STAFF REPORTS

NO. 1022
JUNE 2022

REVISED
AUGUST 2023

When It Rains, It Pours: Cyber Vulnerability and Financial Conditions

Thomas M. Eisenbach | Anna Kovner | Michael Junho Lee

FEDERAL RESERVE BANK *of* NEW YORK

When It Rains, It Pours: Cyber Vulnerability and Financial Conditions
Thomas M. Eisenbach, Anna Kovner, and Michael Junho Lee
Federal Reserve Bank of New York Staff Reports, no. 1022
June 2022; revised August 2023
JEL classification: G12, G21, G28

Abstract

We analyze how systemic cyber risk relates to the financial cycle and show that the potential impact of a cyber attack is systematically greater during stressed financial conditions. This is true over the past two decades and particularly at the onset of the COVID-19 pandemic, when changes in payment activity increased vulnerability by approximately 50 percent relative to the rest of 2020 through more concentration and intraday liquidity stress. We evaluate the effectiveness of policy interventions used to stabilize markets at mitigating cyber vulnerability. We argue that cyber and other financial shocks cannot be treated as uncorrelated vulnerabilities and policy solutions for cyber need to be calibrated for adverse financial conditions.

Keywords: cyber, banks, networks, payments, COVID-19

Eisenbach, Kovner, Lee: Federal Reserve Bank of New York (emails: thomas.eisenbach@ny.frb.org; anna.kovner@ny.frb.org; michael.j.lee@ny.frb.org). The authors thank Danny Brando, Victoria Ivashina, and Stacey Schreft, as well as audience participants at the New York Fed, the Columbia University Cyber Risk to Financial Stability Workshop, and the Bank of Canada Economics of Payments Conference for useful comments. They also thank Ryan Craver for outstanding research assistance.

1 Introduction

Much attention has been paid to the way in which cyber risk may be amplified by the financial system (e.g. Duffie and Younger, 2019; Kashyap and Wetherilt, 2019; Aldasoro et al., 2020a). Existing work tends to treat cyber incidents and financial conditions as independent factors, or only considers how cyber incidents may negatively affect financial conditions. In this paper, we explore how systemic cyber risk varies over the financial cycle, and in particular, whether cyber should be viewed as uncorrelated to other financial vulnerabilities. In other words, when it rains and negative shocks lead to financial market dislocations, does it also pour by increasing the systemic vulnerability to a cyber attack?

We find that cyber risk changes with financial conditions, and thus expand our understanding of the systemic implications of cyber risk on two key dimensions. First, at an individual financial institution level, our paper provides an explicit framework to understand and quantify the implications of an attack over the financial cycle, as well as assess options to increase resiliency. Second, our analysis sheds light on externalities across financial institutions in financial stress. Given that individual institutions are unlikely to internalize externalities associated with shoring up defenses, we estimate that the cyber vulnerability of the financial system will likely be amplified when financial conditions are adverse.

If cyber risk and its amplifications are unrelated to financial risks, financial institutions may gravitate towards considering modal outcomes in order to estimate potential costs. This would place cyber risk with other forms of operational risk unrelated to financial conditions, such as severe weather. However, if cyber risk and its systemic spillovers co-move with financial risks, then cyber resilience should be calibrated to a stressed environment. To the extent that the likelihood of government intervention is higher at a time of financial stress, this may present a further wedge between an individual firm's cost of cyber risk and the societal cost of a significant cyber event.

In this paper we examine how the consequences of a cyber attack evolve over the financial cycle, and in particular, during periods of financial stress. We begin with cyber risk, the risk of loss from computer systems and digital technologies (Brando et al., 2022; Curti et al., 2023). This loss could arise from any of a number of attack methods, including denial of service (DoS) attacks, whereby a firm's website access is impaired, ransomware attacks, whereby an attacker prevent access to data and/or systems, and firm email compromise for fraudulent requests, all of which were identified by the Financial Services Information Sharing and Analysis Center (FS-ISAC) as potential threats.[1] Another method by which

[1] See https://www.fsisac.com/hubfs/NavigatingCyber-2023/NavigatingCyber2023-Final.pdf

a cyber attack could affect the financial system is through an attack on components in the financial supply chain such as to key service providers, for example as seen in the 2023 attack on ION Markets which affected some customers' ability to book and process derivatives trades.

We find that the systemic consequences of a successful cyber attack are higher at times when markets are more volatile and when financial intermediary balance sheets are strained. The onset of the Covid-19 pandemic in March 2020 offers a unique opportunity to study how cyber vulnerability is affected by financial markets volatility, because the shock is unrelated to financial institutions' business models. This period is also unique because it marks the first economic downturn since the global financial crisis, and the first episode of extreme market turmoil in an ample reserves regime.[2] The increased market volatility brought upon by the pandemic shock is also plausibly exogenous to the underlying cyber environment. Further, the episode is instructive with respect to the importance of technological access and resiliency, as many financial institutions shifted to working from home, and have potentially increased the points at which cyber vulnerabilities can be exploited.

