

Repo Collateral Reuse and Liquidity Windfalls

Georgios Angelis Alexiou^{1,2} Sofia Marques Pereira² Victor Rodrigues Gomes¹

¹ Goethe University Frankfurt

² European Central Bank

Abstract

Collateral reuse is central to repo-market liquidity but may amplify fragility through interconnected collateral chains. Theory suggests dealers exploit reuse to obtain liquidity windfalls, pledging collateral received at high haircuts while financing it at lower haircuts. Using confidential transaction-level data on European securities financing markets, we trace complete reuse chains: reused securities account for 11.6% of repo volume, averaging more than €49 billion per day. Haircut differentials on dealer-intermediated chains are negative: dealers forgo liquidity rather than extract it. Evidence indicates relationship-driven intermediation: concessions on sourcing haircuts retain counterparty flow, while windfalls require captive counterparties, absent from bank-to-bank markets.

JEL Codes: G15, G21, G23.

Keywords: repo, collateral reuse, haircuts.

1 Introduction

The repo market is a pivotal pillar of modern finance, anchoring wholesale funding, supporting market making and price discovery in core fixed income markets, allocating scarce safe assets across balance sheets, and assisting monetary policy transmission. Within this system, collateral reuse, the practice of re-delivering securities received in one transaction into another, expands secured funding capacity and improves balance sheet efficiency. It also lengthens collateral chains and tightens linkages among intermediaries and non-banks, creating channels through which operational frictions, settlement failures and shocks can propagate. Understanding how much reuse occurs, who does it, and whether it creates temporary funding wedges is central to market functioning and policy design.

A prominent line of theory predicts that dealers extract temporary liquidity windfalls by setting haircuts asymmetrically along reuse chains. If the haircut on the incoming leg is higher than on the onward leg, the intermediary pockets a temporary free cash wedge. Systematically positive wedges could subsidize matched books, increase effective leverage, compress bid-ask spreads, and dampen the pass-through of policy rates to secured funding. They would also open a distinct channel through which stress moves to funding conditions via abrupt shifts in margin requirements. Despite these implications, empirical evidence has been limited and inconsistent. Previous studies often conflate repos with other securities financing transactions, focus on narrow dealer segments, and lack the security-level mapping required to reconstruct reuse chains. We address these gaps using the Securities Financing Transactions Data Store, a transaction-level source that isolates repos and identifies ultimate cash borrowers and lenders across a broad set of intermediaries, including non-bank financial institutions. Building on the matching approach of [Fuhrer et al. \(2016\)](#) with minor modifications, we trace complete reuse chains and compare haircuts across matched legs within the same collateral cycle.

Two main results emerge. First, reuse in European repos is pervasive and persistent. On average, 11.6 percent of daily volume, roughly 49 billion euros per day in our

sample, is financed with previously received securities, and average chains span close to three intermediaries beyond the original owner. Reuse scales with overall activity and remains present across market conditions, confirming its structural role in secured funding. Second, and most consequential for theory, we find no systematic evidence of haircut-driven liquidity windfalls. Average haircut differentials between the sourcing and onward legs are, if anything, slightly negative. We also do not observe compensating spreads in matched repo rates. Results are robust to the high incidence of zero haircuts.

These findings shift the perspective on repo collateral reuse. The absence of a positive haircut wedge weakens a key mechanism through which models generate temporary liquidity subsidies, leverage pyramids, and acute sensitivity to haircut shocks. In practice, liquidity management appears to adjust through other deal terms rather than through asymmetric margins within reuse chains, suggesting that haircut-induced funding squeezes are a less important amplification mechanism than previously thought. At the same time, the pronounced conditional chain length and the importance of counterparty relationships indicate that reuse remains a material source of interconnectedness. Even without windfalls, settlement frictions and fails can propagate along chains and affect market functioning, which supports continued supervisory focus on the structure and dynamics of reuse networks.

Our contribution is to provide a direct, transaction-level mapping of repo-to-repo reuse and chain length in Europe with broad counterparty coverage, and to deliver an identification-based test that finds no systematic haircut-driven liquidity windfalls. Taken together, these facts point to relationship-driven intermediation: dealers concede on sourcing haircuts to retain counterparty flow, and the windfall configuration emerges only as a limiting case that requires both valuable temporary liquidity and captive counterparties, conditions largely absent from this predominantly bank-to-bank market.

2 Literature Review

2.1 Repo Collateral Reuse

A repurchase agreement (repo) is a two-legged transaction that resembles a collateralized loan: the cash borrower sells securities to the lender and agrees to repurchase them at a pre-specified price, so that the lender earns interest while holding assets that can be liquidated in the event of default. Repos are widely used by dealers to finance market-making and inventory management, and they provide a low-cost vehicle for mutual funds, pension funds, and corporate treasurers to place cash. The outstanding amount of repos and reverse repos in Europe has reached new highs, surpassing EUR 12.5 trillion in 2025 ([ICMA, 2025](#)). Beyond their primary funding function, repos facilitate the circulation and exchange of collateral: cash lenders can obtain specific securities for hedging or short-covering, while securities lenders improve portfolio liquidity without outright sales. These two functions—funding and collateral allocation—are the main motivations for participation in the repo market and underpin its central role in monetary policy transmission, price discovery, and the efficient allocation of safe assets.

One of the noteworthy features of repo arrangements is that the counterparty receiving the collateral gains partial or full legal rights over it. As long as the security can be returned by the repurchase date, the cash lender is free to deploy it in its own transactions, a process known as collateral reuse ([FSB, 2012](#); [Singh, 2014](#)). Our focus is reuse within the repo market, i.e., collateral received via repo and deployed in other repos.

Recent research highlights that collateral reuse tends to boost market efficiency through increased liquidity and additional allocation possibilities, but also increases systemic risk and volatility due to additional interconnectedness (e.g., [Infante, 2019](#); [Gottardi et al., 2019](#); [Accornero, 2020](#); [Brumm et al., 2023](#); [Infante and Saravay, 2025](#)). This trade-off is also emphasized by regulators (e.g., [CGFS, 2013](#); [FSB, 2017a](#); [FSB, 2017b](#)). Re-transferring collateral also has links with asset valuation ([Bottazzi et al.,](#)




Figure 1: Illustration of collateral reuse in a Repurchase Agreement. Collateral goes from A to B in the initial transaction, and B re-pledges it in a repo with C. On the repurchase date of the reuse transaction, C delivers the collateral to B, while B returns the collateral to A on the repurchase date of the initial transaction. RR^i and RR^j represent the repo rates agreed on in each transaction.

Alt text: Two-panel schematic of a repo reuse cycle. The First Leg panel shows boxes A, B, C exchanging collateral rightward and cash leftward, marking the A-to-B trade INITIAL and the B-to-C trade REUSE. The Second Leg panel shows the unwind, cash plus repo rates flowing one way and collateral the other, with a vertical bar labelled TERM between the legs.

2012) and monetary policy (e.g., [Muley, 2016](#); [Andolfatto et al., 2017](#)).

However, despite their importance, these works are all theoretical or normative. Empirical literature on the topic remains limited, mainly due to data availability. [Aitken and Singh \(2010\)](#), [Singh \(2011\)](#), and [Kirk et al. \(2014\)](#) report a trade-off between collateral circulation and the risks intrinsic to reuse, and provide early estimates of collateral velocity, but these studies relied on dealers’ annual reports rather than transaction-level data. More recently, [Infante et al. \(2020\)](#) use U.S. supervisory data on the nine largest primary dealers, finding that 85% of incoming Treasuries are re-transferred, and that reuse is positively related to scarcity and to the central bank’s balance sheet. Extending this line of work, [Infante and Saravay \(2025\)](#) construct a dealer-level “collateral multiplier” from the same supervisory data and show, in a model matched to these facts, that reuse allows intermediaries to reallocate the safety benefits of Treasury securities to a broader investor base: reuse intensifies when the available supply of securities contracts, notably during Federal Reserve asset purchases, thereby alleviating safe asset scarcity. In particular, [Fuhrer et al. \(2016\)](#) develop a bond-level matching algorithm to quantify reuse in the Swiss repo market, reporting an average reuse rate of 5%, peaking at around 20% in 2007. [Issa and Jarnećic \(2024\)](#) apply a similar methodology to Australian data, finding a 3.42% rate, while [Inhoffen and van Lelyveld \(2024\)](#)

use European data from the Money Markets Statistical Report (MMSR) and report a much higher 53% rate. [Jank et al. \(2022\)](#) report 75% for German banks using quarterly data on banks’ bond holdings and securities financing transactions (SFTs).

These divergences partly reflect differences in scope: most earlier work focusing on the European repo market has relied on datasets that combine different types of SFTs (e.g., securities lending, margin lending), making it impossible to disentangle whether reuse occurs within the repo market or across markets, which naturally inflates reuse estimates. This distinction matters because repos follow their own pricing logic, margining practices, and counterparty structures. By restricting our scope to repo-to-repo collateral reuse, our analysis offers a clearer basis for targeted policy responses. In this sense, our work is closer to that of [Fuhrer et al. \(2016\)](#) and [Issa and Jarnecic \(2024\)](#). Moreover, our dataset captures the full network of counterparties active in European repos, including non-bank financial institutions, providing the first detailed mapping of collateral reuse practices that explicitly identifies the participation of these institutions.

2.2 Liquidity Windfalls

Given that the cash lender in a repo receives a security whose market value may fluctuate, repos typically include some degree of over-collateralization in the form of a haircut—the discount on the collateral’s market value relative to its purchase price:

$$haircut = \left(\frac{P_m - P_p}{P_m} \right) \times 100$$

where P_m is the market value of collateral and P_p is the purchase price¹.

Haircut determination is at the center of a significant body of research on repo markets (see [Julliard et al., 2025](#)). Importantly for our purposes, [Infante \(2019\)](#) designs a model where dealers intermediate collateral through reuse, obtaining temporary liquidity by imposing a larger haircut in the initial transaction (dealer as collateral taker)

¹For further details, see: [ICMA ERC Repo Margining Best Practices 2012](#).

relative to the onward transaction (dealer as collateral provider). The resulting “liquidity windfall” means that dealers receive more securities than they need to transfer in order to obtain the same cash they lent out; by re-deploying the full collateral amount, they pocket the cash difference until the repurchase date. This is in line with the model of [Eren \(2014\)](#), where dealers’ demand for funding determines repo rates and haircuts, with dealers charging higher margins when lending to hedge funds than those they face when reusing the sourced collateral.




Figure 2: Illustration of the concept of liquidity windfalls. Counterparty A pledges \$101 in collateral value in exchange for \$100 in cash, facing a positive haircut. The dealer re-pledges the collateral in another repo at a lower haircut (zero, in this example), obtaining more cash than it lent out and retaining this surplus until the repurchase date.

Alt text: Schematic of three boxes labelled A, Dealer, and C. An INITIAL arrow shows A handing 101 dollars of collateral to the dealer for 100 dollars of cash; a REUSE arrow shows the dealer re-pledging that collateral to C for 101 dollars of cash, leaving the dealer holding 1 dollar of surplus.

If systematically present, these windfalls carry consequences for leverage, financial stability, and monetary policy transmission. We highlight three channels through which they would operate:

Leverage and matched-book expansion. Because the windfall frees internal liquidity, it incentivizes dealers to include the same bond in a larger matched book: each fresh reverse-repo is mirrored by an offsetting repo, leaving net exposure small while gross positions escalate. The extra bonds are financed almost entirely with borrowed cash, so assets rise while equity does not, pushing leverage to levels that must adjust sharply the moment haircuts shift. This leverage spiral is spelled out in [Brumm et al. \(2023\)](#). [Wallen and Lu \(2025\)](#) show that dealer balance-sheet capacity is an important driver of leverage supply, which is cyclical and increases US Treasury market fragility. The

FSB (2025) underscores that, especially in the NBFIs sector, leverage through collateral reuse and repo remains a key vulnerability for financial stability.

Exposure to repo and collateral runs. By insulating cash lenders while leaving collateral providers over-collateralized, the windfall mechanism exposes dealers to collateral runs: when a dealer shows signs of stress, collateral providers may refuse to roll their trades; the dealer must then return securities that are no longer in its possession (they have been re-transferred in other repos), forcing it to scramble for replacements or fire-sell assets. This sequence matches the repo runs documented during the 2008 GFC by Copeland et al. (2014), and is formalized in the coordination-failure model of Infante and Vardoulakis (2021). Also based on this concept, Kruttli et al. (2022) argue that distressed prime brokers might not reduce credit to hedge funds due to the benefits of the extra liquidity resulting from collateral reuse.

Bond pricing and policy rate pass-through. At the macro level, haircut-driven extra liquidity can blur the transmission of policy rates. By functioning as a quiet subsidy to the dealer's effective funding cost, liquidity windfalls may weaken the pass-through to repo rates, consistent with evidence that dealer market power muffles the signal central banks aim to send (Eisenschmidt et al., 2024). Cheaper balance-sheet space lets dealers enlarge their matched books and quote tighter bid-ask spreads, a pattern traced in both structural models and data (Huh and Infante, 2021; Adrian et al., 2017). If balance-sheet costs later rise (e.g., because leverage ratios tighten), the entire yield curve shifts between the bounds derived by Du et al. (2023).

Empirical investigation of these predictions is plagued by the same problems as the research on collateral reuse in general, with data availability and methodological issues limiting its scope. Moreover, the few available results contradict each other: Inhoffen and van Lelyveld (2024) find evidence in favor of the hypothesis, while Issa and Jarnećić (2024) find the opposite. The latter show that dealers, in reality, set a negative spread between the haircuts of the initial and reuse repos. They offer a compelling conjecture: it may be that dealers focus on the repo rate spread rather than on the haircut differential. For theoretical models to adapt, one needs to drop the assumption that

dealers strictly prefer immediate cash, as in Infante’s (2019) model, and instead suppose that banks prefer permanent over temporary liquidity, making the rate spread more attractive than the haircut spread.

