which beneficial owners populate the lending side and dealers are the unique cross-market intermediaries, and reports four stylized facts on the joint behaviour of the two markets that motivate the cross-market analysis. Section 4 develops the static two-market search-and-bargaining model. Section 5 presents the empirical evidence. Section 6 concludes.

2 Related Literature

2.1 Securities Lending

The empirical securities-lending literature shows that lending fees reflect demand, supply, and search frictions. D'Avolio (2002) provide the foundational description of the U.S. equity-lending market, where most stocks are cheap to borrow but a small set of specials carries high borrowing costs. Cohen et al. (2007) decompose lending-market movements

into supply and demand shifts and show that demand-driven shorting predicts negative returns. Kolasinski et al. (2013) emphasize search frictions, showing that fees respond more strongly to demand where lending-market search is harder. Saffi and Sigurdsson (2011) connect lending supply to market efficiency and price discovery.

For sovereign bonds, Aggarwal et al. (2021) is the closest benchmark. Using Markit data on European government-bond lending from 2006 to 2017, they document large lendable inventories, substantial on-loan values, and higher safe-bond borrowing demand in collateral-scarcity episodes. They also show that the ECB's PSPP lending program reduced private lending demand. I build on this work by adding the repo side of the same bonds at the ISIN-day level. The question here is not only whether lending fees reflect borrowing demand, but whether the onward repo use of a borrowed sovereign bond is priced in the lending bargain.

2.2 Repo Markets, Collateral Scarcity, and Public Lending Facilities

The repo literature studies how collateral scarcity and funding conditions affect secured money markets. Duffie (1996) provides the canonical theory of repo specialness: a bond in high collateral demand trades at a below-GC repo rate. D'Amico et al. (2018) quantify Treasury specialness and show that security-specific supply and demand matter for repo rates. In Europe, Arrata et al. (2020) show that ECB asset purchases created bond-specific scarcity in repo markets, and Corradin and Maddaloni (2020) document the role of safe-asset scarcity in European sovereign repo during the debt crisis. Pelizzon et al. (2025) connect ECB purchases, repo funding, and futures delivery through euro-area sovereign cash-futures dislocations.

Public securities-lending facilities are especially relevant for this paper. Carrera de Souza and Hudepohl (2024) show that Eurosystem purchases depress repo rates through flow and stock effects, while securities lending partly offsets scarcity. Greppmair and Jank (2026) use the November 2020 Eurosystem SLF pricing reduction as a quasi-natural experiment and show that cheaper public collateral supply expanded repo activity and improved market quality. I use this institutional setting differently. The object of interest is the private securities-lending fee. When the public sector changes the supply of lendable collateral, does that change pass through to private lending fees in the state where repo deployment value matters?

2.3 Collateral Reuse and Intermediation Chains

Collateral reuse is often treated as a contractual or balance-sheet multiplier. Infante (2019) formalizes the liquidity-windfall view, where dealers receive collateral under high

haircuts and redeploy it under lower haircuts. Singh (2014) emphasizes collateral velocity and the macroeconomic importance of collateral chains, while Brumm et al. (2023) studies reuse in a general-equilibrium setting.

The European sovereign setting points to a different active margin. Angelis Alexiou et al. (2025) use transaction-level repo data and find little support for the key haircut-wedge prediction of the liquidity-windfall view. Inbound and outbound haircut differentials are not systematically positive, and haircut and rate differentials move together rather than substituting. This does not mean reuse is irrelevant. It means that the rent-extraction interpretation is not sufficient. In this paper, reuse is an observable trace of dealers bridging securities lending and repo. A dealer borrows a bond from a beneficial owner and can redeploy it in repo when the bond is special. The relevant object is therefore the dealer's repo deployment value for the same bond, not a standalone haircut wedge.

2.4 Intermediary Capital and OTC Search

The mechanism also relates to intermediary asset pricing. He and Krishnamurthy (2013), Brunnermeier and Pedersen (2009), He et al. (2022), and Du et al. (2018) show that dealer balance-sheet capacity shapes prices, liquidity premia, and arbitrage wedges. Hendershott and Menkveld (2014) provides direct microstructure evidence that intermediary inventory state affects prices. These papers study how dealer constraints matter for prices within a market or asset class. I study a same-bond cross-market version: repo conditions affect securities-lending fees because the dealer who borrows the bond also values its repo deployment option.

The model is motivated by OTC search-and-bargaining environments. Duffie et al. (2002, 2005) show how search frictions in decentralized asset markets generate bid-ask spreads, price dispersion, and intermediation rents. Lagos and Rocheteau (2009), Lagos et al. (2011), Vayanos and Weill (2008), Afonso and Lagos (2015), and Atkeson et al. (2015) develop related search models of liquidity, crises, on-the-run premia, federal funds trading, and dealer-customer market structure. The contribution here is to use that search-and-bargaining logic in a two-market same-asset setting. Conditional on a securities-lending match, the lending fee splits the joint surplus from the loan, and the dealer's per-match surplus includes the repo deployment rent earned by re-posting the borrowed bond as collateral. This deployment rent enters the dealer's surplus inside the bargain, not as an outside option, and is the source of the paper's state-dependent pass-through prediction.

3 Institutional Setting and Stylized Facts

The model in Section 4 takes as given that the securities-lending and repo markets for the same euro-area sovereign bond are populated by structurally distinct sets of principals, with the dealer as the unique cross-market intermediary. This section documents that institutional structure, the macro mechanism that drives the high-intermediation regime ($H_{it} = 1$) over the sample, and four stylized facts on the joint behaviour of the two markets. Together the material gives the reader a direct empirical case for studying the two markets jointly, before any theory enters.

3.1 Distinct principals, the dealer as the bridge

The securities-lending and repo markets for the same euro-area sovereign bond are populated by structurally distinct principals. The lending side is dominated by long-only beneficial owners—pension funds, insurers, UCITS and other collective investment vehicles, and official-sector portfolios—which together supply the bulk of government-bond lending inventory (ISLA, 2025); European government-bond lending balances ran at roughly 380–440 billion through 2025, almost entirely against non-cash (bond) collateral. The repo side is the dealer's domain: the ICMA semi-annual survey, covering 10.9 trillion of outstanding European repo, lists *only* dealer banks, debt-management agencies, multilateral institutions, and central securities depositories as direct participants—no pension fund, insurer, mutual fund, or hedge fund appears (ICMA, 2024). This sorting is economic, not statutory: long-only owners lend because securities lending is a low-touch yield overlay on bonds they already hold, whereas direct repo intermediation is a scale-, balance-sheet-, margin-, and clearing-intensive activity that EU fund, insurance, and clearing rules make costly for them rather than prohibit. Even in the September 2022 UK liability-driven-investment stress, when pension-linked vehicles borrowed heavily in repo, they offloaded through primary dealers rather than client-to-client (Pinter, 2023). The dealer is therefore the *unique* cross-market intermediary—the only agent with simultaneous principal access to securities lending (through an agent lender or the Eurosystem securities-lending facility) and to repo (bilateral, triparty, or CCP-cleared). That dual access is the institutional source of the cross-market deployment-value channel: a borrowed bond can be redeployed as repo collateral, so the lending fee can inherit repo specialness even though the beneficial owner never touches the repo leg. The model formalises this through the deployment loading $\Lambda(C_i, H_{it})$, consolidating the beneficial owner and its agent lender into a single lending-side principal.

3.2 Eurosystem holdings and the lending-supply pool

The most relevant macroeconomic driver of the cross-market intermediation regime over the sample is the cycle of Eurosystem asset purchases (PSPP and PEPP) and their subsequent quantitative-tightening runoff. Eurosystem purchases drain the country-level free float (total outstanding net of Eurosystem holdings) from which the Markit lendable inventory is mechanically drawn. The natural structural object is therefore the free float, where total outstanding is taken from the ECB's Centralised Securities Database (CSEC, monthly, in EUR mn), the source consistent with the Eurosystem-holdings denominator. The structural specification is

$$\text{lendable}_{c,t} \;=\; \alpha_c \;+\; \lambda_c \left(\text{CSEC outstanding}_{c,t} \;-\; \text{ECB holdings}_{c,t}\right) \;+\; \varepsilon_{c,t}, \qquad (1)$$

where λ_c is the country-specific fraction of free float that appears in the Markit lendable pool.

Table 1: Structural free-float regression on Markit lendable. Monthly 2020-12 to 2024-11, $N = 48$ per country. λ_c is the fraction of country-c free-float central-government debt that appears in Markit lendable inventory; t-statistics in parentheses. $\hat{\rho}$(free float, lendable) reported for comparison.

	DE	FR	IT	ES	BE	NL	AT
$\hat{\lambda}_c$	+0.056	+0.099	+0.078	+0.105	+0.097	+0.198	+0.245
	(+2.73)	(+6.24)	(+6.62)	(+13.94)	(+16.07)	(+14.43)	(+5.43)
R^2	0.14	0.46	0.49	0.81	0.85	0.82	0.39
$\hat{\rho}$(free float, lendable)	+0.37	+0.68	+0.70	+0.90	+0.92	+0.91	+0.63

Notes: Country-by-country OLS time-series regressions of equation (1); no fixed effects, $N = 48$ monthly observations per country (December 2020–November 2024). Dependent variable: Markit lendable inventory (EUR mn). Cells report $\hat{\lambda}_c$ with t-statistics in parentheses; significance is read from the t-statistics. $\hat{\rho}$ is the Pearson correlation of country free float and Markit lendable. Sources: Markit lendable inventory; ECB Centralised Securities Database outstanding; Eurosystem holdings.

Table 1 reports the estimates. For the major sovereigns the fit is tight (R^2 between 0.46 and 0.85 for FR through NL). The economic interpretation is direct: roughly 5 to 10% of free-float central-government debt shows up as Markit lendable inventory in DE/FR/IT/ES/BE, and 20% to 25% in NL/AT. Germany's lower fraction reflects that Bunds are disproportionately held as HQLA reserve assets by banks and CCPs that do not participate in lending programmes. Holding λ_c fixed, every 100 bn of additional Eurosystem holdings reduces Markit lendable inventory by 5.6 bn (DE), 9.9 bn (FR), 7.8 bn (IT), 10.5 bn (ES), and 9.7 bn (BE). A partial regression that enters Eurosystem holdings, cumulative gross issuance, and cumulative redemption separately (-mn levels, country by country) confirms the supply-drain direction in all three core countries (DE $\hat{\beta}_{\text{ECB}} = -0.135$, $t = -6.69$; FR -0.063, $t = -2.02$; IT -0.067, $t = -3.09$) with cumulative issuance entering at +0.34 to +0.42 and cumulative redemption with the opposite sign.

A within-bond decomposition of cumulative utilisation between January 2018 and November 2024 separates the demand-side component ($\Delta\text{VOL}/\overline{\text{Lend}}$) from the supply-side component ($-\overline{\text{VOL}} \cdot \Delta\text{Lend}/\overline{\text{Lend}}^2$), value-weighted across bonds. In Germany, cumulative within-bond $\Delta\text{Util} = +26.9$ pp decomposes into a demand contribution of -12.3 pp (existing Bunds saw value on loan fall) and a supply contribution of $+17.9$ pp (existing Bunds saw lendable inventory contract): the German within-bond utilisation rise is supply-driven, with the aggregate held pinned by new issuance entering at low utilisation. In Italy, cumulative within-bond $\Delta\text{Util} = +28.0$ pp decomposes into a demand contribution of $+19.1$ pp ($\sim$60%) and a supply contribution of $+12.7$ pp ($\sim$40%): Italian utilisation rose under both a supply contraction and a concurrent demand expansion. In both Germany and Italy, the within-bond utilisation rise pushes bonds into the high-utilisation regime $D_{it}^{U_{75}} = 1$ that activates the cross-market deployment-value channel in the model of Section 4. In the eurozone-core sovereigns where demand was stable or contracting (FR, NL, ES, BE, AT), the same supply contraction was offset at the country level, and aggregate utilisation moved less. The activation regime $H_{it} = 1$ is therefore not an exogenous market-microstructure feature: it is the empirical signature of central-bank asset purchases meeting bond-level collateral demand on the lending market. As the QT runoff continues into 2025 and beyond, the model predicts that utilisation in DE and IT will fall, attenuating the empirical magnitude of the state-dependent fee-demand interaction in subsequent samples.

3.3 Four stylized facts on the joint behaviour of lending and repo

I close the institutional section with four stylized facts on the panel of 1,536 German and Italian sovereign bonds over 2018–2024 that make the cross-market link directly visible in the data, before any model.

Fact 1. Lending fees and repo specialness co-move at low frequency. Figure 1 reports country-month means of the Markit indicative lending fee against repo specialness (the deposit-facility rate minus the bond-specific repo rate). For German Bunds, the unconditional monthly correlation is $\hat{\rho} = +0.78$, with $R^2 = 0.61$ in an OLS regression of fee on specialness on 84 monthly observations. The same scatter for Italian BTPs is positive but flatter; the model-relevant Italian signal sits in the within-bond, state-dependent margin and is averaged out at the country-month aggregate. The joint message is that the two prices are not independent: the cost of borrowing the security and the rent the borrower extracts in repo move together at country-month level, exactly as a joint-bargaining model would predict.

Fact 2. Eurosystem ownership and lending supply move in opposite directions. Figure 2 reports the country-by-country relationship between the Eurosystem share of

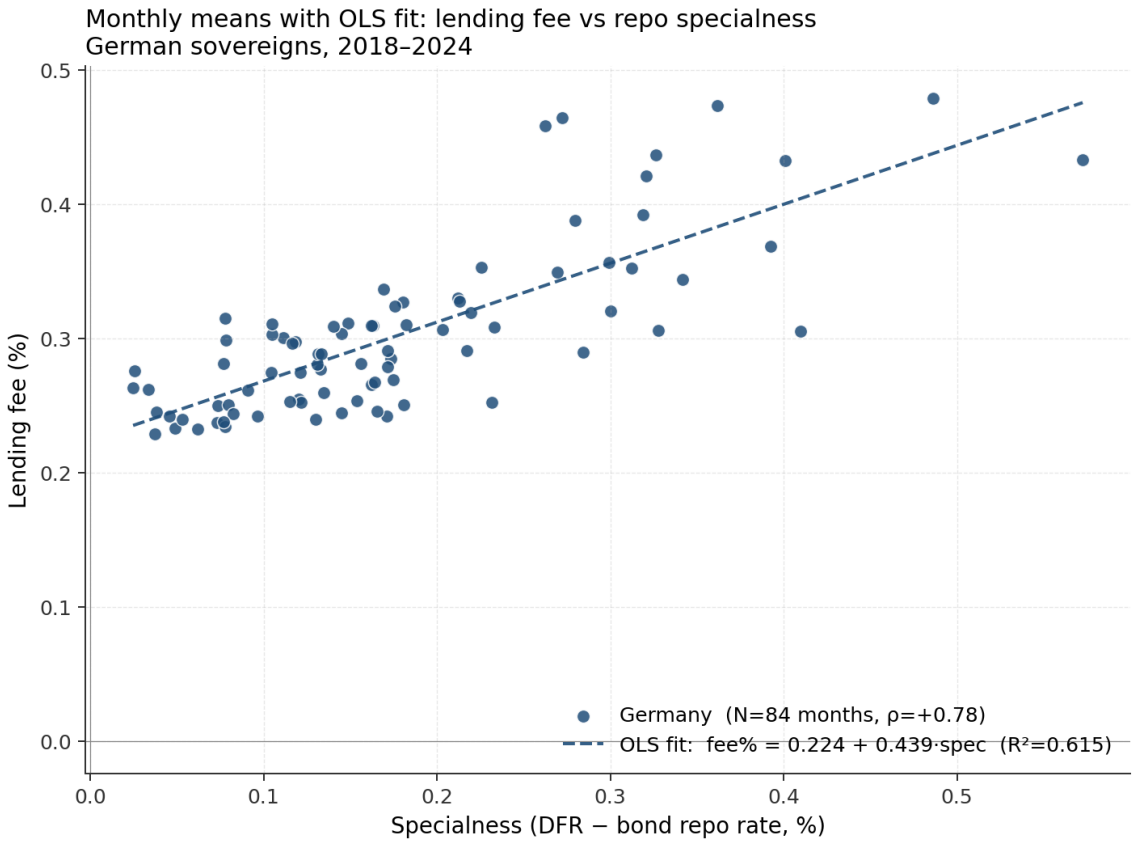


Figure 1: Lending fee versus repo specialness, German Bunds, country-month means, January 2018 to December 2024 (84 months). y-axis: monthly mean of the Markit indicative lending fee (%). x-axis: monthly mean of repo specialness (DFR $- r_{it}^{\text{repo}}$, %). Solid line: OLS fit, slope +0.44, $t = +11.4$, $R^2 = 0.61$. $\hat{\rho} = +0.78$. Quarter-end windows excluded. Source: S&P Global Securities Finance (fees), BrokerTec/MTS (specialness).

central-government outstanding (%, left axis) and the Markit lendable inventory (bn EUR, right axis), monthly, for the ten EU sovereigns with non-trivial Markit coverage. The two series move in opposite directions across all six large issuers, with unconditional cross-country correlations of share and lendable equal to −0.40 (DE), −0.61 (FR), −0.62 (IT), −0.81 (ES), −0.88 (BE), and −0.96 (NL). Phase-conditional correlations sharpen further in the QT-runoff phase from mid-2023 onward, reaching −0.97 in France, −0.93 in Spain, −0.84 in Germany, and −0.82 in Italy. Eurosystem purchases therefore do not simply price collateral demand: they remove bonds from the SL lendable pool and shift the entire intermediation problem, as the structural free-float regression of Section 3.2 quantifies.

Fact 3. In Germany, both pricing measures co-move with repo turnover. Figure 3 reports the monthly bivariate relationships between the two pricing measures (fee, specialness) and three quantity measures (Markit lendable inventory, Markit value on loan, BrokerTec/MTS repo volume). In Germany, both lending fee and repo specialness co-move strongly with monthly repo volume ($\hat{\rho}_{\text{fee,repo-vol}} = +0.78$; $\hat{\rho}_{\text{spec,repo-vol}} = +0.63$). The same holds to a weaker extent with Markit value on loan (+0.51 and +0.52). In Italy the country-month aggregate is uniformly weaker, consistent with Fact 1. When more of the German bond panel is being borrowed and deployed in repo, prices on both sides of the dealer's joint problem rise.

Fact 4. The two markets share a common quantity cycle. Figure 4 reports the three quantity series side by side—lendable inventory, value on loan, and repo volume—monthly, for both countries. The patterns are synchronous. Lendable inventory rises into 2021–2022 as the Eurosystem programmes scale, then collapses sharply through 2022–2023 as PEPP and PSPP enter runoff and dealer inventories reorganise; value on loan and repo volume follow the same broad pattern with somewhat different timing in Germany versus Italy. The three series are not independent flows that happen to share a business cycle; they are tied together through the dealer's intermediation chain.

