Short-Circuiting Short-Term Funding*

R. Jay Kahn, Neth Karunamuni, Mark Paddrik

October 8, 2025

Abstract

This paper examines the potential impact of cyber-induced operational outages in the U.S. repo market. Using transaction-level data and institutional cybersecurity ratings, we simulate disruptions to key cash lenders. Our findings indicate that outages at certain institutions can disrupt over $100 billion in funding and raise repo rates by over 50 basis points. The severity of these disruptions are sensitive to outage timing and duration, with peak settlement times and slower recoveries amplifying stress. The results underscore the importance of both cybersecurity preparedness and institutional resilience in limiting financial market disruption. By linking cyber risk to intraday funding dynamics and rate volatility, this study contributes to the financial stability and operational risk literature, offering a framework for assessing and mitigating cyber threats in core funding markets.

*Kahn, jay.khan@frb.gov, Federal Reserve Board of Governors; Karunamuni, Neth.Karunamuni@ofr.treasury.gov, U.S. Department of the Treasury Office of Financial Research; and Paddrik Mark.Paddrik@ofr.treasury.gov, U.S. Department of the Treasury Office of Financial Research; We would like to thank Mark Carey, Corey Garriott, Stacey Schreft, Stathis Tompaidis, and Peyton Young for their insightful comments in advancing this work. The views expressed in this paper do not necessarily reflect the views of the Office of Financial Research, the U.S. Department of the Treasury, or the Federal Reserve Board and do not establish supervisory policy, requirements, or expectations.

1 Introduction

The U.S. repurchase agreement (repo) market is the largest short-term wholesale funding market in the financial system and serves as a critical medium for liquidity provision and monetary policy implementation. Facilitating over $10 trillion in outstanding agreements daily, the market enables financing by using U.S. Treasuries and other high-quality securities as collateral. As in other over-the-counter (OTC) markets, activity is characterized by persistent bilateral relationships, with volumes evolving in a relatively stable and predictable manner. The centrality of the repo market to intermediation and leverage across both bank and nonbank institutions renders it an essential component of the broader financial infrastructure.

Given the magnitude of flows between counterparties, the functioning of the repo market is sensitive to disruptions in the operational processes that support daily settlement and liquidity redistribution. Repo transactions connect a wide array of institutions, such as dealers, banks, hedge funds, and money market funds, through time-sensitive cash and collateral exchanges. Disruptions at key participants or infrastructure providers can impair access to short-term financing and amplify systemic risk. While no known operational disruptions have previously impacted the market, sudden spikes in rates due to the perceived liquidity shortfalls have underscored the market's vulnerability, as was seen on September 17, 2019, when repo rates surged from approximately 2% to nearly 10%. The episode highlighted the potential for funding stress to propagate through the system rapidly with significant implications for broader financial stability.

The risks posed by operational disruptions have become increasingly salient in assessing cyber threats, which present a growing source of systemic vulnerability. This paper investigates the resilience of the U.S. repo market and specifically of the tri-party segment to operational outages induced by cyberattacks. We combine counterparty-level transaction data with institutional cybersecurity ratings to identify which classes of participants are more susceptible to cyber-induced disruptions. We then quantify the extent to which these vulnerabilities can generate funding shortfalls, alter trading patterns, and affect the pricing of secured funding.

Several institutional features of the tri-party repo market inform our analysis. First, trading relationships are highly persistent and concentrated; market participants transact with a narrow and stable set of counterparties. As such, the sudden unavailability of a key counterparty may result in immediate funding gaps since substitution is not frictionless. Second, participant size is highly heterogeneous, and certain institutions may carry disproportionately more impact on the availability of financing volumes and/or lending rates. Third, trading activity is concentrated during specific time windows, particularly during morning hours, making the timing of disruptions a critical determinant of their

severity. While the Federal Reserve provides repo backstops, access is conditional and usage may be constrained by regulatory or strategic considerations.

We evaluate the market-wide consequences of cyber-induced operational outages through a series of simulation exercises. First, we simulate disruptions at individual institutions, allowing the identity of the affected party to vary, to estimate the range of potential impacts on funding volumes and pricing. Next, we incorporate institution-specific cybersecurity ratings to generate probabilistic estimates of outage occurrence and duration. Using these inputs, we assess the distribution of impacts in multi-institutional disruption scenarios by, evaluating effects on transaction volumes, counterparty reach, and repo rates. Finally, we explore how differences in institutional resiliency modeled through recovery speeds shape the outcomes. Our findings underscore the importance of both preventative cybersecurity measures and recovery capacity in mitigating systemic risk in critical funding markets.

This paper makes three key contributions to the literature on operational risk and financial market resilience. First, it integrates cybersecurity vulnerabilities into the analysis of operational disruptions in the repo market. While cybersecurity ratings identify bank-dealer lenders and bank borrowers as having the highest risk of disruption, our findings emphasize that asset managers, primarily money market funds, are the most systemically important institutions in this context. As the dominant source of secured funding, their sudden unavailability poses a significant threat to market liquidity, particularly for transactions backed by U.S. Treasuries, which are the predominant form of collateral. The structure of the tri-party market amplifies this fragility due to its reliance on a limited number of large institutions to distribute daily funding. Counterfactual simulations demonstrate that enhancing the cybersecurity posture of asset managers yields the most substantial reduction in market-wide risk.

Second, we translate operational outages into economically meaningful measures of market stress. Simulated outages of individual lenders reveal that the average disruption affects seven borrowers, with some disruptions impacting more than twenty borrowers and more than $100 billion in funding. The resulting interest rate impact is substantial with a large lender's outage increasing the aggregate market's interest rates by more than 50 basis points, highlighting the system's sensitivity to the withdrawal of key participants. Particularly due to the stable nature of repo trading relationships where the daily gross volume average change is less than 2%, the effects of a disruption are highly concentrated. Some bank-dealers are among the most impacted, and banks are disproportionately exposed as borrowers. Under the most extreme scenarios, primary dealers face the largest interest rate shocks, highlighting the impact of second-order effects.

Finally, our third contribution is an examination of market resiliency and recovery dy-

namics in the face of cyberattacks. The timing and duration of an outage are critical in shaping the scale of disruption. Outages occurring during peak settlement hours, particularly at 9 a.m. and 2 p.m., have the most immediate and widespread consequences. However, outages that begin earlier in the day (e.g., between 7 and 8 a.m.) can prove even more damaging when institutions exhibit low resiliency and slow recovery times. These findings suggest that regulatory focus should extend beyond baseline cybersecurity to include operation recovery timeliness, with special attention given to timing vulnerabilities and concentration risks in the tri-party repo ecosystem.

Related Literature. The repo market is sensitive to shocks in trust, collateral value, and operational continuity. Prior research has shown that these vulnerabilities are exacerbated in times of market turmoil, where rollover risk and funding fragility become systemically important (López-Espinosa et al. (2012)). Studies of the 2007-09 financial crisis show that investors rapidly withdrew funding from certain borrowers, particularly those perceived to be exposed to credit or liquidity risk, mirroring the dynamics of a traditional bank run but occurring in the shadow banking system (Gorton and Metrick (2012); Copeland et al. (2014)). These runs were driven by increased concerns of a borrower's default and the consequential increased haircuts or outright refusal to roll over repos, especially on relatively lower-quality collateral that counterparties did not want to hold (Gorton et al. (2020); Copeland and Martin (2025)).

Our results build on and extend this literature introducing operational and cybersecurity-driven disruptions as a new channel of systemic risk in the repo market. Our findings suggest that cyber and operational outages could mimic the same disruption dynamics as seen after credit- or collateral-based funding shocks, even in the absence of credit deterioration. Importantly, this study adds a forward-looking dimension to the literature by demonstrating how risk not only resides in asset quality or perceived solvency but rather in the operational integrity and resilience of key market participants, most notably asset managers, whose interruption could disproportionately disrupt market function.

While repo runs have been a less common concern in recent years, research has shown that spikes in U.S. repo rates have become a recurring sign of stress, particularly at quarter-end and year-end. The most notable episode occurred in mid-September 2019, when overnight repo rates surged from around 2% to as high as 10% (Paddrik et al. (2023)), disrupting short-term funding markets and prompting immediate intervention by the Federal Reserve (Afonso et al. (2021)). Despite relatively high aggregate reserve balances in the U.S. banking system at the time, Copeland et al. (2025) argues that liquidity was not sufficiently available when and where it was most needed, leading to delays in outgoing payments by key institutions. In addition, Paddrik et al. (2023) highlights that liquidity constraints such as large Treasury settlements and corporate tax payments, along with

rigid institutional trading patterns, limited the market's ability to reallocate cash effectively. Thus, rather than stemming from a broad-based liquidity shortage, these spikes reflect issues of operational timing and fungibility of relationships, revealing how settlement and funding frictions can generate acute systemic stress even in an otherwise liquid market.

By simulating the impact of lender-specific cyber outages and mapping them to changes in repo volumes and rates, this paper builds on the insight that timing and network location and, not just quantity, of liquidity matters. The finding that outages at peak hours or with longer recovery times lead to disproportionately large disruptions reinforces the view that intraday funding dynamics are central to systemic stability, especially in a market as relationship dependent and time-sensitive as tri-party repo.

These frictions highlight the importance of operational continuity in repo, payments, or clearing systems that rely on daily rollover and timely interdependent payments. Such financial systems are vulnerable to timing mismatches, coordination failures, and erosion of trust, creating stability risk (Duffie and Younger (2019)). Building on the stress-testing approach of Eisenbach et al. (2022), we show how outage timing, firm resilience, and institutional concentration shape the scale and distribution of market disruptions. Our work complements Kotidis and Schreft (2025) by extending the empirical focus from a bank-level natural experiment to a full-market simulation, thereby capturing the cascading effects of funding withdrawal and collateral imbalances across counterparties. We also provide evidence in support of Eisenbach et al. (2023) by showing that cybersecurity vulnerabilities, particularly among key liquidity providers, translate into tangible stress in the pricing and functioning of repo markets under plausible attack scenarios.

Ultimately, this paper adds a new dimension to the literature by integrating cybersecurity preparedness with market microstructure, and by emphasizing the importance of resiliency and not just cybersecurity as a determinant of systemic risk. In doing so, it offers a framework for regulators and market participants to assess not only who is most at risk from a cyber event, but also which disruptions are most dangerous and which mitigation strategies are most effective.

The remainder of this paper is organized as follows. Section 2 provides a background on the importance of the repo market to the financial system, with a focus on the role, make up, and settlement of tri-party repo segment. Section 3 provides an assessment of the consequences of an operational disruption. Section 4 introduces cybersecurity scores and discusses the information these scores reveal about tri-party repo participants. Section 5 applies the cybersecurity scores to estimate the expected impact of a cybersecurity disruption on the market. Section 6 assesses how the resilience of firm recovery influences the implications of a disruption. Section 7 concludes.

2 Background on Short-Term Funding

Money markets and their associated rates play a pivotal role in maintaining financial stability. As a core channel for short-term liquidity, money markets are essential for transmitting monetary policy, meeting daily funding needs, and ensuring the smooth functioning of payment and settlement systems. Conversely, disruptions in these markets can spill over into broader financial conditions, amplifying systemic risk.

Money market rates are indicative of the prevailing liquidity conditions in the financial markets. When these rates are stable, it signifies a healthy level of liquidity in the banking system. Adequate liquidity ensures that financial institutions can meet their short-term funding needs and settle their obligations promptly, preventing potential liquidity crises that can lead to bank failures or distress. Furthermore, money market rates have a cascading effect on various financial instruments and markets. They can influence the yields on short-term government bonds, commercial paper, and corporate debt. This, in turn, affects the cost of borrowing for businesses and consumers, impacting investment decisions and overall economic activity.

Repo markets, which involve the sale of assets together with an agreement to repurchase them on a specified future date at a prearranged price, are used by market participants for many reasons, including financing their portfolios, using cash as collateral to borrow securities, and as a safer alternative to uninsured deposits. The assets underlying a repo are used as collateral to protect cash lenders against the risk that cash borrowers will fail to return the cash. The interest rate on these transactions is calculated based on the difference between the sale price and the repurchase price of the assets underlying the repo.

In principle, the repo rate in the U.S. is governed by the Federal Reserve through its standing lending and borrowing facilities. However, there have been several instances in which overnight rates have jumped outside the target range: The most notable occurring on September 17, 2019, when the average rate was over 3% and as high as 10% for some repo participants (Paddrik et al. (2023)). Figure 1 shows the Federal Funds target rate used as a principal component to the U.S. reference rate, the Secured Overnight Financing Rate (SOFR), along with various rates the Federal Reserve offers banking institutions.

The short-lived nature of the volatility observed in repo rates points to a potential issue regarding the adequacy of funding for financial institutions. In these situations, such as the September 2019 episode, there could be a lack of readily available cash or assets in the financial system. They may be indicative of a sudden surge in demand for short-term funds, which can be driven by various factors like regulatory requirements, unexpected financial stress, or imbalances in the supply and demand for funds (Afonso et al. (2021); Copeland et al. (2025)). However, they also highlight ripple effects resulting from the need

Figure 1: Repo Rates and Federal Funds Rates (percent)

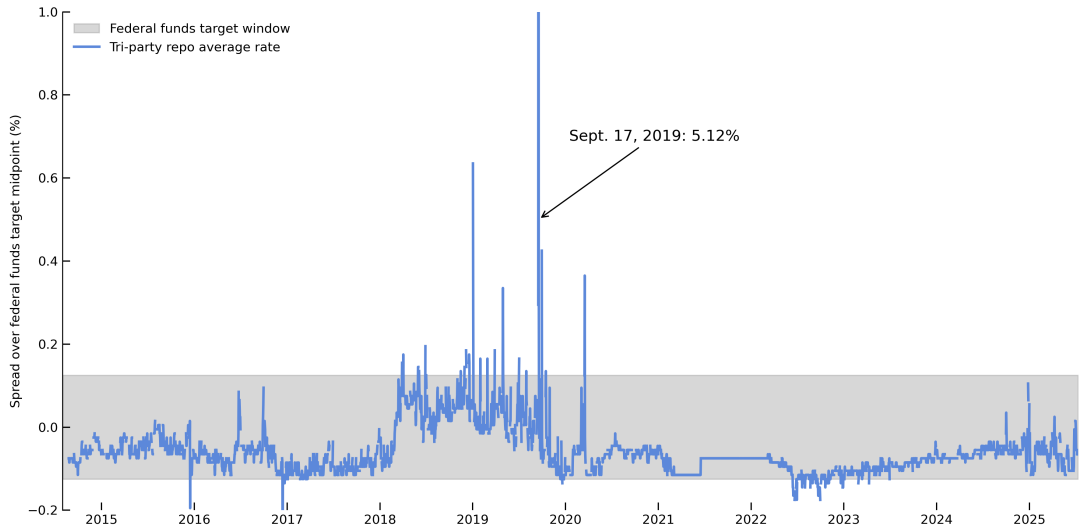

Note: Tri-Party repo average rate is the weighted average daily rate on new overnight Treasury repo transactions from BNY repo data. All rates are spreads over the federal funds target range midpoint.
Source: Federal Reserve's Tri-Party Repo Collection, FRED, Authors' analysis.

to search for funding, causing disruption and exacerbating the liquidity problem.