In order to estimate the impact of a cyber attack we adapt the methodology of Eisenbach, Kovner, and Lee (2022, hereafter EKL), using a scenario based approach to evaluate financial stability risks of cyber attacks amplified through the Fedwire Funds wholesale payment network. To date, there has not been a cyber event with truly systemic consequences on the U.S. financial system. In the absence of actual examples, the wholesale payment network is a natural setting to study cyber vulnerabilities in a financial system. Wholesale payment activity is intimately linked to broad financial system activity, provides a holistic view of liquidity flows between key financial institutions, and offers high-frequency information on aggregate and institution-level liquidity stress. In this approach, we assume that an attack has occurred and then calculate the likely transmission of that attack to other participants in the U.S. financial system, thus quantifying the systemic externalities. Subsequent studies have applied the EKL methodology in other countries (e.g. Kosse and Lu, 2022) and confirmed that key dynamics implied by the methodology played out in an actual cyber attack on a U.S. financial service provider (Kotidis and Schreft, 2022).

March 2020 was marked by sudden severe stress across asset classes and global financial markets (e.g. Federal Reserve Board, 2020; Haddad, Moreira, and Muir, 2021). We show that wholesale payment activity increased along with financial market volatility, became more concentrated, and showed signs of intraday liquidity stress. The financial

[2] Under the ample reserves regime, the aggregate quantity of reserves is intended to be above what is needed for payment purposes, at least during normal times (e.g Logan, 2020).

market stresses in this time period were unusually large, and the speed of the market reaction as well as the global coordination of the financial market deterioration were unprecedented. However, we show that the three indicators of cyber vulnerability observed during this episode are generalizable features of financial market stresses. In particular, we find that the relationship between strains in wholesale payment activity and market uncertainty is a robust feature of the past two decades.

We find that cyber vulnerability, defined as the likely amplification through the interruption of payments flows, was elevated in late February and early March 2020, with the average impact of a cyber attack on one of the largest five banks about 50% greater than the impact of an attack would have been in the rest of 2020. In scenarios where banks hoard liquidity in response to irregular payment flows, forgone payment activity in March 2020 is nearly three times greater than levels outside of March, implying that an attack at a time when financial markets are dislocated would be particularly painful. This additional impact from hoarding emphasizes the way in which externalities could be further amplified by the U.S. payments system's strategic complementaries.[3] Further, we find that delayed recovery from an attack can significantly increase system-level impact: The liquidity shortfall of other banks in the system jumps from $160 billion to roughly $1.5 trillion if an attack lasts for five days instead of one.[4]

Increasing digitization is accompanied by increasing cyber risk which *unconditionally raises* financial stability risks. While there is substantial technical literature on cyber defenses and documenting attacks, this paper abstracts from consideration of the intensity or probability of an attack. That said, it is worth noting that cyber attacks in pursuit of geopolitical goals may coincide with financial volatility, for example as when Russia invaded Ukraine in 2022. To the extent that we document that there is more amplification when financial markets are also stressed, geopolitically motivated attacks timed for maximum damage would also become more likely. In the case of a shock arising from geopolitical conflict, accompanying cyber warfare can be destabilizing.

The March 2020 period showed both increasing potential amplification from a cyber attack and increasing financial market volatility, offering a window for a cyber attack to inflict significant damage. However, we find that official sector interventions to stabilize markets had a mitigating effect on cyber vulnerability, with a decline in hypothetical network impact that starts in the second week of March. This timing corresponds to large

[3]For example, Afonso et al. (2022) find that persistent evidence of intraday strategic payment delays even in normal times.

[4]See Chen et al. (2020), who find strong evidence that payment disruptions could have long-term economic consequences.

liquidity injections by the Federal Reserve. Intuitively, as banks accumulate more liquidity through reserves, they are better able to withstand the unexpected losses in liquidity triggered by a cyber attack on a counterparty. This conclusion, however, may underestimate the impact on markets should a cyber attack impair the trading books and records of a bank and delay or create uncertainty regarding settlement. Specifically, given the significant amount of market transactions cleared and settled within bank holding companies, an attack on a bank holding company with a high concentration of market participants' accounts would directly impact financial market functioning.

These results generalize to other periods of heightened uncertainty. We document that the systemic consequences of a cyber attack rise due to the increased exposure of financial intermediaries and financial market utilities to reliably process transactions in times of increased financial market volatility. In 2020, wholesale payment activity showed a remarkable correlation with the VIX, with a correlation of 0.72 at the daily frequency. This relation holds generally — we document a strong pattern of heightened payment activity in periods of high uncertainty over the past two decades, with a 10-point increase in the VIX corresponding to an increase in payments by about $70 billion per day.

Our paper contributes to a broad literature that studies macroeconomic risks originating from cyber risk. A common theme is the propagation of shocks through interconnected and interdependent systems. These connections can arise through supply chain linkages (Crosignani et al., 2021), or through critical service providers, utilities, and technology infrastructure (Welburn and Strong, 2022). Our study examines this in the context of the financial system, using the complete payments network, which offers a unique andholistic view of full connections in the financial system. Our findings on interactions with financial conditions are likely to generalize to a broader set of industries with interconnections that lend themselves to become increasingly concentrated in times of volatility. Our results have important implications for firms thinking about how to manage both their own cyber risk and the risk of spillovers from other firms.