From facts to question. None of these four facts identifies a causal mechanism on its own. Taken together, however, they imply that the two markets cannot credibly be analysed as independent objects. Prices co-move (Fact 1); central-bank ownership shifts the SL supply schedule (Fact 2); repo activity correlates with both pricing measures (Fact 3); the three quantity series share one cycle (Fact 4); and the institutional structure of Section 3.1 places the dealer as the unique bridge. The empirical question of this paper is therefore not whether the two markets interact, but *how* the dealer's joint optimisation across securities lending and repo shows up in lending fees, and *under what conditions* the cross-market deployment-value channel becomes detectable. Section 4 develops the static search-and-bargaining model that makes this question precise, and Section 5 estimates

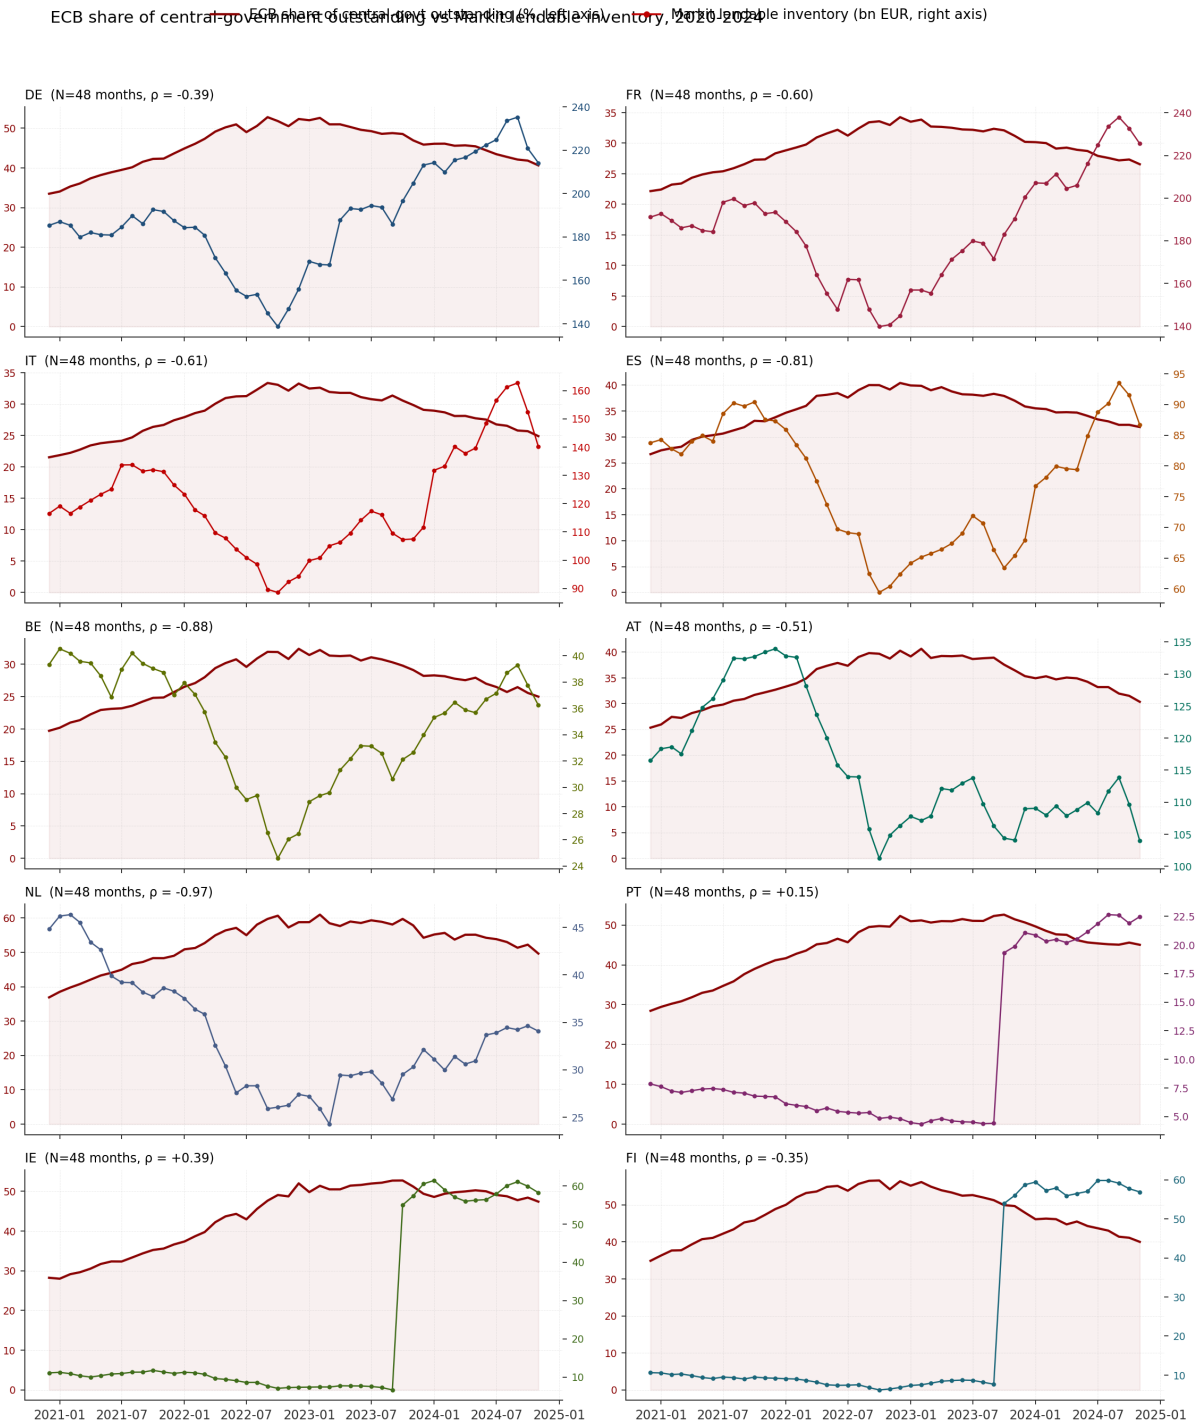


Figure 2: Eurosystem ownership share vs Markit lendable inventory, ten EU sovereigns. One panel per country. Left axis: ECB Eurosystem share of central-government outstanding (%), from monthly ECB country-breakdown releases and CSEC outstanding. Right axis: Markit lendable inventory (bn EUR, monthly mean of daily summed values). Country-by-country Pearson correlations of share and lendable: DE $\hat{\rho} = -0.40$, FR -0.61, IT -0.62, ES -0.81, BE -0.88, AT -0.48, NL -0.96, PT $+0.15$, IE $+0.40$, FI -0.38. Country-specific QT-onset dates from the ECB daily holdings panel: DE 2023-04-30, FR 2023-03-31, IT 2023-07-31, ES 2023-05-31, NL 2023-04-30.

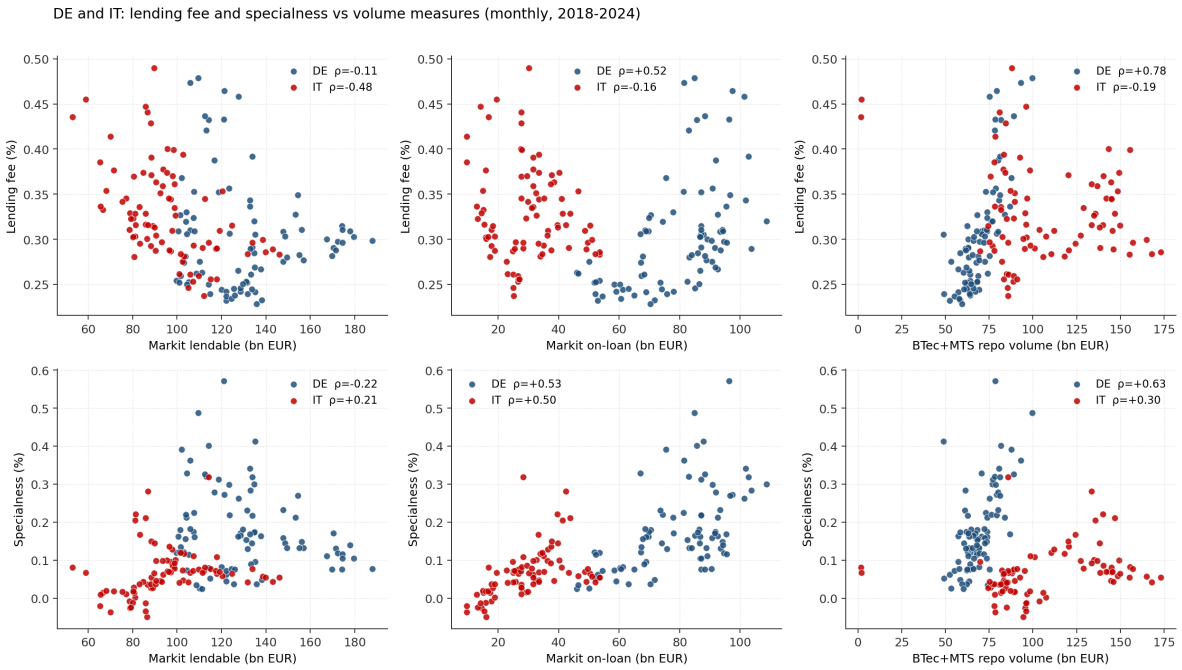


Figure 3: Price-versus-volume grid, monthly country means, 2018–2024 (84 months per country). Rows: lending fee (%), repo specialness (%). Columns: Markit lendable inventory, Markit value on loan, BrokerTec/MTS repo volume (all bn EUR). Each dot is one (country, month). Pearson correlations in each panel. The German cluster delivers visible monthly comovement between the two pricing measures and repo activity; Italy's aggregate is uniformly flatter.

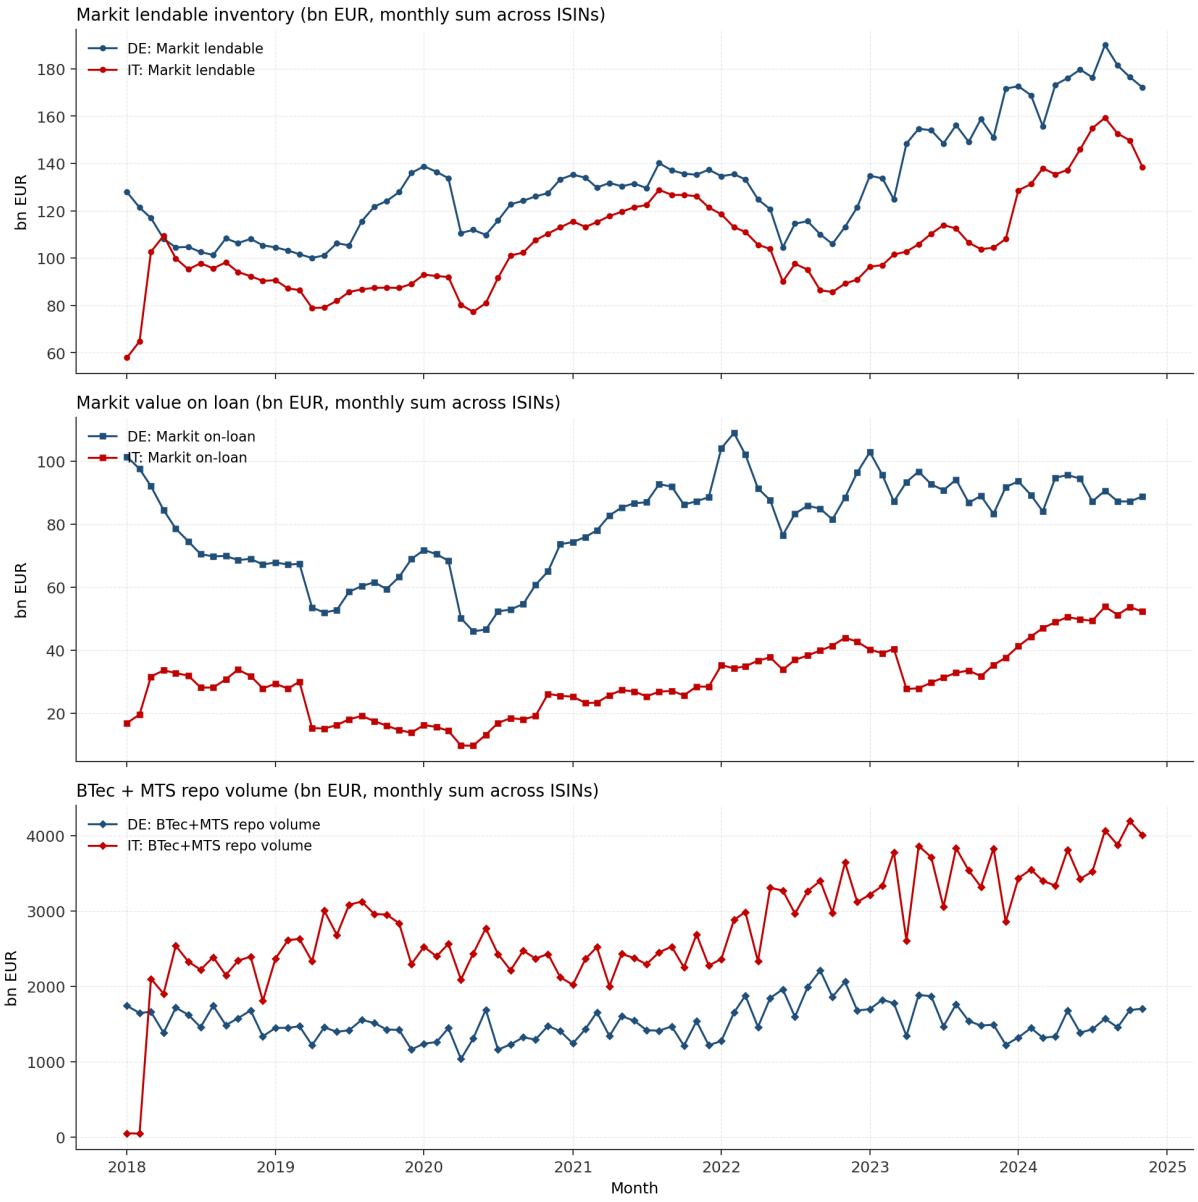


Figure 4: Lendable inventory, value on loan, and repo volume share a common quantity cycle. Monthly mean of daily cross-sectional sums (lendable, value on loan) and monthly total (repo volume), in bn EUR. DE in blue, IT in red. Sample: 2018–2024, quarter-end windows excluded.

the resulting state-dependent pass-through on the same DE+IT panel.

4 Model

This section develops the equilibrium lending fee in a static two-market search-and-bargaining environment along the lines of Duffie et al. (2002) and Duffie et al. (2005), extended to two markets that operate on the same bond simultaneously. Conditional on a securities-lending match between a beneficial owner and a dealer, the lending fee splits the joint surplus from the loan; the dealer's per-match surplus includes the repo deployment rent earned by re-posting the borrowed bond as collateral when the bond is special. Section 4.1 sets out the environment, agents, search-and-matching technology, the repo deployment value, and the Nash bargain. Section 4.2 maps the observed specialness variable to the structural deployment rent. Section 4.3 introduces the deployment-centrality index C_i and the high-intermediation indicator H_{it} that govern the loading on the deployment rent. Section 4.4 derives the fee-pass-through predictions that motivate the empirical analysis.

4.1 Environment and agents

Environment. Time is continuous, $t \in [0, \infty)$. All agents are risk-neutral and discount future payoffs at the common rate $\rho > 0$, which is identified empirically with the ECB deposit-facility rate $\rho = \text{DFR}_t$; this identification is standard in the euro-area sovereign repo literature and reflects the role of the DFR as the relevant short-rate benchmark for collateralised funding. Risk-neutrality is maintained throughout, following the convention of Duffie et al. (2002, 2005) and Lagos and Rocheteau (2009): the holding horizons relevant for securities-lending and repo trades are short relative to the bond's residual maturity, and a risk-aversion specification would add a hedging layer orthogonal to the cross-market deployment channel that is the focus here.

Three agents. The economy has three categories of agents indexed by their access to the two markets.

Beneficial owners. A continuum of risk-neutral beneficial owners (denoted B) holds the inventory of bond i, with mass μ_B per unit of outstanding nominal. Each beneficial owner accesses the securities-lending market through an agent-lender custodian (Section 3.1); throughout the model the beneficial owner and its agent lender are consolidated into a single risk-neutral entity. When holding the bond on its balance sheet at time t, the BO earns the per-period coupon flow c_i and discounts at rate ρ, so the present value of holding is

$$V_{it}^{B,\,\text{hold}} = \frac{c_i}{\rho}, \tag{2}$$

abstracting from the option value of future lending opportunities, which is absorbed into the bond-level intercept α_i in the empirical analogue.[2] When the BO lends the bond to a dealer over a unit time interval at the negotiated fee f_{it}, it earns the fee in addition to the coupon (the GMSLA contract transfers legal title but the BO retains the economic claim on the coupon), so the present value of lending is

$$V_{it}^{B,\,\text{lent}} \;=\; \frac{c_i + f_{it}}{\rho}. \tag{3}$$

The BO's per-match surplus from a successful loan is $V_{it}^{B,\,\text{lent}} - V_{it}^{B,\,\text{hold}} = f_{it}/\rho$.

Dealers. A continuum of risk-neutral dealers (denoted D), with mass μ_D, intermediates between the lending market and the repo market. The relevant aggregate object is the binary high-intermediation indicator $H_{it} \in \{0, 1\}$, which captures whether the dealer's redeployment opportunities bind. Conditional on a successful lending match, the dealer's per-period payoff comprises a lending-side cost $-f_{it}$ paid to the BO and a repo-side deployment rent earned by re-posting the bond as collateral with a repo counterparty.

Repo counterparties. A continuum of risk-neutral repo counterparties (denoted R), with mass μ_R, holds cash inventory and values bond i as specific collateral. The bond's specialness $s_{it} = \text{GC}_t - r_{it}^{\text{repo}}$—its repo rate below general collateral (GC, proxied by the deposit-facility rate DFR_t; Section 5 reports robustness to the GC benchmark)—is the market scarcity price, determined by aggregate repo-market clearing and taken as given by the dealer. Conditional on a bilateral repo match, however, the dealer and the repo counterparty split the specialness surplus by generalised Nash bargaining, and the dealer captures a share $\varphi_R \in (0, 1]$—the *repo rent share*. The dealer's per-deployment surplus is therefore $\varphi_R\, s_{it}^{+}$, where $s_{it}^{+} \equiv \max\{s_{it} - c, 0\}$ nets out a small deployment cost $c \geq 0$, since only specialness above the deployment cost generates a repo rent.