In essence, the observation of sudden, short-lived volatility in repo rates shows that financial institutions might sometimes face challenges in accessing enough short-term funds. The sudden nature of the rate change reflects the difficulty of finding new funding and uncertainty in the degree of short-term funds availability.

2.1 Repo Market Segmentation and Fragmentation

Notably, U.S. repo markets are fragmented with respect to where and with whom repo agreements are made, which can make it more difficult for investors to get a clear picture of the overall market. In a fragmented market is fragmented, market participants have to track multiple liquidity pools. This creates complexity for those searching for short-term funding, and it can make it difficult for participants to understand what is happening in the market.

Second, market fragmentation can lead to increased volatility because there is less liquidity in each trading segment (Anbil et al. (2021)). This means that even small demand or supply shocks can have a big impact on rates and volumes.

The U.S. repo market can be divided into four major segments (see Figure 2), depending on two factors: (1) whether the trades are settled bilaterally or through a tri-party custodian and (2) whether the trades are centrally or non-centrally cleared through the Fixed Income Clearing Corporation (FICC). The focus of this paper will be on the tri-party non-centrally cleared segment, or simply known as tri-party, which is the main venue for

cash supplying entities (e.g., money market funds) and the segment through which the Federal Reserve intervenes in the repo market.

Figure 2: The Four Main Segments of the U.S. Repo Market

		Settlement	
		Tri-Party	**Bilateral**
Clearing	**Centrally**	Fixed Income Clearing Corporation (FICC) GCF Repo · Centrally cleared by FICC · Settled on BNY's Tri-Party platform · General collateral repo only	FICC DVP Service · Centrally cleared by FICC · No central custodian · Specific collateral repo possible
	Non-Centrally	Bank of New York Mellon (BNY) Tri-Party · No central counterparty · Settled on BNY's Tri-Party platform · General collateral repo only	Non-Centrally Cleared Bilateral Repo · No central counterparty · No central custodian · Specific collateral repo possible

Source: Authors' creation.

Relationships are critical in over-the-counter markets, and the repo market is no exception. This is especially true in the non-centrally cleared segments. In these segments, trading requires bilateral agreements, typically governed by Global Master Repurchase Agreements (GMRAs), with key terms such as collateral type and haircut levels negotiated in advance.

Repo participants continually assess the creditworthiness of their counterparties. Lenders are more likely to extend funding to borrowers with strong credit histories and established reputations while those perceived as higher risk may face difficulty accessing liquidity. As a result, participants often build long-term trading relationships with trusted counterparties. Regulatory requirements reinforce this persistence, as firms must conduct due diligence and report transactions to regulators, further discouraging frequent counterparty switching.

Collateral quality also plays a central role. High-quality assets, such as U.S. Treasuries, are preferred due to their liquidity and lower risk. Regulatory constraints may limit which collateral types institutions can accept, encouraging participants to specialize in specific collateral classes and fragmenting the market accordingly. Moreover, establishing a GMRA involves legal and operational overhead that creates additional incentives for participants to maintain durable trading relationships over time.

2.2 Tri-Party Repo

We focus on the tri-party non-centrally cleared segment due to its OTC transaction structure, its role as the venue through which the Federal Reserve conducts operations,

and the availability of regulatory transaction and collateral data. The tri-party repo market segments data is collected by the Federal Reserve Bank of New York, which receives regular reports from the sole tri-party custodian, Bank of New York Mellon (BNY).[1] Our dataset spans from January 2016 through December 2024 and includes detailed information on each transaction, including the identities of the cash lender and borrower, the principal amount of cash borrowed, the collateral posted, its market value, and the total interest due at maturity.

Figure 3 shows the total daily cash borrowed in the tri-party market over our sample period, disaggregated by collateral type. Two features stand out. First, U.S. Treasuries dominate as collateral, comprising about 70% of the total, followed by Agency Mortgage-Backed Securities (MBS) and debt at 20%. Second, there is a notable surge in transaction volume in 2021, largely driven by the Federal Reserve's Overnight Reverse Repo Facility (ON RRP), which was used to support monetary policy implementation and reinforce the lower bound of the federal funds target rate.

Figure 3: Tri-Party Daily Volume by Collateral Type ($ billions)

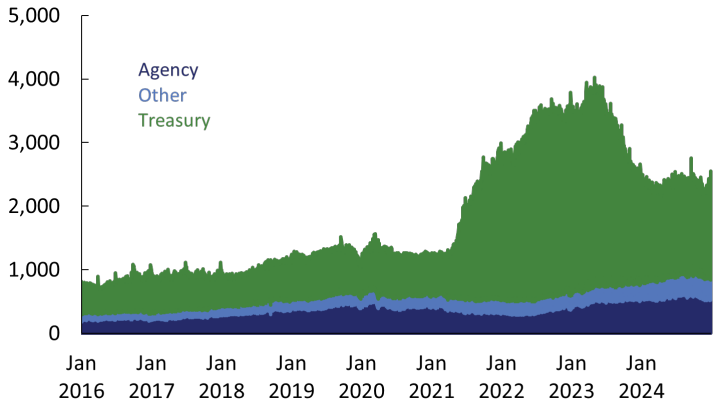

Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

The tri-party segment consists of primary dealers, non-primary dealers, and large banks cash borrowing from asset managers (e.g., money market funds), banks/dealers, government-sponsored entities (GSE), municipalities, and securities lenders/agents. For the purpose of our study, transactions are aggregated up to the holding company level for both cash lenders and borrowers. The result is 139 distinct lenders and 73 distinct borrowers appearing in our sample between January 2016 and December 2024.

Figure 4 shows the distribution of aggregate daily volumes by participant type for both cash lenders and borrowers. Among lenders, asset managers, primarily money market funds (MMFs), dominate the market, followed by bank-dealers and securities lenders

[1]See Erol and Lee (2024) for discussion on technological resiliency of financial system architecture as it relates to the exit of JPMorgan Chase from tri-party and the resulting drop in IT-related investment by the sole provider BNY.

or agents. On the borrowing side, primary dealers are the largest private-sector borrowers. However, in recent years, the Federal Reserve has become the largest overall borrower in the tri-party segment, due to monetary policy operations.

Figure 4: Daily Volume Transacted by Participant Cohort ($ billions)

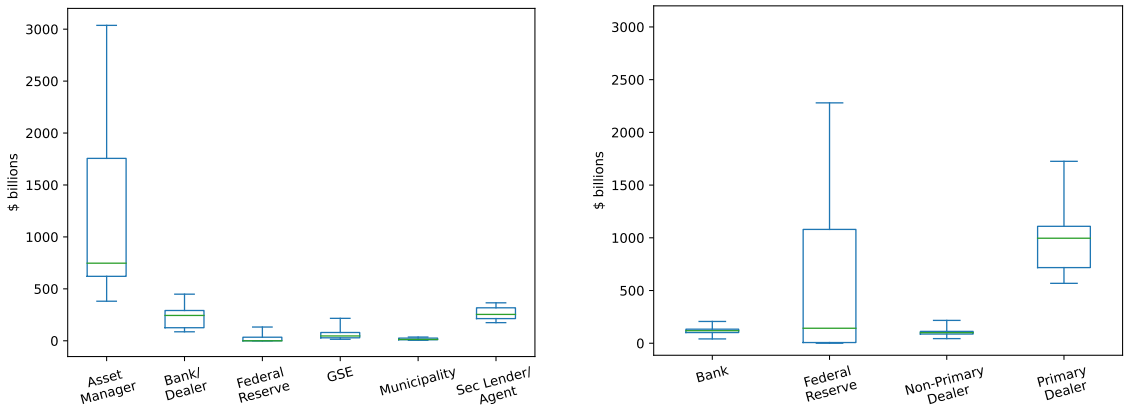

(a) Lenders (b) Borrowers

Note: Data between January 2016 - December 2024. Values represent the distribution of aggregated daily volumes within the respective cohort.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

Table 1 presents the daily trading activity of tri-party borrowers and lenders and the variation in trading relationships we observe over the sample period. The table highlights that the number of borrowers is roughly half that of the number of lenders; consequentially, they perform twice the volume and number of transactions, and manage twice the set of relationships on average.

Despite the potential for up to 10,147 borrower–lender trading pairs, only 1,473 unique pairs are observed over the full sample, representing approximately 14.5% of the theoretical maximum. The volumes and relationship counts presented in Table 1 at the borrower-lender level illustrate how little variation occurs on a typical day. Changes along the intensive margin (the amount of activity within existing relationships) and the extensive margin (the formation or termination of relationships) further emphasize the stability of borrower–lender interactions throughout the period.

On a typical day, about 680 pairs transact, with 97.5% of those same pairs continuing to trade the following day. To quantify the aggregate volume consistency of these relationships, we use cosine similarity, a metric that captures changes in daily gross trading volume between counterparties. On average, we find less than a 2% change in volume from day to day, underscoring the remarkable stability of repo trading relationships.[2]

[2]See Appendix A for a detailed explanation of the cosine similarity metric and its interpretation in this context.

Table 1: Daily Tri-Party Relationship Activity

	Mean	Std Dev	Min	Median	Max
Borrowers #	48.47	5.30	35.00	50.00	58.00
– Volume $B	36.56	96.45	0.02	10.66	620.52
– Trades #	50.54	56.82	1.00	28.56	233.71
– Counterparties #	14.03	12.02	1.00	9.57	42.67
Lenders #	97.10	4.25	83.00	97.00	110.00
– Volume $B	18.38	42.32	0.01	1.67	243.44
– Trades #	25.19	48.04	1.00	6.25	279.96
– Counterparties #	6.99	7.62	1.00	3.96	34.84
Volume $B	1,800.08	959.69	696.51	1,284.39	4027.50
– Δ Net Volume $B	0.77	46.26	-319.38	0.89	353.54
– Δ Gross Volume $B	141.47	56.56	53.16	128.48	508.79
Δ Volume Borrower-Lender Pair $B	0.01	0.81	-8.48	0.00	8.82
– Intensive Increases $B	0.20	0.73	0.00	0.00	8.29
– Intensive Decreases $B	-0.20	0.74	-8.49	-0.00	-0.00
– Extensive Additions $B	0.52	0.81	0.02	0.24	3.08
– Extensive Reductions $B	-0.52	0.79	-3.00	-0.24	-0.02
Borrower-Lender Pairs #	679.93	79.06	478.00	702.00	811.00
– Intensive Increases #	122.86	21.43	63.00	122.00	204.00
– Intensive Decreases #	122.17	20.83	69.00	121.00	205.00
– Extensive Additions #	16.83	7.48	2.00	15.00	78.00
– Extensive Reductions #	16.71	7.35	2.00	15.00	73.00

Note: The first eight rows present borrower and lender statistics, with the *borrower* # and *lender* # rows presenting the distribution of the daily participation counts, followed by average distributional statistic at the daily market participant level for volumes, trades, and counterparties. The subsequent three rows for volume provide distribution statistics for daily aggregate volume changes. The next five rows present Δ *volume borrower-lender pair* change activity, where the average distributional statistic is provided at the daily borrower-lender pair level. The final five rows present the distribution of the count of borrower-lender pairs changes by day, where intensive increases and decreases counts are measured when the change is greater an $1 million.
Source: Federal Reserve Board Tri-Party Repo Collection, Authors' analysis.

Despite transactions being negotiated bilaterally, the structure of trade operations in tri-party repo differs in several ways from bilateral repo markets.[3] Most notably, tri-party repos involve a custodian bank that centrally manages settlement and collateral allocation on behalf of both parties. Settlement occurs both at the initiation of the repo, on a rolling basis throughout the day, and at termination, when the transaction is unwound at 3:30 p.m.. This centralized process enhances operational efficiency but also introduces time-dependent vulnerabilities related to daily funding flows.[4]

[3] In both repo types, the transactions are bankruptcy remote; if a borrower defaults, the lender may liquidate the collateral without entering a bankruptcy proceeding.

[4] Additional facts of note include that tri-party repos involve general collateral, so transactions are not tied to specific securities. Also as the collateral remains with the clearing bank for the duration of the transaction, it cannot be rehypothecated outside the tri-party market, unlike in some bilateral arrangements. While rehypothecation is legally possible within the tri-party system, we do not observe any instances of counterparties switching roles (from borrower to lender or vice versa) to facilitate this practice in our sample.

The majority of funding provided via repos in the U.S. financial system is overnight and, therefore, can be reassessed and negotiated on a daily basis. Figure 5 presents the average intraday settlement cycle for tri-party. The figure on the left highlights the hourly transaction settlement volume distribution. Note that while most trade activity in the repo market occurs during only a few hours in the morning (Clark et al. (2021)), a substantial portion of tri-party is not submitted for settlement at BNY until later in the day. This late settlement is associated with MMF advisors needing to wait until no further withdrawals can be made from the fund before they allocate pre-negotiated trades across their managed funds (McCormick et al. (2021)).

Figure 5: Daily Tri-Party Settlement Cycle

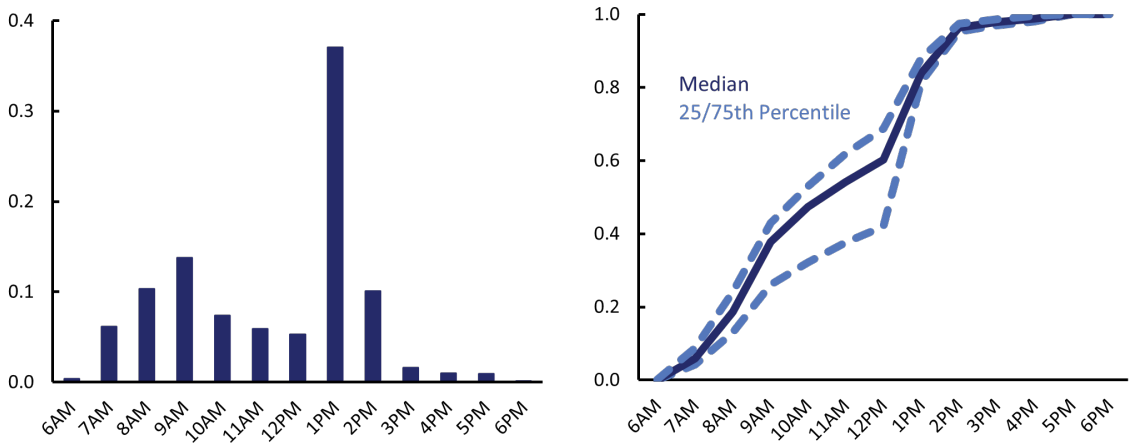

(a) Average Volume (b) CDF

Note: Data between January 2016 to December 2024. The figures show the hourly transaction activity. The figure on the left shows the distribution of average hourly settlement activity at BNY. The figure on the right shows cumulative distributions of volume transacted at hourly intervals, along with the interquartile range, over the sample period.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

The figure on the right provides the median cumulative distribution of daily transaction settlement, along with the interquartile variation, at each hour. The variation highlights that there are some differences in settlement timing, but that generally there are little substantive differences once one excludes the Federal Reserve's ON RRP facility activity.