An important and growing literature studies the financial and economic consequences of cyber risk. While cyber risk is generally recognized by both industry, policy makers, and academics as a significant risk (Brando et al., 2022), there has not been systemic event triggered from a cyber attack as of date. Consequently, ex-post estimates on the cost of cyber risk based on historical cyber incidences have been relatively limited, albeit larger for the financial sector (e.g. Aldasoro et al., 2020a). Researchers have studied the frequency of cyber attacks and how they may be mitigated by bank lending (e.g. Aldasoro et al., 2020b; Crosignani et al., 2021). To overcome this problem, studies have estimated cyber risk by examining scenarios (e.g. Duffie and Younger, 2019, and EKL). A notable exception is

Aldasoro et al. (2020b), which examined how operational costs, including cyber-related costs, changed for financial institutions around the global financial crisis. Our paper addresses a gap in the literature by explicitly studying how systemic cyber vulnerability evolves with the financial cycle. We reaffirm the issue of collective defense against cyber risk illustrated by Anand et al. (2022), and demonstrate the feedback loop between cyber risk and financial conditions that becomes particularly acute during times of high market stress.

The paper proceeds as follows. Section 2 shows the effects of market stress on payment activity in early 2020 and the past two decades more broadly. Section 3 applies cyber scenarios to understand the vulnerabilities during adverse market conditions. Section 4 discusses the mitigating effects of policy responses and Section 5 concludes.

2 Wholesale payment activity and market uncertainty

We document several patterns in wholesale payment activity during adverse financial conditions that relate to the amplification channels of cyber risk. We make use of confidential data on payments sent through Fedwire Funds Service ("Fedwire"), the U.S. wholesale payment system operated by the Federal Reserve which provides detailed information on the accounts and flows between a diverse set of financial institutions.

2.1 Patterns in early 2020

We first illustrate each pattern during the beginning of 2020 and then show with regressions that the patterns hold over a longer sample period using data back through 1997 to 2020.

Level of payment activity. In March 2020, market volatility indices peaked, with the CBOE Volatility Index (VIX) reaching its all-time high of 82.69, above the previous high reached during the financial crisis of 2007–09. Correspondingly, trading volumes were exceptionally high across various markets. Because Fedwire supports the settlement of large-value transactions and trading volumes tend to increase in times of high market uncertainty, we expect a positive relation between market uncertainty and payment system activity. This is what we see in Figure 1, which shows the daily aggregate Fedwire payment value and the VIX in February March and April 2020. Over the full year of 2020, aggregate Fedwire payment value is highly correlated with the VIX, with a correlation of 0.72 at daily frequency. From February to April 2020, during the time in which market stress is most

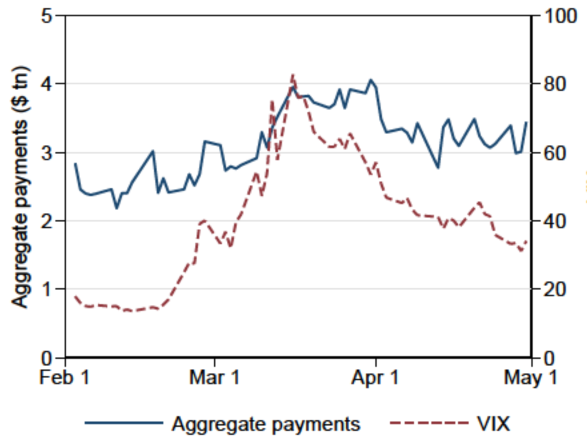

Figure 1: Wholesale payment activity and market uncertainty. The figure shows aggregate payment activity in Fedwire Funds and the CBOE Volatility Index.

acute, aggregate payment activity spikes with the VIX. From the beginning of February to the end of April, the correlation between payment value and the VIX is even greater, at 0.87.

Concentration of payment activity. Historically, activity in the wholesale payment network is highly concentrated, with roughly 50% of payment value accounted for by the top-5 banks (EKL). While this concentration may endogenously arise to facilitate efficient financial transactions, it also increases systemic risk through greater interconnectedness (Erol and Vohra, 2020). In particular, the dependence on core institutions of the network to settle large-value transactions and assist in the flow of liquidity makes the system susceptible to large liquidity dislocations and payment issues if the operations of any key payment bank fail.

Over the course of March 2020, the concentration of payment activity increased. Figure 2 plots the trailing 5-day average share of payment value of the top-5 banks. The top-5 banks' share of daily payment value rises by about 3 percentage points at its peak on March 18 (roughly two times the share's standard deviation in 2020), before falling and stabilizing at levels comparable to the beginning of the year. In sum, not only is there more payment activity, but payment activity becomes more concentrated in times of high market uncertainty.

Risk of coordination failure. A regime with ample reserves should, among other things, satiate liquidity needs associated with payment activity. When liquidity needs are satisfied, banks may send payments asynchronously without concern for their overall liquidity

Figure 2: Concentration of payment activity. The figure shows the trailing 5-day average share of payment value of the top-5 banks.

positions, as the exact timing of payments expected to be received is unlikely to adversely affect overall liquidity positions.