Search and matching. Bilateral meetings between beneficial owners and dealers on the securities-lending market arrive according to a Poisson process with intensity μ_L per BO per unit time. Conditional on a meeting, the BO and dealer negotiate the lending fee f_{it}. The empirical object of interest is the conditional-on-match fee. Bilateral meetings between dealers and repo counterparties arrive at intensity μ_R per dealer per unit time. Conditional on a repo meeting, the dealer posts the borrowed bond as collateral and the general-collateral leg $r_t^{\text{GC}} \approx \text{DFR}_t$ is absorbed into the discount-rate identification, so the relevant per-deployment rent is the dealer's bargained share $\varphi_R s_{it}^{+}$ of the specialness wedge. This rent is scaled by a deployment-intensity loading that depends on bond-level centrality C_i and the dealer's intermediation state H_{it}, as defined in Section 4.3.

[2] The option value of future lending matches is a time-invariant bond-level constant and is absorbed by the bond fixed effect in the regression specifications.

Repo deployment value. The new modelling object is the dealer's deployment value of the borrowed bond. When a dealer borrows bond i from a BO, the dealer can redeploy the bond as collateral in the repo market and capture its bargained share $\varphi_R s_{it}^+$ of the positive specialness per unit time, at an intensity that depends on the bond's deployment centrality and the dealer's regime. Let $C_i \in [0,1]$ denote the latent deployment-centrality rank of bond i (the empirical proxy is constructed as the within-country z-score of average ln repo volume on the first quarter of the bond's observed life; see equation (9) below). Let $H_{it} \in \{0,1\}$ be the high-intermediation-state indicator (proxied empirically by $D_{it}^{U_{75}} = \mathbf{1}[\text{Util}_{it} > p_{75}^{\text{Util}}]$; see Section 4.3). The deployment-intensity loading is

$$\Lambda(C_i, H_{it}) \;=\; \lambda_0 \;+\; \lambda_H\, H_{it} \;+\; \lambda_C\, C_i \;+\; \lambda_{HC}\, C_i H_{it}, \tag{4}$$

with parameter signs $\lambda_0 \geq 0$, $\lambda_H \geq 0$, $\lambda_C \geq 0$, and $\lambda_{HC} > 0$. The dealer's *deployment value* per borrowed bond is $V_{it} = s_{it}^+ \Lambda(C_i, H_{it})$, the gross repo value of redeploying it; the dealer captures the bargained share φ_R, so the deployment-rent component of its flow payoff while holding the bond as inventory is

$$R_{it} \;=\; \varphi_R\, s_{it}^+\, \Lambda(C_i, H_{it}), \tag{5}$$

which enters additively into the dealer's per-match surplus in the conditional-on-match bargaining problem. The strict positivity of λ_{HC} is the structural restriction that delivers the headline centrality cross-partial in Proposition 3 below and is the only restriction that the empirical analysis tests directly. The other loading-function parameters are reduced-form objects that absorb the level of the deployment value and its bond-level and state-level main effects. Equation (5) places the deployment rent in the dealer's per-match surplus inside the lending bargain, so the deployment intensity scales the per-match flow that the beneficial owner and dealer split.

Nash bargain. When a BO and a dealer meet, they split the per-match surplus from a successful loan via generalised Nash bargaining with share $\eta \in (0,1)$ to the lender side (the BO, who receives the fee) and $(1-\eta)$ to the borrower side (the dealer, who pays the fee). This convention is standard for OTC securities-lending bargaining (Duffie et al., 2002, 2005). Let $\tilde{V}_{it}^D$ denote the dealer's outside option, i.e. the value of continuing to search for another BO with the same bond; let B_{it} denote a non-deployment dealer-side borrowing benefit (delivery into a short, hedging, etc.), which absorbs uses of the bond unrelated to repo deployment. The static fee equation that solves the generalised Nash bargain is

$$f_{it}^* \;=\; (1-\eta)\, r_{it} \;+\; \eta\,\big[B_{it} \;+\; R_{it}\big], \tag{6}$$

where r_{it} denotes the BO's outside option (the general-collateral repo rate plus storage costs) and R_{it} is the dealer's repo deployment value defined in equation (5). Equation (6) is the conditional-on-match bargaining block of the OTC search environment of Duffie et al. (2002, 2005): search frictions motivate bilateral negotiation and surplus splitting; the contribution here is to add a repo deployment value for the same bond to the lending bargain.

Foil to the liquidity-windfall view. The deployment-value term R_{it} in equation (6) is not a liquidity-windfall term based on inbound–outbound haircut wedges as in Infante (2019) and Infante and Vardoulakis (2021). Reuse, when it appears in the data, is interpreted as the trace of cross-market bridging—the dealer's borrowing of a bond and its onward deployment as repo collateral—not as a primitive haircut-spread arbitrage motive. The structural object that signs the headline centrality contrast is the deployment loading $\Lambda(C_i, H_{it})$ on the positive deployment rent s^+_{it}, not a haircut wedge.

4.2 Repo Specialness as the Observed Cross-Market Scarcity Price

The static fee bargain (6) carries the deployment rent through $R_{it} = s^+_{it}\,\Lambda(C_i, H_{it})$. To map this object onto bond-day variation in the data, define the empirical specialness

$$s_{it} \equiv r^{GC}_t - r^{\text{repo}}_{it}, \tag{7}$$

where r^{GC}_t is the general-collateral repo rate (proxied by the deposit-facility rate) and r^{repo}_{it} is the bond-specific repo rate at (i,t). As an empirical object, s_{it} can take values of either sign: $s_{it} > 0$ when the bond trades special, $s_{it} \approx 0$ when it trades at general collateral, and $s_{it} < 0$ when its bond-specific repo rate is above general collateral (a regular feature of the data for bonds that are abundant or expensive to fund as repo collateral relative to GC). The structural repo-scarcity rent that enters the dealer's lending surplus is the truncation $s^+_{it} \equiv \max\{s_{it} - c, 0\}$, where $c \geq 0$ is the deployment cost (empirically the 0.05 bps noise floor); the deployment-value channel is silent at $s_{it} \leq c$. The closed-form fee derivation below uses s_{it} in the $s_{it} > c$ regime, and the empirical analysis restricts to $\tilde{s}_{it} > 0.05$ bps when interpreting the centrality contrast structurally. Throughout the paper s_{it} is observed from the BrokerTec/MTS panel rather than solved as the equilibrium of an explicit collateral-demand fixed point; that extension is left for future work.

The dealer's repo deployment rent enters the static bargain (6) through the term

$$R_{it} \;=\; \varphi_R\, s^+_{it}\, \Lambda(C_i, H_{it}), \tag{8}$$

where the deployment loading $\Lambda(C_i, H_{it})$ is defined in Section 4.3. R_{it} is the dealer's bargained share φ_R of the deployment rent s^+_{it} scaled by the bond-state-specific intensity

at which the dealer can redeploy the borrowed bond as repo collateral; in the $s_{it} \leq c$ regime the deployment-value channel is silent because $s_{it}^{+} = 0$.

4.3 Deployment Centrality and Dealer Cross-Market Intermediation Intensity

The deployment loading $\Lambda(C_i, H_{it})$ rises with the bond's deployment centrality C_i: its structural propensity to be redeployed as repo collateral. The natural empirical counterpart is how heavily the bond actually trades in repo. Let $C_i^0 \geq 0$ denote the latent deployment-centrality rank to which the propositions' sign restrictions apply; I use a single transparent monotone proxy for it, the bond's *predetermined repo depth*:

$$C_i^{\text{depth}} \;=\; \text{country-z}\Big(\overline{\ln \text{RepoVol}}_{\,i,\, t \in \mathcal{T}_i^{\text{pre25}}}\Big)\,, \tag{9}$$

the country-demeaned average daily log repo volume of bond i over $\mathcal{T}_i^{\text{pre25}}$, the first quarter of its observed life. Two features make this the right measure. First, it is *predetermined*—computed before the estimation window—so it is a fixed bond characteristic, not a contemporaneous outcome. Second, it is *repo* depth, not lending depth: the model grades the pass-through by how intensively the dealer *deploys* the bond as collateral (the loading Λ), not by how much it is borrowed. Section 5.7 confirms that the contrast appears with repo depth and *not* with the analogous lending-volume measure—exactly as the deployment-value channel predicts; the country demeaning makes the high-/low-centrality split a within-country comparison rather than a cross-country level difference.

The centrality contrast may still be confounded with persistent ISIN-level features correlated with repo depth; the predetermined construction addresses the timing but not all time-invariant unobservables.

The high-intermediation state H_{it}. I define the empirical high-utilization indicator $D_{it}^{U_{75}} \equiv \mathbf{1}[\text{Util}_{it} > p_{75}^{\text{Util}}]$, where p_{75}^{Util} is the sample-matched 75th percentile of utilization (the superscript U_{75} flags this percentile-threshold construction). $D_{it}^{U_{75}}$ measures a bond-level high-dealer-cross-market-intermediation regime. I formalize this by defining

$$H_{it} \;\equiv\; \mathbf{1}[\text{bond } i \text{ is in a high dealer-cross-market-intermediation state at } t], \qquad H_{it} \approx D_{it}^{U_{75}}. \tag{10}$$

$H_{it} = 1$ when beneficial-owner lending inventory is heavily used, sourcing a marginal unit is harder or more valuable, the dealer's cross-market intermediation margin is active, and repo specialness is therefore more likely to be priced into the lending bargain.

Centrality and state shift the deployment loading. In the static formulation (8), C_i and H_{it} enter the bargain through the deployment loading

$$\Lambda(C_i, H_{it}) \;=\; \lambda_0 + \lambda_H H_{it} + \lambda_C C_i + \lambda_{HC}\, C_i H_{it}, \tag{11}$$

with parameter signs

$$\eta \in (0,1), \qquad \lambda_0 \geq 0, \qquad \lambda_C \geq 0, \qquad \lambda_H \geq 0, \qquad \lambda_{HC} > 0, \tag{12}$$

and maintained support conditions on the empirical range of C_i:

$$\Lambda(C_i, H_{it}) \;\geq\; 0 \quad \text{and} \quad \lambda_H + \lambda_{HC}\, C_i \;\geq\; 0 \quad \text{on the empirical support of } C_i^{\text{depth}}. \tag{13}$$

The interpretation is that high-centrality bonds carry higher repo-deployment surplus per unit of positive specialness ($\lambda_C \geq 0$), the high-intermediation state activates additional deployment value ($\lambda_H \geq 0$), and the two complement each other ($\lambda_{HC} > 0$). Centrality matters most when the bond is in a high-intermediation regime. Lemma 1 gives a sufficient primitive matching foundation for this restriction. The support conditions in (13) are additional restrictions, not implications of the weak sign restrictions alone. They matter because C_i^{depth} is country-z-scored and can be negative on the empirical support (on the DE+IT post-25% panel, $p_{10} = -0.34$ and $p_{90} = +1.06$); the maintained restriction is that $\lambda_0 + \lambda_C C_i \geq 0$ and $\lambda_H + \lambda_{HC} C_i \geq 0$ over the observed range of the empirical proxy.

Lemma 1 (Primitive matching foundation for $\lambda_{HC} > 0$)**.** *Let the repo deployment intensity of a borrowed bond be*

$$q_{\mathcal{R}}(C, H) = A(1 + C)^{\alpha} E(H)^{\beta}, \qquad A > 0, \quad \alpha > 0, \quad \beta > 0,$$

where $C \in [0,1]$ is the latent deployment-centrality rank and $E(H) > 0$ is the repo-deployment opportunity set in state $H \in \{0,1\}$ with $E(1) > E(0)$. Then deployment intensity has increasing differences in (C, H):

$$\big[q_{\mathcal{R}}(C_2, 1) - q_{\mathcal{R}}(C_2, 0)\big] - \big[q_{\mathcal{R}}(C_1, 1) - q_{\mathcal{R}}(C_1, 0)\big] > 0 \quad \textit{for } C_2 > C_1.$$

Consequently, the local linear loading $\Lambda(C, H) = \lambda_0 + \lambda_C C + \lambda_H H + \lambda_{HC} CH$ inherits $\lambda_{HC} > 0$ after the positive present-value scaling used in the bargaining equation.

Proof. For $C_2 > C_1$,

$$\big[q_{\mathcal{R}}(C_2, 1) - q_{\mathcal{R}}(C_2, 0)\big] - \big[q_{\mathcal{R}}(C_1, 1) - q_{\mathcal{R}}(C_1, 0)\big] = A\big[(1+C_2)^{\alpha} - (1+C_1)^{\alpha}\big]\big[E(1)^{\beta} - E(0)^{\beta}\big],$$

which is strictly positive because $A > 0$, $\alpha > 0$, $\beta > 0$, $C_2 > C_1$, and $E(1) > E(0)$. A

linear approximation of this increasing-differences object around the empirical support assigns a positive coefficient to the interaction CH. Multiplying by the positive present-value and Nash-bargaining scalars preserves the sign. The lemma is a sufficient primitive structure for the main model's sign restriction; it is not a separate empirical identification of the matching technology. □

Substituting (8) and (11) into (6) delivers the closed-form fee

$$f_{it}^{*} \;=\; (1-\eta)\, r_{it} \;+\; \eta\, B_{it} \;+\; \eta\, \varphi_R\, s_{it}^{+} \big[\lambda_0 + \lambda_H H_{it} + \lambda_C C_i + \lambda_{HC}\, C_i H_{it}\big]. \tag{14}$$

This is linear in s_{it}^{+}, with deployment-loading scaling $\eta\varphi_R\Lambda(C_i, H_{it})$ on the slope: the fee is the product of two bargaining shares (lending η, repo φ_R), the net specialness s_{it}^{+}, and the deployment intensity Λ. The outside option r_{it} and non-deployment benefit B_{it} are absorbed by the bond and time fixed effects in estimation, so only the composite slope $\eta\varphi_R\Lambda$ is identified (and only up to that common scale).

4.4 Headline Prediction: Deployment-Centrality-Graded Pass-Through

Proposition 1 (Repo-to-lending pass-through)**.** *Under* (12) *and* (13)*, on the empirical support of* C_i*,*

$$\frac{\partial f_{it}^{*}}{\partial s_{it}^{+}} \;=\; \eta\, \varphi_R\, \Lambda(C_i, H_{it}) \;\geq\; 0.$$

Proof. Differentiate equation (14) with respect to s_{it}^{+}; the remaining terms are independent of s_{it}^{+}, so the slope is $\eta\varphi_R\Lambda(C_i, H_{it})$. Non-negativity follows from $\eta, \varphi_R > 0$ and the first part of (13). □ □

Proposition 2 (State-dependent pass-through)**.** *Under* (12) *and* (13)*, on the empirical support of* C_i*,*

$$\frac{\partial^2 f_{it}^{*}}{\partial s_{it}^{+}\, \partial H_{it}} \;=\; \eta\, \varphi_R \big(\lambda_H + \lambda_{HC}\, C_i\big) \;\geq\; 0.$$

Proof. Differentiate Proposition 1 with respect to H_{it}; because $H_{it} \in \{0, 1\}$, this derivative is equivalently the finite difference between the high- and low-intermediation regimes in the linear loading. The result follows from $\partial\Lambda/\partial H = \lambda_H + \lambda_{HC} C_i$. Non-negativity is the second part of (13). □ □

Proposition 3 (Deployment-centrality contrast)**.** *Under* (12) *and* (13)*,*

$$\frac{\partial^3 f_{it}^{*}}{\partial s_{it}^{+}\, \partial H_{it}\, \partial C_i} \;=\; \eta\, \varphi_R\, \lambda_{HC} \;>\; 0. \tag{15}$$

The empirical analog, used as the headline, is the between-group contrast *in the state-dependent pass-through across a median split on centrality:*

$$\delta \;:=\; b\big(s^{+}\times D^{U_{75}} \,\big|\, \textit{high-}C\big) - b\big(s^{+}\times D^{U_{75}} \,\big|\, \textit{low-}C\big) \;>\; 0, \tag{16}$$

where $b(s^{+}\times D^{U_{75}} \mid \cdot)$ *is the coefficient on the specialness–state interaction estimated within each centrality group of a within-bond fee regression with ISIN and country×day fixed effects,* $s^{+}_{it} = \max\{\tilde{s}_{it} - c, 0\}$*, and* $\tilde{s}_{it}$ *is the* 20*-trading-day lag of observed specialness. The contrast* δ *identifies* $\eta\varphi_R\lambda_{HC}(\bar{C}_{\text{hi}} - \bar{C}_{\text{lo}}) > 0$*. The low-centrality coefficient is predicted* smaller, not zero*. The continuous triple interaction* $s^{+} \times D^{U_{75}} \times C^{\text{depth}}$ *is reported as corroboration.*

Proof. Differentiate Proposition 2 with respect to C_i. Strict positivity follows from $\eta, \varphi_R > 0$ and $\lambda_{HC} > 0$ in (12). The split contrast δ is the fixed-effect-residualized between-group analog of the structural cross-partial; the structural object signs the empirical projection rather than parameterizing it. □ □

Proposition 3 is the paper's headline structural prediction: the repo-to-lending pass-through is steeper for bonds the dealer deploys more. Within the model, the complementarity arises from the increasing-differences condition in Lemma 1—repo deployment intensity responds more to high intermediation for high-centrality bonds—and the two bargains convert it into the fee at scaling $\eta\varphi_R$. Implementing it as a high-/low-centrality *split* (rather than a fragile continuous triple) is the robust form; the static main-text fee (14) delivers the sign directly from $\eta\varphi_R\lambda_{HC} > 0$.

The static model derives the headline cross-partial $\eta\varphi_R\lambda_{HC} > 0$ from primitives, so the empirical identification in Section 5 is self-contained. On-the-run benchmark status and cheapest-to-deliver status are alternative explanations for why some bonds command a steeper pass-through—benchmark-liquidity rents and futures-delivery specialness, respectively; rather than build either into the model, I rule both out empirically, showing that the repo-depth centrality contrast survives controlling for on-the-run status and absorbs the past-CTD channel (Section 5.9).

5 Empirical Analysis

Notation for Section 5. In the empirical sections, β_k denotes a regression coefficient and $\hat{\beta}_k$ its sample estimate. The static model uses $\eta \in (0, 1)$ for the lender/BO bargaining share, $\varphi_R \in (0, 1]$ for the repo bargaining share, and $\Lambda(C_i, H_{it})$ for the deployment loading. Pearson correlations are denoted $\hat{\rho}$ to distinguish them from the discount rate ρ.