This lack of coherence, pointing to a potential mismatch between theoretical predictions and empirical findings, calls for closer scrutiny. We aim to clarify this discrepancy by identifying the factors that drive haircut differentials and assessing whether they play any meaningful role in dealers’ funding strategies. This question lies at the core of our empirical analysis in Section 5.

3 Data and Reuse Estimation

3.1 The Dataset

The Securities Financing Transactions Data Store (SFTDS) is a joint project of the ECB and seven national central banks (Belgium, Germany, Spain, France, Italy, Luxembourg, and the Netherlands), established under the Securities Financing Transactions Regulation (SFTR, Regulation (EU) 2015/2365). SFTDS covers repos and buy/sell-backs, securities lending, and margin lending, with daily coverage since July 2020, and includes comprehensive information on counterparty identity, transaction dates, trade size, repo rates, haircuts, currencies, counterparty sectors, and collateral ISINs (International Securities Identification Numbers).

The dataset offers key advantages over the Money Market Statistical Reporting (MMSR), a widely used European data source. First, SFTDS contains separate tables for distinct SFTs, making it possible to isolate repo-to-repo reuse rather than conflating it with securities lending or margin lending. Second, its reporting scope extends beyond large banks to include non-bank financial institutions such as hedge funds and pension funds². Third, unlike in MMSR, centrally cleared trades in SFTDS identify the final cash borrower and lender, enabling deeper counterparty-based analysis.

Our sample spans January 2021 to February 2025. We apply a series of filters to the

²For more details, see: [ECB SFTDS Description](#).

initial dataset (detailed in Appendix D), leaving a final sample of around 15 million observations. We further merge in specific bond and counterparty characteristics from commercial data providers and public sources.

3.2 Identifying Collateral Reuse

To identify collateral reuse in SFTDS, we develop a matching algorithm *à la* [Fuhrer et al. \(2016\)](#). It proceeds in two stages. In the first stage, transactions are sorted chronologically and the algorithm iterates through each one, checking whether a subsequent transaction (i) involves the same ISIN, (ii) has a collateral provider equal to the collateral taker of the first, and (iii) has a purchase date that falls within the lifespan of the first. If all conditions hold, the pair is flagged as an initial–reuse match. In the second stage, the algorithm chains these pairs together: an initial transaction in one pair may itself be a reuse transaction in a previous pair, allowing us to trace complete collateral chains from the first provider to the last taker.

Two important deviations from the baseline approach of [Fuhrer et al. \(2016\)](#) should be noted. First, we do not require the repurchase date of the reuse transaction to precede that of the initial transaction, since entities routinely source the same bond elsewhere if timing mismatches arise. This relaxation avoids the significant underestimation arising under the stricter condition. Indeed, when this condition is dropped, the authors’ main reuse metric is almost identical to ours. Second, we do not allow more than one initial transaction to back any single reuse transaction. We recognize that, in practice, a reuse transaction may gather bonds from several prior trades. However, we consider this setup to be more parsimonious given securities’ fungibility and the limited information traders typically have about the provenance of collateral in their inventory³. Together, these choices appear to balance the risks of over- and underestimation, making our estimates more reliable. A more comprehensive treatment of the algorithm can be found in Appendix E.

³This convention reflects the usual approach to reuse from trading desks’ perspective, as clarified in conversations with repo traders.

4 Collateral Reuse in European Repo Markets

4.1 Description of Reuse

We analyze 961 trading days, during which 2174 institutions participated and transferred 5502 different securities. During the sample period, settled repos amounted to around 363 trillion euros, averaging around EUR 378 billion per day and 23.8 million per transaction. We identify 1.7 million trades ($\approx 11\%$) serving as initial transactions and 3.8 million reuse trades ($\approx 25\%$), where 2865 unique securities were re-transferred. In addition, we find 1089 institutions acting as cash lenders in reuse transactions, but only 547 acting as cash borrowers (i.e., collateral providers). This pattern is consistent both for sovereign bonds (1041 lenders vs. 501 borrowers) and corporate bonds (288 lenders vs. 127 borrowers). Altogether, the information we gather suggests that (i) most securities are reused to some extent and (ii) there is much more space or sense for more market participants to serve as collateral suppliers in reuse transactions relative to transacting as cash suppliers, corroborating the need for certain market conditions to fulfill the first role. As shown in Table 1, initial and reuse transactions amount to EUR 29 trillion and EUR 43 trillion, respectively.

Table 1 also shows that reuse activity is heavily concentrated in short maturities, especially one day, for both the number of transactions and trade volume. In contrast, the volume of initial transactions is more spread across the maturity bands, with the 1M–3M and Over 6M tenors also being prominent, although the number of transactions remains focused on day-to-day terms.

Table 1: Distribution of Initial vs. Reuse by Maturity

Maturity	Share of Transactions		Share of Volume	
	Initial	Reuse	Initial	Reuse
Day-to-Day*	86.6%	93.9%	30.2%	89.2%
1W	6.2%	3.8%	9.5%	5.8%
1M–3M	6.7%	2.1%	41.3%	4.8%
3M–6M	0.2%	0.04%	7.5%	0.1%
Over 6M	0.3%	0.1%	11.5%	0.1%
Total	1,705,492	3,848,031	29T	43T

*Day-to-Day includes Overnight, Spot-Next, and Tom-Next.

Figure 3 shows the evolution of the daily value of collateral delivered in repos over the course of our sample. As we can see, there is a noticeable growth in the European repo market during this period, with the total collateral value going from EUR 307 billion to around EUR 495 billion, a 61% increase. The value of reused collateral experiences an even more pronounced expansion in the period, jumping from EUR 35 billion to EUR 58 billion a day, 66% higher. These data highlight that (i) repo markets are becoming increasingly important as a source of wholesale funding for European institutions, and (ii) the value of reused collateral within these markets has been following this growth, and thus so is the exposure of the European financial system to its effects (good and bad).




Figure 3: Transferred Collateral - Total vs. Reuse

Alt text: Two upward-trending lines over the sample period: total transferred collateral rising from about 307 to 495 billion euro, and reused collateral rising from about 35 to 58 billion euro, tracking the overall growth at a lower level.

The set of graphs in Figure 4 shows the counterparty-country decomposition for the three largest euro-area countries (Germany, France, and Italy), depicting the proportion of reused collateral issued in one of these countries that is being reused by entities from the other two, plus British counterparties. We can see that the latter are quite active in the reuse of French and German bonds, while they are basically absent when it comes to Italian collateral. Also noteworthy is the fact that counterparties from the issuer countries are not always the ones reusing domestic securities the most, as can be seen for German bonds at the beginning of the sample and Italian bonds at

the end.




Figure 4: Reuse Volume Composition – Main Bond Issuers

Alt text: Three stacked-area panels tracking how reused collateral splits over time across German, French, Italian, and British counterparties. Panel (a), French bonds, and panel (b), German bonds, show sizeable British reuse; panel (c), Italian bonds, shows British counterparties nearly absent and shares shifting toward domestic holders.

We report two metrics of the extent of collateral reuse in the repo market. The first one is the reuse rate, consisting of:

$$rr_t = \frac{\sum_{i=1}^N Reuse_{it}}{\sum_{i=1}^N Repo_{it}}$$

For each transaction in our sample, we calculate the collateral value corresponding to reuse from past repos, which corresponds to $Reuse_{it}$. $Repo_{it}$ is the traded amount for a given transaction. In simple terms, we calculate the ratio of reused collateral to outstanding collateral per period (in our case, daily). This measure provides insight into

the ability of market participants to generate additional liquidity by re-using collateral (FSB, 2017a).

The second one is the average reuse chain length. As presented in Section 3.2, our algorithm is capable of tracking the path from the first collateral provider to the last collateral taker; Appendix E details the chain-reconstruction procedure.




Figure 5: Illustration of the concept of collateral chains. In the example, collateral transferred by counterparty A is subsequently re-deployed from B to C and finally from C to D. When the time comes to reverse the transaction, the security has to make its way back to A. The collateral chain counts institutions linked by the same collateral, excluding the owner, and is a measure of interconnectedness among market participants.

Alt text: Schematic of four boxes labelled A, B, C, and D connected left to right by arrows marked Collateral, with a dashed return arrow from D back to A and a brace beneath B to D marking a chain of length three. A Collateral Owner node feeds into box A.

$$\overline{Chain}_t = \frac{\sum_{i=1}^N Chain_{it}}{N}$$

where $Chain_{it}$ is the number of nodes in chain i , and N is the total number of chains in period t .

As highlighted by the FSB (2017a), the average chain length is a measure of interconnectedness among market participants, linked by the same collateral being re-transferred in subsequent transactions. Any issue faced by one node in the chain may cause distress along the entire sequence, potentially triggering problems in transaction settlements (e.g., a cascade of failures to deliver reused collateral) that can undermine market confidence and worsen deal terms for those seeking collateral. To the best of our knowledge, our work is the first to directly measure collateral chains, rather than

estimating them through other measures of reuse.⁴

As shown in Figure 6, the reuse rate within the repo market is somewhat volatile, even though its range appears relatively stable. While 1-2 percentage point variations might seem minor, given the average daily volume, they can represent up to EUR 8 billion in additional collateral circulating in the market, something far from negligible. One can also observe a sudden and persistent change in level from July-2022 on. This coincides with the beginning of the ECB's interest rate hiking cycle (July 21, 2022), and might be evidence of a reduction in collateral demand in a market that has just turned cash-driven.

The average reuse rate stands at 11.63%, with a standard deviation of 0.85%. It reaches a minimum of 9.18% in January 2021, and a maximum of 15.25% in November of the same year. This result is considerably higher than those reported by Issa and Jarnećić, 2024 (3.42%) and Fuhrer et al. (2016) (5%). This discrepancy may stem from our algorithm, which, unlike those used by these authors, does not restrict reuse based on repurchase dates. In other words, our approach allows reuse transactions to have a later repurchase date than their corresponding initial transactions. As discussed in Section 3, Fuhrer et al. (2016) find a reuse rate much closer to ours (around 10%) after relaxing this condition.

On the other hand, our estimate is well below those reported by Jank et al. (2022) (75%) and Inhoffen and van Lelyveld (2024) (53%). One possible explanation lies in the scope of the underlying data. Due to the structure of their datasets, both studies are effectively unable to distinguish repos from other securities financing transactions, such as securities lending or margin lending. This would naturally result in a higher reuse rate if other SFTs constantly employ collateral obtained via repo or vice-versa.

Figure 7 outlines the reuse rate profiles of the three main countries in the sample when trading among each other. One can notice that it is not uncommon for institutions located in one country to have a higher reuse rate when dealing with institutions

⁴For instance, FSB (2017b, pp. 7–8) presents an estimate of collateral chain length given by $length_k = \frac{1}{1 - ReuseRate_k}$. This approach is valid only for estimating unconditional chains (i.e., regardless of *de facto* collateral reuse), which is the best one can do when employing aggregate reuse measures instead of tracing how much of the collateral value in each transaction actually originates from previous trades.




Figure 6: Average Reuse Rate

Alt text: Line chart of the daily average value over the sample, fluctuating roughly between 9 and 15 percent around a mean near 11.6 percent, with a visible and persistent downward step in level from around July 2022.


from other countries compared to their own (as an aggregate). Not only that, there are countries, such as Italy, where the “domestic” reuse rate is consistently lower than the “international” ones. This may indicate that the demand profile for securities in other countries creates greater opportunities for reuse gains compared to the domestic Italian market, highlighting the allocation potential of collateral driven by reuse. In plot (d), we can also observe what happens when entities in the aforementioned countries trade with those of a non-EU but highly interconnected market, in the form of the British market. British counterparties also reuse collateral at different rates when trading with German, French, or Italian entities. In particular, the volatility of the reuse rate when facing Italian counterparties relative to the other two origins might signal that factors that turn reuse advantageous in this market for British players (e.g., better deal terms, better allocation of securities) change more frequently in the Italian market.

Not only the counterparties’ origin influences reuse, but also that of the securities. Figure 8 breaks down the reuse rate by the issuing country for the five largest Euro area countries. The reuse rate for French and German bonds is consistently higher and lower than the others, respectively, with the French rate always above average and the German one usually below it. The rate for Spanish bonds is noticeably more volatile in comparison. Interestingly, there seems to be some convergence among most or all


of the rates, especially in the early part of each year.

Figure 9 shows that, aside from a few short periods, short-dated (one year or less) bonds usually have the lowest reuse rates, typically around 8–11%, while the longest maturity buckets tend to show higher rates in the 12–14% range. Reuse rates for bonds with shorter remaining maturities are also more volatile, whereas the >1Y–5Y, >5Y–10Y, and 10Y+ buckets generally follow smoother and often similar trends.


Figure 10 reveals that, across most counterparty relationships, the reuse rate has generally been stable and close to the overall average throughout the sample period, with a few notable excursions. Bank-to-fund transactions exhibit the most pronounced volatility, spiking above 18% during late 2021 and early 2022, potentially reflecting the surge in funding pressures and balance-sheet arbitrage around that time. In contrast, reuse flowing from banks into hedge funds remains consistently subdued, dipping below 8% in recent quarters. This pattern can reflect persistent caution (and perhaps higher operational friction or haircuts) when banks extend reusable collateral to more leveraged trading entities. In terms of reuse volumes, reuse activity is mainly a bank-to-bank phenomenon.