As liquidity becomes scarce, banks more closely manage intraday liquidity by strategically timing payments to better match inflows and outflows, effectively avoiding liquidity shortages (McAndrews and Rajan, 2000). Under intraday liquidity stress, the propensity for banks to delay or halt payment activity in response to irregular payment flows increases (Bech and Garratt, 2003). This form of liquidity hoarding can, in turn, trigger other institutions to hoard liquidity. Individual banks' attempts to preserve their liquidity thus represents a form of coordination failure.

There were several indications that the wholesale payment system became more susceptible to coordination failure in the Covid-19 market turmoil. To start, at an institution level, liquidity needs associated with payment activity grew significantly. One way to see this is to examine a bank's payment activity relative to its reserves. For the top-5 banks, the ratio of daily payments over reserves increased by almost 50%, from about 4 to almost 6 in March 2020 before dropping to about 2 in April. In contrast, the ratio did not change notably for non-top-5 banks (Figure 3, left panel).

Banks typically manage their reserve balance to maintain a desirable level of liquidity. In theory, payments volume could increase but could do so predictably and banks may still be able to manage their reserve balances. However, if greater and more volatile payments made reserve management more difficult, individual banks would have faced greater liquidity risks to processing payments. The increase in payment-related liquidity needs appear to have affected banks' abilities to manage a stable reserve balance. The right panel of Figure 3 plots the trailing 30-day standard deviation in reserve balances for the

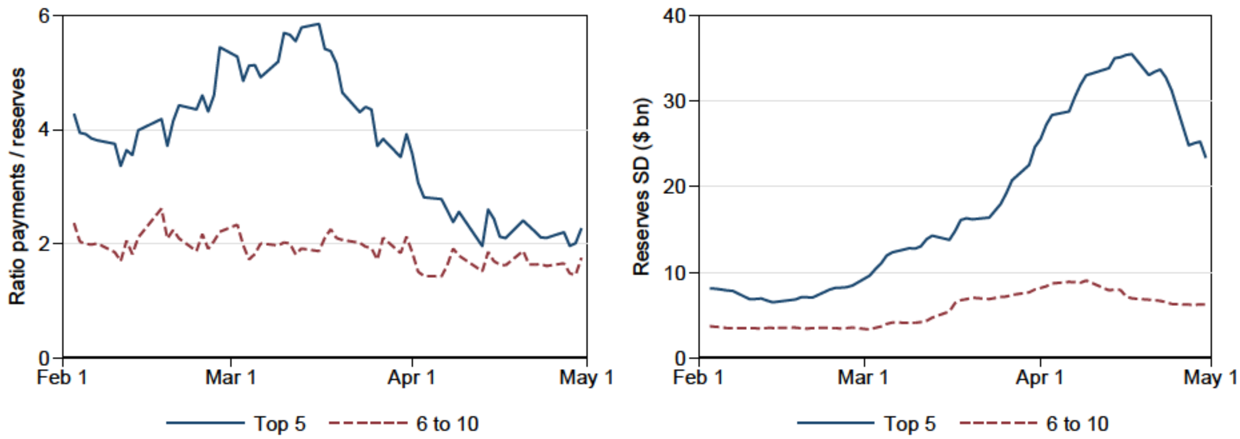

Figure 3: Payment-related liquidity needs and reserve balance volatility. The left panel shows the ratio of average payment value to average reserves for the top-5 banks and for the banks ranked 6 to 10. The right panel shows the trailing 30-day standard deviation of banks' reserve balance, averaged across the top-5banks and for the banks ranked 6 to 10.

most active banks in Fedwire. From mid-February to mid-April, the standard deviation in daily reserves for the top-5 banks increases by roughly a factor of 4.

An informative signal of intraday liquidity stress is delays in settlement times. When banks think their reserves may be scarce, they tend to delay payments until later in the day in order to secure sufficient reserve balances at the end of the day. This coordinated payment behavior was more prevalent pre-2008, when reserves were scarce. These strategic considerations had noticeably diminished post-crisis, due to the dramatic increase in aggregate reserves (Bech, Martin, and McAndrews, 2012) but have reappeared in more recent years (Afonso et al., 2022).

In March 2020, settlement times of late payments noticeably stretched to the end-of-day. Figure 4 shows the timing of the 90th percentile of intraday payment value. Delays begin in late February, around the mark where the VIX increases, and continue to rise until mid-March. In sum, the wholesale payment system is more susceptible to coordination failure in times of high market uncertainty.

2.2 Historical Patterns

We now examine whether the patterns present in payment activity at the beginning of 2020 are a consistent feature over a longer sample. Table 1, shows regression of the key variables characterizing payment activity on the VIX over a more than 20-year sample period from 1997 to 2020. The regressions are at monthly frequency with year fixed effects and control for the level of aggregate reserves, which strongly correlates with payment activity after

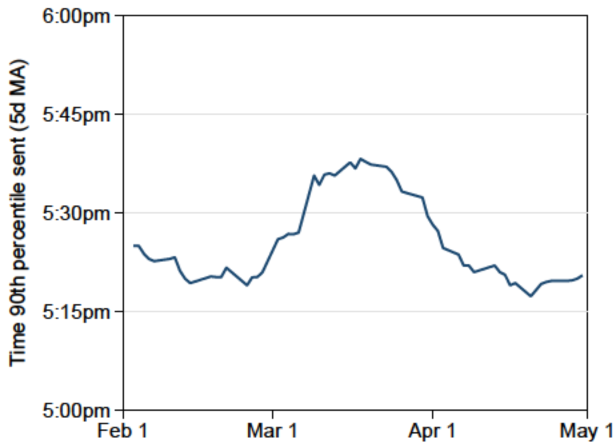

Figure 4: End-of-day payment delay. The figure shows the 5-day moving average of the time of day by which 90% of intraday payment value has been sent.