5.1 Data

The pan-European sovereign bond lending panel combines S&P Global Securities Finance (formerly IHS Markit; lending fees, utilization, short-loan quantity, loan tenure), BrokerTec/MTS Consolidated (repo volume, specialness), interdealer MTS cash-market trades (secondary-market volume, used in the auction first stage of Section 5.14), and the German, French, and Italian debt-management-office auction calendars (reopening dates and allotments). The sample covers 1,536 sovereign bonds across eleven eurozone countries, 2018–2024. The lending fee is the IndicativeFee field in basis points, taken as the same-day quote and used as the primary dependent variable throughout. Utilization is the percentage of lendable inventory on loan (ActiveUtilization), and $D_{it} = \mathbf{1}[\text{Util}_{it} > p_{75}]$ is the state-activation indicator where $p_{75} = 41.87\%$ is the 75th percentile of utilization computed on the pan-European 2018–2024 sample used for the headline Test 3 specification. For robustness specifications estimated on the DE+IT 2021–2024 subsample where matched BrokerTec/MTS repo data are required (Tests 1 and 2, and the repo-side Test 3B in the appendix), the corresponding p_{75} threshold is 50.6%; the Italian-only subsample uses $p_{75} = 38.0\%$. I use the sample-matched p_{75} throughout, and all thresholds are reported in the respective table notes. Short-loan quantity (SLQ_{it}, in face notional) proxies bond-specific borrower demand. Repo specialness is $\text{Specialness}_{it} = \text{DFR}_t - r_{it}^{\text{repo}}$, the DFR-to-bond-specific-repo-rate spread; this is the empirical counterpart of $r_t^{GC} - r_{it}^{\text{repo}}$ in equation (7), with the deposit-facility rate used as the GC proxy.[3][4]

Markit panel selection on lendability. The empirical sample includes only sovereign bonds in the Markit S&P Global Securities Finance lendable universe, which is constructed from beneficial-owner reports across approximately 25,000 institutional accounts globally. A bond enters the universe when at least one beneficial owner reports lendable inventory; bonds excluded entirely from beneficial-owner lending programs (e.g., insurance-mandated holdings held off-program, or bonds with very short remaining maturity outside typical lending windows) are not in the panel. Bonds in the Markit lendable universe are not a random sample of euro-area sovereign debt. The lendability filter is therefore an internal-

[3]During 2018 through mid-2022 market GC traded approximately 8 to 10 bps below DFR. As a robustness check I re-run the headline Test 3 specification with an EURSTR-based specialness $s_{it}^{\text{EURSTR}} = \text{EURSTR}_t - r_{it}^{\text{repo}}$ over the post-October-2019 subsample ($\approx 78\%$ of the baseline panel). The EURSTR construction preserves every sign and significance verdict in the paper: Test 3 $\text{lnSLQ}\times D$ under EURSTR specialness is +1.77 pooled ($t = +4.37$), +2.23 DE ($t = +3.01$), +4.81 IT ($t = +4.14$), +3.65 DE+IT ($t = +5.05$). The EURSTR-5 bp alternative gives identical magnitudes; appendix tables report the full robustness sweep.

[4]Markit lendable-supply variables (`fee_bps`, `short_loan_qty`, `utilisation`, `lendable_qty`, `daily_cost_of_borrow_score`, `specialness_markit`) labelled at date T reflect market state at $T-2$ business days under Markit's reporting convention. I therefore re-run every headline specification with the six Markit columns shifted +2 trading days within each ISIN and confirm that no regression flips sign and no significance verdict changes. The pan-European Test 3 $\text{lnSLQ}\times D$ moves from $t = +5.90$ to $t = +6.06$ under the lag-corrected loader; full results across 30 specifications are reported in the appendix.

validity selection on the population to which the regression coefficients apply: the cross-market deployment-value channel is identified on bonds for which lending is feasible at all, and the empirical coefficients should not be extrapolated to bonds outside the lendable universe (insurance-mandated holdings, short-residual-maturity bonds, and other classes excluded from beneficial-owner lending programs). The pan-European panel of 1,536 ISINs captures the electronically active lendable segment of the eurozone-sovereign universe over 2018–2024.

Coverage of the repo volume measure. The repo volume measure is constructed from BrokerTec and MTS, the two largest electronic interdealer repo automatic trading systems (ATS) in Europe. The ICMA European Repo Market Survey reports combined monthly turnover on BrokerTec and MTS in the EUR 5–8 trillion range per platform in recent years, with Italian government securities accounting for 40.3% of ATS-traded repo collateral by outstanding value and German government securities for 15.4% (ICMA, 2026). Total European repo activity is larger: the ICMA survey reports outstanding repo and reverse repo of EUR 13.7 trillion on the books of 59 participating entities at December 2025, and aggregate SFTR-reported outstanding for the EU plus the UK was EUR 25.9 trillion on the nearest reporting date. These figures are not used to form a coverage ratio because the available public aggregates mix transaction turnover (the BrokerTec/MTS measure) and outstanding positions (ICMA, SFTR); their role is only to document that the electronic interdealer segment is a partial view of the broader repo market. The BrokerTec/MTS measure covers the CCP-cleared electronic interdealer segment only, omitting CCP general-collateral-financing facilities (Eurex GC Pooling, LCH eurGCPlus), triparty repo, and the rapidly growing dealer-to-customer platforms (Tradeweb, GLMX). ATS share of outstanding interdealer repo has trended down to 25.1% of the ICMA survey sample in December 2025, an eight-year low (ICMA, 2026). For the purposes of this paper, the BrokerTec/MTS measure captures *bond-specific* electronic repo activity at high frequency on the segment where repo specialness is price-discovered and where Italian and German sovereigns are the dominant collateral, the right measurement domain for the state-dependent mechanism tested here.[5] The omission of other segments introduces measurement error in the repo-volume regressor; under classical measurement error the within-bond regression coefficient is a lower bound on the true elasticity, so the within-bond estimates reported below should be read as such. The classical-measurement-

[5] A coverage diagnostic on the analysed panel rejects a differential ATS-coverage interpretation of the cross-country results: ISIN-day coverage with positive ATS-reported repo activity is 100% in both Germany and Italy, with 38.6 trades per ISIN-day in DE versus 37.3 in IT and volume-per-lendable ratios of 0.25% (DE) versus 0.21% (IT). Coverage is uniform across the 2018–2024 sample years. The within-panel cross-country comparison is unbiased given the equal coverage within the panel, and the IT/DE ratio of $\sim$1.7$\times$ in the cross-market simultaneity coefficient reflects a genuine cross-country difference rather than measurement attenuation. The panel *ex ante* excludes bonds with no ATS repo activity (about 20% of nominal Bund outstandings traded purely OTC, by ECB-internal estimates), which is an external-validity restriction on the lendable-and-ATS-active subset of the eurozone sovereign universe.

error reading does not apply if reporting noise is correlated with the true regressor (a non-classical structure that the publicly available data do not let me characterise directly), in which case the direction of bias is not determined.

BO-side utilization as the empirical state variable. The empirical state variable $D_{it}^{U_{75}}$ is BO-side utilization: the fraction of lendable inventory currently on loan, thresholded at the within-sample 75th percentile. I do not interpret this indicator as literal dealer inventory. It proxies a high cross-market intermediation regime H_{it}: a regime in which sourcing a marginal unit is harder or more valuable, the dealer's cross-market intermediation margin is active, and repo specialness is more likely to be priced into the lending bargain. The empirical strategy uses BO-side utilization directly as the regime indicator and does not rely on dealer-position data.

Regression design. I exploit within-bond, within-day variation via ISIN and day fixed effects. The ISIN fixed effect absorbs time-invariant bond-level heterogeneity (issuer credit class, coupon, original-issue size, indexation status, persistent dealer-network features); the day fixed effect absorbs cross-sectional average shocks at daily frequency (euro-area-wide funding conditions, calendar events, common monetary-policy news). Time-varying bond-level confounders that move within bond and within day are not absorbed by either fixed effect; the predetermined-state robustness battery in Section 5.5 addresses one form of this concern by holding the state indicator $D^{U_{75}}$ fixed at k trading days before the fee observation, while the post-25%-of-life cohort restriction in Section 5.7 addresses time-varying confounders that operate through the bond's initial pricing window. Standard errors are two-way clustered (ISIN and day) in headline specifications. Quarter-end windows are excluded to avoid window-dressing distortions.

5.2 Within-window timing

Two within-sample timing patterns sharpen the empirical specification choices and complement the broader stylized facts of Section 3. First, the gap between aggregate Markit lendable supply and value on loan for German Bunds reaches its tightest configuration in late 2022 and early 2023 (Figure 5), the same window in which Bund repo specialness peaks. Second, weekly aggregate German ShortLoanQuantity co-moves positively with weekly median Bund specialness across the sample, with baseline correlation $\hat{\rho} = +0.36$ rising to $\hat{\rho} = +0.46$ when quarter-end windows are excluded at the $\pm$15-business-day margin (Figure 6). The pan-European pooled analogue gives $\hat{\rho} = +0.30$ baseline, rising to $\hat{\rho} = +0.45$ on the same exclusion margin. The joint peak of demand and scarcity at the aggregate level motivates the within-bond, state-dependent specifications of Tests 1–3 below: the same demand shocks that move repo activity in the high-utilization regime should also move lending fees there, and the time-series alignment of the two markets sharpens

around the QT-onset window where the policy-shock identification in Section 5.11 also sits.

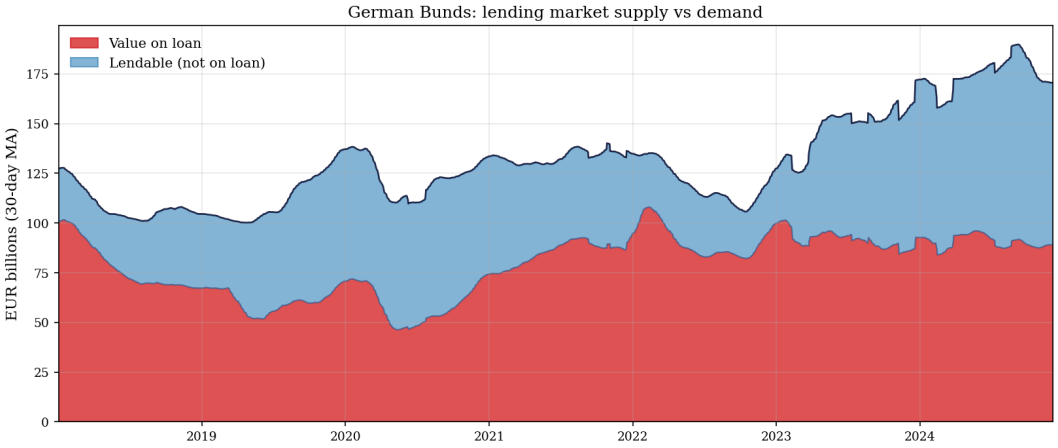


Figure 5: German Bunds: aggregate lendable value (blue area) and value on loan (red area), EUR billions, 30-day moving average. The tightening episode in late 2022 / early 2023 coincides with the peak in Bund specialness. Source: S&P Global Securities Finance (Markit). Sample: weekly observations, January 2018–December 2024.

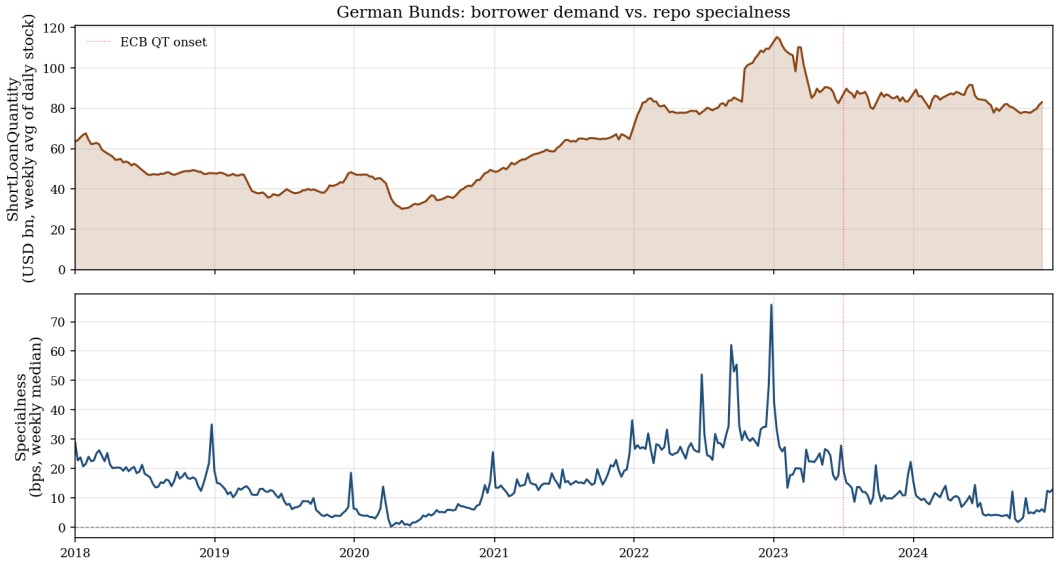


Figure 6: German Bunds: weekly aggregate ShortLoanQuantity (top, brown area) and weekly median repo specialness (bottom, blue line). Both series peak jointly in late 2022/early 2023. Weekly level correlation $\hat{\rho} = +0.36$ (baseline ± 5 business-day quarter-end exclusion) rising to $\hat{\rho} = +0.46$ at ± 15 days. Top panel: aggregate ShortLoanQuantity (face notional, EUR billion). Bottom panel: median repo specialness (bps). Source: S&P Global Securities Finance (SLQ), BrokerTec/MTS (specialness).

5.3 Test 1: Scarcity and Intermediation Activity

The model predicts that lending market scarcity is associated with higher repo intermediation activity for the same bond, consistent with a cross-market bridging role for dealers. The specification is:

$$\ln(\text{RepoVol}_{it}) = \alpha_i + \alpha_t + \beta_1 \, \text{Util}_{it} + \beta_2 \, \text{Fee}_{it} + \varepsilon_{it} \tag{17}$$

Table 2 reports the results across six specifications. The coefficient on utilization is $\hat{\beta}_1 = 0.0056$ ($t = 7.53$, $p < 0.001$): a one-percentage-point increase in active utilization is associated with a 0.56% increase in daily repo volume for the same bond on the same day. The estimate is unchanged under double-clustering ($t = 7.49$), under $t-2$ matching ($t = 7.51$), and in the nonlinear specification where crossing the 75th percentile of utilization is associated with 17.0% higher repo volume ($t = 6.25$). The p_{90} threshold adds 10.5% and is individually significant ($t = 3.54$), consistent with the same association continuing at moderate scarcity levels; reverse causality (repo-side dealer activity raising both observed utilization and observed repo volume) is acknowledged below and not separately ruled out. The lending fee is insignificant in all specifications. This is not a rejection of the model, which predicts a state-dependent fee response rather than an unconditional fee-volume slope. The within-bond specification also absorbs persistent bond-level fee differences through ISIN fixed effects, leaving only high-frequency fee variation in this first descriptive

test.

Table 2: Test 1: Lending Scarcity and Repo Activity

	Linear		**Threshold**		**Double-cluster**	
	$t-1$	$t-2$	$t-1$	$t-2$	$t-1$	$t-2$
	(1)	(2)	(3)	(4)	(5)	(6)
Util	0.0056***	0.0056***			0.0056***	0.0056***
	(0.0007)	(0.0007)			(0.0007)	(0.0008)
$\mathbf{1}[\text{Util} > p_{75}]$			0.170***	0.170***		
			(0.027)	(0.027)		
$\mathbf{1}[\text{Util} > p_{90}]$			0.105***	0.105***		
			(0.030)	(0.030)		
Fee (bps)	0.0002	0.0001	0.0002	0.0002	0.0002	0.0001
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
ISIN FE	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes	Yes
SE cluster	ISIN	ISIN	ISIN	ISIN	ISIN×Day	ISIN×Day
Within-R^2	0.032	0.032	0.015	0.018	0.032	0.032
N	176,382	176,380	176,382	176,382	176,382	176,380
ISINs	423	423	423	423	423	423

Notes: Standard errors in parentheses. Sample: January 2021–December 2024; quarter-end windows excluded per Bassi et al. (2024). Dependent variable is $\ln(\text{RepoVol}_{it})$. Columns (1)–(4): standard errors clustered at the ISIN level. Columns (5)–(6): double-clustered (ISIN × date). $t-1$ and $t-2$ refer to the Markit settlement-lag matching assumption. $p_{75} = 50.6\%$ is the 75th percentile of utilization computed on the DE+IT 2021–2024 sample used for this specification; p_{90} is the corresponding 90th percentile. The threshold differs from the pan-European $p_{75} = 41.87\%$ used in the headline Test 3 specification because Tests 1 and 2 require matched BrokerTec/MTS repo data, which is available for German and Italian sovereigns. The positive association is consistent with the scarcity-to-intermediation prediction but cannot be separated from reverse causality (repo-side dealer cross-market intermediation can mechanically raise both utilization and volume). All continuous variables are winsorized at the 1st and 99th percentiles. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

These results are consistent with the model's first prediction: lending-market scarcity is associated with higher repo activity for the same bond. The effect is economically meaningful and is estimated from within-bond, within-day variation after absorbing permanent bond characteristics and aggregate market conditions. A direct cross-market bridge regression of $\ln(\text{RepoVol}_{it})$ on $\ln(\text{SLQ}_{it})$ on the same panel (the appendix bridge regression) provides a complementary reduced-form link from borrower demand to repo activity through dealer intermediation.

Reverse-causality caveat. The Test 1 regression identifies a positive within-bond association between lending utilization and repo volume, but the direction of causality cannot be determined from this specification alone. Two channels can produce the same correlation. First, as the model predicts, scarcity in the lending market induces dealers to shift activity toward repo to source and redeploy the scarce bond. Second, repo-side dealer cross-market intermediation, where dealers reuse received collateral as a by-product

of cross-market positioning, can mechanically raise both utilization and repo volume simultaneously. The two channels are not mutually exclusive, and the Test 1 specification does not distinguish them. The identification burden is therefore carried by Tests 2 and 3: Test 2 documents that unconditional repo volume does not move fees (contra what a pure reverse-causality story of repo-side dealer activity driving utilization would generate), and Test 3 documents a state-dependent price response consistent with the deployment-value channel.

5.4 Test 2: Repo Market Activity and Lending Fees

The model predicts that higher repo market activity for a given bond is associated with a higher lending fee for that bond: more frequent repo deployment increases the bond's value to the dealer in the model, making the dealer willing to bid more aggressively for it in the lending market. The specification is:

$$\text{Fee}_{it} = \alpha_i + \alpha_t + \beta_1 \ln(\text{RepoVol}_{it}) + \beta_2 \,\text{Specialness}_{it} + \varepsilon_{it} \tag{18}$$

Table 3 reports the results across five specifications. The unconditional coefficient on log repo volume is statistically indistinguishable from zero in every within-bond specification: $\hat{\beta}_1 = 0.12$ ($t = 0.22$) contemporaneously, 0.11 ($t = 0.20$) in the lead, and 0.00 in the reverse. Repo specialness, by contrast, is precisely estimated across all specifications ($\hat{\beta}_2 \approx 0.29$, $t > 6$, well above conventional significance thresholds): a one basis point increase in specialness is associated with a 0.29 bps higher lending fee, consistent with bond-specific repo conditions feeding into lending prices.