(a) Reuse Rates by Counterparty Country Pair - German Lender



(b) Reuse Rates by Counterparty Country Pair - French Lender



(c) Reuse Rates by Counterparty Country Pair - Italian Lender



(d) Reuse Rates by Counterparty Country Pair - British Borrower

Figure 7: Reuse Rates by Counterparty Country Pair (% , MA-30). Country A $\rightarrow$ Country B means that an entity located in A lends cash to an entity in B. For instance, the data represented by DE $\rightarrow$ IT is the average daily reuse rate for when Italian counterparties borrow cash (and, therefore, transfer securities) from a German counterparty.

Alt text: Four line-chart panels of 30-day moving-average reuse rates. Panel (a) is a German cash lender, panel (b) a French cash lender, panel (c) an Italian cash lender, and panel (d) a British borrower, each plotting rates against domestic and foreign counterparties. Rates with foreign counterparties frequently exceed domestic ones, and those involving Italian counterparties are the most volatile.




Figure 8: Reuse Rate by Issuer Country (MA-30)

Alt text: Five lines for the largest euro-area issuers. French bonds trace the highest reuse rate throughout and German bonds the lowest; Spanish bonds are noticeably more volatile, and all the lines partly converge early in each year.

Finally, the average chain length sits at around 2.95. As shown in Figure 11, it is




Figure 9: Reuse Rate by Bond Remaining Maturity (MA-30)

Alt text: Several lines grouped by remaining maturity. Short-dated bonds, one year or less, carry the lowest and most volatile rates, around 8 to 11 percent, while the longer-maturity buckets run higher and smoother, around 12 to 14 percent.




Figure 10: Reuse Rate by Counterparty Sector Pair (MA-30)

Alt text: Several lines by sector pair. Most stay close to the overall average, but the bank-to-fund line spikes above 18 percent in late 2021 and early 2022, while bank-to-hedge-fund reuse stays subdued, dipping below 8 percent in recent quarters.

also quite stable throughout the analyzed period, consistently above 2.5 and never dropping below 2. [Singh \(2011\)](#) introduced the concept of “collateral velocity”, which is essentially the ratio of total pledged collateral to incoming collateral. This gives an idea of how many market transactions are simultaneously secured by the same security, on average. In other words, it is an estimate of the reuse chain length, something we are now able to observe more directly. Singh found that collateral velocity (or chain length) was around 3 before the collapse of Lehman Brothers at the end of 2007, falling to 2.4 by the end of 2010. Our results suggest that collateral chains have since rebounded to pre-crisis levels.




Figure 11: Average Reuse Chain Length

Alt text: Line chart over the sample period showing a stable, roughly flat series that hovers around 2.95, stays consistently above 2.5, and never drops below 2.

4.2 Drivers of Reuse

We proceed to identify the determinants of collateral reuse in the repo market. For this, we employ different specifications under the following general form:

$$Reuse_i = \beta_0 + \Phi X + \beta_1 OISSpread_i + \beta_2 Hike_i + \beta_3 March23_i + \gamma_{t,b,cp} + u_i$$

where *OISSpread* is the spread between the 3-month €STR-based Overnight Index Swap (OIS) and the realized €STR for a given date, capturing agents' expectations on future market developments; *Hike* and *March23* are dummies flagging the dates after July 2022 and the whole month of March 2023, isolating the effects of the new ECB monetary policy cycle and the financial turmoil in 2023 (e.g., Credit Suisse downfall, Silicon Valley Bank bankruptcy); $\gamma_{t,b,cp}$ are combinations of time (week), bond (security ISIN) and counterparty fixed effects; and Φ and X are, respectively, vectors of coefficients and independent variables corresponding to each group that we test for. The groups refer to different aspects of trade, namely transaction, counterparty, and collateral characteristics.

$Reuse_i$ is a binary variable taking the value 1 if transaction i is reusing collateral, and 0 otherwise. This turns our model into a Linear Probability Model (LPM),

where coefficients can be directly interpreted as changes in the probability of collateral reuse. Importantly, while nonlinear models such as logit and probit are theoretically preferable for binary outcomes due to their bounded predictions and efficient estimation, they are computationally intensive and often incompatible with high-dimensional fixed effects structures. In other words, employing any of those would prevent us from taking full advantage of our vast data availability to make use of more rigorous specifications. In this setting, we opt for a LPM, which allows us to flexibly incorporate strict fixed effects that capture unobserved heterogeneity across several dimensions. Although LPM has known limitations, including heteroskedasticity and potential predictions outside the $[0,1]$ range, the use of robust standard errors and the consistency of the estimates under standard assumptions make it a pragmatic and credible choice for estimating conditional associations. Full descriptive statistics of the variables can be found in Appendix [A](#).

GROUP 1: TRANSACTION CHARACTERISTICS

To analyze this group, we opt to specify our regressions as:

$$\begin{aligned} Reuse_i = & \beta_0 + \phi_1 \ln(Notional_i) + \phi_2 \ln(Term_i) + \phi_3 Haircut_i + \phi_4 RepoRate_i \\ & + \beta_1 OISSpread_i + \beta_2 Hike_i + \beta_3 March23_i + \gamma_{t,b,bw} + u_i \end{aligned}$$

where *Notional* is the transaction amount and *Term* is the transaction tenor in days, both in natural logarithm; *Haircut* and *RepoRate* are the remaining transaction terms; and $\gamma_{t,b,bw}$ are *time* $\times$ *bond* $\times$ *borrower* fixed effects. The choice of the borrower, rather than the more strict counterparty-pair fixed effects, is justified by the fact that the lender generally does not know the source of the collateral; therefore, the decision and capacity to reuse it lie with the borrower.

As we clarify in Section [5](#), we replace CCP “haircuts” with initial margins from a major European central clearing counterparty. These are defined at the ISIN level; that is, each bond has a specific margin parameter. We observe margins on a quarterly basis

and extend each parameter throughout the corresponding quarter. For instance, the same ISIN transferred in January 2022 and March 2022 would carry the same margin, namely that corresponding to Q1 2022. This frequency is not an issue, as we do not expect CCP margins to change frequently, and certainly not drastically. This does mean, however, that any time fixed effects at a higher frequency would absorb the margin variation entirely, which is why we include only *bond* $\times$ *borrower* fixed effects in Group 1 regressions for CCP trades, foregoing the trade-week dimension.

GROUP 2: COUNTERPARTY CHARACTERISTICS

Our specification to test for the attributes of borrowers and lenders is:

$$\begin{aligned} Reuse_i = & \beta_0 + \phi_1 DB_i + \phi_2 HFB_i + \phi_3 DEB_i + \phi_4 FRB_i + \phi_5 ITB_i + \phi_6 RATB_{bwi} \\ & + \phi_7 EXPOB_{bwi} + \phi_8 EXPOL_{li} + \beta_1 OISSpread_i + \beta_2 Hike_i + \beta_3 March23_i + \gamma_{t,b} + u_i \end{aligned}$$

where *DB* and *HFB* are dummies flagging if the borrower is, respectively, a dealer⁵ or a hedge fund; *DEB*, *FRB*, and *ITB* are dummies indicating if the borrower is located in Germany, France, or Italy; *RATB* is the borrower's rating in a numeric scale that goes from 1 (equivalent to D) to 22 (equivalent to AAA)⁶; and *EXPOB* and *EXPL* represent, respectively, the exposure of the borrower to the lender and vice-versa in a given transaction, quantified as fraction of the cash amount, in that quarter, that the borrower (lender) received (lent out) from a given lender (to a given borrower). The subscripts *bwi* and *li* refer, respectively, to borrower and lender characteristics when transaction *i* is sealed. $\gamma_{t,b}$ are *time* $\times$ *bond* fixed effects.

GROUP 3: COLLATERAL CHARACTERISTICS

Finally, we proceed to capture the influence of collateral specifics on reuse with:

$$Reuse_i = \beta_0 + \phi_1 REMMAT_{bi} + \phi_2 \ln(Outstanding_{bi}) + \phi_3 IssuerDE_{bi} + \phi_4 IssuerFR_{bi}$$

⁵Further explanation on our classification of dealers is available in the next section.

⁶Ratings follow Standard and Poor's model.

$$\begin{aligned}
& + \phi_5 \text{IssuerIT}_{bi} + \phi_6 \text{ASKBID}_{bi} + \phi_7 \text{Specialness}_{bi} + \phi_8 \text{OTR}_{bi} + \phi_9 \text{CTD}_{bi} + \\
& + \phi_{10} \text{ECBHC}_{bi} + \beta_1 \text{OISSpread}_i + \beta_2 \text{Hike}_i + \beta_3 \text{March23}_i + \gamma_{t,bw} + u_i
\end{aligned}$$

where *REMMAT* is the bond remaining maturity in years; *Outstanding* is the natural logarithm of the security's outstanding amount in the trade date; *IssuerDE*, *IssuerFR*, and *IssuerIT* are dummies indicating if the issuer is located in Germany, France, or Italy; *ASKBID* is the difference between the ask and the bid prices for the security on a given date, capturing liquidity conditions; *Specialness* is the difference between the repo rate for a given transaction and the ongoing ECB Deposit Facility Rate (DFR)⁷; *OTR* is a dummy equal to one if the security is the most recently issued (on-the-run) benchmark within its maturity segment on the trade date, typically associated with higher liquidity and tighter pricing; *CTD* is a dummy equal to one if the security is the cheapest-to-deliver into the relevant futures contract on the trade date; *ECBHC* is the haircut attributed to the security in the ECB collateral framework on the trade date, which we understand to be a suitable quantitative substitute for bond ratings as a measure of credit risk. The subscript *bi* refers to the bond *b* serving as collateral for transaction *i*. Finally, $\gamma_{t,bw}$ are *time* $\times$ *borrower* fixed effects.

Table 2 presents the results for transaction characteristics, starting with two intuitive findings: smaller and shorter transactions are more likely to reuse collateral. This appeals to considerations of risk exposure on the notional side and to the repurchase date dynamics on the term side, meaning that reusing collateral from longer maturities in shorter-maturity repos makes more sense if one aims to constrain exposure. Despite being statistically significant in some setups, the haircut/margin and repo rates are not economically meaningful for determining reuse. The statistical or economic insignificance of the OIS spread and the dummy flagging the turbulent month of March 2023 signal that greater financial instability has a weaker effect on collateral reuse in the repo market than one would guess. The change in this pattern when we relax the

⁷It might be argued that specialness in this definition fits better as a transaction characteristic. Our understanding, however, is that this is an attribute closely tied to bond specifics, and significant variation between transactions with the same underlying collateral should rarely be observed over short periods of time unless there is significant distress in the market, something we already control for.

bond fixed effects in Table 4, at least for OTC trades, indicates that participants' concerns about heightened risks when entangled in collateral chains during turmoil are related to collateral quality. Finally, the ECB-hike dummy is, at most, economically insignificant, indicating that the decision to reuse collateral is somehow detached from the monetary policy environment. Most of these results apply to both over-the-counter and centrally cleared transactions, with the exception of a positive effect of term on the probability of reuse for the latter. Although somewhat surprising, this result may reflect a relative lack of concern among counterparties about tenor, given the additional layer of security provided by the CCP.

Table 3 shows that dropping borrower fixed effects drastically reduces explanatory power (as measured by the R-squared). That said, even with a comprehensive set of counterparty characteristics (country, sector, rating, connections, etc.), we still fall far short of achieving comparable explanatory power to the stricter fixed-effects specification when it comes to explaining reuse probability. This suggests that what really matters is not the institution's observable characteristics, but rather *who they are*. In essence, this points to borrower-specific market stance or behavior that is difficult to capture with standard statistical tools—though recognizing this limitation is itself informative.

Dealer activity does not appear to be a determinant of the probability of reuse. By contrast, hedge fund participation notably reduces this probability when those funds act as borrowers and face non-zero haircuts. Italian borrowers are also much less likely to reuse than their counterparts in other jurisdictions. Given that Italian banks are often perceived as riskier than other major nationalities in the sample, both results suggest that reuse is more feasible for counterparties with a stronger market stance. Furthermore, the importance of this factor for the broader repo market is reflected in the coefficients for borrower and lender exposure: greater borrower dependence on a lender for transaction volume significantly decreases the probability of reuse, whereas greater lender dependence on a borrower has the opposite effect, further underscoring the importance of counterparties' competitive position.

Finally, Table 4 shows that only specific collateral features appear to matter for the decision to reuse. This may reflect the fact that most collateral in euro-area repos is already constrained to high-quality securities. One can confirm this by checking, in Table A1, that the average ECB haircut in our sample is 4.23%, which is typically feasible only for bonds rated AAA to A-⁸. This conclusion is reinforced by the finding that the haircut measures themselves are neither statistically significant nor economically meaningful. The bond issuer, however, does seem to matter: German and French bonds are less likely to be reused. One explanation is that, viewed as safer assets issued by the EU's largest economies, they may have higher-value uses than being consecutively transferred in repos. The coefficients on outstanding amount and specialness provide conflicting evidence regarding the hypothesis that bond scarcity prompts collateral reuse, with the latter supporting the premise and the former contradicting it. Finally, the results for cheapest-to-deliver bonds indicate that those who obtain these securities prefer to hold them rather than repo them out again, given their intrinsic cost-saving features.

⁸See: [ECB Haircut Schedule](#).

Table 2: Results Group 1 – Transaction Characteristics

	Dependent Variable: <i>Reuse</i>		
	OTC	CCP	OTC non-zero
ln(Notional)	−0.034*** (−2.97)	−0.031*** (−12.75)	−0.028*** (−3.24)
ln(Term)	−0.024*** (−6.12)	0.022*** (8.50)	−0.024*** (−7.12)
Haircut/Margin	0.002*** (6.28)	0.002 (1.47)	0.002*** (6.69)
Repo Rate	−0.007 (−0.66)	0.004** (2.18)	−0.010 (−1.05)
OIS Spread	−0.024 (−1.64)	−0.005 (−0.74)	−0.009 (−0.77)
March 23	−0.006 (−0.61)	0.009** (2.11)	0.001 (0.10)
ECB Hike	−0.005 (−0.59)	−0.010* (−1.69)	0.001 (0.05)
Constant	0.82*** (4.44)	0.70*** (17.32)	0.68*** (4.80)
Time × Bond × Borrower FE	Yes	No	Yes
Bond × Borrower FE	No	Yes	No
Adj. R-squared	0.60	0.24	0.62
Number of obs	1,499,713	12,953,080	836,874

Notes: t-statistics in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Standard errors are clustered at the borrower level. The OTC non-zero column restricts the OTC sample to transactions with a non-zero haircut.