2008 (Eisenbach, Frye, and Hall, 2019).

Column (1) shows that the relation between the aggregate level of payments and the VIX is highly significant and the coefficient of about 0.007 implies that for a 10-point increase in the VIX, payments increase by about $70 billion per day. Column (2) shows that concentration of payment activity measured as the top-5 banks' share of payment value is significantly decreasing in aggregate reserves and increasing in the VIX, consistent with times of reserve scarcity and market uncertainty leading to greater concentration of payment activity among the largest banks. Column (3) shows that the risk of coordination failure measured by the time of day at which 90% of payment value has been sent is significantly decreasing in aggregate reserves and increasing in the VIX, consistent with times of reserve scarcity and market uncertainty leading banks to strategically delay payments.

Heightened payment activity, concentration of payments, and intraday liquidity stress during market turmoil have implications for the amplification of a cyber attack through the financial system. In the context of the wholesale payment system, a cyber attack could be timed at periods where payment activity is heightened. The system-level impact of an attack varies over time, and increases when payment activity is greater. An attacker could view periods of high financial market uncertainty as a proxy for greater impact to the system as a whole and use it to time attacks. The greater concentration in payment activity could mean that a pointed attack on a key institution could have a greater impact on the network as well. The shock could be further exacerbated by other banks' reactions, especially with greater intraday liquidity stress. This stress historically occurs when payments volumes are high and volatile and banks may be incentivized to conserve reserves or think strategically about payments timing. For at least the last two decades, when it

Table 1: Wholesale payment activity and market uncertainty. The table shows linear regressions of aggregate Fedwire payment value, the top-5 banks' share of payment value and the time by which 90% of payment value is sent on the VIX and aggregate reserves. All variables are averaged from daily to monthly frequency and all regressions include year fixed effects. Heteroskedasticity-consistent standard errors are reported in parentheses. Sample is April 1997 to December 2020.

	(1) Agg. payments	(2) Top-5 share	(3) Time 90th pctl.
VIX	0.0069***	0.0529***	0.1787***
	(0.0026)	(0.0150)	(0.0425)
Agg. reserves	0.1617*	-1.5744***	-5.5857***
	(0.0867)	(0.4134)	(1.4650)
Year FEs	Yes	Yes	Yes
Observations	285	285	285
Adj. within-R^2	0.157	0.068	0.164

rains, it pours, as financial market uncertainty is associated with cyber vulnerabilities in the payment system.

3 Cyber vulnerability during adverse market conditions

We adopt the cyber scenario approach used in EKL in a form modified for the analysis of adverse market conditions in 2020 in order to gain more insights into the impact of financial uncertaingy on vulnerability to a cyber attack. All our scenarios assume that a cyber attack compromises the normal functioning of a targeted institution's systems such that it is unable to send any payments from the beginning of a Fedwire day. The scenarios we employ vary in terms of (i) the target institution, (ii) the reaction function of other banks, and (iii) the time it takes to recover.

The scenarios are meant to represent attacks that affect the availability or integrity of the attacked institution's systems or data. For example, a cyber attack may impair the availability of relevant data or communication systems of an institution, or may compromise the integrity of its operations either by manipulating or corrupting the data. In both instances, the attacked institution's ability or willingness to perform large-value payments would be stifled, as assumed by our scenarios.

While the target institution is assumed to be unable to send any payments, the institution is still able to receive payments due to the institutional features of Fedwire, where payments are actualized when Fedwire receives a payment request from the sender. The balance in an institution's reserve account thus increases with incoming payments, even if the institution is unable to observe or interact with the Fedwire network due to a cyber incident. For the duration of the impairment, an attacked institution soaks up liquidity without releasing payments, restricting the flow of liquidity — a problem which was observed following the attacks on September 11, 2001 (Lacker, 2004). Our scenarios therefore calculate counterfactual end-of-day reserve balances for all institutions, i.e. what their liquidity position would have been if they had not received any payments from the attacked institution and had responded as specified by their reaction function.