Demand-side unconditional effect. Replacing $\ln(\text{RepoVol}_{it})$ with $\ln(\text{SLQ}_{it})$ in the unconditional specification, estimated on the pan-European 2018–2024 sample used for Test 3 (Table 4), yields a strongly negative coefficient: $\hat{\beta}_1 = -5.96$ ($t = -22.17$) pooled, -6.12 ($t = -13.84$) for DE, -6.85 ($t = -15.03$) for IT. Higher borrower demand for a specific bond is associated with *lower* within-bond, within-day lending fees. This is consistent with the model: when demand is met by willing lenders, more bonds are deployed, utilization may remain moderate, and fees do not need to rise. The state-dependent interaction (Test 3) shows that this relationship reverses at high utilization, when dealer capacity is constrained, a pattern consistent with the deployment-value channel.

The cross-sectional specification (column 5) reveals a different pattern: the average log repo volume positively predicts the average lending fee ($\hat{\beta}_1 = 4.82$, $t = 3.14$, $p = 0.002$), controlling for issuer. Bonds with structurally deeper repo markets are associated with higher lending fees on average, though this may reflect omitted bond-level characteristics. The contrast, positive cross-sectional effect, zero within-bond time-series effect, motivates

Test 3.

Table 3: Test 2: Repo Market Thickness and Lending Fees

	Same-day (primary) (1)	**Multi-day** (robustness) (2)	**Lead** (Fee_{t+1}) (3)	**Reverse** (Fee_t on Vol_{t-1}) (4)	**Cross-sect.** (OLS) (5)
ln(RepoVol)	0.119	−0.099	0.107	0.000	4.819***
	(0.554)	(0.563)	(0.545)	(0.0002)	(1.536)
Specialness	0.289***	0.334***	0.294***	0.005***	0.031
	(0.042)	(0.050)	(0.043)	(0.001)	(0.210)
DE dummy					−4.602
					(2.694)
ISIN FE	Yes	Yes	Yes	Yes	No
Day FE	Yes	Yes	Yes	Yes	No
SE cluster	ISIN	ISIN	ISIN	ISIN	HC3
Within-R^2 / R^2	0.031	0.037	0.031	0.015	0.034
N	176,382	176,506	176,047	176,382	424[a]

Notes: Standard errors in parentheses. Sample: January 2021–December 2024; quarter-end windows excluded. Dependent variable is Fee_{it} (bps) in columns (1)–(4) and the bond-level time-average $\overline{\text{Fee}}_i$ in column (5). Columns (1)–(4): ISIN and day FE with ISIN-clustered SEs; robust to double-clustering (ISIN × date). [a]The pooled sample includes 424 ISINs; Tables 2–3 report 423 because one ISIN enters the sample after the first available Markit observation date and is dropped by the $t-1$ lag constraint. Column (3): tomorrow's fee on today's repo volume. Column (4): today's fee on yesterday's repo volume. Column (5): cross-sectional OLS on 424 ISINs (cross-sectional averages), HC3 robust SEs. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5.5 Test 3a: State-Dependent Cross-Market Spillover (High Utilization)

The paper's structural headline is the repo-side pass-through of Tests 5.6–5.7: the lending fee's sensitivity to repo specialness, graded by deployment centrality. I introduce the channel through its demand-side counterpart—a transparent state-dependence in lending-market demand—and then show that this counterpart is the reduced form of the specialness pass-through and is dominated by it in a head-to-head. This demand-side test maps to Proposition 2 under the lending-market demand proxy $\ln \text{SLQ}_{it}$ in place of the model's repo-side scarcity price s_{it}^{+}: lending fees respond to demand shocks primarily in the high-intermediation regime $H_{it} = 1$, with empirical proxy $D_{it}^{U_{75}}$. I test this with a within-bond panel regression:

$$\text{Fee}_{it} = \alpha_i + \alpha_t + \beta_1 \ln \text{SLQ}_{it} + \beta_2 \ln \text{SLQ}_{it} \times D_{it} + \delta_D D_{it} + \gamma_{\text{Spec}} \cdot \text{Spec}_{it} + \varepsilon_{it}, \qquad (19)$$

where $D_{it} = \mathbf{1}[\text{Util}_{it} > p_{75}]$ and $p_{75} = 41.87\%$ is the 75th percentile of utilization computed on the pooled pan-European sample. The static model predicts the state-dependent sign $\beta_2 > 0$. I treat the unconditional slope β_1 as descriptive rather than as a primitive sign restriction.

ln SLQ as a demand proxy. SLQ_{it} is constructed by Markit as $\text{Util}_{it} \times \text{Lendable}_{it}$, so the algebraic identity

$$\ln \text{SLQ}_{it} \;=\; \ln \text{Util}_{it} \;+\; \ln \text{Lendable}_{it}$$

holds within ISIN-day exactly (median absolute identity error = 0.0000 across the 728,502 ISIN-days of the pan-European panel). The Test 3a interaction $\ln \text{SLQ} \times D$ therefore decomposes additively into a demand component $\ln \text{Util} \times D$ and a supply component $\ln \text{Lendable} \times D$. The identifying assumption for $\ln \text{SLQ}$ as a demand proxy is that the variance of $\ln \text{Util}$ dominates the variance of $\ln \text{Lendable}$ within ISIN and that the demand component carries the state-dependent fee response. The assumption is testable: a joint horse-race regression that includes both interactions simultaneously delivers $\hat{\beta}^{\text{demand}}_{\text{Pooled}} = +12.26$ ($t = +10.74$) and $\hat{\beta}^{\text{supply}}_{\text{Pooled}} = +1.38$ ($t = +3.88$), with the demand interaction 3 to 5 times larger than the supply interaction in every country sample. For Germany specifically, the supply interaction $\ln \text{Lendable} \times D$ becomes statistically insignificant ($t = +1.24$, $p \approx 0.21$) once the demand interaction is controlled for, while the demand interaction remains highly significant ($t = +6.23$). The headline $\ln \text{SLQ} \times D$ coefficient is therefore dominantly a demand-state interaction, with a small but non-zero supply component that the model also predicts (through the matching-rate channel: higher lendable supply raises the equilibrium fee in the loaded state by raising the dealer's matching intensity). The full decomposition table is reported in the appendix.

The demand-side interaction. Table 4 reports the panel estimates. On the pan-European sample of 1,536 bonds across eleven eurozone countries (728,502 ISIN-days), the interaction $\hat{\beta}_2 = +2.15$ is significant at the 1% level ($t = 5.90$). Country-by-country: $t = 3.44$ (DE), $t = 4.60$ (IT), with the interaction significant at the 10% level in seven of eleven countries; under a Bonferroni correction for eleven country tests, four countries survive at the 5% familywise level (IT, BE, DE, and IE). France ($t = -0.56$) and the Netherlands ($t = -0.13$) show no effect. The model does not pre-register country-specific predictions about which markets activate the channel; the cross-country heterogeneity is reported as empirical description rather than as a test of a structural hypothesis. A continuous specification interacting $\ln \text{SLQ}_{it}$ with Util_{it} directly (no threshold) yields $t = 9.58$ pooled. The unconditional slope is strongly negative ($\hat{\beta}_1 = -5.96$, $t = -22.17$). I do not treat that negative level slope as a primitive sign restriction of the static model; it is consistent with borrower demand being accommodated in low-utilization states, while the direct model prediction is the positive state-dependent slope.

Robustness and interpretation. The demand-side $\ln \text{SLQ} \times D$ result is a state-dependent demand association, not a claim that every high-utilization bond is literally dealer-inventory constrained. The appendix reports the threshold sweep, cash-collateralized variant, borrower-concentration split, collateral-transformation split, bid-ask placebo, and

Table 4: Test 3a: The Demand-Side SLQ × D State-Dependent Interaction

	Pan-European	Germany	Italy
$\hat{\beta}_1$ (unconditional ln SLQ)	−5.96	−6.12	−6.85
	($t = -22.17$)	($t = -13.84$)	($t = -15.03$)
$\hat{\beta}_2$ (ln SLQ × D, threshold)	+2.15	+2.14	+5.10
	($t = 5.90$)	($t = 3.44$)	($t = 4.60$)
ln SLQ × Util (continuous)	+0.072	+0.059	+0.082
	($t = 9.58$)	($t = 4.88$)	($t = 4.81$)
Specialness control	Yes	Yes	Yes
ISIN FE, day FE	Yes	Yes	Yes
N (ISIN-days)	728,502	123,692	161,905

Notes: Dependent variable: Fee_{it} in basis points. t-statistics in parentheses use two-way clustered standard errors by ISIN and day. $D_{it} = \mathbf{1}[\text{Util}_{it} > 41.87\%]$ uses the pan-European 75th-percentile utilization threshold. Quarter-end windows excluded. The continuous specification replaces ln SLQ × D with ln SLQ × Util. Appendix tables report threshold sweeps, collateral splits, borrower concentration, bid-ask placebo tests, and country details. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

nonparametric SLQ and utilization tests. These checks leave the same basic reading: the interaction is not concentrated at a single threshold cell and is difficult to explain as a simple utilization nonlinearity. The economic magnitude is modest relative to total fee volatility, but the sign and incidence line up with the deployment-value mechanism after absorbing ISIN and day fixed effects.

Economic magnitude of the demand-side interaction. The demand-side coefficient $\hat{\beta}_2 = +2.15$ bps is small in absolute basis points but economically meaningful when benchmarked against the relevant frictions in the lending and repo markets. At a +1-standard-deviation ln SLQ shock—a doubling of demand for a typical bond—the realised fee response in the high-utilisation regime is +1.93 bps for Germany and +5.64 bps for Italy, equal to 22% and 44% of the median bid-ask spread on the same bond (8.87 bps DE, 12.69 bps IT), 11% and 62% of the mean specialness premium, and 7% and 16% of the mean lending fee. Table 5 reports the country-by-country benchmark. For a dealer holding a typical lendable position of EUR 1–4 bn at the country mean utilisation rate, this translates into incremental annual lending revenue of approximately 172k per dealer-bond in Germany and 125k in Italy from a single +1-sd demand realisation persisting at the mean utilisation rate; across a portfolio of ∼100 bonds the demand-state heterogeneity moves dealer revenue by 1.5–17 million annually, a magnitude in the same order as the dealer balance-sheet repo costs documented by D'Avernas and Vandeweyer (2024) and the liquidity-windfall margins of Infante (2019). The marginal fee response is therefore sufficient to influence the dealer's cross-market allocation decision, as the deployment-value channel requires.

Table 5: Economic magnitude of the demand-side $\ln \mathrm{SLQ} \times D$ interaction

Country	$\hat{\beta}_2$ (bps)	Realised at +1sd (bps)	Median BA (bps)	Mean spec (bps)	Realised / BA
DE	+2.14	+1.93	8.87	19.57	22%
IT	+5.10	+5.64	12.69	8.28	44%
DE+IT	+3.91	—	—	—	—
FR	−0.37	−0.26	12.02	13.11	−2%
NL	+2.15	+1.92	9.55	14.21	20%
ES	+2.15	+2.23	17.19	9.60	13%
BE	+2.15	+1.83	12.23	12.50	15%
AT	+2.15	+1.61	25.62	12.56	6%
FI	+2.15	+2.03	16.37	14.61	12%
PT	+2.15	+2.94	20.91	13.39	14%
IE	+2.15	+2.88	31.40	11.61	9%
Pooled	+2.15	—	—	—	—

Notes: "Realised at +1sd" is $\hat{\beta}_2$ times the within-bond standard deviation of lnSLQ for that country, expressed in basis points and measuring the bps fee impact of a +1-sd lnSLQ demand shock on a bond in the high-utilisation regime. "Median BA" is the median Markit/Bloomberg bid-ask spread of the bond on the same trading day, in basis points. "Mean spec" is the mean repo specialness DFR $- r_{it}^{\mathrm{repo}}$ for the same panel of bonds in basis points. "Realised / BA" reports the realised fee impact as a share of the country median bid-ask. DE+IT and Pooled aggregate samples lack a single bond-level benchmark and are omitted from the ratio columns.

Predetermined-state and nonlinearity checks. Two robustness families test whether $\ln \mathrm{SLQ} \times D$ simply picks up a mechanical utilization nonlinearity. First, appendix non-parametric tests replace the binary high-utilization indicator with utilization-quintile interactions. For the depth-centrality mechanism, the response is concentrated in the top utilization quintile rather than forming a smooth monotone gradient, which is the activation pattern predicted by the state-dependent mechanism. Second, lagged-state specifications replace contemporaneous D_{it} with predetermined utilization states. Table 6 reports the compact version. The pan-European Test 3a interaction remains positive and significant at lags of 20, 40, 60, and 120 trading days, and the DE+IT depth-centrality triple remains positive over the same horizons. These checks do not make the panel estimate timing-causal, but they show that the demand-side interaction is not only a same-day utilization convexity.

Heterogeneity decomposition. A within-bond residual specification shows that much of the pooled $\ln \mathrm{SLQ} \times D$ signal reflects persistent cross-bond heterogeneity in fee-demand sensitivity. The within-bond, time-varying component survives in Italy but is close to zero in Germany. I therefore use the pan-European and German results mainly as broad state-dependence and cross-sectional centrality evidence, while the Italian sample carries the within-bond policy-shock identification. The appendix reports the full residual decomposition across lags.

Table 6: Predetermined high-utilization states

	Test 3a			Test 3c		
k	$\hat{\beta}$	t	N	$\hat{\beta}$	t	N
20	+2.45	+8.39	699,817	+12.40	+3.12	212,907
40	+2.25	+7.80	673,951	+16.78	+3.63	210,848
60	+2.08	+6.96	650,742	+14.26	+3.14	207,965
120	+1.64	+4.69	592,200	+13.57	+2.30	198,344

Notes: The table replaces the contemporaneous high-utilization indicator with $D_{i,t-k}$ for $k \in \{20, 40, 60, 120\}$ trading days. Test 3a is estimated on the pan-European panel with ISIN and day fixed effects. Test 3c is estimated on the DE+IT post-25%-of-life panel with ISIN and country×day fixed effects and the model-aligned $\tilde{s}^{+,c}_{it}$ object. Standard errors are two-way clustered by ISIN and day.

5.6 Test 3b: The Repo-to-Lending Specialness Pass-Through

Test 3b is the paper's headline structural test: it replaces the lending-side demand proxy $\ln \mathrm{SLQ}_{it}$ of Test 3a with the empirical analog of the model's repo-market scarcity price—the object the fee equation (14) says the lending fee actually prices. That it is the repo price, and not the lending-demand quantity, that the fee tracks is shown directly by a head-to-head on the pooled panel: entering the specialness state interaction $s^+ \times D^{U_{75}}$ and the demand state interaction $\ln \mathrm{SLQ} \times D^{U_{75}}$ in the same regression, the specialness interaction is large and significant (+10.9, $t = +7.7$) while the demand interaction is an order of magnitude smaller (+1.0, $t = +3.1$). The demand interaction of Test 3a is thus the reduced form of the specialness pass-through. The model-aligned object is $\tilde{s}^{+,c}_{it} \equiv \max\{\tilde{s}_{it} - c, 0\}$ with $c = 0.05$ bps and $\tilde{s}_{it} \equiv \mathrm{Specialness}_{i,t-20}$, the 20-trading-day lag of observed specialness. Directional timing tests do not identify a clean daily ordering between the two markets. In lead-lag horse races with ISIN and country-by-day fixed effects, fee regressions on lagged and future $\tilde{s}^{+,c}_{it}$ reject both blocks. The lag block is larger in the full panel (+19.76 versus +10.06, equality $p = 0.010$) and in the high-utilization cell (+40.52 versus +19.31, $p = 0.0004$), but the high-depth split is essentially symmetric (+16.31 versus +16.29, $p = 0.997$). Reverse regressions do not support a repo-leading interpretation: the residualized panel VAR has fee residuals forecasting specialness residuals ($p = 0.001$), while specialness residuals do not forecast fees ($p = 0.239$), with economically tiny coefficient sums. I therefore read the centrality block as cross-sectional mechanism evidence on which bonds load the cross-market channel, not as a within-bond timing-causal estimator; the identified policy-shock evidence is provided by the Italian SLF DiD in Section 5.11. The empirical state-dependence test is

$$\mathrm{Fee}_{it} \;=\; \alpha_i + \alpha^c_t + \beta_1\, \tilde{s}^{+,c}_{it} + \beta_2\, D^{U_{75}}_{it} + \beta_3\, \tilde{s}^{+,c}_{it} \times D^{U_{75}}_{it} + \varepsilon_{it}. \tag{20}$$

The model predicts $\beta_3 > 0$ in the regime $s_{it}^+ > 0$. The relevant structural object is Proposition 2's state-dependent pass-through:

$$\frac{\partial^2 f_{it}^*}{\partial s_{it}^+ \, \partial H_{it}} \;=\; \eta \, \varphi_R \big(\lambda_H + \lambda_{HC} \, C_i\big) \;\geq\; 0,$$

under the parameter restrictions $\eta \in (0,1)$, $\varphi_R \in (0,1]$, $\lambda_H \geq 0$, $\lambda_{HC} > 0$ and the support condition $\lambda_H + \lambda_{HC} C_i \geq 0$ on the empirical range of C_i. Test 3b without centrality is the state-pass-through analog; Test 3c (Section 5.7) adds centrality and tests the cross-partial $\partial^3 f_{it}^* / (\partial s_{it}^+ \, \partial H_{it} \, \partial C_i) = \eta \varphi_R \lambda_{HC} > 0$ from Proposition 3. On the DE+IT post-25% panel, the coefficient on $\tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}}$ is +13.88 ($t = +2.70$), and the length-balanced version is +9.44 ($t = +2.31$). This is consistent with high-intermediation-state activation of the cross-market channel.

5.7 Test 3c: The Deployment-Centrality Contrast

The central test follows Proposition 3. Using the depth-only headline proxy C_i^{depth} defined in equation (9) and the noise-floor-thresholded positive-specialness object $\tilde{s}_{it}^{+,c} \equiv \max\{\tilde{s}_{it} - c, \, 0\}$ with $c = 0.05$ bps, I estimate the model-aligned regression

$$\text{Fee}_{it} \;=\; \alpha_i + \alpha_t^c + \beta_1 \, \tilde{s}_{it}^{+,c} + \beta_2 \, \tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}} + \beta_3 \, \tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}} \times C_i^{\text{depth}} + \text{controls} + \varepsilon_{it}, \quad (21)$$

where $\tilde{s}_{it} \equiv \text{Specialness}_{i,t-20}$ is the 20-trading-day lag of observed specialness, α_t^c is country-by-day fixed effects, $\beta_3 > 0$ is the structural prediction $\beta_3 \approx \eta \varphi_R \lambda_{HC}$ from Proposition 3 (active in the regime $s_{it}^+ > 0$), and the saturated controls term comprises the three two-way interactions $D_{it}^{U_{75}} \times C_i^{\text{depth}}$, $\tilde{s}_{it}^{+,c} \times C_i^{\text{depth}}$, and $\tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}}$, plus the three main effects $D_{it}^{U_{75}}$, C_i^{depth}, and $\tilde{s}_{it}^{+,c}$ entered as separate regressors (C_i^{depth} is collinear with the ISIN fixed effect and absorbed by it; $\tilde{s}_{it}^{+,c}$ has within-bond, within-day variation). The companion continuity specification with raw lagged specialness $\tilde{s}_{it}$ in place of $\tilde{s}_{it}^{+,c}$ is retained as a transparency check; raw $\tilde{s}_{it}$ is no longer the model-facing object. The regression sample is the post-25%-of-life DE+IT panel restricted to lagged-specialness non-missing observations ($N = 212{,}907$, 501 ISINs).