3 Disruption Risk

During episodes of disruption in the repo market, the Federal Reserve may intervene to preserve market liquidity and sustain credit intermediation. Through open market operations, the Federal Reserve acts as a counterparty, supplying cash to primary dealers against high-quality collateral, such as U.S. Treasury securities, through overnight or

term repo. These interventions help stabilize short-term interest rates and maintain the functioning of funding markets. However, while such measures are essential tools for mitigating broader systemic risk, they entail implementation challenges and are not friction less, as evidenced by the market dislocations observed in September 2019.

To better understand the potential impact of operational disruptions in the tri-party repo market, we begin with a simple stress-testing exercise. In this baseline scenario, we assume that all cash lenders are equally likely to experience an outage, drawing from a uniform distribution across the full set of participants. This stylized approach abstracts away from heterogeneity in institution size, trading volume, or cybersecurity posture, providing a neutral benchmark for assessing systemic vulnerability under uniform stress conditions. These benchmark results allow for an informed comparison with subsequent analysis shown in Section 5, where we incorporate implied probabilities of a disruption for each lender based on their cybersecurity ratings and remeasure the expected impacts from an outage.

To implement this stress test, we conduct a choose-one simulation to examine the distribution of impacts resulting from a single lender outage. We assume that a cyberattack disables one cash lender on a given day for the entire day, rendering it unable to transact for the entire trading session.[5] Iterating through each of the 137 unique cash lenders in our sample, we evaluate the consequences of their hypothetical outage. The results, detailed below, quantify the stress placed on the market in terms of disrupted volumes, affected borrowers, and resulting changes in repo rates.

3.1 Disruption Estimates

Figure 6 displays the distribution of impacts across all cash lenders from a full-day operational outage, measured in terms of daily transaction volumes (left panel) and the number of affected cash borrowers (right panel) over the sample period from November 2019 through November 2024. This timeframe corresponds to the availability of cybersecurity data, which is discussed in later sections. To preserve institutional confidentiality, we restrict the upper bound of the figures to the 90th percentile.

In the left panel, we observe the distribution of daily volume disrupted by a single outage. On this scale, the median impact is relatively small and indistinguishable from zero. The 75th percentile reaches just under $10 billion, while the mean impact lies above

[5]Some firms do include clauses or operational provisions to address rollover risk due to operational outages. These protections are not standard in the core GMRA text but, typically, are handled through annexes, bilateral side letters, or internal operational protocols. These 'force majeure' clauses may allow for temporary suspension of obligations during a severe operational disruption, such as cyberattack, communication failures. or Infrastructure outages, allowing for 'fail forward' for one day unless explicitly terminated. For the purposes of this exercise, we will assume no such clauses are in places for any repo transactions intended to take place on the day of the disruption.

Figure 6: Disruption Distribution on Borrowers

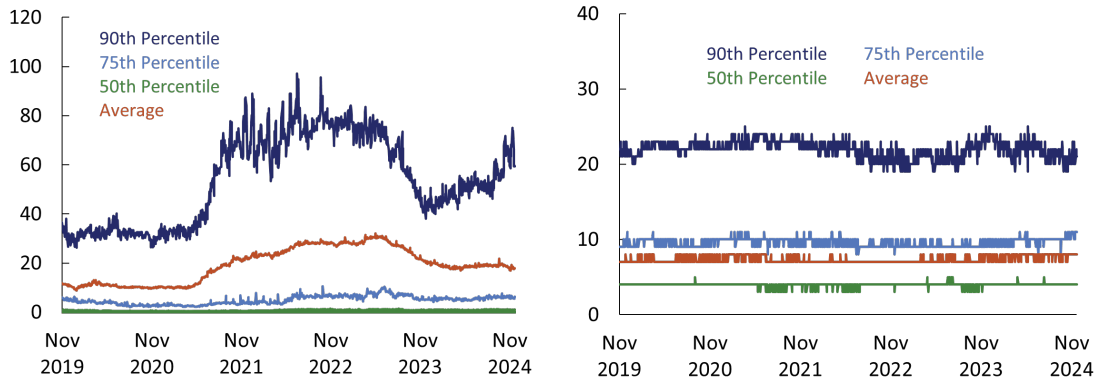

(a) Amounts ($ billions) (b) Number of Borrowers

Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

this percentile, highlighting a long right tail in the distribution. This is further emphasized by the 90th percentile that ranges from approximately $30 billion to $100 billion over the sample period, indicating that a small number of large lenders account for disproportionately high daily volumes relative to their peers.

In the right panel, we observe the distribution of borrowers affected by a single lender outage. The median number of borrowers impacted is around five even though the average is about eight. At the 90th percentile, just over 20 borrowers are affected. While there is still evidence of a long right tail, this distribution is less skewed than that of volume impacts. One reason for this lower skewness is that borrowers typically maintain relationships with multiple cash lenders, which increases the likelihood of being affected by any given outage, but also evenly distributes exposure across the market.

Figure 7 displays lenders' average disrupted collateral volumes by type. As expected, and consistent with earlier discussions on collateral composition, Treasury collateral accounts for the largest share of disrupted volume, exceeding that of both Agency and other collateral types. The average disrupted Treasury volume began rising in early 2021 and remained elevated through early 2024, with typical values just under $40 billion in average daily Treasury collateral disrupted.

Figure 8 presents the distribution of disrupted daily volumes by lender (left panel) and borrower cohorts (right panel). Among lender cohorts, median disrupted volumes are relatively similar across groups; however, asset managers, bank-dealers, and securities lenders/agents exhibit long right tails in their distributions, excluding the Federal Reserve. This skewness is consistent with the patterns observed earlier in Figure 6.

In the right panel, we examine the distribution of disrupted daily volumes for cash borrowers. The Federal Reserve stands out as the largest cash borrower, with its up-

Figure 7: Average Disrupted Collateral Volume by Type ($ billions)

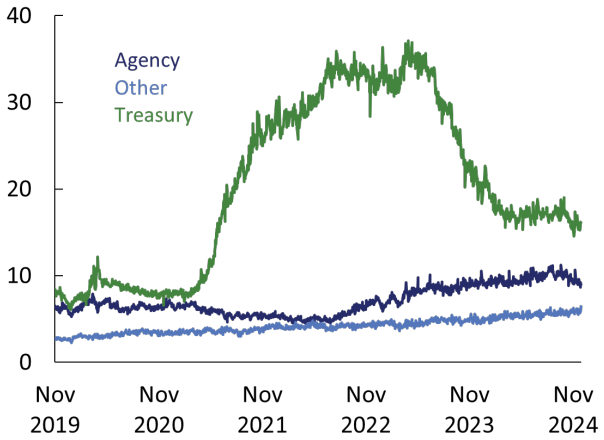

Note: Averages represent the amount of daily volumes transacted, collateralized by type, for lender's transacting non-zero volumes of that respective collateral.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

Figure 8: Daily Disrupted Volume by Market Participant Type

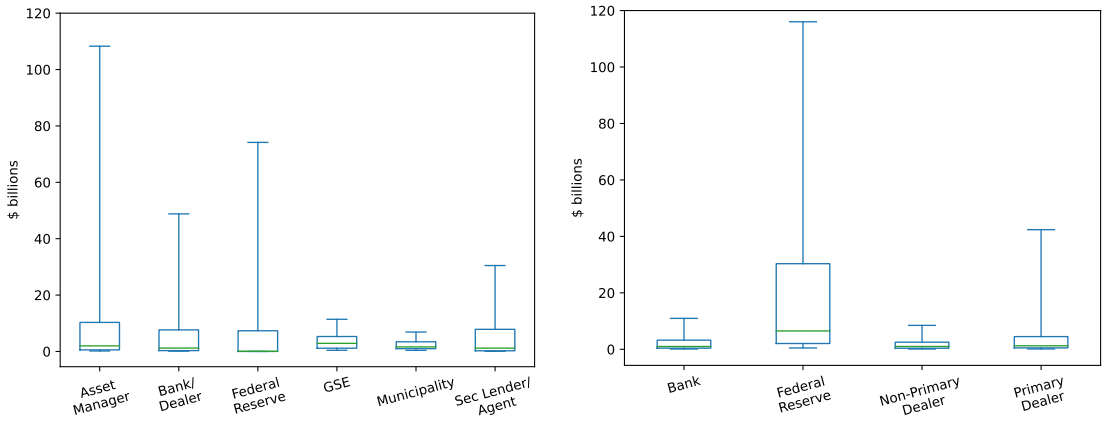

(a) Lenders (b) Borrowers

Note: Values represent the distribution of daily entity-level volumes within the respective cohort. Whiskers are set to the 5th and 95th percentiles to preserve confidentiality.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

per quartiles significantly higher than those of the other groups. While the remaining three borrower types have similar median disrupted volumes, primary dealers experience larger disruptions on worse days compared to banks and non-primary dealers.

3.2 Disruption Impact on Rates

While stable bilateral relationships among repo market participants can facilitate efficient day-to-day liquidity allocation, they may also engender systemic fragility by constraining the adaptability of funding searches. Regulatory interventions, such as those

implemented by the Federal Reserve, can mitigate disruptions by ensuring continued access to liquidity for borrowers. Nevertheless, the necessity of re-routing funding flows in response to counterparty withdrawal can exert upward pressure on rates. For instance, the failure of a key participant to transact may necessitate rapid reallocation across a concentrated network of counterparties, amplifying stress and volatility within the repo market.

3.2.1 Bartik regressions

To examine the effect of losing a lender with whom a borrower has an existing relationship on the interest rates that the borrower must pay, we estimate a two-stage least squares (2SLS) regression using a Bartik-style instrument. The objective is to isolate quasi-exogenous variation in the availability of funding from specific money market funds (MMFs) to specific dealers. This instrument is designed to isolate the effect of increases in cash available to a borrower (which should decrease rates) from increases in the demand for cash from dealers (which should increase rates). We focus on MMFs as the lender group in this analysis because they are relatively homogeneous, and much of the variation in their aggregate lending is driven by plausibly exogenous factors, such as Treasury bill supply or investor redemptions.

The first-stage regression estimates the change in a borrower's repo volume as a function of shifts in the overall funding availability from MMFs that have historically lent to that borrower. Specifically, we weigh each MMF's change in repo supply by its historical propensity to lend to a given borrower:

$$\Delta\text{volume}_{j,t} = \alpha_j + \delta_t + \beta \sum_j w_{i,j,t}\text{volume}_{i,t} + \epsilon_{j,t} \tag{1}$$

where i indexes MMFs, j indexes borrowers, and $w_{i,j,t}$ represents the weight placed on each MMF's activity based on prior lending relationships and measured using outstanding shares from either the previous day or two weeks. The weighted sum captures the idea that if MMFs that typically fund borrower j simultaneously reduce lending, the borrower will face a constrained funding environment, resulting in an aggregate volume shortfall.

In the second stage, we regress the repo rate paid on overnight Treasury transactions by each borrower on the instrumented change in volume:

$$\text{rate}_{j,t} = \alpha_j + \delta_t + \Delta\text{volume}_{j,t} + \epsilon_{j,t} \tag{2}$$

These results illustrate that disruptions in relationships in the repo market can meaningfully affect borrowers' marginal cost of funding. In doing so, they highlight the potential for relationship-driven segmentation to amplify stress from a cyber event.

Table 2 presents the results. The OLS specification in Column 1 shows a positive relationship between changes in volume and rates, which is consistent with borrower demand driving both higher rates and larger quantities. Columns 2 and 3 report 2SLS estimates using alternative weighting schemes for the Bartik instrument.[6] These specifications, based on historical weights from one day and two weeks, respectively, show strong first-stage relevance, with F-statistics exceeding 2,900, and negative second-stage coefficients. These findings indicate that reductions in MMF funding availability lead to higher borrowing costs, underscoring the importance of stable lender relationships in the repo market.

Table 2: Changes in MMF Lending on Dealer Group Rates

Dependent variable	**Weighted average rate**								
	OLS	2SLS							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Δ Principal Amount	0.668**	-3.274**	-6.109***	-4.785***	-6.706***	-4.348**	-7.595***	-2.684*	-2.692*
	(0.312)	(1.406)	(1.329)	(1.490)	(1.352)	(1.951)	(1.858)	(1.460)	(1.407)
Δ Principal Amount^2				-0.292**	-0.585***				
				(0.124)	(0.127)				
× FR ON RRP Volume						1.815	2.421*		
						(1.477)	(1.408)		
× Top Quartile TGCR-ON RRP								-1.625	-10.013***
								(3.259)	(3.127)
First Stage Instrument									
Bartik instrument		0.661**	0.761***	0.681***	0.761***	0.506*	0.585***	0.962*	1.031***
Robust standard error		0.011	0.012	0.012	0.012	0.012	0.013	0.017	0.017
F statistic for IV in first stage		3608.074	4077.367	1823.443	2038.658	2248.838	2472.099	2115.230	2296.964
Anderson-Rubin test		0.02	0.00	0.00	0.00	0.06	0.00	0.06	0.00
Observations	71,902	71,902	71,902	71,902	71,902	71,902	71,902	71,902	71,902

Note: Principal amount is scaled in billions. Column 2 uses the Bartik instrument derived using weights for the prior business day, whereas column 3 uses the average weight over the prior two weeks. Standard errors in parenthesis. Beyond the direct change in principal borrowed, the squared change in principal borrowed, and interactions with Federal Reserve's Overnight Reverse Repo Facility (FR ON RRP) volume levels and periods of more higher rate levels (defined as the spread in Tri-Party General Collateral Rate (TGCR) to ON RRP) are presented in columns 4-9. t-statistics in parenthesis, and * p<0.10, ** p<0.05, *** p<0.01
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

These results illustrate that disruptions in relationships in the repo market can meaningfully affect borrowers' marginal cost of funding. In doing so, they highlight the potential for relationship-driven segmentation to amplify stress from a cyber event.