Table 3: Results Group 2 – Counterparty Characteristics

	Dependent Variable: <i>Reuse</i>		
	OTC	CCP	OTC non-zero
Dealer Borrower	0.005 (0.19)	0.029 (1.53)	−0.019 (−0.85)
Hedge Fund Borrower	−0.082 (−1.23)	— —	−0.172*** (−3.09)
German Borrower	−0.011 (−0.24)	0.028 (0.75)	0.047 (0.92)
French Borrower	−0.036 (−0.92)	−0.007 (−0.32)	−0.031 (−0.65)
Italian Borrower	−0.081* (−1.84)	−0.115*** (−4.11)	−0.187*** (−3.53)
Borrower Rating	−0.003 (−0.77)	0.004 (1.02)	0.001 (0.22)
Borrower Exposure	−0.182*** (−4.49)	−0.681*** (−4.68)	−0.154*** (−3.05)
Lender Exposure	0.086** (2.44)	0.408** (2.57)	0.039 (1.25)
OIS Spread	−0.003 (−0.25)	−0.005 (−0.82)	0.012 (1.02)
March 2023	−0.009 (−0.47)	0.010 (1.42)	−0.003 (−0.15)
ECB Hike	0.017* (1.87)	−0.004 (−0.69)	0.003 (0.24)
Constant	0.283*** (4.09)	0.149** (2.04)	0.204*** (3.49)
Time × Bond FE	Yes	Yes	Yes
Adj. R-squared	0.27	0.09	0.35
Number of obs	1,077,077	11,122,017	629,756

Notes: t-statistics in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Standard errors are clustered at the borrower level. The Hedge Fund Borrower indicator is omitted in the CCP column due to collinearity. The OTC non-zero column restricts the OTC sample to transactions with a non-zero haircut.

Table 4: Results Group 3 – Collateral Characteristics

	Dependent Variable: <i>Reuse</i>		
	OTC	CCP	OTC non-zero
Remaining Maturity	0.001** (2.36)	−0.001* (−1.68)	0.000 (0.98)
ln(Outstanding)	0.053*** (6.17)	0.019 (1.58)	0.055*** (7.84)
German Issuer	−0.051*** (−3.86)	−0.032*** (−3.30)	−0.053*** (−4.04)
French Issuer	−0.059*** (−5.91)	−0.016* (−1.71)	−0.045*** (−5.75)
Italian Issuer	−0.003 (−0.12)	0.014 (0.43)	−0.004 (−0.16)
Ask–Bid	−0.030* (−1.71)	−0.011 (−0.49)	−0.010 (−0.49)
Specialness	0.039*** (3.53)	0.003 (0.17)	0.048*** (3.89)
Cheapest-to-Deliver	−0.037** (−2.44)	−0.010 (−1.28)	−0.036*** (−2.60)
On-the-Run	−0.005 (−0.52)	−0.001 (−0.20)	0.004 (0.53)
ECB Haircut	0.001 (0.35)	0.005 (1.50)	0.002 (1.09)
OIS Spread	−0.020*** (−3.10)	−0.009 (−1.31)	−0.013** (−1.99)
March 2023	−0.002 (−0.18)	0.012 (1.53)	0.002 (0.22)
ECB Hike	0.004 (0.34)	−0.002 (−0.39)	−0.005 (−0.36)
Constant	−0.269*** (−3.06)	0.028 (0.24)	−0.320*** (−4.49)
Time × Borrower FE	Yes	Yes	Yes
Adj. R-squared	0.15	0.08	0.16
Number of obs	1,430,461	10,316,955	873,777

Notes: t-statistics in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Standard errors are clustered at the borrower level. The OTC non-zero column restricts the OTC sample to transactions with a non-zero haircut.

5 Investigating Liquidity Windfalls

We now proceed to test for the existence of liquidity windfalls via collateral reuse in repo transactions. As discussed in Section 2, theory predicts that dealers intermediate collateral by running a reverse repo to source securities and a subsequent repo to source liquidity using the same ISIN. The twist is on the deal terms: dealers would impose a larger haircut in the initial transaction than they face in the reuse one, pocketing a temporary free cash buffer—the liquidity windfall.

Given the context, a logical first step is to distinguish the institutions that are central to this analysis. Our identification of dealers relies on the official list of Primary Dealers (PDs) for the reporting countries. A dealer is considered active in a transaction when (i) it acts as the cash borrower or lender and (ii) the issuer of the security is a country where the bank has PD status. The list of active dealers in our sample can be found in Appendix C. We observe 52 distinct dealers in our sample, of whom 47 trade sovereign bonds at least once and 45 trade corporate bonds. Dealers borrow cash in 47% of transactions and lend cash in 46%, while dealer-to-dealer transactions account for 21%. In reuse transactions, 51 dealers act as cash providers, while 43 act as collateral providers; dealers are collateral providers in more than 50% of reuse trades, consistent with their role as collateral intermediaries.

We then compute our variable of interest: haircut differentials between matched initial–reuse pairs:

$$\Delta H_j = \text{Haircut}_{\text{initial}} - \text{Haircut}_{\text{reuse}}$$

An entity intermediating collateral via reuse receives a liquidity windfall whenever $\Delta H_j > 0$. Consistently positive differentials would corroborate the hypothesis that this extra cash cushion is systematic.

Before proceeding, we note that haircut differentials are computed exclusively for pairs in which both transactions are conducted over-the-counter. The rationale lies in basic repo microstructure: centrally cleared transactions are not subject to a haircut.

Instead, the central counterparty (CCP) requires margin to be posted according to its own internal framework. While the margining process is intricate and contingent on the CCP's risk management practices, its central component is the *initial margin*, which may be loosely interpreted as a haircut analogue.⁹

At first glance, the availability of data on initial margins might suggest that margin-haircut differentials could serve as a measure of additional liquidity. A critical asymmetry, however, complicates this interpretation: the CCP seeks to mitigate the risk posed not only by the cash borrower but also by the lender, and margins are therefore levied on *both counterparties*¹⁰. As a result, reusing collateral originating from a centrally cleared transaction entails a substantial cost from the outset. This gives rise to a configuration in which haircut differentials can virtually never act as a driver of collateral reuse whenever a CCP is involved in either leg: (i) pairs of centrally cleared initial-and-reuse transactions generate cumulative margin requirements; (ii) when a security flows from an OTC trade into a centrally cleared one, the haircut in the former would need to depart markedly from prevailing market practice to offset the typically higher margins; and (iii) conversely, reusing collateral originating from a centrally cleared trade in a bilateral transaction obliges the reusing institution to bear both margin and haircut.

Drawing on publicly available margin parameters from a major central counterparty, we illustrate these points. Figure 12 shows that initial margins consistently exceed bilateral haircuts by a large amount. The difference in the means is even more striking: 3.9% (margins) against -0.10% (haircuts). Figure 13 further reinforces that any haircut differential computed outside of OTC-OTC pairs would invariably be negative. Accordingly, only purely bilateral initial-and-reuse pairs are economically meaningful for the present analysis.

⁹Unlike a haircut, which represents a discount applied to the market value of collateral, the initial margin constitutes an additional payment posted in cash or securities. It is therefore not strictly equivalent to a bilateral haircut, although the two may be regarded as comparable from a *cost management perspective*.

¹⁰Due essentially to liquidity considerations, the margins for cash borrower and lender are not necessarily symmetric.




Figure 12: Initial Margins and Haircuts, 10-day Moving Average

Alt text: Two lines comparing central counterparty initial margins with bilateral haircuts. The initial-margin line lies well above the haircut line throughout the sample, with mean margins near 3.9 percent against mean haircuts near minus 0.10 percent.




Figure 13: Haircut Differentials With CCP Participation, 10-day Moving Average

Alt text: Single line for reuse pairs in which a central counterparty takes part in at least one leg. The differential stays consistently and substantially below zero across the whole sample.

5.1 Descriptive Evidence

Figure 14 presents the evolution of the average haircut differentials between matched OTC pairs. The full-sample average sits around -11 basis points (bps). When weight-

ing the differences by the reused cash amount, the mean moves to -7 bps, while most of the time-series structure remains unchanged. This is the first piece of evidence against the existence of liquidity windfalls: the prediction of a positive differential whenever a dealer pockets the spread between a reverse repo and a subsequent repo is not borne out. Table 5 sharpens this point. The negative differential comes essentially from the dealer intermediated chains, reaching -33.5 bps (-26.8 when weighted) on average relative to +3.8 bps (+7.2 weighted) for the remaining links. If dealers were extracting windfalls, the opposite pattern should hold.

Table 5: Summary Statistics – OTC-OTC Haircut Differentials

Sample	Obs	Mean	Std. Dev.	Min	Max
All observations	282,418	-0.1162	0.8469	-7.7311	7.4500
Dealer-intermediated	116,787	-0.3355	0.9002	-6.1293	7.2910
Non-dealer-intermediated	165,631	0.0384	0.7705	-7.7311	7.4500
Weighted**					
All	282,418	-0.0702	0.2954	-1.9048	1.6022
Dealer-intermediated	116,787	-0.2687	0.3312	-2.8490	1.7229
Non-dealer-intermediated	165,631	0.0721	0.3098	-2.0320	1.9488

** Weighted by reused value.



Figure 14: Haircut Differentials, 10-day Moving Average

Alt text: Single line for matched over-the-counter initial and reuse pairs, fluctuating around a mildly negative full-sample average of roughly minus 11 basis points, or about minus 7 basis points when weighted by reused value.

5.2 Econometric Assessment

To move beyond descriptive evidence, we econometrically assess haircut differentials. As observed by [Issa and Jarnecic \(2024\)](#), this estimation is likely to suffer from self-selection bias: factors may influence both the decision to reuse collateral and the haircut differential between the initial and reuse transactions. We follow their approach in estimating a two-stage [Heckman \(1979\)](#) model.

In the first stage, we estimate a probit model for the probability of selection into collateral reuse. For this, we further exploit our granular evidence on the drivers of reuse by including variables from the three groups of characteristics — transaction, counterparty, and collateral — that best predict reuse:

STAGE I

$$\Pr(\text{reuse}_i = 1) = \Phi\left(\beta_0 + \beta_1 \ln(\text{Notional}_i) + \beta_2 \text{Specialness}_i + \beta_3 \text{HFB}_i + \beta_4 \text{ITB}_i\right)$$

$$\begin{aligned}
& +\beta_5 \text{EXPOB}_{bwi} + \beta_6 \text{EXPOL}_{li} + \beta_7 \text{IssuerDE}_{bi} + \beta_8 \text{IssuerFR}_{bi} \\
& +\beta_9 \ln(\text{Oustanding}_{bi}) + \beta_{10} \text{CTD}_{bi} + \beta_{11} \text{OISSpread}_i + \gamma_t
\end{aligned}$$

Our main goal in the first stage is to obtain the inverse Mills ratio (IMR), which controls for the non-random selection into reuse: it represents the conditional expectation of the outcome disturbance given that an observation enters the sample, and including it in the outcome equation removes the resulting selection bias. We estimate the first stage as a probit with time (week) fixed effects and standard errors clustered at the borrower level. We deliberately avoid the more saturated fixed effects structures used in our second stage: probit estimation with high-dimensional fixed effects is subject to the incidental parameters problem, which biases slope coefficients (and therefore the constructed IMR) when within-cell sample sizes are small. Week fixed effects absorb aggregate time-varying conditions while preserving the large within-cell sample sizes needed for consistent probit estimation, and the richer heterogeneity is absorbed by the fixed effects in the second stage. Table 6 presents the Stage I results.

An important methodological point is that, while transactions vary considerably in size, ordinary least squares would assign equal weight to all trades; the resulting coefficients would describe the haircut differential for the *average trade* rather than for the *average euro* of collateral. To focus the fit on economic exposure, we estimate Stage II by weighted least squares (WLS), assigning each observation a weight equal to the reuse transaction's collateral value,

$$w_i = C_i > 0, \quad \hat{\beta}_{\text{WLS}} = (X^\top W X)^{-1} X^\top W d,$$

where d_i is the haircut differential, X_i its covariate vector, and $W = \text{diag}(w_1, \dots, w_n)$. The resulting coefficients quantify how explanatory variables affect the haircut differential *per euro of collateral*, preventing small transactions from disproportionately driving results.

Table 6: Stage I – Probit for selection into collateral reuse

	Dependent variable: <i>Reuse</i>	
	OTC	OTC non-zero haircut
ln(Notional)	-0.077*** (-4.28)	-0.060*** (-3.00)
Specialness	0.037 (0.44)	0.109 (1.27)
Hedge Fund Borrower	-0.032 (-0.31)	-0.166 (-1.44)
Italian Borrower	0.404*** (2.77)	0.555*** (2.62)
Borrower Exposure	-0.455*** (-4.25)	-0.487*** (-3.87)
Lender Exposure	0.155 (1.17)	-0.059 (-0.32)
German Issuer	-0.118** (-2.31)	-0.167** (-2.39)
French Issuer	-0.193*** (-4.02)	-0.160** (-2.35)
ln(Outstanding)	0.309*** (7.92)	0.325*** (7.56)
Cheapest-to-deliver	-0.102** (-2.08)	-0.106** (-2.04)
OIS spread	-0.041 (-1.37)	-0.026 (-0.55)
Week fixed effects	Yes	Yes
Observations	1,812,226	1,044,797
Pseudo R^2	0.063	0.074

Notes: t-statistics in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Standard errors are clustered at the borrower level.