The deployment-centrality contrast (headline). The model's headline prediction (16) is implemented as a between-group contrast rather than a continuous triple, which is more robust to functional form and to the weighting of long-lived bonds. Splitting the DE+IT post-25% panel at the median of C_i^{depth} and estimating the state-dependent specialness pass-through $b(\tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}})$ within each group, the pass-through is large and significant for repo-central bonds (+15.4, $t = +3.0$) and statistically indistinguishable from zero for the rest (+4.2, $t = +1.5$); the between-group contrast is $\delta = +13.9$ ($t = +2.5$), which identifies $\eta \varphi_R \lambda_{HC} (\bar{C}_{\text{hi}} - \bar{C}_{\text{lo}}) > 0$. The contrast is robust across a median split, a top-

versus-bottom-tercile split ($\delta = +12.6$, $t = +2.4$), the positive-specialness cell ($\delta = +14.4$, $t = +2.4$), and the sample-matched utilization threshold; the low-centrality coefficient is *smaller, not zero*, as the model predicts (it is an activation gradient, not an on–off switch). The continuous-triple form of the same cross-partial is positive and significant as a corroboration (reported in the appendix). The contrast is specific to *repo* depth: ranking bonds instead by their predetermined *lending* volume collapses it to a small, insignificant $\delta = +4.1$ ($t = +0.8$), confirming that the grading runs through repo deployment rather than borrowing demand—as the deployment-value model requires. It is also robust to the centrality construction: the raw (un-demeaned) repo-depth level gives $\delta = +13.1$, and a full-sample (rather than predetermined) repo-depth measure gives $\delta = +11.5$.

Robustness of the contrast. The split δ is robust to the concern that high-centrality bonds are mechanically observed for more days. Re-weighting each ISIN to roughly equal total weight (weights $\propto$ 1/post-25% length, normalized to mean one) leaves the point estimate stable—it does not fall, moving from +12.6 to +14.3—while precision declines to $t = +1.5$, the mechanical cost of down-weighting the long-lived, high-centrality bonds that carry the channel. A length-bias-free construction that recomputes centrality on a fixed pre-2022 window and estimates only on 2022–2024, so that no ISIN's centrality calculation uses more observations than another's, returns a contrast of the same sign and comparable magnitude. The result is also insensitive to the noise-floor cutoff $c \in \{0, 0.025, 0.05, 0.10\}$ and, as Section 5.1 discusses, to the EURSTR and EURSTR−5 bp general-collateral benchmarks; full sweeps are in the appendix.

The censoring restriction. The deployment loading is active only where the bond is special: below the cost threshold ($s \leq c$) there is no scarcity rent to deploy across markets, so the fee–specialness pass-through should switch off. It does, cleanly, for the bonds that carry the channel. Estimating the simple fee–specialness slope—without any triple—separately by centrality group and specialness cell (Table 7), repo-central (high-C) bonds load a pass-through of +20.5 ($t = +4.1$) where the bond is special and an indistinguishable-from-zero +1.2 ($t = +0.1$) below the threshold: a textbook switch-off exactly where the model says the channel turns off. Low-centrality bonds load no positive pass-through in either cell (+2.7, $t = +0.8$ when special), consistent with their never deploying the bond; the noisy negative slope they display below the threshold lies outside the model's deployment domain (non-central, non-special bonds) and is not interpreted.

A model-aligned utilization-quintile state-shape reinforces the reading: the pass-through is flat across the bottom four utilization quintiles and carries entirely in the top quintile ($Q_4 = +15.6$, $t = +3.1$; $Q_4 - Q_0$ Wald contrast +14.5, $t = +2.3$; Table 8). The appendix repeats the specialness-based specifications using EURSTR and EURSTR−5 bps as alternative general-collateral benchmarks.

Table 7: Censoring: simple fee–specialness pass-through by centrality and specialness cell (DE+IT, post-25%)

Centrality group	Specialness cell	∂Fee/∂s	t	N
High-C (repo-central)	positive ($\tilde{s} > 0.05$)	+20.5	+4.1	98,029
High-C (repo-central)	non-positive ($\tilde{s} \leq 0.05$)	+1.2	+0.1	43,920
Low-C	positive ($\tilde{s} > 0.05$)	+2.7	+0.8	41,818
Low-C	non-positive ($\tilde{s} \leq 0.05$)	−55.0	−3.5	29,140

Notes: Dependent variable: Fee_{it} in basis points; regressor is the 20-trading-day lag of specialness $\tilde{s}_{it}$ (in percentage points, so the slope is bps per pp). Sample: DE+IT post-25%-of-life panel, 2018–2024, quarter-end excluded; high-/low-C is the median split on C_i^{depth}; ISIN and day fixed effects, two-way clustered standard errors by ISIN and day. The channel-carrying high-C bonds show the predicted switch-off (positive when special, zero below the threshold); the low-C non-positive cell is non-central, non-special bonds outside the model's deployment domain.

The model-aligned specialness object. The structural deployment rent in Section 4 is the net specialness $s_{it}^{+} = \max\{s_{it} - c, 0\}$, and the model's cross-partial is $\partial^3 f^*/(\partial s^{+} \partial H \partial C) = \eta \varphi_R \lambda_{HC} > 0$. The empirical analog used throughout this section is

$$\tilde{s}_{it}^{+,c} \equiv \max\{\tilde{s}_{it} - c,\, 0\}, \qquad c = 0.05 \text{ bps}, \tag{22}$$

the noise-floor-thresholded 20-trading-day lag of specialness; the raw lag $\tilde{s}_{it}$ enters only as a transparency check. The split δ and the censoring test above are estimated on $\tilde{s}_{it}^{+,c}$.

A common-scarcity saturation check confirms the deployment-centrality grading is not a generic scarcity proxy: in the continuous-triple specification, adding lending scarcity, utilization, on-the-run status, repo volume, repo transaction count, and their high-utilization interactions moves the centrality interaction only marginally (from +15.1, $t = +3.4$, to +14.7, $t = +3.4$). The deployment-depth pattern is therefore not absorbed by common-scarcity controls, and the appendix threshold sweep shows it is not driven by the 0.05 bps noise-floor choice.

The state-shape is concentrated at the top utilization quintile. Bin slopes $\hat{\phi}_0, \ldots, \hat{\phi}_3$ are individually indistinguishable from zero on the depth-only proxy; only $\hat{\phi}_4 = +14.42$ ($t = +3.05$) is significant at conventional levels, and the $Q_4 - Q_0$ Wald contrast is +13.14 ($t = +2.08$, significant at 5%). This is consistent with the model's prediction that the cross-market channel activates in the high-intermediation regime $H_{it} = 1$ rather than scaling smoothly with utilization across the support. I read this as top-utilization activation evidence and do not interpret the lower-quintile null as a power statement on a continuous gradient.

Model-aligned utilization-quintile state-shape. Replacing $\tilde{s}_{it}$ with $\tilde{s}_{it}^{+,c} = \max\{\tilde{s}_{it} - 0.05, 0\}$ in the bin-specific specification produces a comparable top-quintile activation pattern: $Q_0, \ldots, Q_3$ slopes remain individually indistinguishable from zero, Q_4 carries the

Table 8: Depth-only centrality by utilization quintile: top-utilization activation (DE+IT, post-25%)

Utilization quintile	$\hat{\phi}_q$ (bin slope)	SE	t	N
Panel A. Bin-specific $\tilde{s} \times C^{\text{depth}}$ slope by utilization quintile				
Q_0 (lowest)	+1.28	4.19	+0.30	212,907
Q_1	−3.04	4.40	−0.69	212,907
Q_2	+1.97	2.71	+0.73	212,907
Q_3	+7.22	4.52	+1.60	212,907
Q_4 (highest)	+14.42	4.72	+3.05	212,907
Panel B. Wald contrast: $Q_4 - Q_0$				
$Q_4 - Q_0$	+13.14	6.31	+2.08	212,907

Notes: Dependent variable: Fee_{it} in basis points. Sample: DE+IT post-25%-of-life panel, 2018–2024, quarter-end excluded. Utilization quintiles $Q_0, \ldots, Q_4$ are constructed within country (country-quantile cuts on Util_{it}). Each row of Panel A reports the bin-specific slope $\hat{\phi}_q$ on $\mathbf{1}\{\text{Util_qbin}_{it} = q\} \times \tilde{s}_{it} \times C_i^{\text{depth}}$ from a single regression that includes all five bin slopes simultaneously (no global $\tilde{s} \times C$ term; the Q_0 dummy is absorbed by the ISIN fixed effect). The specification includes ISIN and country×day fixed effects with bin-specific $\tilde{s}$ and C^{depth} doubles for $Q_1, \ldots, Q_4$; standard errors are two-way clustered by ISIN and day. Panel B reports the linear Wald contrast $\hat{\phi}_4 - \hat{\phi}_0$ with covariance-based standard error from the same fitted regression. The pattern is consistent with top-utilization activation rather than a smooth monotone gradient: $Q_0, \ldots, Q_3$ are individually indistinguishable from zero, while Q_4 carries the channel and the $Q_4 - Q_0$ contrast is significant at the 5% level. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

channel at +15.55 ($t = +3.06$), and the $Q_4 - Q_0$ Wald contrast strengthens slightly to +14.54 ($t = +2.27$). Full results are in the appendix; the existing raw-$\tilde{s}$ table is retained as the continuity diagnostic. The model-aligned state-shape therefore reinforces rather than weakens the top-utilization activation reading.

5.8 Cross-Market Simultaneity Test

Tests 1–3 establish separate channel components: lending-market scarcity is associated with higher repo activity (Test 1), unconditional repo volume does not move fees (Test 2), and the interaction $\ln \text{SLQ} \times D$ moves fees at high utilization (Test 3). The model predicts that the same demand shock should *simultaneously* move both markets. To test this directly I estimate the state-dependent $\ln \text{SLQ} \times D^{U_{75}}$ interaction separately on four cross-market outcomes, lending fee, repo specialness, $\ln \text{RepoVol}$, and $\ln(1 + \text{RepoCount})$, on the DE+IT panel with two-way (ISIN, day) clustered standard errors. Table 9 reports the result.

Across four separately estimated market outcomes, the same high-intermediation lending-demand state is associated with higher securities-lending fees, higher repo specialness, greater repo volume, and more repo transactions. The coordinated signs survive two-way clustering by ISIN and day. The specialness response is the marginal cell of the four ($\hat{\theta} = +0.0117$, $t = +2.20$, just above the 5% threshold), while the fee, $\ln \text{RepoVol}$, and

Table 9: Cross-market simultaneity, marginal regressions (DE+IT)

Outcome	$\hat{\theta}$	SE	t	N
Fee_{it} (bps)	+4.56	0.66	+6.94	281,110
Specialness_{it} (bps)	+0.0117	0.0053	+2.20	281,110
$\ln \text{RepoVol}_{it}$	+0.158	0.020	+7.84	281,110
$\ln(1 + \text{RepoCount}_{it})$	+0.082	0.011	+7.56	281,110
ISIN FE	Yes			
Country×day FE	Yes			
SE cluster	ISIN × day (two-way)			

Notes: Dependent-variable units: Fee_{it} in basis points; Specialness_{it} in basis points; $\ln \text{RepoVol}_{it}$ and $\ln(1 + \text{RepoCount}_{it})$ in natural logs. Each row is a separate two-way fixed-effects regression of the listed outcome on $\ln \text{SLQ}_{it} \times D_{it}^{U_{75}}$, with $\ln \text{SLQ}_{it}$ and $D_{it}^{U_{75}}$ as controls, ISIN and country×day fixed effects, and standard errors two-way clustered by ISIN and day. Threshold: $D_{it}^{U_{75}} = \mathbf{1}[\text{Util}_{it} > 41.87\%]$. Coefficient $\hat{\theta}$ is the state-dependent shock loading on each outcome. Sample: DE+IT, 2018–2024, quarter-end excluded. The summed-squared-t statistic across the four outcomes equals 171.6; this statistic is reported as a coordination diagnostic only (the four marginal regressions share an RHS regressor and have correlated residuals across LHS variables, so a χ^2_4 critical value would overstate joint significance), and inference is based on the four marginal t-statistics. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

ln(1+RepoCount) rows are all at the $|t| > 6.9$ level; the joint cross-market signal therefore rests primarily on the three high-precision rows, with specialness providing additional but weaker corroboration.

Joint cross-market test with proper clustering. The four marginal regressions in Table 9 share the same right-hand side, so their residuals are correlated across outcomes for the same ISIN-day. To test the joint null that all four outcome-specific coefficients are zero with a covariance matrix that respects this dependence, I stack the four outcomes into a single long-format panel of length $4N$ (one observation per ISIN-day-outcome cell), standardize each outcome (subtract its mean, divide by its standard deviation), and estimate a single regression with outcome-specific shock coefficients and outcome×ISIN and outcome×day fixed effects, with standard errors two-way clustered by ISIN and day across the stacked system. Table 10 reports the result.

The four standardized outcome-specific coefficients are individually significant at $|t| > 3.7$ and the joint Wald test rejects the null of zero coefficients at $\chi^2(4) = 100.15$ ($p < 10^{-20}$). The high-intermediation lending-demand state is associated with a coordinated state-dependent response across the four outcomes on the same ISIN-day, and this conclusion does not rest on inflated standard errors from outcome-by-outcome regressions ignoring the outcome residual covariance.

A within-country shuffle of $\ln \text{SLQ}_{it} \times D_{it}^{U_{75}}$ across ISINs (200 replications, DE+IT panel) produces a placebo distribution with χ^2-95th-percentile of 8.95 on a tightened specification (no outcome-specific main effects) whose empirical χ^2 is 163.54; the plus-one adjusted

Table 10: Stacked-system joint cross-market test, properly clustered

Outcome (standardized)	$\hat{\theta}_m^{\text{std}}$	SE	t
Fee_{it} (bps)	+0.211	0.028	+7.41
Specialness_{it} (bps)	+0.099	0.026	+3.76
$\ln \text{RepoVol}_{it}$	+0.213	0.030	+7.11
$\ln(1 + \text{RepoCount}_{it})$	+0.188	0.025	+7.65
Joint test $H_0 : \theta^{\text{std}}_{\text{Fee}} = \theta^{\text{std}}_{\text{Spec}} = \theta^{\text{std}}_{\text{RepoVol}} = \theta^{\text{std}}_{\text{RepoCount}} = 0$			
$\chi^2(4)$ statistic			100.15
p-value			$< 10^{-20}$

Notes: Sample: DE+IT panel, 2018–2024, quarter-end excluded ($N = 281{,}110$ ISIN-days, stacked into $4N = 1{,}124{,}440$ outcome-ISIN-day cells). Each outcome is standardized to mean zero and unit variance before stacking, so $\hat{\theta}_m^{\text{std}}$ is the standardized response of outcome m to a one-unit increase in $\ln \text{SLQ}_{it} \times D_{it}^{U_{75}}$. Stacked regression includes outcome-specific shock coefficients (θ_m^{std} for $m \in \{\text{Fee}, \text{Spec}, \ln \text{RepoVol}, \ln \text{RepoCount}\}$), outcome-specific main effects on $\ln \text{SLQ}_{it}$ and $D_{it}^{U_{75}}$, outcome×ISIN fixed effects, and outcome×day fixed effects. Standard errors are two-way clustered by (un-stacked) ISIN and day across the full $4N$ panel. The joint $\chi^2(4)$ is computed from the stacked-system covariance matrix on $\widehat{\boldsymbol{\theta}}^{\text{std}} = (\hat{\theta}^{\text{std}}_{\text{Fee}}, \hat{\theta}^{\text{std}}_{\text{Spec}}, \hat{\theta}^{\text{std}}_{\text{RepoVol}}, \hat{\theta}^{\text{std}}_{\text{RepoCount}})$ with the matrix $R = I_4$ and $q = \mathbf{0}$. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

placebo p-value is 0.005, rejecting the null that random reassignment of the shock label across same-country ISINs would generate the joint cross-market signature. The empirical χ^2 in Table 10 (100.15) is also far above the placebo 95th percentile, so the joint cross-market response is not a label-permutation artifact.

5.9 Discriminating the Deployment-Value Channel from Within-Market Alternatives

The headline specialness pass-through and the centrality contrast are consistent with the deployment-value mechanism, but they also admit competing within-market data-generating processes. The decisive discriminators are those a purely within-market mechanism cannot reproduce: the pass-through vanishes where the deployment loading $\Lambda(C_i, H_{it})$ is structurally zero (the negative-specialness test below), the same demand state moves the repo market as well as the lending fee (the cross-market simultaneity test), and the pass-through is graded by repo-deployment centrality (Test 3c).

Negative-specialness discriminating test. The deployment-value loading $\Lambda(C_i, H_{it})$ vanishes whenever specialness is non-positive: at $s_{it} \leq 0$ the bond is not scarce in repo, no deployment rent exists, and the model predicts the state-dependent demand interaction to shut down on the boundary. A pure haircut-arbitrage mechanism in which the dealer earns reuse rents regardless of specialness sign would predict the interaction to persist into the non-positive cell. Estimating the headline Test 3 specification separately on the $s > 0$ and $s \leq 0$ cells gives the discrimination cleanly (Table 11). In the positive-specialness

cell, the interaction reproduces the full-sample headline ($\hat{\beta}_2 = +2.10$, $t = +5.65$ pooled; +2.15 DE, $t = +3.43$; +5.10 IT, $t = +4.34$). In the non-positive cell, the interaction collapses to insignificance across every sample ($t = +1.35$ pooled, +0.79 DE, +0.33 IT, +1.29 DE+IT); the point estimate falls to roughly 45% of the positive-cell value with 1.9× the standard error. The pattern is exactly what the deployment-value model predicts and is inconsistent with a generic haircut-arbitrage mechanism.