Evaluating the severity of a cyber event requires pinning down conditions under which the liquidity positions of other banks, which are not directly attacked, should be considered as materially impaired. Our scenarios for an attack on day t consider a bank i impaired if its counterfactual end-of-day reserve balance r_t^i is more than two standard deviations below the bank's historical average reserve balance. Specifically, we calculate a time-varying threshold b_t^i given by

$$b_t^i = \bar{r}_t^i - \frac{2\sigma_{\text{ref}}^i}{\bar{r}_{\text{ref}}^i}\,\bar{r}_t^i,$$

where $\bar{r}_t^i$ is the trailing 30-day average reserve balance of bank i on day t, and σ_{ref}^i and $\bar{r}_{\text{ref}}^i$ are the trailing 30-day standard deviation and average of bank i's reserve balance at a reference date. Here, $\bar{r}_t^i$ is meant to capture a time-varying target reserve balance of bank i, and the ratio $2\sigma_{\text{ref}}^i/\bar{r}_{\text{ref}}^i$ represents a liquidity buffer ratio of two standard deviations during normal times. We therefore set the reference date to February 19, the point at which the VIX begins to rise but results are robust to choosing a different reference date. Because the effective liquidity buffer scales with the trailing average balance, the threshold adjusts to the changing quantity of reserves observed in the latter part of the sample. Results are not sensitive to the details of the impairment threshold.[5]

3.1 Baseline scenario

The baseline scenario examines the impact of an attack on a single top-5 bank, assuming no reaction by other banks and focusing only on the first day. Figure 5 summarizes

[5]The threshold differs in two ways from that used in EKL, which is given by $b_t^i = \bar{r}_t^i - 2\sigma_t^i$, where σ_t^i represents the standard deviation in the past 30-day reserve balance of bank i at time t. First, a reference date is used to pin down the buffer for all dates. This is because the variation in end-of-day balance during a period of severe market turmoil is unlikely to reflect a bank's tolerance toward reserve volatility but rather reflects intraday liquidity stress. Second, the buffer is taken to be proportional to the trailing average balance at time t, to account for changes in the quantity of reserves held by banks.

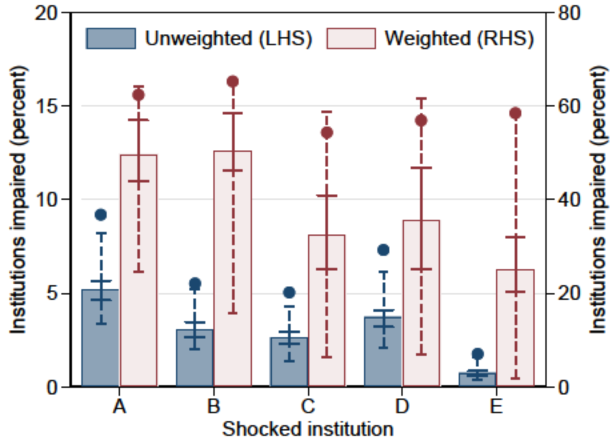

Figure 5: Impact of an attack on a top-5 bank. The figure shows the distribution across days of the unweighted share of institutions impaired by a shock to each of the top-5 institutions and of the share weighted by payments (excluding the attacked institution). Bars represent the average impact; solid whiskers represent the p25/p75 range; dashed whiskers the p1/p99 range; dots show the maximum impact days in March.

the impact of attacks on each of the five institutions, showing the average across all days in 2020 (bars) as well as percentiles of the distribution across days (whiskers) and the maximal impact during the month of March 2020 (dots). The unweighted share is the raw fraction of impaired institutions, and the weighted share is the fraction of impaired institutions weighted by their payments in 2020 (not including the attached institution itself).[6] The weighted shares are considerably larger than the unweighted shares, reflecting the concentration of payment activity, and the variation across days is at least as large as the variation across attacked institution. Our focus is on the maximal impact during the month of March 2020, as represented by the dots in the figure. As anticipated, the worst impact in the time of market turmoil is close to or even above the 99th percentile across all of 2020, both in terms of raw share and weighted share.

Figure 6 shows the time series of the impact of an attack, averaged across the top-5 banks. The weighted impact starts out fairly high and, after a dip in mid-February it steadily climbs until March 3, when the first cut of the Federal Funds target rate of 50 basis points was announced. At the peak, around 60% of institutions by payment share would have been impaired in an attack on a top-5 bank. Compared to around an average of 40% across all of 2020, vulnerability is therefore about 50% higher. While the raw share of impaired banks increased throughout March and peaks on March 30, the weighted share

[6]Similar results are obtained when institutions are weighted by assets. Weighting by payment share enables the analysis to take into account the impact on branches of FBOs, which account for a significant fraction of both payments and reserves.

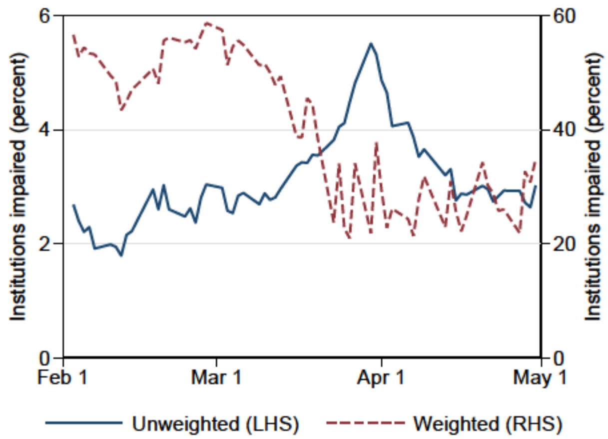

Figure 6: Average impact of an attack on a top-5 bank. The figure shows the daily time series of the impact of an attack, averaged across the top-5 banks.

decreases dramatically through the end of March, which we explore in greater detail when discussing the mitigating effects of policy interventions in Section 4.