3.2.2 Rate Impact Distribution

Building on the results above, we incorporate the estimated rate impact into our choose one simulation to evaluate how an individual lender's outage affects repo rates. We utilize the coefficient from column (3) specification, which finds that for each $1 billion disruption (a negative outcome), rates are increased by 6 basis points. Therefore, we can compute a given lender's impact on the overnight rate paid by a non-Federal Reserve

[6]Many Bartik instruments use growth in the "shift" variable rather than the level. We use the level since we are interested in the effect of dollar losses and do not want to exaggerate the effects of small deviations from small baselines.

borrower using Equation 3, where $\text{volume}_{i,j,t}$ represents the total overnight volume lost (in billions) by borrower j due to the outage of lender i on day t.

$$\text{rate impact}_{i,j,t} = -0.06109 \cdot \text{volume}_{i,j,t} \tag{3}$$

We then extend this calculation to estimate each lender's average impact on market rates, using a weighted average across affected borrowers as shown in Equation 4. This allows us to assess not only the borrower-specific effects, but also the broader implications for overall market pricing in the event of a lender-specific disruption.

$$\text{market rate impact}_{i,t} = \sum_{j} \left[(-0.06109) \cdot (\text{volume}_{i,j,t}) \cdot \left(\frac{\text{volume}_{j,t}}{\text{volume}_{t}} \right) \right] \tag{4}$$

Figure 9 displays the distribution of the average market rate impact in the choose one simulation. Similar to the earlier results on disrupted volumes, the distribution exhibits notable right skewness. The median rate impact is close to 0%, with the 75th percentile slightly above that. The mean exceeds the 75th percentile, again indicating a long right tail. The 90th percentile is several times larger than the 75th percentile and has shown a steady upward trend over time, closely approaching 0.4% towards the end of the sample period.

Figure 9: Change in Market Weighted Average Rate (%)

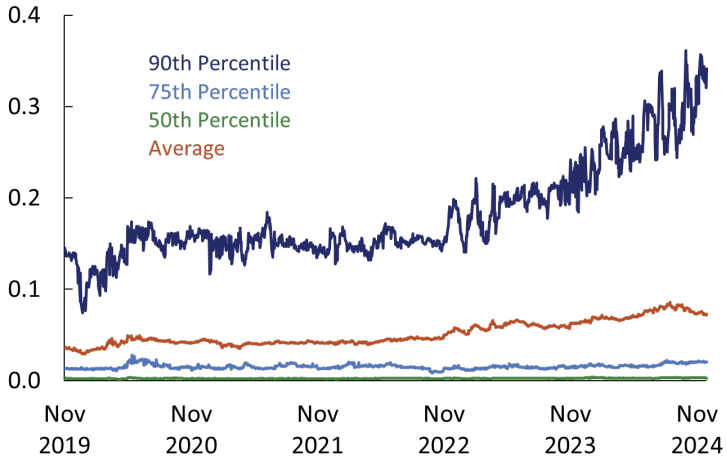

Note: Data between November 2019 - November 2024. Whiskers are set to the 5th and 95th percentiles to preserve confidentiality.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

This rising 90th percentile contrasts with the simulated disrupted volumes in the choose-one exercise, which increased during the Federal Reserve's period of elevated repo borrowing but declined once the Fed began tapering later in the sample. The continued rise in rate impact may reflect growing borrowing concentration, where large volumes are

increasingly channeled through a small number of borrower-lender pairs. Empirically, concentration tends to be higher in inter-affiliate transactions, borrowing and lending between subsidiaries of the same parent institution, which can amplify the rate sensitivity to disruptions in these concentrated relationships.

Figure 10 shows the distribution of daily average market rate impacts from the choose one simulation, disaggregated by lender (left panel) and borrower cohorts (right panel). To preserve confidentiality, the upper percentiles are capped at the 95th percentile. In the left panel, we observe that while the median impact among non-Federal Reserve lenders is effectively 0%, asset managers, bank-dealers, and securities lenders/agents exhibit long right tails. This aligns with prior figures showing that these same cohorts had right-skewed distributions in disrupted volumes. Among them, the bank-dealer group contains the lender with the potential to cause the largest rate impact, though the interquartile range is marginally wider for the securities lenders/agents cohort, possibly indicating more concentrated lending activity within that group.

Figure 10: Change in Rate by Market Participant Type

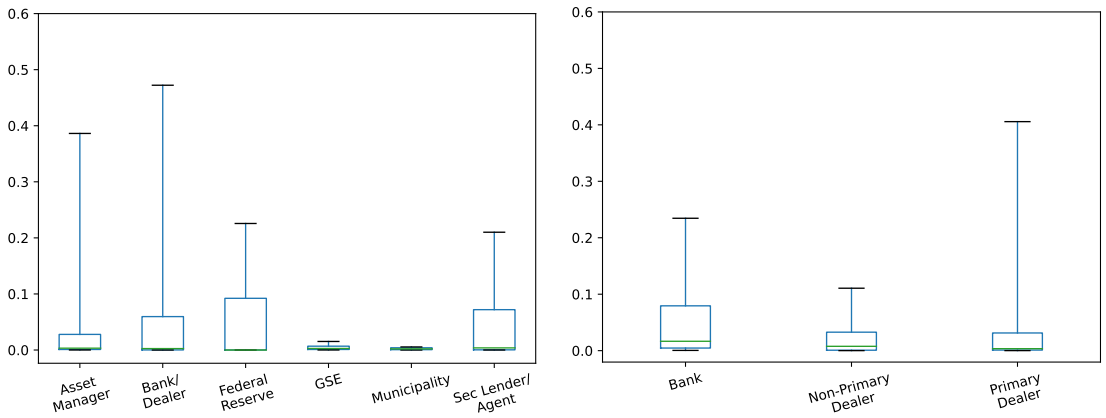

(a) Lenders (b) Borrowers

Note: Values represent the distribution of daily entity-level rate changes within the respective cohort. Whiskers are set to the 5th and 95th percentiles to preserve confidentiality.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

In the right panel, we examine the distribution of daily average rate impacts to borrower cohorts resulting from a single lender outage. Banks experience the highest median impact, as well as a notably wider interquartile range, reflecting greater variability in outcomes. In contrast, primary dealers show the least variability—their interquartile range is narrow, and the median impact is close to 0%. However, they exhibit the highest maximum rate impact among all borrower groups, with at least one lender's outage causing an average rate increase of more than 0.4% for primary dealers. This suggests that while disruptions to primary dealers are less frequent, the consequences can be severe when

they do occur.

4 Cybersecurity

The ability to assess cybersecurity posture can significantly enhance the value of financial system stress testing. As in the exercise in Section 3, traditional stress tests often assume uniform vulnerability across institutions or rely on hypothetical scenarios, which can obscure the distinct operational risks posed by cyber threats. Cybersecurity ratings provide an empirical foundation for differentiating institutions based on their relative exposure to cyber risk, enabling scenario design that reflects heterogeneous vulnerabilities.

Historically, cybersecurity rating data has been used primarily for risk assessment and vendor management, allowing firms to evaluate the security posture of third-party providers and counterparties. Financial institutions, insurers, and asset managers have leveraged these scores to inform due diligence, procurement decisions, and underwriting practices. However, these applications have largely remained operational or governance-focused, with limited integration into broader financial risk models or market-level systemic risk assessments.

4.1 Cybersecurity Ratings and Risk

To assess the cybersecurity risk of financial institutions, we use cybersecurity ratings from BitSight Technologies, a provider of externally observed cybersecurity performance data. BitSight ratings aim to quantify an institution's security posture using publicly observable information on security configurations and security events. Ratings are updated daily, with new observations typically incorporated within 48 hours. The service is used across industries, including by cyber insurers, who collectively underwrite over $5 billion in cyber insurance premiums.

BitSight provides an overall security rating ranging from 250 to 900, where a higher relative security rating represents better cybersecurity posture. In practice, the effective range spans 300 to 820, with an average around 720. In addition to the overall rating, BitSight offers scores across specific "risk vector" categories, such as Patching Cadence, Web Application Headers, and Open Ports. These sub-ratings capture more granular aspects of an institution's cybersecurity posture. For example, the Open Ports vector identifies internet-exposed ports that could serve as potential attack surfaces. Risk vector ratings are typically reported using a letter grade scale from A to F.

To translate these ratings into daily cyber outage probabilities, we use estimates from Rincon and Ordóñez (2023), which empirically links BitSight ratings to reported cybersecurity incidents. The report finds correlations between higher ratings and lower incident

frequencies: firms with poor cybersecurity ratings are 2.6 times more likely to experience an incident than highly rated firms. This relationship holds for both overall BitSight scores and risk vector grades.

We calibrate our model's disruption probabilities using this data. For the overall BitSight rating, we construct a stepwise function that maps each rating bucket to a scaling factor, which adjusts a baseline average daily disruption probability of 1 in 1,000 (or 0.1%). Each firm's disruption probability is calculated by multiplying this baseline by the appropriate rating-based scaling factor.[7]

Similarly, we apply Rincon and Ordóñez (2023) relative attack probabilities for risk vector grades using the same baseline approach. Since risk vector ratings are already bucketed by letter grades, we directly apply the corresponding relative probabilities from the report. For each institution and risk vector, we then multiply this baseline probability by the scaling factor, based on the BitSight risk vector grade's relative attack frequency from the report.

4.2 Cybersecurity Ratings and Tri-Party

As noted in prior sections, individual tri-party transactions are aggregated at the holding company level. We then match institutions in the tri-party dataset to the BitSight database for the period in which cybersecurity data is available, spanning from November 2019 through November 2024. Figure 11 shows the extent of BitSight coverage for tri-party institutions. In terms of volume, coverage is substantial: throughout the sample period, approximately 95% or more of total institutional tri-party volumes correspond to institutions with available BitSight ratings.

The next set of figures combines BitSight cybersecurity ratings with tri-party repo data to illustrate how cyber risk is distributed across different types of market participants. Figure 12 presents the distribution of daily cybersecurity ratings by cohort, with lenders shown on the left and borrowers on the right. Among lenders, there is greater variation in security ratings. The lowest-rated lenders are banks and dealers, while the highest-rated that exclude the Federal Reserve are the government-sponsored enterprises (GSEs). On the borrower side, non-Federal Reserve institutions exhibit less variation with most firm types clustering around similar median ratings. However, the lower tail of the distribution is slightly worse for banks, indicating a subset of institutions with comparatively weaker cybersecurity postures.

Figure 13 displays the average non-zero daily volumes by cybersecurity score for both lenders and borrowers. Most daily volume is concentrated among institutions in the

[7]The baseline disruption probability used here is an assumption used to scale the simulations in the preceding sections, and used for demonstrative purposes in viewing the results. We use this assumption as a tool, rather than presenting a relative ratio based set of results.

Figure 11: BitSight Coverage of Tri-Party Institutions

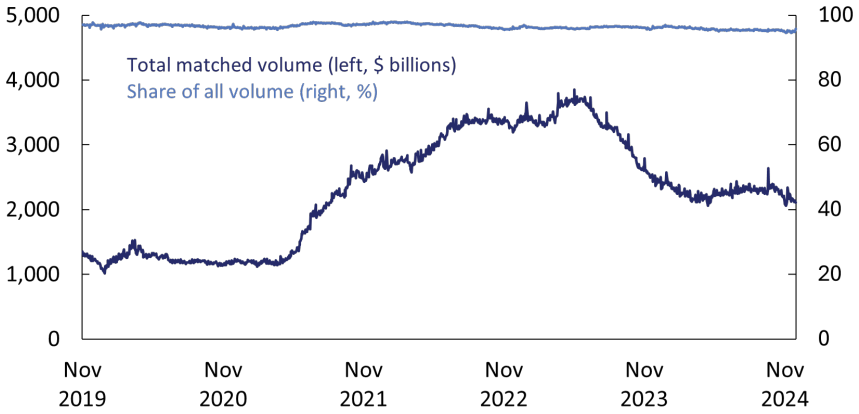

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

Figure 12: Range of Daily Cybersecurity Ratings

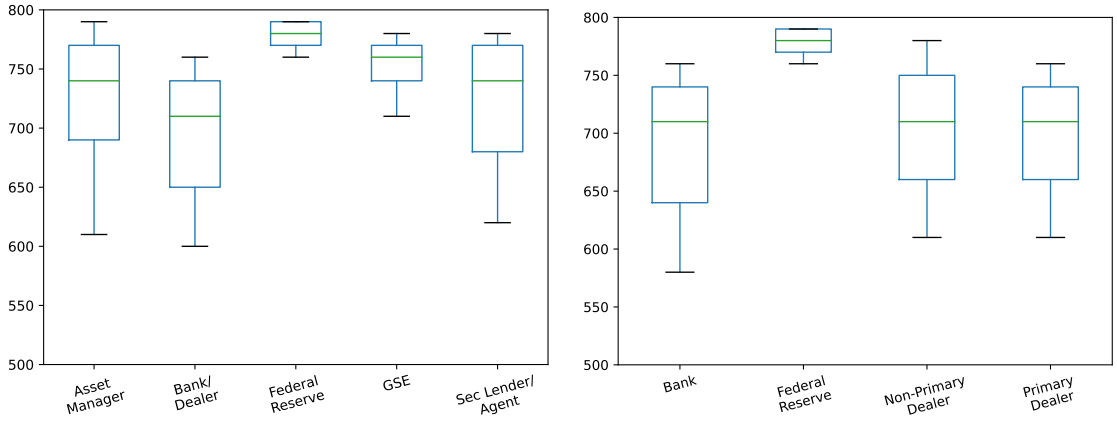

(a) Lender (b) Borrower

Note: Values represent the distribution of daily entity-level cybersecurity ratings within the respective cohort. Whiskers are set at the 10th and 90th percentiles to preserve confidentiality. Municipality type is excluded from the lender figure for confidentiality.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

highest-rated cybersecurity buckets. However, on the lending side, there remains a substantial share of daily volume associated with firms that have cybersecurity scores below 720. This indicates that a meaningful portion of market activity is conducted by institutions that are relatively more vulnerable to cyberattacks.