STAGE II

$$\begin{aligned}\Delta H_j = & \phi_0 + \phi_1 \ln(\text{Notional}_{j2}) + \phi_2 \ln(\text{Term}_{j2}) + \phi_3 \text{DEALERINT}_j + \phi_4 \Delta R_j + \phi_5 \text{Specialness}_{j2} \\ & + \phi_6 \text{CTD}_{j2} + \phi_7 \text{Hike}_{j2} + \phi_8 \text{March23}_{j2} + \phi_9 \text{IMR}_{j2} + \gamma_{t,b,bw} + \varepsilon_j\end{aligned}$$

where ΔH_j is the haircut differential between the initial and reuse transaction for pair j ; *DEALERINT* is a dummy equal to 1 if a dealer is wedged between two non-dealers in the chain, our main variable of interest; ΔR_j is the repo rate differential between initial and reuse transactions; and IMR is the inverse Mills ratio from Stage I. The subscript $j2$ indicates that attributes refer to the reuse transaction. We impose *time* $\times$ *bond* $\times$ *borrower* fixed effects, so that time effects control for risk and liquidity factors, while bond and borrower effects absorb collateral and counterparty heterogeneity.

Table 7 reports our results across different specifications, ensuring robustness to distinct fixed effect and standard error clustering structures. The Heckman selection correction performs as anticipated. The inverse Mills ratio has a positive sign independently of the model, confirming that selection into collateral reuse is non-random and positively correlated with unobserved determinants of the haircut differential. In the less-saturated specifications (1)–(3), the IMR is statistically significant at the 1% level, with point estimates between 0.29 and 0.34. Once saturated *week* $\times$ *ISIN* $\times$ *borrower* fixed effects are absorbed in (4)–(5), and in the non-zero-haircut subsample of (6), the IMR remains positive and of comparable magnitude (0.32–0.43) but is imprecisely estimated, consistent with the large reduction in within-cell variation following the absorption of singleton observations and reflected in the low within- R^2 of these columns. Overall, the pattern supports the presence of selection bias in the economically meaningful specifications and confirms that the second-stage coefficients are estimated free of selection contamination, with the saturated columns serving more as conservative stress tests than as efficient estimators.

The remaining controls reveal that the haircut differential is largely orthogonal to standard transaction- and collateral-level determinants. Notional and term enter with intuitive signs but are economically negligible: a doubling of trade size or tenor shifts the haircut differential by at most a few basis points, and the effect is statistically indistinguishable from zero in most columns. Among collateral characteristics, specialness is negative but achieves significance only under *time* $\times$ *bond* fixed effects, while the cheapest-to-deliver indicator changes sign across fixed effect structures—positive when only week $\times$ ISIN heterogeneity is absorbed, negative and large once borrower fixed effects are added—suggesting that any genuine CTD effect is confounded by borrower-specific reuse patterns. The aggregate-environment dummies for March 2023 and the post-July-2022 ECB hiking cycle are similarly small and insignificant. Taken together, these results echo the broader message of Section 4: most observable transaction and collateral characteristics fail to systematically shape reuse outcomes once richer heterogeneity is properly absorbed, leaving the dealer-intermediation indicator and the IMR as the only structurally informative coefficients in the second stage.

The most informative auxiliary result concerns the coefficient on the repo-rate differential ΔR . If dealers were managing temporary liquidity through pricing rather than through margins—conceding on the haircut side while extracting an offsetting wedge on the reuse rate—we would expect ΔR to enter the haircut-differential equation with a systematically negative coefficient, indicating that the two pricing dimensions act as substitutes. The data show no such pattern: across all six specifications the coefficient is small, statistically indistinguishable from zero, and switches sign between the unsaturated and saturated columns. Reduced haircut differentials are therefore not being compensated by wider rate spreads on the reuse leg, and the absence of a trade-off between the two pricing dimensions weighs against any reformulation of the windfall mechanism in which dealers preserve a funding subsidy through a different pricing channel. The result is instead consistent with practical frictions that constrain pricing on both margins simultaneously: balance-sheet and leverage-ratio

costs make asymmetric pricing unattractive on either dimension, competitive pressure from money-market funds compresses the scope for widening either spread, and operational risk from daylight exposures favors matching, rather than stacking, of terms across the legs of a reuse chain.

Column (6) addresses the concern that the high incidence of zero haircuts in both initial and reuse transactions could be masking results. While zero-zero pairs are consistent with our conclusion—no haircut differential means no windfall—we show that excluding them does not weaken our findings. In fact, the results become stronger, with a more negative constant and larger coefficients, evidence that the absence of positive haircut differentials is not an artifact of the zero-haircut phenomenon recently documented in the literature (e.g., [FSB, 2025](#); [Hermes et al., 2025](#)).

Because the regression constant absorbs different objects across specifications, the estimated coefficients alone do not deliver a level that can be read as the average haircut differential. To recover a quantity in the same units and with direct economic meaning as the dependent variable, Figure 15 reports the sample average of the model's conditional predictions when dealer intermediation is set to one (Dealer=1) and to zero (Dealer=0), with all remaining covariates and absorbed fixed effects held at their observed values.

The figure delivers a clear and uniform message: once the full battery of econometric controls is applied, the average haircut differential along dealer-intermediated chains remains negative in every specification, ranging from -30.9 bps to -15.8 bps. The shaded confidence bands are constructed based on the standard errors of the estimated counterfactual means. These standard errors are calculated analytically using the delta method, which approximates the variance of a function of estimated parameters.

The point estimate conditional on dealer intermediation never crosses zero. When a dealer intermediates between two non-dealers the haircut on the initial leg is on average *lower* than the haircut on the onward leg, precisely the opposite of the asymmetric wedge predicted by theory on liquidity windfalls. Even in the most demand-

ing specifications (4)–(5), where saturated fixed effects strip away essentially all cross-sectional variation and the implied $\hat{\beta}_{\text{dealer}}$ is statistically indistinguishable from zero, the dealer-intermediated counterfactual remains comfortably below zero, indicating that dealer chains continue to embed a negative differential even within the same borrower–bond–week cell. Far from extracting a temporary liquidity windfall, dealer intermediaries appear to systematically surrender collateral value as securities pass through their books—a result consistent with the descriptive evidence in Table 5.




Figure 15: Dealer Counterfactual Means Across Specifications

Alt text: Point plot with shaded confidence bands across six specifications. The Dealer set to one series (red circles) stays below zero throughout, from minus 30.9 to minus 15.8 basis points, while the Dealer set to zero series (blue squares) is near or above zero in less-saturated columns and falls below zero under saturated fixed effects. A dashed line marks zero.

6 Discussion

Our results paint a picture of European repo markets in which collateral reuse is structurally pervasive but liquidity windfalls are not. Reuse accounts for around 11.6 percent of daily volume — close to 49 billion euros a day on average — and chains routinely span close to three intermediaries beyond the original owner. Yet the haircut wedge that theory predicts dealers should extract from these chains does not appear in the data: average differentials are mildly negative overall and more clearly nega-

Table 7: Stage II – Haircut differential analysis

	No FE (1)	Week FE (2)	Week $\times$ ISIN FE (3)	Week $\times$ ISIN $\times$ Borrower FE (4)	Week $\times$ ISIN $\times$ Borrower FE (5)	Non-zero haircut (pair) (6)
ln(Notional)	0.001 (0.083)	0.001 (0.100)	-0.017* (-1.700)	-0.022 (-1.467)	-0.022 (-1.467)	-0.010 (-0.833)
ln(Term)	-0.042 (-1.273)	-0.043 (-1.387)	-0.020 (-1.111)	-0.019 (-0.679)	-0.019 (-0.704)	-0.007 (-0.250)
Dealer intermediary	-0.314*** (-8.722)	-0.303*** (-8.912)	-0.344*** (-8.821)	-0.034 (-0.291)	-0.034 (-0.293)	-0.080 (-0.615)
ΔR	0.052 (0.703)	0.033 (0.446)	-0.025 (-0.694)	-0.029 (-0.906)	-0.029 (-0.906)	-0.021 (-0.750)
Specialness	-0.028 (-0.500)	-0.046 (-1.278)	-0.116*** (-2.829)	-0.030 (-0.682)	-0.030 (-0.682)	-0.082 (-0.965)
Cheapest-to-deliver	-0.006 (-0.200)	-0.012 (-0.444)	0.092** (2.300)	-0.144*** (-2.880)	-0.144*** (-2.880)	– –
ECB Hike	0.046 (0.541)	-0.045 (-0.429)	-0.040 (-0.690)	-0.047 (-0.810)	-0.047 (-0.810)	-0.032 (-0.516)
March 2023	-0.005 (-0.096)	0.054 (0.730)	0.038* (1.727)	-0.009 (-0.563)	-0.009 (-0.474)	-0.018 (-1.000)
IMR	0.294*** (4.200)	0.288*** (4.056)	0.343*** (2.744)	0.429 (1.521)	0.429 (1.516)	0.323 (1.099)
Constant	-0.335* (-1.872)	-0.270 (-1.607)	-0.077 (-0.939)	-0.218 (-1.298)	-0.218 (-1.290)	-0.405 (-1.468)
Observations	259,450	259,450	246,552	197,875	197,875	75,784
Adjusted R^2	0.054	0.070	0.428	0.764	0.764	0.802
Within R^2	0.054	0.050	0.058	0.003	0.003	0.001
Time FE	No	Yes	No	No	No	No
Time $\times$ Bond FE	No	No	Yes	No	No	No
Time $\times$ Bond $\times$ Borrower FE	No	No	No	Yes	Yes	Yes
Clustered SE	Borrower	Borrower	Borrower	Borrower	Borrower–Week	Borrower

Notes: t-statistics in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Columns (1)–(5) use the full OTC-OTC reuse pair sample and include the inverse Mills ratio from the full-sample first-stage probit. Column (6) restricts to observations by excluding pairs in which both transactions have zero zero haircuts; it uses the inverse Mills ratio from the corresponding non-zero haircut first stage. The CTD indicator is omitted in Column (6) because it is collinear with the saturated $Time \times Bond \times Borrower$ fixed effects.

tive on the dealer-intermediated links that should, if anything, exhibit the most pronounced positive spreads. We discuss below what this configuration implies for the mechanisms emphasized in the literature, for financial stability and monetary policy, and where its main limitations lie.

6.1 Reconciling the evidence with theory

Discussions with market participants suggest a straightforward reading of why the windfall mechanism is not detected. Within financial institutions, securities remaining after necessary coverage has been finalized are swept into a collateral pool, and allocation decisions depend primarily on the achievable repo rate.¹¹ Desks therefore optimize against the rate at which a security can be repoed out, not against margin asymmetries that would generate temporary cash buffers. On this account, the borrowing side of each link is the constrained side, while abundant reserves keep the intermediary's shadow value of temporary cash low. This is consistent with the conjecture of [Issa and Jarnecic \(2024\)](#) that dealers prefer permanent over temporary liquidity, and that pricing, not margining, is the active dimension along reuse chains.

A natural extension of the windfall hypothesis is that dealers might still be extracting a wedge, only through repo rates rather than haircuts. Our second-stage results push back against this reformulation as well: the coefficient on the rate differential is small, statistically indistinguishable from zero, and switches sign across specifications. Reduced haircut differentials are not being compensated by wider rate spreads on the reuse leg. Several practical frictions help rationalize this. Balance-sheet and leverage-ratio costs make asymmetric pricing unattractive on either dimension; competitive pressure from money market funds and other cash providers compresses the scope for widening spreads on the onward leg; and operational risk from daylight exposures favors matching rather than stacking deal terms across the legs of a chain. In this environment, the matched-book economics that windfall models rely on simply

¹¹This account draws on detail about institutions' daily repo operations provided by repo traders and specialists.

do not hold up at the level of individual reuse pairs.

This reading does not invalidate the broader logic of dealer intermediation in repo. Models such as [Infante \(2019\)](#), [Eren \(2014\)](#) and [Brumm et al. \(2023\)](#) capture important features of how collateral circulates and how risk concentrates on dealer balance sheets. What our evidence challenges is the specific assumption that dealers strictly prefer immediate cash and accordingly engineer haircut asymmetries to obtain it. Relaxing that assumption reconciles the logic of intermediation with the data. When temporary liquidity has a low shadow value and intermediation rents accrue through durable counterparty relationships, the natural policy is not a windfall-generating margin asymmetry but a haircut concession on the sourcing leg, offered to retain counterparty flow. The windfall configuration requires, in addition, captive counterparties that can be charged an asymmetric margin without defecting—a condition absent from a predominantly bank-to-bank market in which links form between counterparties with symmetric access to alternative sources of funding.

6.2 Implications for financial stability

The absence of haircut-driven windfalls weakens one of the more acute amplification channels in the literature. Without a systematic positive wedge, the leverage-pyramiding logic that would otherwise inflate gross positions on the back of reused collateral loses much of its bite, and so does the sensitivity to abrupt margin shifts that [Brumm et al. \(2023\)](#) emphasize. By extension, the coordination-failure mechanics in [Infante and Vardoulakis \(2021\)](#) are dampened: if dealers are not over-collateralized on the sourcing leg relative to the onward leg, the asymmetric exposure that drives collateral runs in those models is, on average, not present in our data.