Table 11: Negative-specialness discriminating test: deployment value vs. haircut arbitrage

Cell	Sample	$\hat{\beta}_2$ ($\ln \mathrm{SLQ} \times D$)	SE	t
Full panel (Table 4 headline reproduction)				
all	Pooled	+2.154	0.365	+5.90
all	DE	+2.140	0.621	+3.44
all	IT	+5.097	1.108	+4.60
all	DE+IT	+3.909	0.633	+6.17
Positive-specialness cell ($s_{it} > 0$)				
pos_spec	Pooled	+2.103	0.372	+5.65
pos_spec	DE	+2.152	0.628	+3.43
pos_spec	IT	+5.104	1.176	+4.34
pos_spec	DE+IT	+3.652	0.652	+5.60
Non-positive cell ($s_{it} \leq 0$, $\Lambda \equiv 0$)				
neg_spec	Pooled	+0.948	0.704	+1.35
neg_spec	DE	+0.938	1.183	+0.79
neg_spec	IT	+0.596	1.810	+0.33
neg_spec	DE+IT	+1.743	1.351	+1.29

Notes: Sample: pan-European panel, 2018–2024, quarter-end excluded. Specifications match Table 4 (ISIN and day fixed effects, two-way clustered SE by ISIN and day). The positive- and non-positive-specialness cells are mutually exclusive partitions of the panel on contemporaneous $s_{it} = \mathrm{DFR}_t - r_{it}^{\mathrm{repo}}$. The non-positive cell is the cell in which the deployment-value loading $\Lambda(C_i, H_{it})$ is structurally zero by the model.

Futures delivery does not subsume the channel. Cheapest-to-deliver status is the natural alternative grading variable, since CTD bonds carry the strongest futures-delivery specialness. Constructing $\mathrm{PastCTD}_{it}$ as ever-CTD strictly before t from the Eurex contract feed and racing the past-CTD triple $\tilde{s}_{it}^{+,c} \times D_{it}^{U_{75}} \times \mathrm{PastCTD}_{it}$ against the repo-depth centrality triple, the centrality triple survives (+11.0, $t = +2.3$) while the CTD triple collapses to insignificance (+4.9, $t = +1.0$), though each is significant standalone (+15.1 and +11.8). Repo-deployment centrality absorbs the notified-CTD channel, not the other way around: delivery history is not a separate state-dependent moderator beyond how heavily the bond is deployed in repo.

Identification. A generic within-market inventory or balance-sheet model can reproduce the within-market state-dependence of the fee, but not the cross-market moments: the cross-market simultaneity, the deployment-centrality contrast, the unconditional de-

mand sign, and the country-specific signs of the SLF/QT difference-in-differences each require the cross-market deployment channel that a within-market model lacks. The appendix reports the full identification mapping across competing mechanisms.

5.10 Liquidity as a Moderator: Pointer to Appendix

A fourth-margin test interacts the headline state-dependent regression with bid-ask quartile dummies. On the Italian subsample, the state-dependent interaction is largest for the most-illiquid bonds ($\hat{\beta}_2^{IT,Q_1} = 10.07$, $t = 4.99$) and attenuates to the most-liquid quartile ($\hat{\beta}_2^{IT,Q_4} = 5.55$, $t = 3.19$), a $\sim 1.8\times$ ratio consistent with binding dealer-capacity constraints when intermediation is more costly. On the DE-only subsample, a Wald test of $\hat{\beta}_2^{DE,Q_1} = \hat{\beta}_2^{DE,Q_4}$ fails to reject equality ($p = 0.222$), consistent with bid-ask compression in the Bund benchmark. This is a separate moderator from the borrower-concentration and collateral-transformation splits already reported. Full results and per-quartile t-statistics in the appendix.

5.11 Difference-in-Differences: ECB Securities Lending Facility

The ECB publishes a list of government bond ISINs available for lending through its Securities Lending Facility (SLF) to eligible Eurosystem counterparties. Two recent papers establish the SLF as a collateral-supply shifter on the public side: Carrera de Souza and Hudepohl (2024) document that Eurosystem purchases depress repo rates through both flow and nonlinear stock effects and that securities lending partly offsets this scarcity, and Greppmair and Jank (2026) use the November 2020 SLF pricing reduction as a quasi-natural experiment, finding that cheaper public-collateral supply nearly doubled facility usage and expanded private repo activity through a collateral-multiplier mechanism rather than substituting for it. I exploit the same institutional object on a different shock, the July 2023 QT onset, to test the private-side transmission to securities-lending fees, which neither paper studies. During QT the ECB reduces its holdings of SLF-listed bonds, shrinking the latent supply buffer that the SLF provides and tightening lending-market conditions for SLF-listed bonds relative to non-SLF bonds. At the same time, the unwinding of Eurosystem holdings returns bonds to private repo markets, affecting bond-specific repo conditions. Both channels operate on SLF-listed bonds after QT onset, and the DiD exposure therefore captures a composite treatment; the interpretation-and-limitations paragraph below returns to this point. The exposure is:

$$Z_{it} = \text{SLF}_i \times \text{PostQT}_t, \tag{23}$$

where SLF_i indicates whether ISIN i is on the SLF list and $\text{PostQT}_t = \mathbf{1}[t \geq \text{July 2023}]$. I estimate a triple-difference[6] that interacts Z_{it} with the high-utilization indicator D_{it}:

$$\text{Fee}_{it} = \alpha_i + \alpha_t + \tau_1 Z_{it} + \tau_2 Z_{it} \times D_{it} + \mathbf{X}'_{it}\boldsymbol{\gamma} + \varepsilon_{it}, \tag{24}$$

where $\mathbf{X}_{it}$ stacks the high-utilization indicator D_{it}, the specialness control Spec_{it}, and the unabsorbed lower-order interactions $D_{it} \times \text{PostQT}_t$ and $\text{SLF}_i \times D_{it}$. The level terms SLF_i and PostQT_t are absorbed by ISIN and day fixed effects, respectively. The sign prediction is country-specific. For Italian BTPs, the fee response to the QT-induced supply contraction should be amplified at high utilization when the deployment-value channel is active ($\tau_2^{IT} > 0$). The German DiD is sign-consistent but is not a second identified causal cell, since its pre-trends are not flat.

Results. Table 12 reports the Italian and German estimates. For Italian BTPs, the simple DiD is small and statistically indistinguishable from zero ($\hat{\tau}_1^{IT} = +4.19$, $t = +0.68$), while the triple interaction is positive at the 1% level ($\hat{\tau}_2^{IT} = +6.45$, $t = 3.56$). The estimate indicates a state-dependent response to the composite SLF-listing × QT treatment, with the caveats in Section 5.11 below. Italian event-study pre-trends are flat (none of the 23 pre-QT months significant; mean $|t| = 0.70$), which supports the parallel-trends assumption; Figure 7 plots the event-study coefficients γ_k around the QT onset. The Italian estimate is the paper's identified causal estimate of a state-dependent lending-fee response—the fee-side counterpart to the auction evidence on the repo side. For German Bunds, the simple DiD is positive and marginal ($\hat{\tau}_1^{DE} = +10.57$, $t = 1.83$), but this estimate is confounded by a pre-existing upward trend (five of 23 pre-QT months significant, mean $|t| = 1.40$). The triple interaction is negative for Germany ($\hat{\tau}_2^{DE} = -4.15$, $t = -3.29$), consistent with a ceiling effect at high utilization where Bunds were already scarce pre-QT. Germany is therefore sign-consistent with the high-scarcity ceiling interpretation rather than an independent identifying test. I do not report a pooled euro-area DiD as a headline estimate: country-level effects are heterogeneous and the pooled coefficient is sensitive to the country mix. A pooled DE+IT diagnostic with country interactions gives $\hat{\tau}_2^{IT} = +6.77$ ($t = +3.89$) and $\hat{\tau}_2^{DE} = -5.58$ ($t = -4.01$); a Wald test rejects equality with $\chi^2(1) = 30.35$, $p < 0.0001$, formalizing the sign asymmetry.

[6]The triple difference-in-differences design exploiting (i) time variation, (ii) variation between treated and control units in the relevant cross-section, and (iii) variation between treated and control outcomes within the treated unit is used by Vuillemey (2020) to identify the effect of an institutional innovation (central clearing in 19th-century Le Havre coffee futures) on trade flows in cleared versus uncleared commodities. My setting is structurally analogous: the ECB's QT onset varies over time, SLF-listed versus non-SLF bonds provide the treated-versus-control cross-section, and the high-utilization indicator D_{it} provides the within-bond dimension over which the mechanism predicts heterogeneous response.

Table 12: ECB SLF Difference-in-Differences: State-Dependent Fee Response

	(1) Italy	(2) Germany
$\hat{\tau}_1$ (Simple DiD: Z)	+4.19	+10.57*
	(+0.68)	(+1.83)
$\hat{\tau}_2$ (Triple DiD: $Z \times D$)	+6.45***	−4.15***
	(+3.56)	(−3.29)
Specialness control	Yes	Yes
ISIN, day FE	Yes	Yes
Pre-QT parallel trends	0/23 sig.	5/23 sig.†
N (ISIN-days)	161,905	123,692

Notes: Dependent variable: Fee_{it} in basis points. t-statistics in parentheses, computed from two-way clustered standard errors (ISIN, day). $Z_{it} = \text{SLF}_i \times \text{PostQT}_t$ where $\text{PostQT}_t = \mathbf{1}[t \geq \text{July 2023}]$ and SLF_i indicates inclusion on the ECB's SLF list. $D_{it} = \mathbf{1}[\text{Util}_{it} > p_{75}]$ using sample-matched p_{75}: 38.0% for Italy (column 1) and 50.6% for Germany (column 2). Specifications include ISIN fixed effects, day fixed effects, repo specialness control, and unabsorbed lower-order interactions. An ISIN-level wild-cluster bootstrap of $\hat{\tau}_2^{IT}$ on the column (1) sample ($B = 999$) yields Rademacher $p = 0.001$ and Webb $p = 0.003$, yielding the same statistical conclusion under finite-cluster inference. †The German pre-QT trend is statistically detectable (five of 23 pre-period months significant, mean $|t| = 1.40$), so the German estimates should be interpreted as sign-consistent corroboration rather than as an independently identified causal cell; the appendix reports matched-control and donor-pool-placebo diagnostics. A pooled euro-area DiD is not reported as a headline estimate: country-level effects are heterogeneous and the pooled coefficient is sensitive to the country mix; per-country DiDs and the DE/IT sign-asymmetry test are reported in the appendix. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

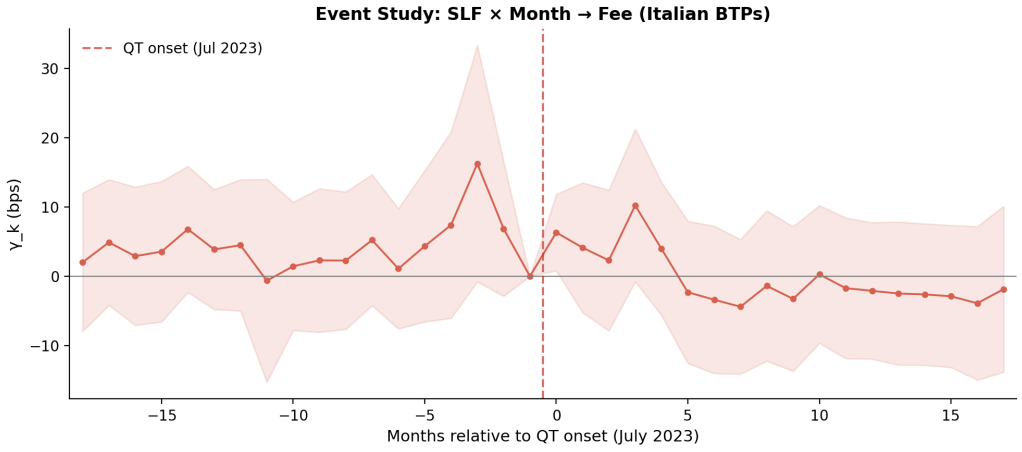


Figure 7: Event study: SLF × month → lending fee (Italian BTPs). Each point is the coefficient γ_k from $\text{Fee}_{it} = \alpha_i + \alpha_t + \sum_k \gamma_k \cdot \text{SLF}_i \cdot \mathbf{1}[\text{month}_t = k] + \phi \cdot \text{Spec}_{it} + \varepsilon_{it}$, with $k = -1$ (June 2023) as the base period. Shaded band: 95% CI. Dashed vertical line: QT onset (July 2023). Pre-QT coefficients are individually indistinguishable from zero at conventional levels (mean $|t| = 0.70$ across $k < -1$). Post-QT fee effect peaks at $k = +3$ and decays thereafter, consistent with a gradually absorbed supply shock. 106 SLF versus 217 non-SLF Italian BTPs; $N = 161{,}905$ ISIN-days.

Interpretation and scope. The Italian SLF/QT cell is the paper's identified causal estimate of a state-dependent lending-fee response, and the fee-side counterpart to the auction evidence on the repo side. The German estimates are useful only as sign-consistent corroboration because their pre-trends are not flat. A pooled DE+IT sign-asymmetry test confirms that the Italian and German state-dependent coefficients differ sharply, but the causal weight remains on Italy. Appendix tables report matched-control diagnostics for Germany, within-Italy heterogeneity, centrality-moderated DiD, and the data limitation that the panel does not contain corporate-bond, equity-lending, or non-eurozone-sovereign off-mechanism placebo assets.

5.12 Cross-Country SLF/QT Evidence

To test whether the SLF/QT fee response is specific to the German and Italian cells, I extend the difference-in-differences to every eurozone sovereign with non-trivial Markit coverage and report the country-by-country result.

Country-stratified ECB SLF DiD. I run the headline DiD specification on each country's subsample, with the Markit-reporting-lag correction (footnote 4) and country×day fixed effects, and report the result in Table 13.[7]

Table 13: Country-stratified ECB SLF DiD by eurozone country

Country	$\hat{\tau}_1$ (t)	$\hat{\tau}_2$ (t)	sig. at 5%?	N
Significant state-dependent response				
DE	+7.94 (+4.00)	−4.02 (−2.45)**	yes	123,692
IT	−5.73 (−2.03)	+7.40 (+3.30)***	yes	161,905
FR	+1.97 (+0.99)	−7.31 (−4.69)***	yes	132,892
No significant response				
NL	−1.26 (−0.37)	+0.01 (+0.00)	no	41,183
ES	−1.64 (−0.56)	+2.00 (+1.23)	no	88,615
BE	+6.78 (+1.26)	−4.33 (−1.94)*	no†	54,875
AT	−4.46 (−1.99)	+1.56 (+0.62)	no	42,617
FI	−16.17 (−1.29)	−4.23 (−1.03)	no	33,401
PT	+1.22 (+0.27)	+2.27 (+0.43)	no	23,732

Notes: $\hat{\tau}_2$ is the state-dependent coefficient $\text{SLF}_i \times \text{PostQT}_t \times D_{it}$ on the country's panel, with ISIN and country×day fixed effects and two-way clustered SE by ISIN and day. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. †Belgium is the marginal case: $\hat{\tau}_2 = -4.33$ ($t = -1.94$, significant at 10% but not 5%). Sources: country panels of the matched Markit–BrokerTec/MTS data.

[7] The country-stratified specification uses a slightly tightened panel construction relative to Table 12: the lag-2 Markit-reporting correction is applied within country, and country×day fixed effects replace day fixed effects to absorb residual country-level common shocks. The DE and IT point estimates in Table 13 ($\hat{\tau}_2^{\text{DE}} = -4.02$, $\hat{\tau}_2^{\text{IT}} = +7.40$) therefore differ slightly from the DE/IT-only headline DiD coefficients in Table 12 (−4.15, +6.45), but the sign and significance pattern is preserved across the two specifications. The cross-country interpretation rests on the Table 13 pattern; Table 12 is the primary DE+IT identified policy-shock cell.

The SLF/QT response is detectable in three countries—DE, IT, and FR, each significant at the 5% level—and statistically indistinguishable from zero in the remaining six. Among the three, $\hat{\tau}_2$ is positive in IT and negative in DE and FR, and a Wald test rejects equality of $\hat{\tau}_2^{\mathrm{DE}}$ and $\hat{\tau}_2^{\mathrm{IT}}$ at $p < 0.0001$. The opposite signs are informative in a limited but robust way: a single common Europe-wide shock cannot move German and Italian fees in opposite directions, so the pattern rules out a pan-European confound. They do not, however, identify the deployment channel on their own. The SLF/QT shock bundles a public lending-supply margin and a repo-deployment margin, and I do not claim to predict the country-by-country sign from primitives; the cross-country determinant of the sign is left as an open pattern. The clean, sign-stable repo-side evidence comes instead from the auction reopenings of Section 5.14, which relieve specialness with the *same* sign in all three countries.

5.13 Alternative Mechanisms for the Cross-Country Sign Pattern

The cross-country sign asymmetry in the SLF/QT DiD ($\hat{\tau}_2^{\mathrm{IT}} > 0$, $\hat{\tau}_2^{\mathrm{DE}} < 0$, $\hat{\tau}_2^{\mathrm{FR}} < 0$, with the remaining six eurozone sovereigns indistinguishable from zero) is a sharp test of competing mechanisms. Four candidate alternatives fail to predict the opposite-sign pattern across Germany and Italy under any single shock.

Alternative 1: liquidity-premium dynamics under QT. A QT-induced reduction in central-bank balance-sheet support could differentially raise liquidity premia in less-liquid sovereigns: BTPs (lower liquidity) would experience the largest fee increase, Bunds (most liquid) the smallest. This predicts $\hat{\tau}_2^{\mathrm{IT}} > \hat{\tau}_2^{\mathrm{DE}}$ (consistent with the data in magnitude order), but it predicts both to be positive, not opposite signs. The opposite-sign data therefore reject this mechanism.

Alternative 2: PSPP / PEPP rebalancing. Differential portfolio rebalancing under QT would predict $\hat{\tau}_2$ signs correlated with the country's ECB holdings share or with the PSPP unwind speed. PSPP unwind is roughly proportional to GDP share, which predicts similar (not opposite) signs across DE and IT. PEPP unwind is not active in the sample. Rejected.

Alternative 3: MTS-versus-Eurex routing artifact. Italian repo trades dominantly on MTS, German repo on Eurex and BrokerTec; differential routing-of-trade responses to QT could mechanically alter measured cross-country repo volumes. The empirical ATS coverage diagnostic on the analysed panel rejects this: 100% of ISIN-days in both DE and IT have positive ATS-reported repo activity, with 38.6 trades/ISIN-day in DE versus 37.3 in IT and volume-to-lendable ratios of 0.25% and 0.21% respectively.

Coverage is uniform across the 2018–2024 sample. The routing alternative is empirically rejected.

Alternative 4: country-specific SLF supply scale. The SLF withdrawal scale is roughly proportional to each country's prior SLF usage, giving quantitative magnitudes but not opposite signs. The opposite-sign requirement is what defeats this alternative.

Interpreting the cross-country signs. The model can rationalize opposite signs through country-specific bargaining and matching primitives, but this is a consistency check rather than an ex-ante prediction: the mapping from primitives to the sign is not sharp enough in the data to forecast which country goes which way. I therefore treat the cross-country signs as ruling out a common Europe-wide confound and as consistent with country-specific deployment primitives, and rely on the auction reopenings (Section 5.14) for the clean, sign-stable repo-side evidence.