Figure 14 shows the average non-zero daily volume by risk vector grade for Patching Cadence and Open Ports. In both cases, the majority of volume is associated with institutions receiving the highest grade of A. However, for Patching Cadence, there is still a notable amount of volume tied to lower grades such as C and D, significantly more than for Open Ports. This suggests that vulnerabilities related to patching practices may be a

Figure 13: Average Daily Volume by Cybersecurity Score ($ billions)

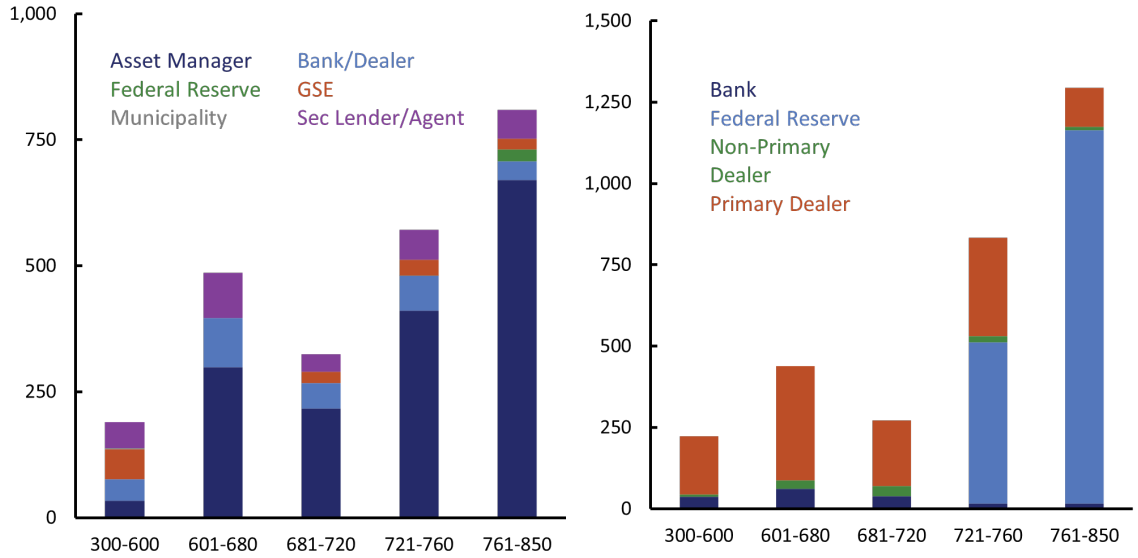

(a) Lender (b) Borrower

Note: A higher score corresponds to better cybersecurity posture.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

more prevalent and potentially problematic issue in the tri-party repo market compared to open port exposures.

Figure 14: Average Daily Volume by Cybersecurity Grade ($ billions)

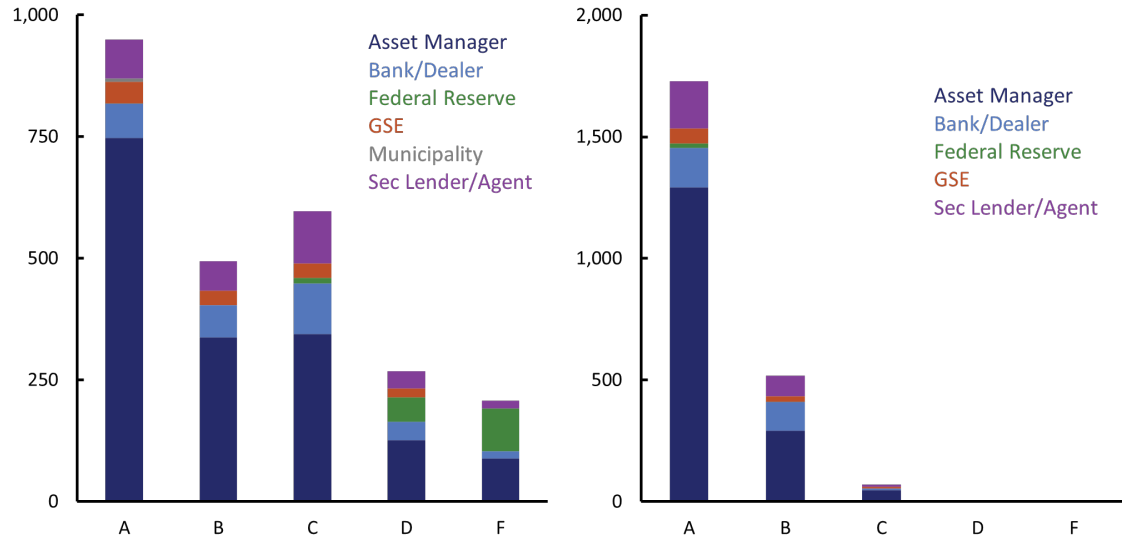

(a) Patching Cadence (b) Open Ports

Note: A (F) corresponds to the highest (lowest) cybersecurity grade.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

Table 3 provides a detailed breakdown of the average non-zero daily collateral trans-

acted across the 16 BitSight risk vectors. For the vast majority of these categories, trading volume is heavily concentrated among institutions with the highest grades, A and B, indicating that most collateral in the repo market is transacted by relatively secure counterparties. However, there are a few notable exceptions where a significant share of collateral is associated with lower-rated institutions. In particular, the SSL Configurations and Mobile Application Security risk vectors show substantial concentrations in the C grade, suggesting persistent vulnerabilities in these specific areas among a non-trivial segment of market participants.

Table 3: Average Daily Volume Transacted by Cybersecurity Vector Grade ($ billions)

Risk Vector	A	B	C	D	F
Botnet Infections	1,946	117	214	33	-
Desktop Software	952	319	283	175	167
DKIM	989	1,021	287	-	-
Insecure Systems	1,908	138	190	91	3
Malware Servers	2,297	-	-	-	-
Mobile Application Security	413	798	692	106	56
Open Ports	1,724	510	57	4	3
Patching Cadence	950	474	571	195	108
Potentially Exploited	1,884	215	144	57	-
Server Software	2,204	61	35	4	-
Spam Propagation	2,248	97	45	24	-
SPF	2081	93	115	22	9
SSL Certificates	1,025	1,012	220	22	20
SSL Configurations	554	625	726	274	120
Unsolicited Communications	2,295	16	-	-	-
Web Application Headers	345	232	306	530	886

Note: A grade of A (F) corresponds to the highest (lowest) cybersecurity grade.
Source: Federal Reserve Board Tri-Party Repo Collection, BitSight, Authors' analysis.

5 Cybersecurity and Disruption Risk

Pairing cybersecurity ratings with market transaction data offers a powerful and timely approach to assessing the threat of operational disruptions in financial markets. While cybersecurity ratings and vulnerability metrics provide insight into which institutions are more likely to experience an attack, they offer limited information about the systemic importance of those institutions or the consequences of their failure. By integrating these security indicators with granular market data like trading volumes, counterparty relationships, and collateral flows regulators and analysts can more precisely identify which vulnerabilities pose the greatest systemic risk. This combined approach enables the as-

sessment of not only the likelihood of a cyber incident, but also its potential economic impact, including disrupted liquidity, elevated funding costs, and contagion across market participants.

5.1 Estimating Security Impact

We extend our analysis to estimate the potential impact of outages based on the predicted likelihood of a cyber disruption affecting a cash lender. As discussed in Section 4, we assume a baseline daily disruption probability of 1 -in 1,000 (or 0.1%), which is then scaled by each institution's relative cyber risk using its BitSight rating. We denote this institution-specific disruption probability as θ. Using this risk-adjusted probability, we compute the daily expected impact of a disruption using Equation 5, where $\nu_{i,t}$ corresponds to either (i) volumes or (ii) rate impact associated with lender i on day t:

$$E[\nu_t] = \sum_i \theta_{i,t} \cdot \nu_{i,t} \tag{5}$$

Figure 15 illustrates the expected impact across three dimensions: total collateral affected, the average overnight borrowing rate, and the number of institutions impacted. In the top-left panel, the trend in disrupted collateral closely mirrors the overall transaction volume, with affected collateral averaging around $3 billion per day and peaking above $5 billion in some periods. The top-right panel shows the expected impact on the market-weighted average overnight rate, which is less closely tied to total transaction volumes. Across the sample, this impact ranges from approximately 0.004% to 0.011%. The bottom panel reports the expected number of impacted lenders and borrowers. We can see that these values are relatively low, as the expected number of disrupted cash lenders is roughly one-tenth of a lender, with just one borrower expected to be impacted.

The key takeaway is that while the likelihood of a cyber disruption on any given day is low, the expected amount of collateral at risk is non-trivial. Moreover, the rate impact is shaped by the concentration of lending and the distribution of disruption probabilities. In scenarios where the baseline probability increases due to heightened threat environments or elevated sector-specific risks, the expected impacts on collateral, rates, and participant exposure would correspondingly rise.

Figure 16 further disaggregates the expected collateral disruption by cash lender type. To preserve confidentiality, we report values for only three major lender cohorts: asset managers, bank-dealers, and securities lenders/agents. These groups account for the vast majority of cash lending activity on any given day. The top-left panel displays the expected total collateral volume disrupted. Asset managers clearly dominate in terms of expected disruption, with recent values around $2 billion and peaks exceeding $4 billion.

Figure 15: Expected Impact of a Cyber Disruption

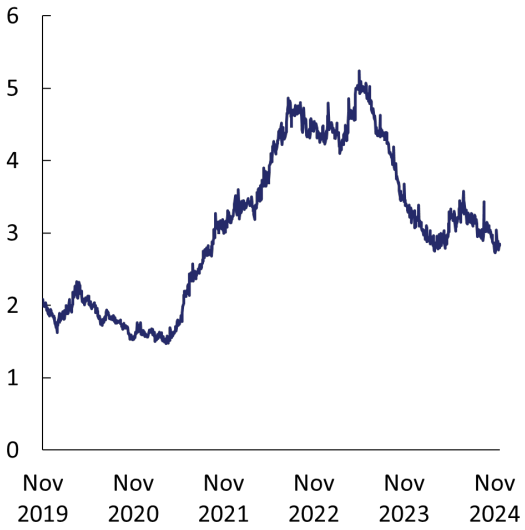

(a) Total Collateral ($ billions)

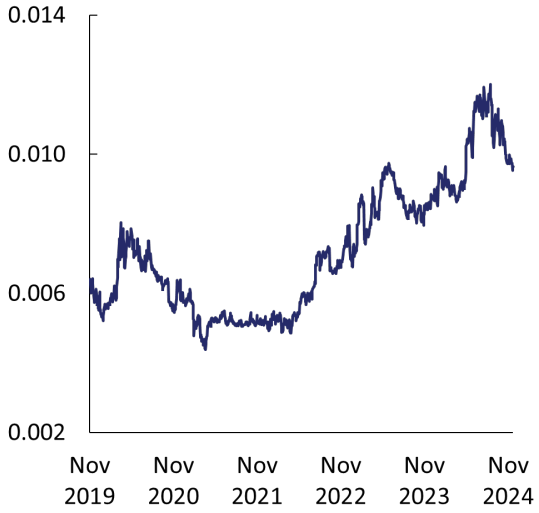

(b) Change in Market Weighted Average Rate (%)

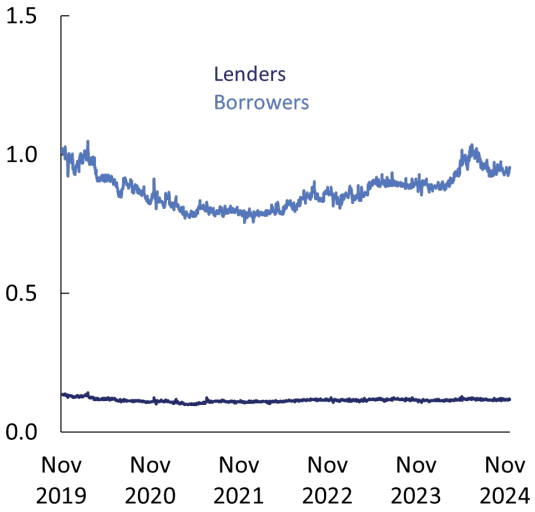

(c) Impacted Institutions (count)

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

In contrast, bank-dealers and securities lenders/agents show much smaller and relatively stable impacts, generally hovering around $0.5 billion or slightly less throughout the sample period.

The top-right panel focuses on disruptions to Treasury collateral, which closely mirrors the total volume trend. As noted in earlier sections, Treasuries represent approximately 70% of all pledged collateral, making this pattern expected. Asset managers again exhibit the largest expected impact, with recent Treasury-specific disruptions slightly over

Figure 16: Expected Cyber Disruption Impact by Collateral and Lender Type ($ billion)

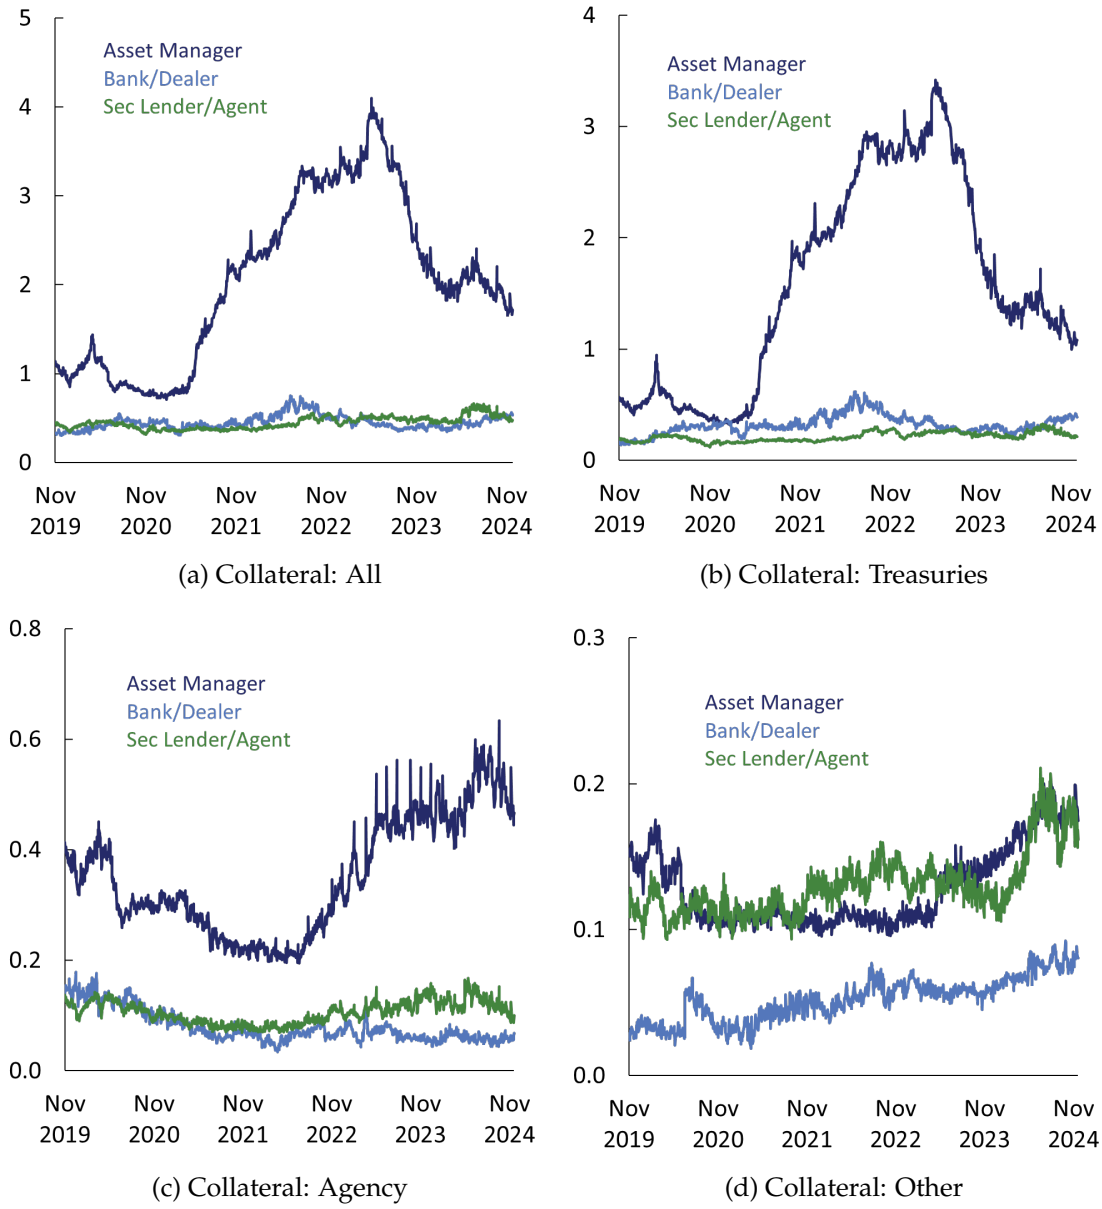

(a) Collateral: All

(b) Collateral: Treasuries

(c) Collateral: Agency

(d) Collateral: Other

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

$1 billion. The bottom panels break down the expected impact for Agency and Other collateral types. For Agency collateral, trends remain consistent, though recent disruptions among asset managers have increased to nearly $600 million. For Other collateral, both asset managers and securities lenders/agents exhibit nearly equal expected impacts, just under $200 million, while the impact associated with bank-dealers is about half that amount.