This should not be read as a clean bill of health for collateral reuse. Two qualifications matter. First, reuse remains a substantial source of interconnectedness. The average chain length close to three, broadly stable across the sample and consistent with the pre-2008 levels reported by [Singh \(2011\)](#), means that operational frictions, settlement fails, and idiosyncratic distress at any node can still propagate along the

chain. Even without windfalls, a failure to deliver reused collateral forces replacement or unwinding further down the chain, with potential effects on market functioning. Second, our findings concern central tendencies. Stress episodes are precisely the moments at which haircut asymmetries might emerge, and our sample, while including the March 2023 turmoil and the ECB hiking cycle, does not cover a systemic event of the magnitude of 2008. The structural orientation of liquidity management toward rates rather than margins may itself shift under acute stress, and this is a margin on which continued supervisory attention is warranted.

The findings also align with the recent [FSB \(2025\)](#) emphasis on leverage in the non-bank financial sector. Our results suggest that, at least for the euro-area repo segment, NBFI leverage built on reused collateral is unlikely to be amplified by haircut windfalls of the kind that early models flagged. Instead, the relevant exposures appear to run through the network structure of intermediation — who is connected to whom, through which collateral, and at what chain depth — rather than through margin asymmetries within individual reuse pairs.

6.3 Implications for monetary policy

For monetary policy transmission, the practical message is that the rate channel of repos is doing more of the work than the literature on collateral reuse has typically allowed. If liquidity along reuse chains is managed through repo rates rather than through asymmetric haircuts, then policy rates pass through to secured funding without being significantly muffled by a haircut subsidy that would otherwise cushion dealer funding costs. This is consistent with the segmented-markets evidence in [Eisen-schmidt et al. \(2024\)](#), which already pointed to dealer market power as a friction in transmission, but suggests that the friction operates through pricing rather than through the kind of margin-driven balance-sheet expansion described in [Huh and Infante \(2021\)](#) and [Adrian et al. \(2017\)](#).

By the same token, central bank actions that affect collateral valuations — through asset purchases, eligibility changes, or shifts in the official haircut schedule — should

be expected to transmit primarily through scarcity and pricing rather than through induced changes in private haircut wedges along reuse chains. The risk that policy shocks inadvertently trigger a margin-induced liquidity crunch in sovereign-bond markets, while not negligible, appears smaller than windfall-based models would imply.

7 Conclusion

This paper provides new empirical insights into collateral reuse and liquidity dynamics in European repo markets. Using a novel, transaction-level dataset from the Securities Financing Transactions Data Store, we directly measure collateral chains and test the liquidity windfalls hypothesis.

Two main results emerge. First, reuse in European repos is pervasive and structural rather than marginal. On average, 11.6 percent of daily volume, roughly 49 billion euros, is financed with previously received securities, chains span close to three intermediaries beyond the original owner, and activity is predominantly bank-to-bank. Reuse scales with overall activity and persists across market conditions, supporting its role as a core feature of secured funding rather than an episodic phenomenon. Second, we find no systematic evidence of haircut-driven liquidity windfalls. Average haircut differentials between sourcing and onward legs are mildly negative overall, and more clearly negative on the dealer-intermediated chains where windfall theories would predict the strongest positive wedge; the pattern is not compensated by repo rate spreads, which weighs against reformulations of the mechanism through alternative pricing channels. Beyond these two results, the pattern of differentials is consistent with relationship-driven intermediation. When temporary cash has a low shadow value and intermediation rents accrue through durable counterparty relationships, a haircut concession on the sourcing leg is the natural policy of a franchise intermediary: the negative differential is the price of maintaining its position in the reuse network, and its magnitude, 16 to 31 basis points in our counterfactual estimates, provides a

lower bound on the per-period value of that position. On this reading, the windfall hypothesis is a limiting case that requires both valuable temporary liquidity and captive counterparties, conditions associated with the dealer-to-client financing structures studied in earlier work and absent from a network in which links form predominantly between banks with symmetric market access.

These findings shift the perspective on repo collateral reuse. Without a systematic haircut wedge, the leverage-pyramiding and acute haircut-shift sensitivity that windfall models generate lose much of their force. Liquidity management along reuse chains operates through repo rates and term, and the two pricing dimensions do not substitute for one another: transfers to the cash-borrowing side of a link are cheapest through the haircut, while competition pins the rate at the market benchmark. For monetary policy, transmission to secured funding accordingly runs more through the rate channel than through margin-induced balance-sheet effects. For financial stability, the principal vulnerability associated with reuse lies less in haircut subsidies and more in the network structure of intermediation: chain length, counterparty concentration, and the propagation of operational frictions and settlement fails along maintained links.

By providing a direct, transaction-level mapping of repo-to-repo reuse with broad counterparty coverage, including non-bank financial institutions, and by delivering an identification-based rejection of the windfall hypothesis, this paper reorients empirical and policy attention toward the margins that matter. The relationship-intermediation reading organizes the road ahead, yielding testable conjectures: stress episodes and scarce-reserves regimes raise the shadow value of temporary cash and should compress the concession; leverage-ratio reporting dates shift balance-sheet costs and should move dealer-intermediated reuse volumes and differentials together; and any windfalls that do emerge should concentrate on captive counterparties. Together with regulatory-threshold, collateral-framework, and settlement-fail designs, these tests should further sharpen our understanding of how reuse shapes funding conditions in European repo markets.

References

- Accornero, M. (2020). Collateral reuse, liquidity and financial stability. Technical Report 10/2020, DiSSE Working Paper Series.
- Adrian, T., Boyarchenko, N., and Shachar, O. (2017). Dealer balance sheets and bond liquidity provision. *Journal of Monetary Economics*, 89:92–109.
- Aitken, J. and Singh, M. M. (2010). The (sizable) role of rehypothecation in the shadow banking system. Technical report, International Monetary Fund.
- Andolfatto, D., Martin, F. M., and Zhang, S. (2017). Rehypothecation and liquidity. *European Economic Review*, 100:488–505.
- Bassi, C., Behn, M., Grill, M., and Waibel, M. (2024). Window dressing of regulatory metrics: evidence from repo markets. *Journal of Financial Intermediation*, 58:101086.
- BCBS (2014). Basel iii leverage ratio framework and disclosure requirements. Technical report, Basel Committee on Banking Supervision.
- Bottazzi, J. M., Luque, J., and Páscua, M. R. (2012). Securities market theory: Possession, repo and rehypothecation. *Journal of Economic Theory*, 147(2):477–500.
- Brumm, J., Grill, M., Kubler, F., and Schmedders, K. (2023). Reuse of collateral: Leverage, volatility, and welfare. *Review of Economic Dynamics*, 47:19–46.
- CGFS (2013). Asset encumbrance, financial reform and the demand for collateral assets. Technical Report 49, Committee on the Global Financial System.
- Copeland, A., Martin, A., and Walker, M. (2014). Repo runs: Evidence from the tri-party repo market. *The Journal of Finance*, 69(6):2343–2380.
- Du, W., Hébert, B., and Li, W. (2023). Intermediary balance sheets and the treasury yield curve. *Journal of Financial Economics*, 150(3):103722.
- Eisenschmidt, J., Ma, Y., and Zhang, A. L. (2024). Monetary policy transmission in segmented markets. *Journal of Financial Economics*, 151:103738.
- Eren, E. (2014). Intermediary funding liquidity and rehypothecation as determinants of repo haircuts and interest rates. In *27th Australasian Finance and Banking Conference*.

- FSB (2012). Strengthening oversight and regulation of shadow banking. a policy framework for addressing shadow banking risks in securities lending and repos. Technical report, Financial Stability Board.
- FSB (2017a). Non-cash collateral reuse: Measure and metrics. Technical report, Financial Stability Board.
- FSB (2017b). Re-hypothecation and collateral reuse: Potential financial stability issues, market evolution and regulatory approaches. Technical report, Financial Stability Board.
- FSB (2025). Leverage in nonbank financial intermediation: Final report. Technical report, Financial Stability Board.
- Fuhrer, L. M., Guggenheim, B., and Schumacher, S. (2016). Re-use of collateral in the repo market. *Journal of Money, Credit and Banking*, 48(6):1169–1193.
- Gottardi, P., Maurin, V., and Monnet, C. (2019). A theory of repurchase agreements, collateral reuse, and repo intermediation. *Review of Economic Dynamics*, 33:30–56.
- Griffiths, M. D. and Winters, D. B. (1997). On a preferred habitat for liquidity at the turn-of-the-year: Evidence from the term-repo market. *Journal of Financial Services Research*, 12(1):21–38.
- Hermes, F., Schmeling, M., and Schrimpf, A. (2025). The international dimension of repo: five new facts. Technical Report 3065, ECB Working Paper Series.
- Huh, Y. and Infante, S. (2021). Bond market intermediation and the role of repo. *Journal of Banking & Finance*, 122:105999.
- ICMA (2025). European repo market survey number 49. Technical report, International Capital Market Association.
- Infante, S. (2019). Liquidity windfalls: The consequences of repo rehypothecation. *Journal of Financial Economics*, 133(1):42–63.
- Infante, S., Press, C., and Saravay, Z. (2020). Understanding collateral reuse in the us financial system. In *AEA Papers and Proceedings*, volume 110, pages 482–486.
- Infante, S. and Saravay, Z. (2025). Treasury reuse and the demand for safe assets. *The Review of Financial Studies*. Advance article, hhaf094.

- Infante, S. and Vardoulakis, A. P. (2021). Collateral runs. *The Review of Financial Studies*, 34(6):2949–2992.
- Inhoffen, J. and van Lelyveld, I. (2024). Safe asset scarcity and reuse in the european repo market. Technical Report TI 2024-018/IV, Tinbergen Institute Discussion Paper.
- Issa, G. and Jarnecic, E. (2024). Collateral reuse as a direct funding mechanism in repo markets. *Pacific-Basin Finance Journal*, 86:102449.
- Jaafar, A., Polizzi, S., and Reghezza, A. (2023). Persistency of window dressing practices in the us repo markets after the gfc: The unexplored role of the deposit insurance premium. *European Financial Management*, 29(2):634–663.
- Jank, S., Moench, E., and Schneider, M. (2022). Safe asset shortage and collateral reuse. Technical Report 16439, CEPR Discussion Papers.
- Julliard, C., Pinter, G., Todorov, K., Wijnandts, J.-C., and Yuan, K. (2025). What drives repo haircuts? evidence from the uk market. *Management Science*.
- Kirk, A., McAndrews, J., Sastry, P., and Weed, P. (2014). Matching collateral supply and financing demands in dealer banks. *FRBNY Economic Policy Review*, 20(2).
- Kruttli, M. S., Monin, P. J., and Watugala, S. W. (2022). The life of the counterparty: Shock propagation in hedge fund-prime broker credit networks. *Journal of Financial Economics*, 146(3):965–988.
- Muley, A. (2016). Rehypothecation and monetary policy. Technical report, Massachusetts Institute of Technology.
- Munyan, B. (2017). Regulatory arbitrage in repo markets. Technical report, SSRN.
- Singh, M. (2011). Velocity of pledged collateral: analysis and implications. Technical Report Working Paper 11/256, International Monetary Fund.
- Singh, M. (2014). *Collateral and Financial Plumbing*. Risk Books, London, UK.
- Wallen, J. and Lu, L. (2025). Negative treasury haircuts. Technical report, SSRN Working Paper.

Appendices

The following appendices are part of the submitted manuscript.

A Descriptive Statistics

Table A1: Descriptive Statistics (Alphabetically Ordered)

Variable name	Unit	Freq.	Obs.	Mean	Std. dev.	Min	Max
Ask-Bid	euro	d	11,767,249	0.10	0.13	0.00	3.67
Bond Outstanding Amount	billions	d	15,248,207	20.4	11.7	6.9	67.3
Bond Remaining Maturity	years	t	15,249,327	8.10	8.86	0.002	99.9
Borrower Rating	integer	d	12,248,563	17.66 (A+)	1.93	8 (B)	22 (AAA)
Borrower-Lender Exposure	ratio	q	15,256,046	0.06	0.13	0.00	1.00
Cheapest-to-deliver	pp	d	15,256,046	0.14	0.35	0.00	1.00
Dealer Borrower (DB)	dummy	t	15,256,046	0.19	0.39	0.00	1.00
Dealer Lender (DL)	dummy	t	15,256,046	0.46	0.50	0.00	1.00
ECB Haircut	pp	d	15,246,679	4.23	3.48	0.40	40
ECB Hike	dummy	d	15,256,046	0.70	0.46	0.00	1.00
French Borrower	dummy	t	15,256,046	0.21	0.40	0.00	1.00
French Issuer	dummy	t	15,256,046	0.21	0.40	0.00	1.00
German Borrower	dummy	t	15,256,046	0.27	0.44	0.00	1.00
German Issuer	dummy	t	15,256,046	0.27	0.44	0.00	1.00
Haircut	pp	t	15,256,046	0.14	1.14	-6.00	13.01
Hedge Fund Borrower	dummy	t	15,256,046	0.03	0.16	0.00	1.00
Italian Borrower	dummy	t	15,256,046	0.15	0.36	0.00	1.00
Italian Issuer	dummy	t	15,256,046	0.32	0.47	0.00	1.00
Lender-Borrower Exposure	ratio	q	15,256,046	0.07	0.16	0.00	1.00
March 2023	dummy	d	15,256,046	0.02	0.16	0.00	1.00
Notional	millions	t	15,256,046	23.8	32.3	0.4	216
OIS Spread	pp	d	15,256,046	0.09	0.28	-0.44	1.05
On-the-run	pp	d	15,256,046	0.22	0.42	0.00	1.00
OTC Trade	dummy	t	15,256,046	0.12	0.32	0.00	1.00
Repo Rate	pp	t	15,256,046	1.84	1.92	-0.96	4.00
Reuse	dummy	t	15,256,046	0.21	0.41	0.00	1.00
Specialness	pp	t	15,256,034	0.14	1.71	-4.50	5.75
Term	days	t	15,254,810	2.18	9.74	1	720

Notes: Freq. = Frequency of variables: t = for each transaction, d = daily, q = quarterly. Obs. = Number of observations for each variable. Unit: pp = percentage points.