5.14 Test 4: Auction Reopenings as an Exogenous Repo-Supply Shock

The observational pass-through of Tests 3b–3c shares a common-factor problem: repo specialness and the lending fee are both driven by collateral demand, so the cross-sectional association cannot by itself separate the deployment channel from common demand. The model's repo side, however, makes a sharp prediction that an exogenous supply shock can test directly. Specialness is the scarcity price of the specific bond, so an increase in the bond's available supply must lower it. Treasury auction reopenings provide that shock: a reopening—a tap of an already-outstanding line—is an exogenous, calendar-set increase in the bond's collateral supply, announced well in advance and uncorrelated with the bond's contemporaneous lending-fee innovations.

Using the harmonized auction calendars of the German, French, and Italian debt-management offices (2018–2024), I estimate the dynamic effect of a reopening with local projections (Jordà, 2005). For each horizon h, I regress the bond's specialness h trading days from the reopening on the log allotment size, controlling for lagged specialness—which absorbs the pre-auction scarcity cycle—with ISIN and day fixed effects and standard errors clustered by ISIN. Figure 8(a) plots the impulse response. The pre-shock horizons ($h < 0$) are flat and insignificant, so the relief is not the continuation of a trend; specialness falls on impact, reaches -0.055 pp at $h = 2$ ($t = -8.5$), and mean-reverts over roughly two weeks. The relief scales with the exogenous shock—a one-log-unit larger allotment adds -0.012 pp of relief ($t = -2.4$)—and carries the *same* sign in all three countries in country-level event-window estimates (-0.039 pp DE, -0.021 pp FR, -0.034 pp IT, each $t < -6$), with null placebos at pseudo-event dates 60 and 90 business days before each reopening.

Specialness is therefore a genuine, supply-driven repo price, not a passive correlate of the fee.

The supply increase is directly observable in the secondary cash market. Re-estimating the same local projection on interdealer cash-market volume for these bonds, trading volume jumps by about one log unit on the auction day ($t = +19$) and the probability that the bond trades at all rises by 27 percentage points, both decaying within days (Figure 8(b)). The reopening visibly lands new collateral on dealer balance sheets—a transparent first stage for the supply-to-specialness channel.

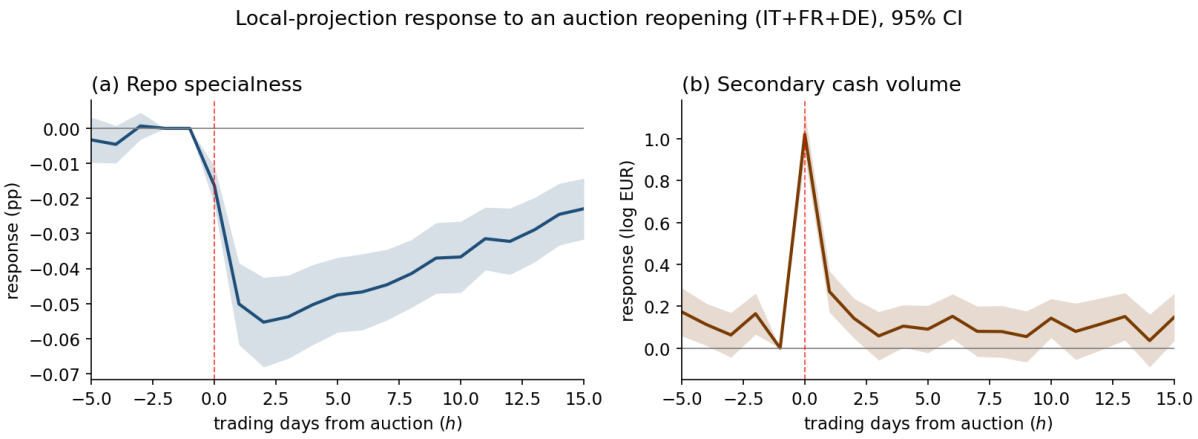


Figure 8: Local-projection response to an auction reopening, Germany, France, and Italy pooled, 2018–2024. Jordà (2005) local projections of (a) repo specialness (pp) and (b) interdealer secondary-market cash volume (log EUR) on the log reopening allotment, by trading-day horizon h relative to the auction. Each horizon-h regression controls for lagged specialness and includes ISIN and day fixed effects; shaded bands are 95% confidence intervals from ISIN-clustered standard errors. Horizons $h < 0$ serve as a placebo. Source: S&P Global Securities Finance and BrokerTec/MTS (specialness), interdealer MTS cash trades (volume), and the DE/FR/IT debt-management-office auction calendars.

The lending fee does not respond, and this null is a feature of the design. The same local projection run on the fee is flat and insignificant at every horizon, as is its interaction with the high-utilization state. A reopening is a *transitory* supply shock—specialness reverts within two weeks—whereas the fee prices the bond's *persistent* deployment value, so it correctly ignores a transient specialness dip. The auction therefore identifies the repo-side mechanism (supply → specialness) cleanly and exogenously but cannot identify the fee pass-through; that requires a *persistent* shift in deployment value, which is precisely what the SLF/QT regime change of Section 5.11 supplies. The two quasi-experiments are complementary by construction: auctions move the transitory repo price, the SLF/QT regime moves the persistent fee.

5.15 Empirical Synthesis

The main empirical evidence has six parts. First, lending-market demand and repo activity are linked at the same ISIN-day level (Test 1). Second, the paper's headline is the repo-to-lending specialness pass-through, graded by deployment centrality (Tests 3b–3c): the lending fee's sensitivity to repo specialness is large for repo-central bonds and small for the rest, with a between-group contrast of $\delta = +13.9$ ($t = +2.5$) that identifies the model's cross-partial $\eta\varphi_R\lambda_{HC}$; in a head-to-head the specialness state interaction dominates the lending-demand interaction by an order of magnitude. The contrast is robust to length-balancing (the point estimate is stable, with precision falling as the long-lived high-centrality bonds are down-weighted) and to a length-bias-free fixed-pre-period construction. Third, the demand-side counterpart—the lnSLQ$\times D$ interaction (+2.15 bps, $t = +5.90$), the reduced form of the specialness pass-through—is itself economically meaningful, equal to 22–44% of the bond's median bid-ask spread and of the order of 10^5 per dealer-bond in annual lending revenue (125k–172k in Germany and Italy). Fourth, the deployment channel is discriminated from within-market alternatives: the pass-through vanishes off the positive-specialness domain where Λ is inactive, the fee tracks the repo *price* rather than repo volume, the cross-market simultaneity test moves fees, specialness, and repo volume together, and cheapest-to-deliver bonds carry no separate futures-delivery channel. Fifth, in local projections Treasury auction reopenings—exogenous, calendar-set increases in collateral supply—relieve specialness with the same sign in DE, FR, and IT and with flat pre-trends, while leaving the fee unmoved: this transitory supply shock prices into specialness but not the fee, confirming that specialness is a genuine supply-driven repo price and that the fee tracks persistent rather than transitory deployment value (Test 4). Sixth, the ECB SLF/QT episode delivers the paper's identified causal estimate of a state-dependent lending-fee response: in Italy, SLF-listed BTPs show a +6.45 bp higher fee response in the loaded state after QT, with flat pre-trends (wild-cluster-bootstrap $p = 0.001$); the opposite-signed responses across DE, IT, and FR additionally rule out a common Europe-wide confound, though the bundled nature of the shock means the cross-country sign is not claimed as a structural prediction. Together with the auction evidence, the two episodes identify the channel's two halves—the repo price (auctions) and the fee response (the Italian SLF/QT cell).

The appendix contains the remaining diagnostics: utilization-nonlinearity tests, lagged-state checks, threshold sweeps, the general-collateral-benchmark robustness, March 2023 flight-to-quality descriptives, the identification mapping across competing mechanisms, the non-additive OnRun specification check that leaves the OffRun centrality contrast unchanged, the supply/demand decomposition of lnSLQ confirming that the demand-side interaction is dominantly demand-driven, and the ATS coverage diagnostic on the DE-vs-IT cross-market comparison.

6 Conclusion

This paper studies how repo-market conditions pass through to securities-lending fees for the same sovereign bond. The central idea is simple: when a dealer borrows a bond in the lending market, the bond may also have value as collateral in repo. That repo deployment value should matter most when repo specialness is positive, when the bond is central to repo intermediation, and when lending-market utilisation is high.

I formalise this mechanism in a static two-market search-and-bargaining model along the lines of Duffie et al. (2002, 2005), in which the lending fee prices the dealer's repo deployment value: the bond's repo specialness, scaled by the intensity with which the dealer redeploys it as collateral, and split across a securities-lending bargain (lender share η) and a repo bargain (dealer share φ_R). The model delivers a small set of sign predictions on the fee–specialness pass-through—a non-negative pass-through, steeper in the high cross-market intermediation state, and steeper still on bonds the dealer deploys more—summarised in the deployment-centrality contrast $\eta\varphi_R\lambda_{HC} > 0$.

The evidence supports the deployment-value channel in the dealer-active parts of the euro-area sovereign collateral market. The headline is the repo-to-lending specialness pass-through, graded by deployment centrality: on the German and Italian post-issuance panel, the lending fee's sensitivity to repo specialness in the high-utilisation state is large for repo-central bonds and indistinguishable from zero for the rest, a between-group contrast of $\delta = +13.9$ ($t = +2.5$) that identifies the model's cross-partial; the contrast is robust to length-balancing and to a length-bias-free construction, and the pass-through concentrates in the positive-specialness domain where the channel operates and switches off below it. In a head-to-head the specialness state interaction dominates its demand-side counterpart, the lnSLQ$\times D$ interaction (+2.15 bps, $t = +5.90$, equal to 20–45% of the bond's bid-ask spread); the latter is the reduced form of the same channel and is itself economically meaningful, of the order of 10^5 per dealer-bond in annual lending revenue. The deployment channel is separated from within-market alternatives by the vanishing of the pass-through off the positive-specialness domain (where Λ is inactive), the cross-market simultaneity of fees, specialness, and repo volume, and the absence of a separate cheapest-to-deliver channel.

Two quasi-experiments identify the channel's two halves. On the repo side, Treasury auction reopenings—exogenous, calendar-set increases in a bond's collateral supply—relieve its repo specialness with the same sign in Germany, France, and Italy, confirming that specialness is a genuine supply-driven repo price rather than a passive correlate of the fee. On the fee side, the ECB SLF/QT episode delivers the paper's identified causal estimate of a state-dependent lending-fee response: in Italy, SLF-listed BTPs show a +6.45 bp higher fee response in the loaded state after QT begins ($t = 3.56$, wild-cluster-bootstrap

$p = 0.001$), with flat pre-trends. Across countries the response is opposite-signed—positive in Italy, negative in Germany and France—which rules out a common Europe-wide confound, since no single pan-European shock can move German and Italian fees in opposite directions; but because the SLF/QT shock bundles a public lending-supply and a repo-deployment margin, I do not claim to predict the country-by-country sign from primitives.

The results also have limits. The centrality evidence is best read as cross-sectional mechanism evidence on which bonds load the deployment-value channel, not a timing-causal effect of daily lagged specialness. The lendability filter on the Markit panel is an internal-validity selection: the coefficients apply to the lendable-and-ATS-active subset of the eurozone sovereign universe and should not be extrapolated outside it. The classical-measurement-error reading of the within-bond regression coefficient applies only under classical noise; with non-classical reporting noise the direction of bias is not determined. The German SLF estimate is sign-consistent with the high-scarcity ceiling but fails the parallel-trends diagnostics, so I treat it as corroborative rather than causal.

Taken together, the findings show that repo and securities-lending markets cannot always be understood separately. For sovereign bonds that are important in repo intermediation, repo deployment value is reflected in securities-lending fees when lending markets are tight: the lending fee is, in part, a repo price. This cross-market pricing channel helps explain why collateral scarcity, public lending facilities, and dealer intermediation capacity transmit across markets that are institutionally separate but linked by the same bond.

References

Afonso, G., & Lagos, R. (2015). Trade dynamics in the market for federal funds. *Econometrica*, *83*(1), 263–313. https://doi.org/10.3982/ECTA10586

Aggarwal, R., Bai, J., & Laeven, L. (2021). Safe-asset shortages: Evidence from the European government bond lending market. *Journal of Financial and Quantitative Analysis*, *56*(8), 2689–2719. https://doi.org/10.1017/S0022109020000770

Angelis Alexiou, G., Pereira, S. M., & Rodrigues-Gomes, V. (2025). *Repo collateral reuse and liquidity windfalls* (Working Paper No. 3147). European Central Bank.

Arrata, W., Nguyen, B., Rahmouni-Rousseau, I., & Vari, M. (2020). The scarcity effect of QE on repo rates: Evidence from the euro area. *Journal of Financial Economics*, *137*(3), 837–856. https://doi.org/10.1016/j.jfineco.2020.04.009

Atkeson, A. G., Eisfeldt, A. L., & Weill, P.-O. (2015). Entry and exit in OTC derivatives markets. *Econometrica*, *83*(6), 2231–2292. https://doi.org/10.3982/ECTA11477

Bassi, C., Behn, M., Grill, M., & Waibel, M. (2024). Window dressing of regulatory metrics: Evidence from repo markets. *Journal of Financial Intermediation*, *58*, 101086. https://doi.org/10.1016/j.jfi.2024.101086

Brumm, J., Grill, M., Kubler, F., & Schmedders, K. (2023). Re-use of collateral: Leverage, volatility, and welfare. *Review of Economic Dynamics*, *47*, 19–46. https://doi.org/10.1016/j.red.2022.03.003

Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, *22*(6), 2201–2238. https://doi.org/10.1093/rfs/hhn098

Carrera de Souza, T., & Hudepohl, T. (2024). Frictions in scaling up central bank balance sheet policies: How eurosystem asset purchases impact the repo market. *Journal of Banking & Finance*, *158*, 107037. https://doi.org/10.1016/j.jbankfin.2023.107037

Cohen, L., Diether, K. B., & Malloy, C. J. (2007). Supply and demand shifts in the shorting market. *Journal of Finance*, *62*(5), 2061–2096. https://doi.org/10.1111/j.1540-6261.2007.01269.x

Corradin, S., & Maddaloni, A. (2020). The importance of being special: Repo markets during the crisis. *Journal of Financial Economics*, *137*(2), 392–429. https://doi.org/10.1016/j.jfineco.2020.02.006

D'Amico, S., Fan, R., & Kitsul, Y. (2018). The scarcity value of Treasury collateral: Repo-market effects of security-specific supply and demand factors. *Journal of Financial and Quantitative Analysis*, *53*(5), 2103–2129. https://doi.org/10.1017/S0022109018000790

D'Avernas, A., & Vandeweyer, Q. (2024). *Treasury market liquidity and dealer balance sheet costs* (tech. rep.). Working Paper.

D'Avolio, G. (2002). The market for borrowing stock. *Journal of Financial Economics*, *66*(2–3), 271–306. https://doi.org/10.1016/S0304-405X(02)00206-4

Du, W., Tepper, A., & Verdelhan, A. (2018). Deviations from covered interest rate parity. *Journal of Finance*, *73*(3), 915–957. https://doi.org/10.1111/jofi.12620

Duffie, D. (1996). Special repo rates. *Journal of Finance*, *51*(2), 493–526. https://doi.org/10.1111/j.1540-6261.1996.tb02692.x

Duffie, D., Gârleanu, N., & Pedersen, L. H. (2002). Securities lending, shorting, and pricing. *Journal of Financial Economics*, *66*(2–3), 307–339. https://doi.org/10.1016/S0304-405X(02)00226-X

Duffie, D., Gârleanu, N., & Pedersen, L. H. (2005). Over-the-counter markets. *Econometrica*, *73*(6), 1815–1847. https://doi.org/10.1111/j.1468-0262.2005.00639.x

Greppmair, S., & Jank, S. (2026). Collateral scarcity and market functioning: Insights from the eurosystem securities lending facilities. *Review of Finance*, rfag004. https://doi.org/10.1093/rof/rfag004

He, Z., & Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, *103*(2), 732–770. https://doi.org/10.1257/aer.103.2.732

He, Z., Nagel, S., & Song, Z. (2022). Treasury inconvenience yields during the covid-19 crisis. *Journal of Financial Economics*, *143*(1), 57–79. https://doi.org/10.1016/j.jfineco.2021.05.044

Hendershott, T., & Menkveld, A. J. (2014). Price pressures. *Journal of Financial Economics*, *114*(3), 405–423. https://doi.org/10.1016/j.jfineco.2014.08.001

ICMA. (2024, December). *European repo market survey, number 48* (tech. rep.). International Capital Market Association.

ICMA. (2026, March). *European repo market survey, number 50* (tech. rep.). International Capital Market Association. https://www.icmagroup.org/assets/documents/Regulatory/Repo/Surveys/ICMA-European-Repo-Market-Survey-number-50-conducted-December-2025-published-March-2026-260326.pdf

Infante, S. (2019). Liquidity windfalls: The consequences of repo rehypothecation. *Journal of Financial Economics*, *133*(1), 42–63. https://doi.org/10.1016/j.jfineco.2019.02.004

Infante, S., & Vardoulakis, A. (2021). Collateral runs. *Review of Financial Studies*, *34*(6), 2949–2992. https://doi.org/10.1093/rfs/hhaa139

ISLA. (2025). *Securities lending market report, h2 2025* (tech. rep.). International Securities Lending Association.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182.

Kolasinski, A. C., Reed, A. V., & Ringgenberg, M. C. (2013). A multiple lender approach to understanding supply and search in the equity lending market. *Journal of Finance*, *68*(2), 559–595. https://doi.org/10.1111/jofi.12007

Lagos, R., & Rocheteau, G. (2009). Liquidity in asset markets with search frictions. *Econometrica*, *77*(2), 403–426. https://doi.org/10.3982/ECTA7250

Lagos, R., Rocheteau, G., & Weill, P.-O. (2011). Crises and liquidity in over-the-counter markets. *Journal of Economic Theory*, *146*(6), 2169–2205. https://doi.org/10.1016/j.jet.2011.10.001

Pelizzon, L., Subrahmanyam, M. G., & Tomio, D. (2025). Central bank-driven mispricing. *Journal of Financial Economics*, *166*, 104004. https://doi.org/10.1016/j.jfineco.2025.104004

Pinter, G. (2023). *An anatomy of the 2022 gilt market crisis* (Staff Working Paper No. 1019). Bank of England. https://www.bankofengland.co.uk/working-paper/2023/an-anatomy-of-the-2022-gilt-market-crisis

Saffi, P. A. C., & Sigurdsson, K. (2011). Price efficiency and short selling. *Review of Financial Studies*, *24*(3), 821–852. https://doi.org/10.1093/rfs/hhq124

Singh, M. (2014). *Collateral and financial plumbing*. International Monetary Fund.

Vayanos, D., & Weill, P.-O. (2008). A search-based theory of the on-the-run phenomenon. *Journal of Finance*, *63*(3), 1361–1398. https://doi.org/10.1111/j.1540-6261.2008.01360.x

Vuillemey, G. (2020). The value of central clearing. *Journal of Finance*, *75*(4), 2021–2053. https://doi.org/10.1111/jofi.12902