5.2 Distribution of Impacts

Having established the expected impact of a disruption, we now turn to examining the distribution of possible outcomes. To investigate this, we implement a Monte Carlo simulation over the sample period, simulating outage scenarios and tracking their corresponding impact. While we allow for multiple institutions to experience outages on the same day, we assume that cyber disruptions occur independently across institutions.[8] Furthermore, we assume that for each institution, the occurrence of a disruption on a given day is independent of both historical and future outcomes. Lastly, we assume that any disruption results in a full-day outage, halting all of the affected lender's transactions on that day.

For each trading day t and each institution i active on that day, we generate a binary outage outcome by comparing the institution's derived disruption probability, $\theta_{i,t}$, to a random number Z, drawn from a uniform distribution between 0 and 1. If $Z \leq \theta_{i,t}$, the institution is designated as experiencing an outage. The total disrupted volume for each day is then calculated using Equation 7. This simulation is repeated 10,000 times to capture the full distribution of potential disruption outcomes.

$$\text{outage}_{i,t} = \begin{cases} 1, & \text{if } Z_{i,t} \leq \theta_{i,t} \\ 0, & \text{if } Z_{i,t} > \theta_{i,t} \end{cases} \tag{6}$$

$$\text{disrupted volumes}_t = \sum_i \text{outage}_{i,t} \cdot \text{volumes}_{i,t} \tag{7}$$

Figure 17 presents the results of the simulated disruption outcomes. The left panel shows the distribution of aggregate daily volumes disrupted by institutions experiencing an outage. The distribution is notably right-skewed, with total disrupted volumes remaining relatively modest up to the 95th percentile. Beyond that point, the 99th percentile reveals a sharp increase, with daily disrupted volumes ranging from $50 billion to over $250 billion. The right panel displays the distribution of disrupted cash borrowers. A similar pattern emerges, with a long right tail indicating a substantial rise in the number of affected borrowers beyond the 95th percentile. In the 99th percentile scenario, approximately 25 to 30 borrowers are impacted by simulated outages. This could be thought to being akin to a shared third-party service which forced Ion Trading to shut down a key futures trading service.

[8] An extension of the simulation framework would involve modeling joint outages resulting from shared reliance on service providers, such as cloud platforms, custodians, or network infrastructure firms. To incorporate this dependency, the simulation would introduce correlated outage probabilities among groups of institutions linked to the same vendor. Specifically, one could define clusters of institutions that share a common provider using BitSight and simulate vendor-level outage events, which, when triggered, would cause all dependent institutions to experience simultaneous disruptions.

Figure 17: Cyber Disruption Distribution

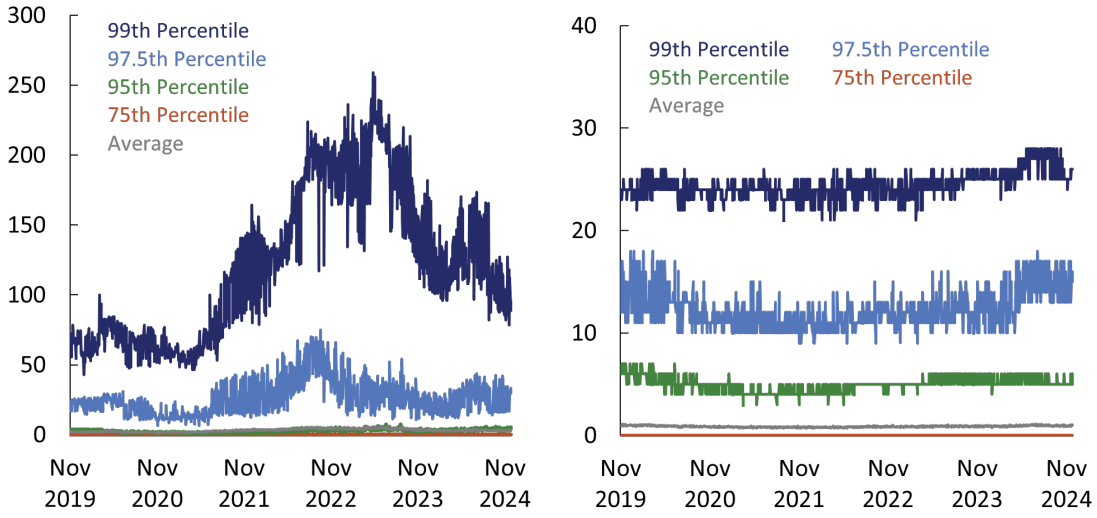

(a) Amounts ($billions) (b) Number of Borrowers

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

Figure 18 illustrates the distribution of rate increases resulting from the simulated outages. The upper percentiles exhibit substantially larger rate impacts. At the 99th percentile, the rate increase ranges from 0.25% to 0.90%. Notably, the spread between the 97.5th and 99th percentiles is relatively narrow. This may suggest that beginning at the 97.5th percentile, scenarios increasingly involve outages at highly concentrated cash lenders, with additional disruption beyond that driven by large but less systemically important lenders.

Overall, these results highlight the pronounced right-skewness in the distribution of disruption outcomes. While the expected values of impacted volumes and rates are already non-trivial, the right tail of the distribution is an order of magnitude larger. Note that these outcomes are based on the assumption of independent disruptions across institutions. If this assumption were relaxed to allow for correlated outages, whether due to shared service providers or time-linked vulnerabilities, both the frequency and severity of extreme outcomes would rise significantly, shifting the entire distribution further to the right and deepening tail risk.

5.3 Counterfactual Impact Assessment

To illustrate the importance of an institution's cybersecurity rating, we conduct a counterfactual exercise that evaluates the relative impact of improving an institution's BitSight cybersecurity score, representing a strengthening of its cyber risk posture. We begin by

Figure 18: Change in Market Weighted Average Rate (%)

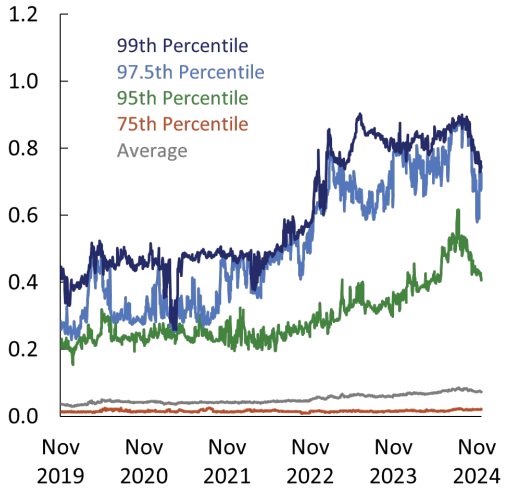

Note: Data between November 2019 - November 2024.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

examining the expected amount of disrupted collateral under current cybersecurity conditions and then quantify how this impact changes following a hypothetical score upgrade.

We divide the sample of lending institutions into quintile cohorts, based on each institution's median cybersecurity rating over the sample period. The first quintile includes institutions with median cybersecurity ratings in the bottom 20%, while the fifth quintile includes those in the top 20%. The expected value of disrupted collateral for each quintile is computed and shown in the left panel of Figure 19. As expected, quintile one exhibits the highest expected disruption values across the entire series. This is intuitive, as institutions in this cohort by construction have the highest disruption probabilities relative to their peers.

Interestingly, quintile four registers the second-highest disruption values for most of the period, despite having lower disruption probabilities than the bottom three quintiles. This can be attributed to the substantially higher transaction volumes associated with institutions in this group. A similar trend is observed for quintile five: although these institutions have the lowest disruption probabilities, they consistently rank third in expected collateral impact, again due to the scale of their lending activity. These results highlight that both cyber risk exposure and market activity levels jointly determine the potential systemic impact of a disruption.

We next aim to isolate the specific effect of improvement in a cybersecurity vector grade. The right panel of Figure 19 displays the change in expected amount of collateral disruption within each quintile cohort after all institutions in the cohort receive a positive one-step cybersecurity rating upgrade. As discussed in prior sections, an institution's

Figure 19: Counterfactual based on Cybersecurity Score ($ billions)

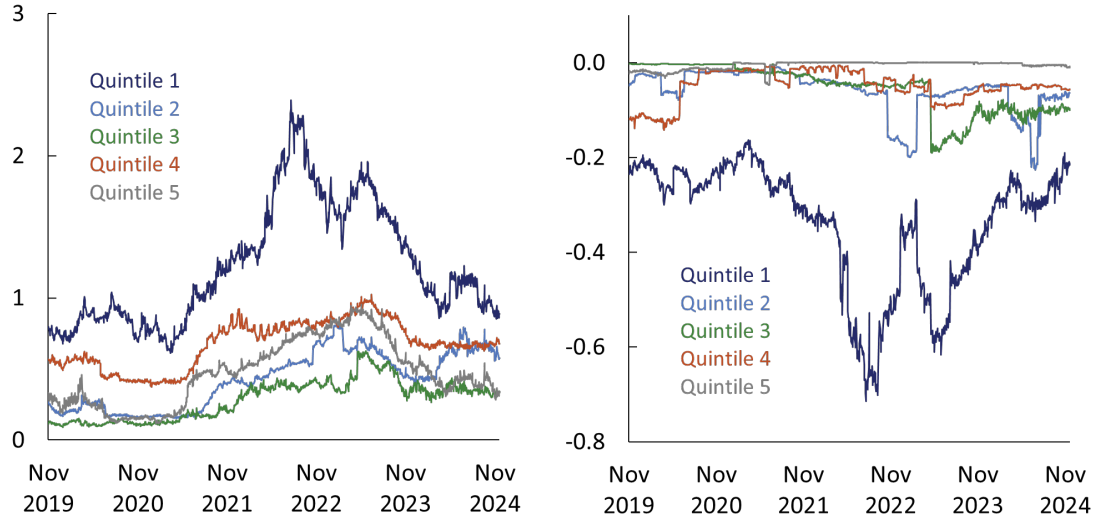

(a) Expected Amounts

(b) Change in Disrupted Volume

Note: Quintile 1 (5) corresponds to the lenders within the lowest (highest) quintile of cybersecurity ratings.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

relative probability of cyber disruption is modeled using a step-wise function based on its BitSight rating. Therefore, a one-step upgrade results in a lower disruption probability, unless the institution is already in the highest rating category, in which case its score remains unchanged.

The results show that quintile 1, which includes institutions with the lowest median cybersecurity scores, experiences the largest reduction in expected collateral disruption of nearly one-third its original value. This outcome is intuitive: The contribution of disruption probability to the expected collateral impact is greatest in this group, so improvements in cybersecurity yield the largest marginal benefit. Although the remaining quintiles also experience reductions in expected impact following the upgrade, these effects are smaller in magnitude, consistent with their lower baseline disruption probabilities. Overall, this exercise underscores the outsized benefit of improving cybersecurity among the most vulnerable institutions.

While focusing on institutions with the lowest cybersecurity ratings is intuitive, addressing this group can be challenging, as its composition spans nearly all observed institution classes. However, if one were to prioritize a specific group, Figure 16 suggests that asset managers should be the focus, as they account for the largest share of volumes expected to be impacted by a disruption. To test this, we perform a counterfactual analysis of institution classes, simulating a one-step rating upgrade across each group. Figure 20 displays the resulting decrease in expected disrupted collateral by cohort.

Figure 20: Counterfactual Change in Disruption Volume by Type ($ billions)

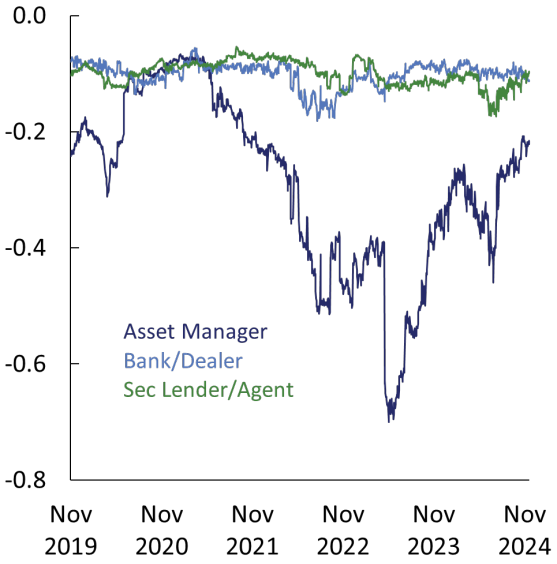

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

The results show that the reduction in expected impact for asset managers is indeed substantial. As a share of their pre-upgrade expected impact, a one-step improvement in cybersecurity rating leads to a nearly 20% reduction in expected disrupted collateral. Interestingly, bank-dealers and securities lenders/agents exhibit similar relative reductions of approximately 20%, though the absolute magnitudes are smaller due to their lower transaction volumes.

This exercise highlights the significance of cybersecurity ratings in determining expected disruption outcomes. Even a modest, one-step rating upgrade produces meaningful absolute and relative improvements in systemic risk exposure across multiple cohorts. The importance of such upgrades is especially pronounced for lenders with large transaction volumes, like asset managers, or for those operating with relatively low cybersecurity ratings.