3.2 Cascade scenario and coordination failure

The analysis in the previous section assumes that all institutions, other than the directly attacked institution continue to make payments as usual. This non-reaction of banks assumes that banks may not be sensitive to intraday liquidity conditions, and hence, may not react to abnormal conditions experienced throughout the course of a day. However, in Section 2 we show evidence of intraday liquidity stress in late February and March which suggests that other banks are likely to react to large deficits in intraday liquidity positions by delaying or halting payment activity. Furthermore, relative to a typical operational outage, a suspected cyber attack may be accompanied by greater uncertainty and a lack of common knowledge regarding the source, magnitude, and recovery. This uncertainty could be exacerbated by attacked banks, who may be reluctant to disclose to counterparties and clients the exact state of their internal systems or data.

To evaluate the potential impact when banks react, this section considers a cascade scenario where banks react to a lack of incoming payments by suspending their own payments and hoarding liquidity. Specifically, banks' reaction function is assumed to be based on a threshold: Whenever the counterfactual net payment deficit (intraday) passes some liquidity-hoarding threshold, the bank is assumed to halt payments for the remainder of the Fedwire day. The liquidity-hoarding threshold is set to equal the maximum realized

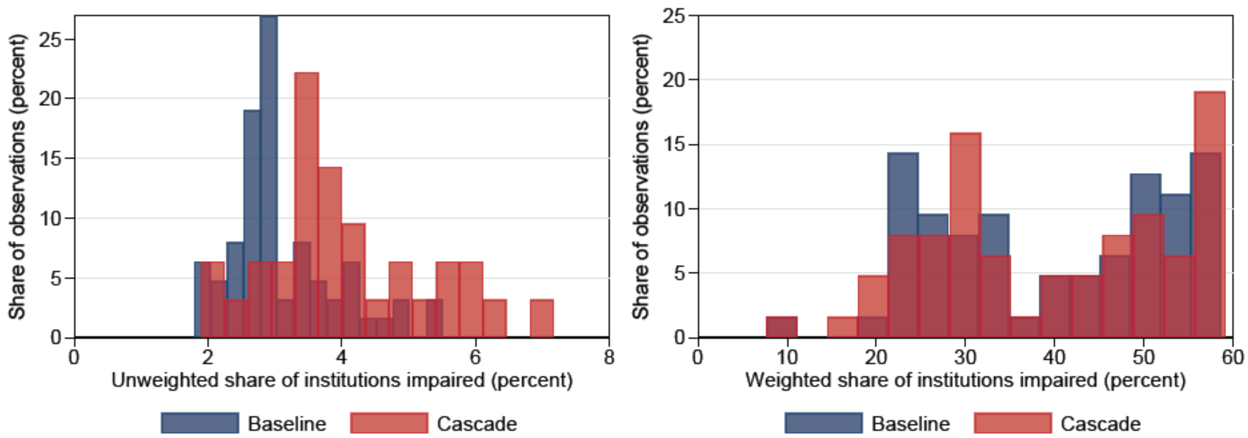

Figure 7: Comparison of simple and cascade scenario for February to April 2020. The figure shows the distribution across days of the impact for the baseline scenario and the cascade scenario for February to April 2020, averaged across the top-5 institutions. The left panel shows the unweighted share of impaired institutions. The right panel shows the share of impaired institutions weighted by payments (excluding the attacked institution).

net payment deficit of the institution in the entire year of 2020.[7]

A priori, it is not clear if the impact should be greater in the cascade scenario as it involves some banks actively preserving their liquidity position which helps them but hurts others. While the bank targeted in the attack is exogenously specified as in the baseline scenario, the set of banks that are triggered to hoard liquidity in the cascade is endogenous to the payment network structure. Figure 7 compares the impact in the simple and cascade scenario involving top-5 banks, for the period of February to March. Impact is slightly greater under the cascade scenario with a more notable shift in the distribution of the raw share of institutions. This is consistent with the core periphery structure of the payments network and suggests that the large core banks' hoarding at the expense of more periphery banks becoming impaired.

An additional risk in the cascade scenario pertains to the payments that are not made as a result of hoarding behavior. In contrast to the simple scenario, systemic risk sprouts not only from the compromised liquidity positions of banks, but also from system-level disruptions in payment activity that supports financial markets and the broader economy. Figure 8 shows the average daily forgone payment value in the cascade scenario, both the payments foregone by the attacked institution and the payments foregone by other institutions due to the cascade. Both increase considerably in March but the foregone payments from the cascade notably more so, reaching close to the level of the foregone payments of

[7]Given the high value of payments that occurred, particularly in March 2020, the liquidity-hoarding threshold is a conservative cutoff.

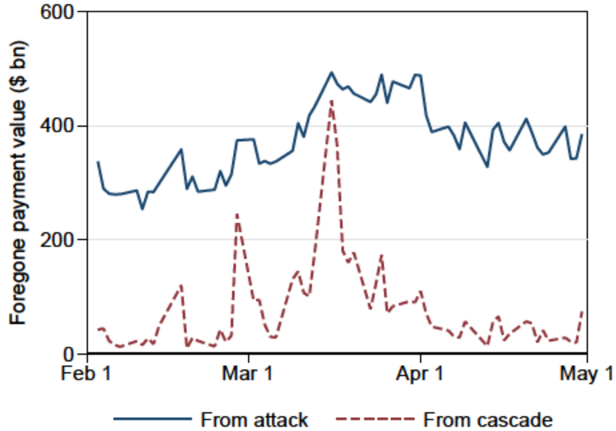

Figure 8: Forgone payment value in cascade scenario. The figure shows the payments foregone by the attacked institution and the payments foregone by other institutions due to the cascade., averaged across the top-5 institutions.

the attacked institution.