B Haircut Variability

In analyzing haircuts as drivers of reuse and eventually haircut differentials as a variable of interest, some may point out that there is low variability in haircuts for bilateral transactions for a more precise statistical inference, thanks to current academic consensus. In this section, we show that this does not hold.

Table B1 shows that, even though the fraction of zero haircuts in the sample is significant (around 42%), our number of observations remains robust after discounting those. The mean also shows a shift of only -8 basis points.

Table B1: Summary Statistics: Haircuts

Sample	Observations	Mean	Std. dev.	Min	Max
Full OTC sample	1,819,118	-0.104	1.798	-6.000	13.010
Non-zero OTC sample	1,049,218	-0.180	2.365	-6.000	13.010

Figure B1 shows that this ratio remains consistent through time, rarely varying beyond the 55-60% levels, undermining concerns that haircut related analysis might be biased by time-series changes in these levels.




Figure B1: Share of non-zero haircuts by quarter

Alt text: Marker line from first quarter 2021 to first quarter 2025, ranging mostly between about 49 and 62 percent and staying close to a dashed horizontal line at the overall average of about 57.7 percent.

Finally, Figure B2 shows that there is considerable variability in haircuts, both as

an mean absolute value (to avoid positive and negative haircuts canceling each other out) and by a volume-weighted average.




Figure B2: Haircuts - Daily Average

Alt text: Two lines of over-the-counter haircuts, one a mean absolute value and one a volume-weighted average. Both display considerable day-to-day variability across the sample, with wide swings rather than a smooth trend.

C Dealer List

Institution	AT	BE	DE	DK	ES	FI	FR	GR	IT	LU	NL	PT	SK	EU
ABN AMRO Bank N.V.			✓							✓	✓			
Alpha Bank S.A.								✓						
Banco Bilbao Vizcaya Argentaria S.A.					✓				✓	✓		✓		✓
Banco Cooperativo Español S.A.					✓									
Banco Monte dei Paschi di Siena S.p.A.									✓					✓
Banco Santander S.A.					✓				✓			✓		✓
Barclays Bank Ireland PLC	✓		✓		✓	✓	✓	✓	✓		✓	✓		✓
BofA Securities Europe S.A.	✓	✓				✓	✓	✓	✓		✓	✓		✓
BNP Paribas	✓			✓	✓	✓	✓	✓	✓			✓		✓
BNP Paribas Fortis		✓												
Bred Banque Populaire	✓										✓			✓
CaixaBank S.A.					✓									✓
Cecabank S.A.					✓									✓
Citibank Europe PLC	✓	✓			✓	✓	✓		✓		✓	✓	✓	
Citibank Global Markets Europe AG								✓						✓
Commerzbank AG	✓		✓		✓			✓						✓
Cooperatieve Rabobank U.A.			✓											
Credit Agricole CIB S.A.		✓			✓	✓	✓		✓			✓		✓
Danske Bank A/S			✓	✓		✓								✓
DekaBank			✓											
Deutsche Bank AG	✓	✓	✓		✓	✓	✓	✓	✓			✓	✓	✓
DZ Bank AG			✓											✓
Erste Group Bank AG	✓									✓			✓	✓
Eurobank S.A.								✓						✓
Goldman Sachs Bank Europe SE	✓		✓		✓	✓	✓	✓	✓		✓	✓		✓
HSBC Continental Europe S.A.	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓
ING Bank N.V.											✓			
Intesa Sanpaolo S.p.A.			✓					✓	✓	✓			✓	✓
Jefferies GmbH			✓					✓			✓	✓		
J.P. Morgan SE	✓	✓		✓	✓	✓	✓	✓	✓		✓	✓	✓	✓
KBC Bank N.V.		✓											✓	✓
Landesbank Baden-Württemberg	✓		✓											✓
Landesbank Hessen-Thüringen			✓											
Mediobanca Banca di Credito Finanziario S.p.A.									✓					
Mizuho Securities Europe GmbH			✓						✓					
Morgan Stanley Europe SE	✓	✓		✓	✓		✓	✓	✓		✓	✓		✓
National Bank of Greece S.A.								✓						✓
Natixis S.A.	✓	✓	✓		✓		✓		✓	✓	✓			✓
NatWest Markets N.V.		✓					✓		✓		✓			✓
Norddeutsche Landesbank			✓											
Nordea Bank Abp			✓	✓		✓								✓
Nomura Financial Products Europe GmbH	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓		✓
Piraeus Bank S.A.								✓						✓
Raiffeisen Bank International AG	✓													✓
Scotiabank Designated Activity Company			✓											
Société Générale S.A.	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓			✓
SEB				✓										✓
TD Global Finance Unlimited Company														✓
UBS Europe SE			✓					✓	✓					✓
UniCredit Bank AG (Munich)			✓											
UniCredit S.p.A.	✓								✓					✓
Volksbank Wien AG	✓													

The list of potential dealers is based on registered primary dealers (PD) in any of the reporting countries. A dealer is registered as active in a transaction if: (i) it acts as the cash borrower or cash lender and (ii) the security's issuer is in a country where it holds PD status. XS ISINs have their PDs equalized to EU ones.

D Data Cleaning Details

This appendix provides additional detail on the data filtering and sample construction decisions summarized in Section 3.1.

D.1 Intragroup Transactions

We drop intragroup transactions (20.54% of the original sample) due to the distinct nature of these repos. When a bank moves cash or collateral between its own branches or subsidiaries, the volume and terms of that transfer are often standardized or set in advance. If some of these intragroup repos exist solely to shift assets within the same institution rather than to fund external trades, they serve a fundamentally different purpose from ordinary repo transactions. Consequently, treating any subsequent trades as collateral reuse could stem from a misinterpretation of their true intent.

D.2 Window Dressing

The practice of window dressing by banks has been a topic of debate for decades (e.g., [Griffiths and Winters, 1997](#); [Munyan, 2017](#); [Bassi et al., 2024](#)). Window dressing involves the use of short-term financial transactions to improve accounting numbers around the quarter-end reporting date. Repos lead to an expansion in banks' balance sheets by allowing borrowing against their assets, and the incoming cash increases Leverage Ratio exposure as defined in Basel III ([BCBS, 2014](#)). Banks seeking to report a better version of this metric therefore reduce their repo positions at quarter-ends, a pattern that has been widely documented. This could even involve closing reverse repo positions and using the released cash to decrease repo exposure ([Jaafar et al., 2023](#)), further reducing outstanding volumes.

Window dressing in repo is thus observed through sharp volume reductions at quarter-ends, which tend to affect reuse both in size and motivation, adding confounding effects to the analysis. We follow the analysis in ([Bassi et al., 2024](#), p. 27) to determine an adequate exclusion window: $T - 2$ to $T + 2$ for the first three quarters

and $T - 10$ to $T + 3$ for the last quarter, amounting to around 9.5% of our sample.

D.3 Trimming Outliers

We trim the data at the 1% level for repo rate (1.41% of observations affected), volume (1.31%), and haircut (1.28%) to mitigate potential data quality issues that could result in artificial outliers. The whole process reduces our sample to 43% of the original aggregate volume (entire sample). Notably, 12% of the excluded value is due solely to the 1% trimming process on volumes, corroborating the problems affecting extreme values and substantiating our decision to clean it. The final sample consists of around 15 million observations.

D.4 General Collateral Treatment

In terms of collateral setting, General Collateral (GC) repos allow the cash borrower to choose from a predetermined set of securities, while Specific Collateral (SC) repos allow the cash lender to select the security. Some argue that GC transactions should be excluded from a collateral reuse analysis. [Inhoffen and van Lelyveld \(2024\)](#) point out that in centrally cleared transactions the CCP usually does not allow collateral reuse in GC repos, apart from transactions with the Eurosystem.

Our view is that excluding GC is not necessarily optimal. In bilateral transactions, even though the security is not selected by the cash lender, it can still be re-transferred if the occasion arises, especially since ISINs in GC baskets tend to be high-quality collateral. For centrally cleared operations, the security, although encumbered post-transaction, could have been sourced via reuse initially; excluding it would mean removing at least one initial–reuse pair. Given the small representation of GC repos in our sample, both in number of transactions (5.45%) and total volume (5.1%), we keep them in our baseline analysis.

E The Reuse Identification Algorithm

This appendix provides the complete specification of the algorithm we use to identify collateral reuse and to reconstruct collateral chains in SFTDS. A summary is provided in Section 3. The algorithm builds on the matching approach of [Fuhrer et al. \(2016\)](#) and proceeds in three stages: (i) it matches every repo transaction against subsequent transactions that could plausibly redeploy its collateral (*initial–reuse pairs*); (ii) it assigns a euro value of reuse to each matched pair, resolving conflicts that arise when several reuse transactions draw on the same initial transaction; and (iii) it links matched pairs into complete collateral chains from the first provider to the last taker. Throughout this appendix we describe each stage in turn, state explicitly how conflicts between competing matches are resolved, and illustrate the mechanics with a stylized worked example and a multi-generation example of the algorithm’s transaction-level output.

The algorithm is run separately for every security (ISIN), since reuse by definition involves re-delivery of the same security. It operates on fixed-term repos only: open-term transactions are excluded (see Appendix D) because both the matching window and the conflict-resolution logic require a defined repurchase date. Counterparties are identified at the level of the *final* cash borrower and lender. This look-through matters for centrally cleared trades: matching on the direct counterparty would interpose the CCP on every leg and mechanically break any cleared chain at the clearing house.

E.1 Notation and Inputs

For a given ISIN, the input is the set of repo transactions $n = 1, \dots, N$, each described by the tuple

$$(P_n, T_n, t_n^p, t_n^r, v_n),$$

where P_n is the collateral provider (the final cash borrower), T_n is the collateral taker (the final cash lender), t_n^p is the purchase date, t_n^r is the repurchase date, and v_n is the market value of the collateral in euro. Dates are observed at daily frequency.

Before matching, transactions are sorted in ascending order by (1) purchase date, (2) repurchase date, and (3) collateral value, and indexed in that order. The sort is not innocuous: it determines both which transaction is examined first as a potential initial transaction and the order in which competing reuse candidates claim an initial transaction's collateral (Stage 2). Sorting by earliest start, then shortest maturity, then smallest size makes the algorithm fully deterministic and implements a first-in-first-out convention for collateral, consistent with how trading desks treat fungible inventory.

E.2 Stage 1: Identifying Initial–Reuse Pairs

The algorithm iterates over transactions in index order. When transaction n is examined as a potential initial transaction, a subsequent transaction m qualifies as a candidate reuse transaction if the following four conditions hold simultaneously:

- (C1) both transactions involve the same security (guaranteed by construction, since the algorithm runs ISIN by ISIN);
- (C2) the collateral provider of the second transaction is the collateral taker of the first, $P_m = T_n$: the entity that received the security must be the one re-delivering it;
- (C3) the purchase date of the second transaction falls within the lifespan of the first, $t_n^p \leq t_m^p \leq t_n^r$: the collateral can only be redeployed while it sits on the taker's balance sheet;
- (C4) transaction m has not already been claimed as the reuse transaction of another initial transaction.

All candidates satisfying (C1)–(C4) are matched to n at once and flagged as claimed. Four features of this stage deserve emphasis, as they govern how the algorithm behaves when several matches compete for the same transactions.

Exclusive backing. Condition (C4) enforces that *each reuse transaction is backed by at most one initial transaction*. We recognize that, in practice, a trader may assemble the collateral for one repo from several prior trades. However, given the fungibility of securities and the limited information desks typically retain about the provenance of bonds in their inventory, a one-initial-per-reuse convention is the more parsimonious representation; any part of the reuse transaction’s value not covered by its initial transaction is attributed to the trader’s own inventory (Stage 2).

Greedy chronological priority. Because transactions are examined in sort order and claim all their candidates immediately, earlier initial transactions have priority: a transaction claimed by initial transaction n is no longer available to any later initial transaction, even one that would also satisfy (C1)–(C3). Matching is therefore a transparent greedy rule rather than a global optimization; combined with the deterministic sort, it makes the output fully reproducible.

The drain rule. Once a transaction has been examined as an initial transaction, it is marked as drained and can no longer be claimed as a reuse transaction by a later initial. Together with the sort order, this prevents a specific double-counting problem: two same-day transactions in the same ISIN between the same two entities but in opposite directions would otherwise be identified as reuse cases *of each other*, which is implausible.

Chains remain possible. The drain rule does not prevent chains: a transaction that has been claimed as the reuse of an earlier transaction still acts, in its own turn, as a potential initial transaction for further redeployment. This is precisely what allows Stage 3 to trace collateral across more than two links.

Finally, note that no condition is imposed on the *repurchase* date of the reuse transaction: the reuse leg may mature after the initial leg. This is a deliberate deviation from [Fuhrer et al. \(2016\)](#), discussed in Section [E.6](#) below.

E.3 Stage 2: Value Assignment and Conflict Resolution

Stage 1 establishes *which* transactions are linked; Stage 2 determines *how much* collateral value flows through each link. Let $M_n = \{m_1, m_2, \dots, m_K\}$ denote the reuse transactions matched to initial transaction n , ordered by the dataset sort. The assignment applies three rules.

Rule 1: the cap. The value of reuse assigned to any single match can never exceed the value of the initial transaction:

$$\bar{v}_{m_k} = \min(v_{m_k}, v_n).$$

If the reuse transaction is larger than its initial transaction, only the covered portion is classified as reuse; the remainder is attributed to the taker's own inventory. Partial reuse is thus explicitly allowed.