6 Cyber Resilience and Disruption Risk

Assessing cyber *resilience*, the ability to recover following an attack, is a critical complement to traditional cybersecurity evaluations, which often focus narrowly on preventative controls and threat exposure. While identifying vulnerabilities and breach likelihood is essential, it offers only a partial view of systemic risk. Even in financial markets like the tri-party repo system, well-secured institutions can suffer disruptions. The key differentiator in these cases becomes not whether an entity is breached, but how quickly and effectively it can recover. A high-resilience institution will limit the duration and

scope of disruptions, preserving market confidence and continuity. Without assessments of resilience, regulators and market participants may underestimate the systemic consequences of operational outages, particularly during peak liquidity windows.

The daily settlement cycle of the tri-party repo market is highly time-sensitive, with most transaction activity occurring early in the day. The majority of repo trades are submitted to the clearing bank before 9 a.m., reflecting the market's need to establish funding positions well in advance of the operational day (see Figure 21). This front-loaded pattern ensures that participants have sufficient time for collateral allocation and liquidity management. Notably, this timing excludes Federal Reserve ON-RRP transactions and money market fund (MMF) advisor complexes, which tend to submit trades later, typically closer to their 1 p.m. redemption deadlines.

Figure 21: Average Hourly Volume by Participant Type ($billions)

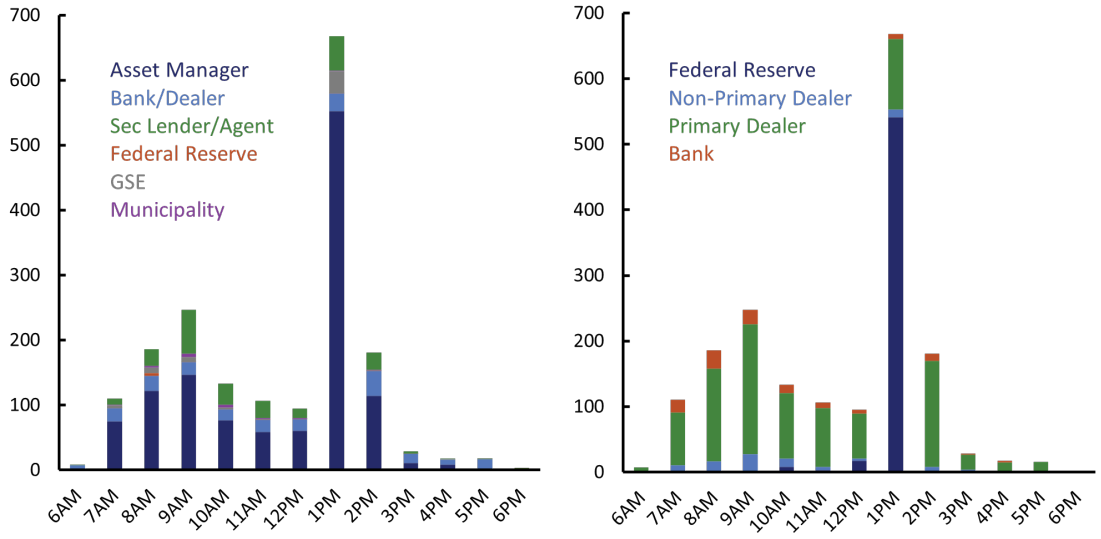

(a) Lenders (b) Borrowers

Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

Following this early activity, most of the daily volume concludes prior to the 3:30 p.m. unwind, when the clearing bank processes the return of cash and collateral. Very little trading occurs after this point, underscoring the importance of early-day execution. This concentrated cycle highlights the market's dependence on timely submission and settlement, making it especially vulnerable to operational or cyber disruptions during peak hours.

After removing the Federal Reserve from the transaction data and focusing on collateral types, the importance of settlement timing in the tri-party repo market becomes even more apparent, as shown in Figure 22. Less liquid collateral tends to be settled earlier in the day, reflecting its more time-sensitive nature in securing funding.

Figure 22: Average Collateral and Cybersecurity Score ($billions, Cybersecurity Score)

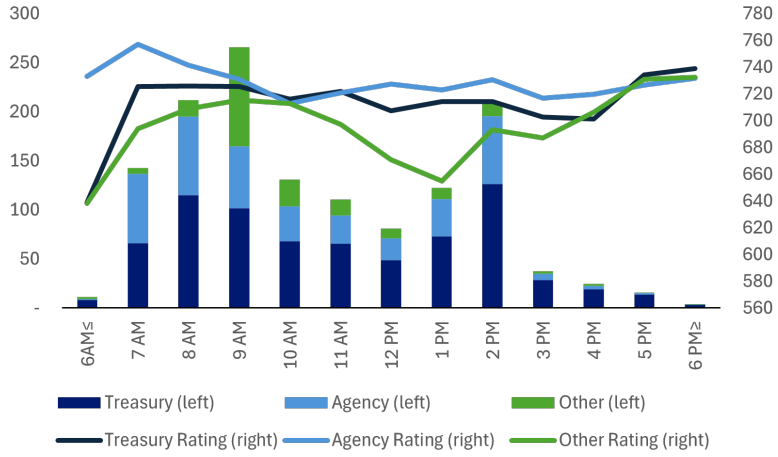

Note: Data between November 2019 - November 2024. The Federal Reserve transaction activity are removed.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

Also plotted are the weighted average cybersecurity ratings of cash lenders by collateral segment. While cybersecurity ratings are generally comparable across groups, we observe a notable dip around noon among lenders accepting other collateral. This suggests that institutions active later in the settlement cycle and willing to accept lower-quality or more heterogeneous collateral may also have weaker cybersecurity postures. Combined with the lower scores observed among 6 a.m. participants, these findings point to two distinct windows of elevated cyber vulnerability. Given the critical role of timing and collateral in intraday liquidity flows, the presence of weaker cybersecurity during these periods could significantly amplify the consequences of an operational disruption, particularly among less-regulated or non-bank financial institutions.

6.1 Estimating Resilience Impact

To model the recovery dynamics following a cyber-related disruption in the tri-party repo market, we use a Poisson distribution to estimate the probability of an institution recovering at various times throughout the trading day. This approach captures the stochastic nature of operational recovery, where institutions may regain functionality at any point in time, but the likelihood of recovery varies over the day.

In our framework, we assume that recovery events follow a Poisson process with a specified average recovery rate. To estimate the likelihood that an institution has recovered at any point during the trading day, we model the time to recovery as an exponentially distributed random variable, which corresponds to the waiting time until a firm recovers. The cumulative distribution function of the exponential distribution gives the probability that an institution has recovered by time h, as shown in Equation 8. This

probabilistic structure allows us to simulate partial-day outages, where institutions may re-enter the market at different times, rather than assuming uniform full-day disruptions.

$$P(\text{cyber recovery by } h) = 1 - e^{-\lambda h} \tag{8}$$

By adjusting the Poisson parameter λ, we can reflect different levels of operational resiliency. As the expected time to recover from a cyberattack follows the form $1/\lambda$, a higher λ implies a faster expected recovery, corresponding to institutions with more robust contingency planning and quicker response capabilities. This offers a flexible and interpretable way to introduce recovery uncertainty into disruption simulations, allowing us to analyze how the timing of re-entry affects intraday funding flows in the repo settlement cycle.

To estimate the expected volume still disrupted at time h, we adjust Equation 5 to incorporate the probability that a recovery has not yet occurred, based on the exponential distribution. The result is the expected volume remaining affected by the outage given h:

$$E[\text{disrupted volume by } h] = \sum_{h} \sum_{i} (\theta_{i,t} \cdot \text{volume}_{i,t,h}) \cdot e^{-\lambda h} \tag{9}$$

Applying this simple model, we can explore a range of cyber disruption scenarios by varying both the start time of the disruption and the expected recovery time, $1/\lambda$, for the affected lenders. For example, Figure 23 plots the expected volume of disruption at time h, assuming the disruption begins at 6 AM, across four different average recovery time scenarios. As the figure illustrates, the consequences vary sharply depend on the assumed recovery speed. For instance, the difference in the maximum expected volume disrupted between a one-hour and a four-hour recovery time exceeds a factor of ten.[9] This result underscores how sensitive market functioning can be to institutional resiliency.

While the magnitude of a disruption is a key determinant of its impact, the timing of recovery plays an equally consequential role in shaping financial system outcomes. Delays in restoration can propagate through adjacent markets, exacerbating liquidity shortfalls and amplifying systemic risk. The simulations reported in Figure 23 demonstrate that recovery assumptions materially alter the temporal profile of disruption. Specifically, increasing the average recovery time from one hour to four hours shifts the peak of expected disruption volume by approximately two hours and substantially prolongs the period before stabilization. These findings highlight that resilience to cyber risk must be assessed not solely in terms of breach prevention, but also in terms of post-disruption recovery capacity.

[9] Appendix B presents expected disruption volumes across various disruption start times (6 a.m.–3 p.m.) and recovery time assumptions (1–4 hours).

Figure 23: Expected Disruption Volume at 6 a.m. by Average Recovery Time

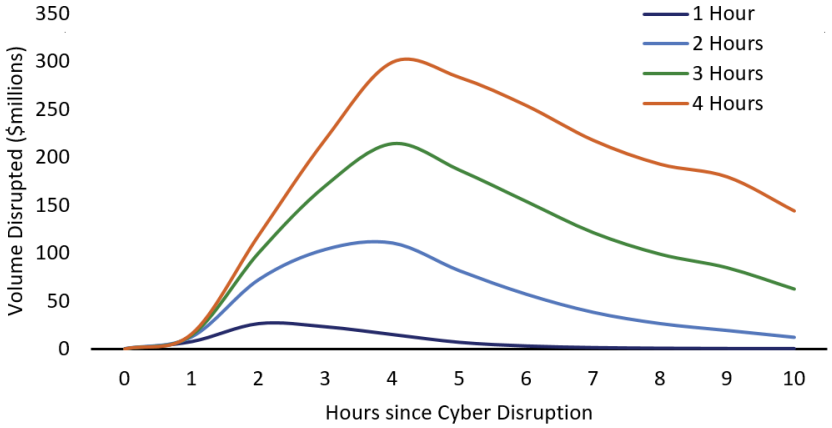

Note: Data between November 2019 - November 2024. The Federal Reserve transaction activity are removed.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

6.2 Counterfactual Impact Assessment

To quantify the joint effects of disruption timing and recovery dynamics on systemic vulnerability, we conduct a counterfactual simulation exercise that varies both the start time of a cyber event and the expected recovery horizon across a set of plausible scenarios. For each scenario, we estimate the maximum expected volume of operational disruption across the trading day. This framework allows us to characterize how the interaction between outage timing and recovery speed shapes overall market impact.

As shown in Figure 24, the relationship between disruption characteristics and market impact is highly nonlinear. Outages that begin early in the trading day (i.e., between 6 a.m. and 9 a.m.) and exhibit slow recovery profiles (e.g., 3–4 hours) generate markedly higher peaks in disrupted volume. By contrast, disruptions that occur later in the day or that are resolved more swiftly tend to produce substantially more limited effects. These results underscore the dual importance of enhancing recovery capabilities and ensuring operational continuity during periods of heightened liquidity demand. Together, they point to the need for systemic safeguards that account not only for the severity of cyber events, but also for their temporal alignment with key market functions.

We extend the counterfactual framework by shifting focus from peak intraday disruption to the volume of disruption remaining at the tri-party unwind. This alternative metric captures the risk that operational failures lead to unresolved funding mismatches persisting into the end-of-day settlement cycle, when cash and collateral positions must be reconciled. Such risks are particularly acute for repo transactions that are expected to rollover the prior day's overnight funding when delays in settlement may compromise counterparties' ability to meet liquidity needs or secure collateral.

Figure 24: Maximum Disruption by Recovery Time

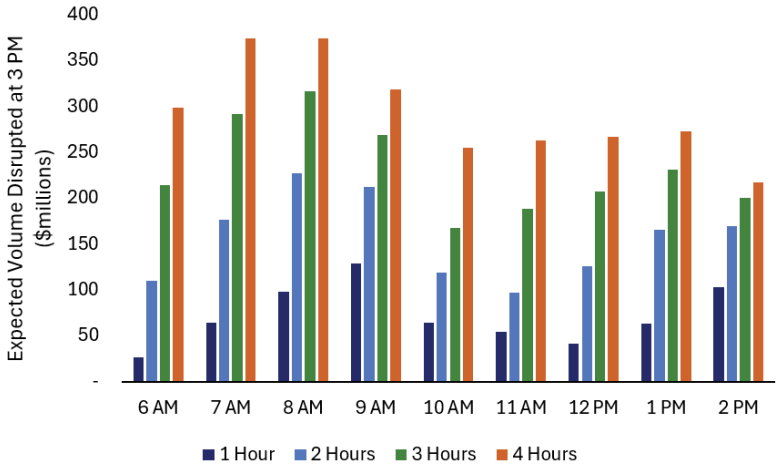

Note: Data between November 2019 - November 2024. The Federal Reserve transaction activity are removed.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

The results, presented in Figure 25, indicate that even moderate delays in recovery can result in a significant portion of daily volume remaining unsettled by the unwind deadline, particularly when disruptions occur later in the trading day. In these cases, relatively rapid recovery on the order of two to three hours may still prove insufficient to avoid settlement failures. For instance, a disruption that begins in the morning and resolves after four hours produces a comparable volume of residual disruption to a late-day outage with the same recovery duration. These findings underscore the time-critical nature of tri-party settlement and the elevated systemic risk posed by cyber incidents that compress or eliminate the window for intraday remediation. In particular, disruptions occurring in the afternoon interact with reduced market flexibility, leaving less scope for reallocation or recovery and thereby increasing the likelihood of end-of-day settlement gridlock.

Figure 25: Disruption Remaining at Tri-Party Unwind by Recovery Time

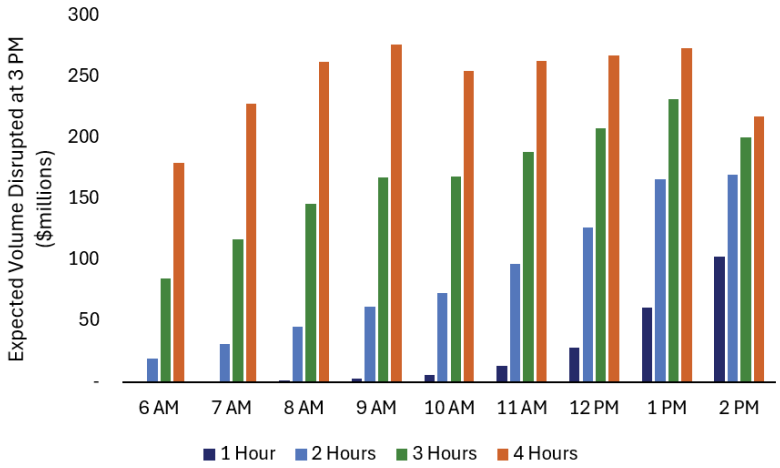

Note: Data between November 2019 - November 2024. The Federal Reserve transaction activity are removed.
Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.