Rule 2: overlapping matches share the initial transaction's capacity. When an initial transaction backs several reuse transactions, the algorithm checks whether they overlap in time. Two matched transactions overlap when the later one starts before the earlier one matures, $t_{m_k}^p \leq t_{m_j}^r$ for $j < k$: in that case, both would draw on the initial transaction's collateral *simultaneously*, and their combined reuse value must be disciplined by the collateral actually available. The earliest match, m_1 , always receives its full capped value $\bar{v}_{m_1}$. Each subsequent match that overlaps earlier ones can be assigned at most the initial transaction's *residual capacity* — the value of the initial transaction net of what is already deployed in the overlapping matches — with the checks performed pairwise in chronological order:

$$r_{m_k} = \min(\bar{v}_{m_k}, \text{residual capacity of } n \text{ at } t_{m_k}^p).$$

When the residual capacity reaches zero, subsequent overlapping candidates are assigned a value of zero: they are *not* classified as reuse (their collateral is attributed en-

tirely to own inventory), and — consistent with the exclusive-backing rule — they are not returned to the pool to be re-matched with other initial transactions. This choice is conservative: it curbs overestimation at the risk of discarding some genuine reuse.

Rule 3: maturing matches restore capacity. If a later match does *not* overlap any earlier match — that is, it starts only after all previously matched reuse legs have matured — the collateral has returned to the taker’s balance sheet and can be redeployed anew. Such a match receives its full capped value, regardless of how much was assigned to earlier, already-matured matches. An important corollary is that the *total* reuse recorded against one initial transaction over its lifespan may legitimately exceed its collateral value, because the same security can be reused sequentially; what the algorithm restricts is only the amount reused *at any one time*.

For every transaction, Stage 2 records four variables. On the reuse leg: the *value of reuse* r_m (the euro amount of its collateral identified as redeployed from its initial transaction) and a *reuse dummy* equal to one whenever $r_m > 0$. On the initial leg: the *total reuse* $R_n = \sum_k r_{m_k}$ (the cumulative euro amount of its collateral redeployed by its taker) and the *count of reuse* C_n (the number of reuse transactions with positive assigned value that it backs). The reuse dummy underlies the transaction-level regressions of Section 4, while the value of reuse aggregates to the daily reuse volumes and the 11.6% reuse rate reported in the main text.

One further remark concerns onward reuse. When a reuse transaction m later acts as an initial transaction for the next link of a chain, the capacity available for onward redeployment is its full collateral value v_m , not merely the portion r_m classified as reused. This again reflects fungibility: the taker in m receives v_m of the security and can redeploy all of it, irrespective of where the provider sourced it.

Conflict resolution at a glance

Since the treatment of competing matches is the substantive heart of the algorithm, we summarize the complete set of tie-breaking and conflict rules:

1. **One-to-many is allowed.** A single initial transaction may back multiple reuse transactions — concurrently, up to its collateral value, and sequentially, without limit, as long as each reuse leg starts within its lifespan.
2. **Many-to-one is not allowed.** Each reuse transaction is backed by at most one initial transaction; any shortfall is attributed to own inventory rather than to a second initial transaction.
3. **Each transaction is claimed at most once,** with priority given by the sort order (earliest purchase date, then earliest repurchase date, then smallest value). Earlier initial transactions claim first.
4. **Concurrent reuse is capacity-constrained.** Overlapping matches divide the initial transaction's value first-come-first-served: the earliest claims in full, later ones receive at most the residual.
5. **Sequential reuse is unconstrained.** Once a matched reuse leg matures, the collateral is available again in full; total reuse over the initial transaction's life may exceed its value.
6. **Exhausted matches are dropped conservatively.** A match arriving after capacity is exhausted is not counted as reuse at all and is not reassigned elsewhere.
7. **Mutual matches are ruled out.** The drain rule prevents two offsetting same-day transactions between the same entities from being classified as reuses of each other.
8. **Chains are permitted.** A claimed reuse transaction can itself serve as the initial transaction of the next link.

E.4 A Worked Example

Table [E1](#) illustrates Stage 2 with a stylized example. Entity A delivers €100 million of a bond to entity B against cash (transaction n , running from 10 to 20 January). Stage 1

matches four subsequent transactions in which B re-delivers the same ISIN within that window.

Table E1: Stylized Example: Value Assignment and Conflict Resolution

	Leg	Purchase	Repurchase	Value	Assigned reuse	Rule applied
n (initial)	A $\rightarrow$ B	10 Jan	20 Jan	100	—	backs m_1 – m_4 ; total reuse = 180; count = 3
m_1	B $\rightarrow$ C	11 Jan	14 Jan	60	60	earliest match: full value
m_2	B $\rightarrow$ D	12 Jan	13 Jan	55	40	overlaps m_1 : residual $100 - 60 = 40$; €15m own inventory
m_3	B $\rightarrow$ E	13 Jan	14 Jan	30	0	overlaps m_1, m_2 : capacity exhausted; not classified as reuse
m_4	B $\rightarrow$ F	15 Jan	19 Jan	80	80	no overlap: collateral returned, full capacity restored

Notes: Values in € million. All five transactions involve the same ISIN; each m_k satisfies conditions (C1)–(C4) with respect to n . “Assigned reuse” is the value of reuse r_{m_k} recorded on the reuse leg; the reuse dummy equals one for m_1, m_2 , and m_4 , and zero for m_3 .

The example activates every conflict rule. The earliest match, m_1 , claims its full €60 million (Rule 2, first claimant). Transaction m_2 starts while m_1 is still outstanding, so only the residual $100 - 60 = 40$ is available: €40 million is classified as reuse and the remaining €15 million as B’s own inventory (Rules 1 and 2). Transaction m_3 also overlaps both outstanding matches, but capacity is now exhausted ($60 + 40 = 100$), so it is assigned zero, carries a reuse dummy of zero, and remains in the dataset as an ordinary transaction — it is not re-matched to any other initial transaction (Rule 6, conservative exclusion). By 15 January, however, m_1, m_2 , and m_3 have all matured and the collateral is back with B; transaction m_4 therefore overlaps nothing and receives its full €80 million (Rule 3). The initial transaction ends with a total reuse of $60 + 40 + 80 = 180 > 100$ — entirely legitimate, since the collateral was reused sequentially, never simultaneously beyond the €100 million available — and a count of reuse of three. If C, D, or F subsequently re-delivers the bond within the lifespan of its own leg, Stage 1 will match those transactions in turn, and Stage 3 will link them into a chain.

Table E2 complements the stylized example by showing what the algorithm’s transaction-level *output* looks like when redeployment spans several generations. Since SFTDS data are confidential and individual transactions cannot be disclosed, the table is populated with fictitious entities, dates, and amounts, constructed to be exactly consistent with the rules above. Entity B receives €50 million of the bond from entity A (transaction 30575, running from 4 to 11 August) and re-delivers it in four subsequent trans-

actions. The first two, 35354 (€20 million) and 35373 (€25 million), are covered in full. Transaction 35402 overlaps both, so only the residual $50 - 20 - 25 = 5$ million of its €15 million is classified as reuse (Rule 2). By 10 August all three legs have matured, and 35602 therefore starts against restored capacity (Rule 3), but is capped at the initial transaction’s €50 million (Rule 1), with the remaining €15 million of its €65 million attributed to B’s own inventory. Transaction 30575 accordingly ends with a count of reuse of four and a total reuse of €100 million — twice its own value, reflecting sequential redeployment.

The later generations illustrate the remaining rules. Entity D redeploys the €25 million received in 35373 into transactions 35418 (€10 million, in full) and 35423 (€15 million of €18 million, capped by the residual); both mature on 9 August, *after* their initial transaction 35373 itself matures — permitted by the relaxed repurchase-date condition discussed in Section E.6. Entity E redeploys €12 million in transaction 35456 even though only €5 million of its incoming leg 35402 was classified as reuse: onward capacity is the full collateral amount received, consistent with fungibility. Finally, entity H passes €14 million onward in transaction 35540, completing the fourth generation.

Table E2: Illustrative Example: Transaction-Level Output of the Algorithm

ISIN	Index	Purchase date	Repurchase date	Provider	Taker	Amount	Value of reuse	Total reuse	Reuse dummy	Count of reuse
DE0001	30575	04/08/2023	11/08/2023	Entity A	Entity B	50,000,000	0	100,000,000	0	4
DE0001	35354	07/08/2023	08/08/2023	Entity B	Entity C	20,000,000	20,000,000	0	1	0
DE0001	35373	07/08/2023	08/08/2023	Entity B	Entity D	25,000,000	25,000,000	25,000,000	1	2
DE0001	35402	07/08/2023	09/08/2023	Entity B	Entity E	15,000,000	5,000,000	12,000,000	1	1
DE0001	35418	08/08/2023	09/08/2023	Entity D	Entity G	10,000,000	10,000,000	0	1	0
DE0001	35423	08/08/2023	09/08/2023	Entity D	Entity H	18,000,000	15,000,000	14,000,000	1	1
DE0001	35456	09/08/2023	10/08/2023	Entity E	Entity I	12,000,000	12,000,000	0	1	0
DE0001	35540	09/08/2023	10/08/2023	Entity H	Entity J	14,000,000	14,000,000	0	1	0
DE0001	35602	10/08/2023	11/08/2023	Entity B	Entity F	65,000,000	50,000,000	0	1	0

Notes: Illustrative example with fictitious entities, dates, and amounts (SFTDS data are confidential); all figures are exactly consistent with the matching and conflict-resolution rules of Appendix E. Amounts in euro. *Index* is the transaction’s position in the sorted dataset (gaps correspond to unrelated transactions in the same ISIN). *Amount* is the market value of the collateral. *Value of reuse* is the portion of the transaction’s collateral identified as redeployed from its (unique) initial transaction; any remainder is attributed to the provider’s own inventory. *Total reuse* is the cumulative amount of the transaction’s collateral redeployed by its taker in subsequent transactions. *Reuse dummy* flags transactions with a positive value of reuse; it is zero for transaction 30575, which opens the chain. *Count of reuse* is the number of reuse transactions the transaction backs.

E.5 Stage 3: Chain Reconstruction

Stages 1 and 2 produce a mapping from every initial transaction to the reuse transactions it backs (with positive assigned value). Because each reuse transaction has at most one initial transaction, this mapping defines a forest: a collection of directed trees whose nodes are transactions and whose edges are initial–reuse links. Stage 3 reads chains off this forest as follows:

1. *Roots.* A root is an initial transaction that is not itself the reuse transaction of any earlier trade — the point where the collateral first enters the market from its original owner.
2. *Traversal.* Starting from each root, a depth-first search follows initial–reuse links until it reaches a transaction that backs no further reuse (a leaf).
3. *Chains.* Every maximal root-to-leaf path with at least one link is recorded as a chain. When a transaction backs several reuse transactions, the tree branches, and one chain is recorded per leaf; branching chains therefore share their upstream links.

For each chain, the algorithm stores the ISIN, the ordered list of transaction indices and repo identifiers, the purchase and repurchase dates of the first and last links, and the chain length, defined as the number of reuse links. A chain with k transactions connects $k + 1$ institutions — that is, k institutions beyond the original collateral owner — which is the interconnectedness measure analyzed in Section 4.

Figure E1 displays the forest implied by the example output of Table E2. From the single root 30575, the depth-first search records five maximal chains: (30575, 35354), (30575, 35602), and (30575, 35402, 35456) alongside the two branches through 35373. The highlighted path is the longest: transaction 30575 ($A \rightarrow B$) is redeployed in 35373 ($B \rightarrow D$), again in 35423 ($D \rightarrow H$), and once more in 35540 ($H \rightarrow J$) — a chain of length three, spanning five institutions from the original owner A to the final taker J. Branching chains share their upstream links: the chain ending in 35418 traverses the same first two transactions as the highlighted one.




Figure E1: Reuse chains implied by the example output of Table E2. Each arrow links an initial transaction to a reuse transaction it backs; every transaction has at most one incoming arrow (exclusive backing), so the links form a tree rooted at transaction 30575. The highlighted path is the longest of the five maximal chains, 30575 → 35373 → 35423 → 35540: a chain of length three passing the collateral from entity A through B, D, and H to J, five institutions in total.

Alt text: Tree diagram in four generation columns: transaction 30575 sits as the single root on the left and branches through four second-generation nodes, then third- and fourth-generation nodes. Grey arrows mark initial-to-reuse links, and a highlighted blue path traces the longest chain, 30575 to 35373 to 35423 to 35540.

E.6 Design Choices and Relation to the Baseline Approach

Two deliberate deviations from the baseline approach of [Fuhrer et al. \(2016\)](#) deserve restating alongside the full algorithm.

First, we impose no condition on the repurchase date of the reuse transaction: the onward leg may mature after the incoming leg. Requiring the reuse transaction to be fully nested within the initial transaction’s lifespan presumes that the re-user must have the security back in hand before returning it, whereas in practice entities routinely source the same (fungible) bond elsewhere when timing mismatches arise. The stricter nesting condition leads to significant underestimation: when [Fuhrer et al. \(2016\)](#) drop it, their headline reuse rate rises to approximately 10%, close to the 11.6% we estimate (Section 4).

Second, as described above, we do not allow more than one initial transaction to back a single reuse transaction, attributing any uncovered portion to own inventory instead. This convention is conservative with respect to the liquidity-windfall analysis in Section 5: it can only understate the amount of matched reuse, never fabricate links between unrelated trades.

More broadly, the algorithm is a deterministic greedy procedure, not a global match-

ing optimization. We view this as a feature: the first-in–first-out claiming convention mirrors desk practice for fungible collateral, every match is traceable to an explicit rule, and the output is exactly reproducible from the sorted transaction record. Together, the capacity constraints of Stage 2 (which curb overestimation from overlapping matches) and the relaxed repurchase-date condition (which curbs underestimation from timing mismatches) balance the two error margins, which we believe makes the resulting reuse estimates more reliable than either polar design.