7 Conclusion

This paper highlights the critical intersection of cybersecurity and market functioning in the U.S. tri-party repo market. Although this market is often viewed as operationally stable due to its persistent trading relationships and high-quality collateral, our findings demonstrate that cyber-induced operational outages can lead to severe disruptions. The analysis shows that the sudden absence of major lenders, particularly cash lending asset managers, can create liquidity gaps, raise interest rates, and propagate shocks through the market's interconnected network of borrowers and lenders. These risks are material, and they are amplified by the structural reliance on a small number of institutions and tight daily settlement windows.

Our results underscore that operational and cyber risks are systemic in nature, not just technical concerns. Timing and recovery play a decisive role in the scale of disruption, as outages occurring during peak settlement hours or lasting for extended periods are significantly more damaging. This adds a new layer to our understanding of intraday funding stress, aligning cyber risk with existing evidence on settlement frictions and market rigidity observed during events like the September 2019 repo rate spike.

By simulating real-world outage scenarios and integrating firm-level cybersecurity data, we contribute a novel empirical approach to assessing systemic vulnerabilities. The findings support regulatory focus on both preventive cybersecurity measures and post-attack resilience, particularly for institutions that serve as primary liquidity providers. Counterfactual simulations suggest that strengthening the cybersecurity of cash lending asset managers offers the highest potential to reduce market-wide fragility, given their

outsized role in distributing short-term funding.

Looking ahead, our framework provides a basis for further inquiry into how cyber risk interacts with funding markets, payment systems, and central bank policy. Regulators and market participants alike must treat cybersecurity not as a siloed IT function but as a core component of financial stability planning. In an increasingly digitized and interdependent financial system, the ability to prevent, absorb, and recover from cyber disruptions will be as important as the capital and liquidity buffers that underpin conventional risk management.

References

Afonso, G., Cipriani, M., Copeland, A. M., Kovner, A., La Spada, G., and Martin, A. (2021). The market events of mid-september 2019. *Economic Policy Review*, 27(2).

Anbil, S., Anderson, A., and Senyuz, Z. (2021). Are repo markets fragile? evidence from september 2019. *Federal Research Board Finance and Economics Discussion Series*, 21-028.

Clark, K., Copeland, A., Kahn, R. J., Martin, A., Paddrik, M. E., and Taylor, B. (2021). Intraday timing of general collateral repo markets. *Liberty Street Economics*, (20210714).

Copeland, A., Duffie, D., and Yang, Y. (2025). Reserves were not so ample after all. *The Quarterly Journal of Economics*, 140(1):239–281.

Copeland, A. and Martin, A. (2025). Repo over the financial crisis. *The Journal of Finance*, 80(2):911–936.

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.

Duffie, D. and Younger, J. (2019). Cyber runs. *Hutchins Center Working Paper*, (51).

Eisenbach, T. M., Kovner, A., and Lee, M. J. (2022). Cyber risk and the us financial system: A pre-mortem analysis. *Journal of Financial Economics*, 145(3):802–826.

Eisenbach, T. M., Kovner, A., and Lee, M. J. (2023). When it rains, it pours: Cyber vulnerability and financial conditions. *Federal Reserve Bank of New York Staff Report*, (1022).

Erol, S. and Lee, M. J. (2024). Financial system architecture and technological vulnerability. *Federal Reserve Bank of New York Staff Report*, (1122).

Gorton, G. and Metrick, A. (2012). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451.

Gorton, G. B., Metrick, A., and Ross, C. P. (2020). Who ran on repo? *AEA Papers and Proceedings*, 110:487–492.

Kotidis, A. and Schreft, S. (2025). The propagation of cyberattacks through the financial system: Evidence from an actual event. *The Journal of Finance*. forthcoming.

López-Espinosa, G., Moreno, A., Rubia, A., and Valderrama, L. (2012). Short-term wholesale funding and systemic risk: A global covar approach. *Journal of Banking & Finance*, 36(12):3150–3162.

McCormick, M. J., Paddrik, M. E., and Ramírez, C. A. (2021). The dynamics of the u.s. overnight triparty repo market. FEDS Note 2021-08-02, Board of Governors of the Federal Reserve System.

Paddrik, M. E., Young, H. P., Kahn, R. J., McCormick, M., and Nguyen, V. (2023). Anatomy of the repo rate spikes in September 2019. *Journal of Financial Crises*, 5(4).

Rincon, A. C. and Ordóñez, G. (2023). The impact of cyber security management practices on the likelihood of cyber events and its effect on financial risk. Research report, Moody's Analytics. Prepared in collaboration with BitSight.

A Measuring Relationship Persistence

Suppose the repo market is made up **I** lenders and **J** borrowers. Let $\mathbf{X}_t$ represent the set of repo agreements outstanding in the market on day t, where x_{ij} represent the amount of principal cash held in repo between lender i and borrower j. Given the OTC structure of trade, this creates a noncentralized allocation problem that the set of lenders and borrowers must solve to arrive at $\mathbf{X}_t$.

To assess how dynamic the market is, we look at the allocation of daily funding between $\mathbf{X}_t$ and $\mathbf{X}_{t+1}$. Specifically, we can compute the *cosine similarity* which measures the similarity between two vectors of an inner product space, the descriptive the amount of repo done across all participants, reflected in Equation (A.1). Such that a value closer to zero reflects high similarity whereas a value closer to one reflects low similarity.

$$\mathrm{X}^{t+1} = 1 - \frac{\mathbf{X}^t \cdot \mathbf{X}^{t+1}}{\|\mathbf{X}^t\|\|\mathbf{X}^{t+1}\|} \tag{A.1}$$

For example, let us consider two end-of-day repo settlements in which three lenders and three borrowers repo can \$1 dollar between themselves daily. Figure A.1 represents the settlement networks, where each link represents a repo between the lender and borrower nodes. Between each pair of days, we compute the cosine similarity measure. The difference between t and $t+1$ reflects one repo transaction change, which results in the measure of 0.11.

Figure A.1: Repo Cosine Similarity Measure Example

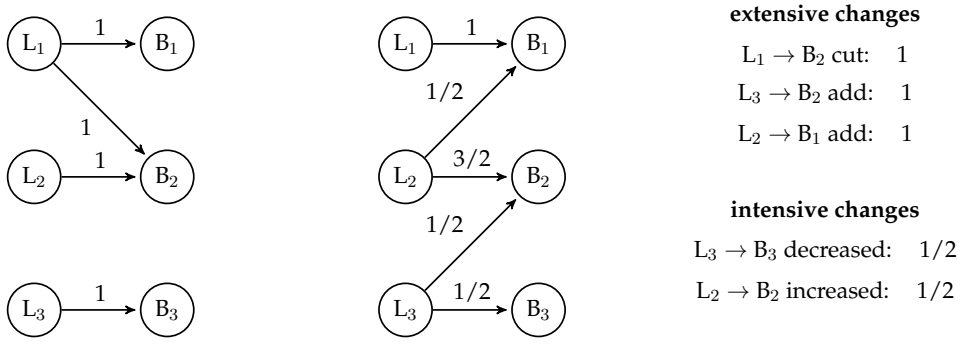

(a) Day t (b) Day $t+1$

Note: Figures (a) and (b) present two example repo settlement networks where borrowers, B_i, and lenders, L_j, are depicted as nodes, and links represent a repo agreement. We compute the variation across the networks by applying the cosine similarity measure between pairs of days.
Source: Authors' creation.

Figure A.2 presents the application of the cosine similarity measure to the tri-party market segment. Figure A.2 presents the daily measure value (dark blue) presented in Equation A.1, along with the seven day moving average (light blue). Figure A.2 highlights that the degree of search and volume volatility has steadily declined over time, starting in 2016 when the average was 2.2% and the highest observation was 16.6%, versus in 2022 where the average was 0.1% and the highest observation was 1.5%. Notably the trend in volume volatility has also declined over this period, as demand and supply shocks have become less frequent.

Figure A.2: Tri-Party Daily Cosine Similarity

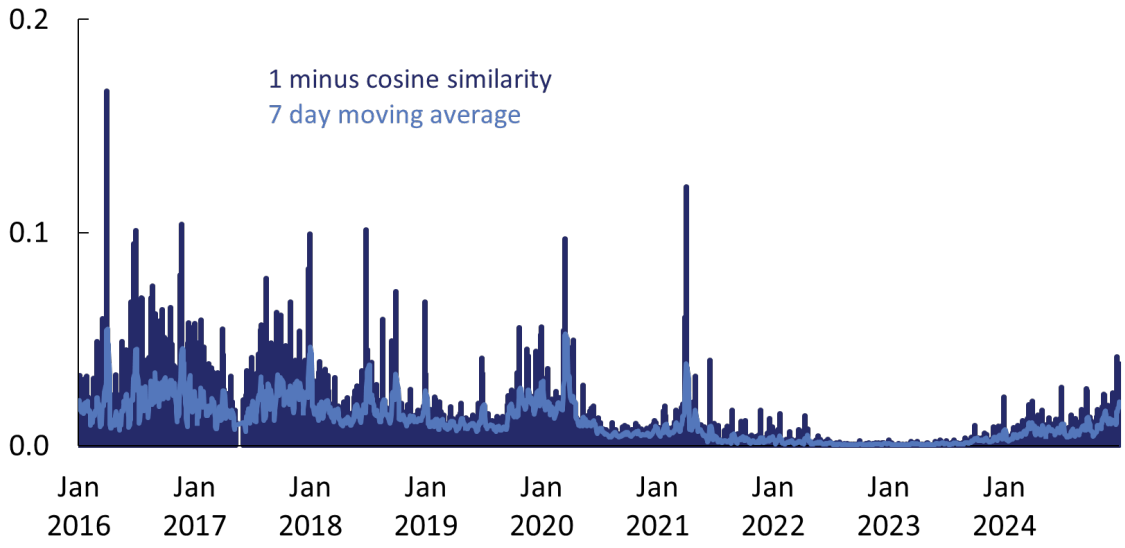

Note: Plot illustrates the daily cosine similarity measure.
Source: Federal Reserve's Tri-Party Repo Collection, Authors' analysis.

B Cyber Resilience to Hourly Disruption

Figure B.1: Expected Disruption since Time of Cyberattack by Resilience

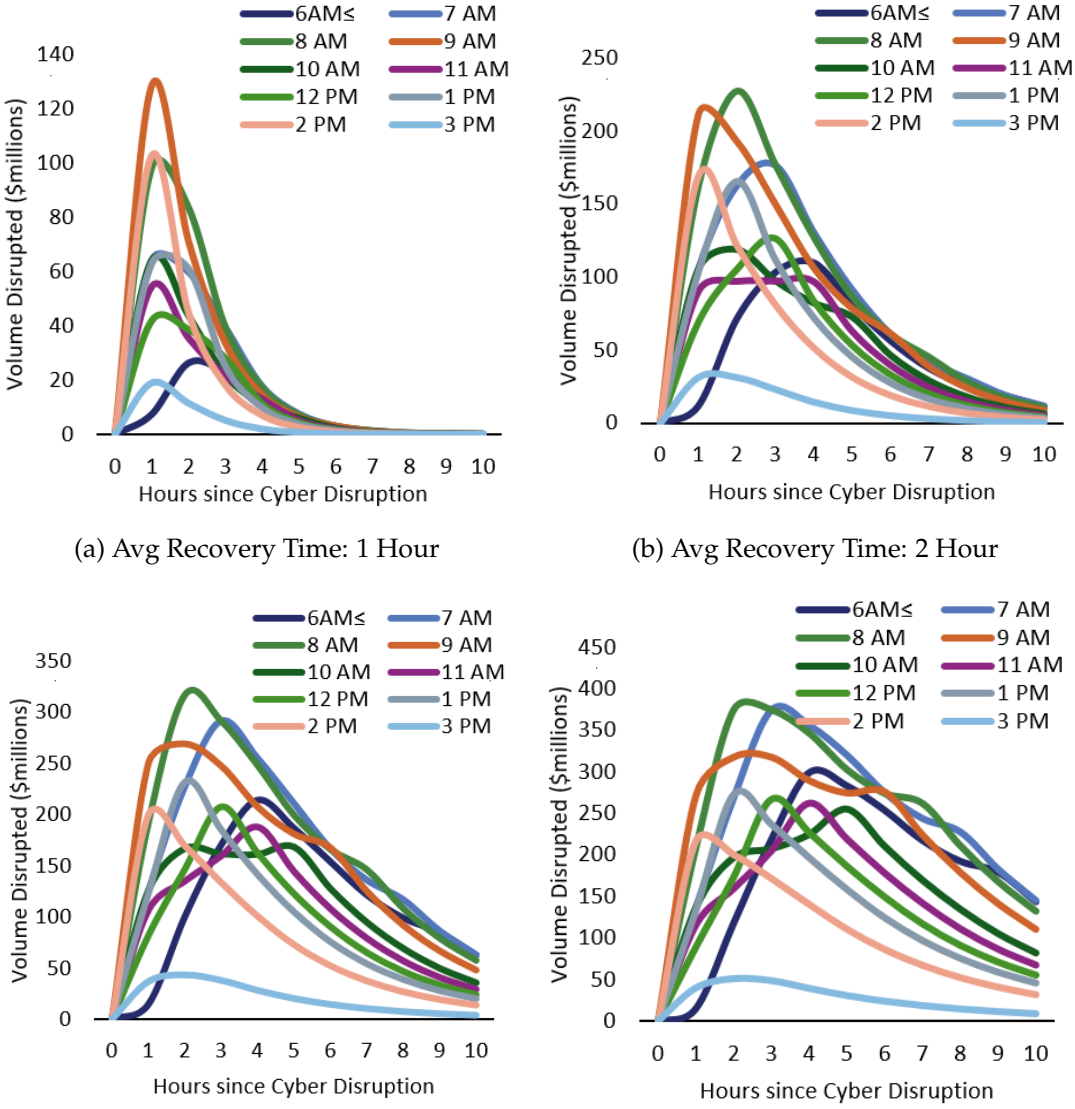

(a) Avg Recovery Time: 1 Hour

(b) Avg Recovery Time: 2 Hour

(c) Avg Recovery Time: 3 Hour

(d) Avg Recovery Time: 4 Hours

Source: Federal Reserve's Tri-Party Repo Collection, BitSight, Authors' analysis.