3.3 The consequences of delayed recovery

So far, the analysis focused on the single-day impact of a cyber attack. An extended cyber incident, due to delays in operational recovery, however, may result in deeper consequences for the system. Recovery is one of five functions of the NIST cybersecurity framework, along with Identify, Protect, Detect, and Respond, and pertains to timely recover to normal operations to reduce the impact from a cybersecurity incident. The extent to which an attacked institution is able to restore its capabilities and services that were impaired depends on the strengths of its recovery function.

We examine the impact of delayed recovery from a cyber attack by extending the baseline one-day scenario to consider a multi-day scenario. The multi-day scenario maintains the assumption that banks other than the attacked institution continue to make payments as usual. This allows us to analyze the severity of liquidity dislocations grows with each day, and the emergency liquidity support that may be required.

Starting with the day of the attack, we cumulate the set of impaired institutions across additional days, such that the n-day share impaired is equal to the share of institutions that become impaired as a result of payment deficits arising from day 1 to day n. The results are summarized in the left panel of Figure 9. The impact of an attack averaged across the top-5 banks substantially increases with a delayed recovery in late February to mid March, with the net weighted share increasing from about 45% on day one, to over 70% by day five. In other words, by the fifth day, the vast majority of the network (by payment share)

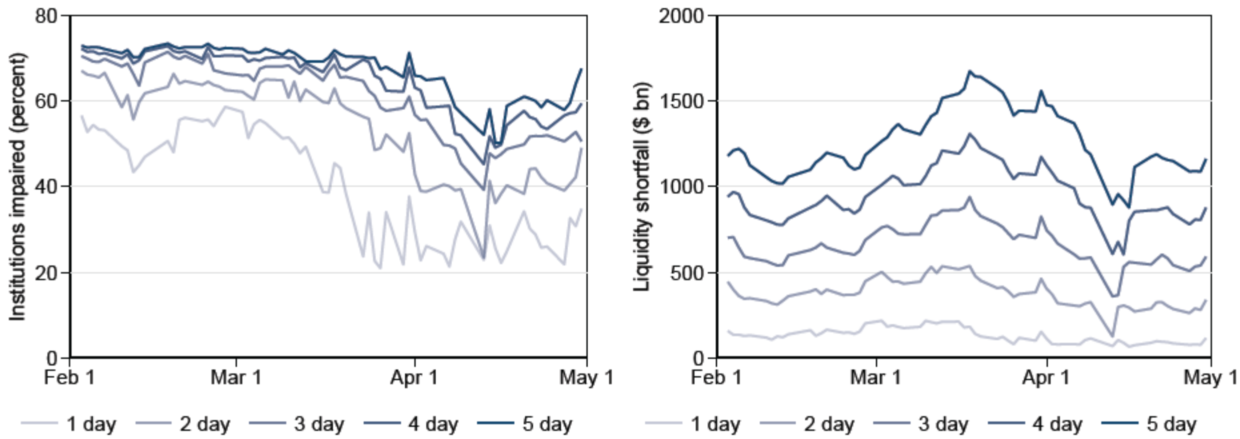

Figure 9: Multi-day scenario. The left panel shows the weighted share of impaired institutions for each multi-day scenario, averaged across the top-5 institutions. The right panel shows the total liquidity shortfall of impaired institutions for each multi-day scenario, averaged across the top-5 institutions.

is put in a compromised liquidity position as a result of the cyber attack.

However, the increase in the share of impaired institutions provides only a partial picture of the severity of a prolonged disruption to a top-5 bank, especially during the period of adverse market conditions. In particular, delayed recovery increases the severity of liquidity dislocations as many institutions' reserve balance can drop below zero in the scenario. It is therefore instructive to quantify the short-term liquidity support that would be required to restore the reserve balances of impaired banks back to the impairment threshold. The right panel of Figure 9 shows the aggregate liquidity shortfall, defined as the gap between institutions' impairment thresholds and their counterfactual reserve balance, aggregated across all impaired institutions. Averaged across 2020, the liquidity shortfall grows from $120 billion on day one to $1.1 trillion on day five. By comparison, in March, the liquidity shortfall increases from $164 billion to almost $1.5 trillion over the five days, reaching a peak of almost $1.7 trillion.

3.4 Attack on DFMUs

Finally, we consider a scenario involving an attack on designated financial market utilities (DFMUs) that impairs the systems of the attacked institution, as in the baseline scenario. We focus on two DFMUs: CHIPS, a wholesale payment system that offers multilateral netting benefits, and CLS, which provides settlement across different currencies' payment systems and is a key part of the infrastructure for global foreign exchange markets. In the event that either DFMU were to be rendered inoperable, banks could attempt to divert

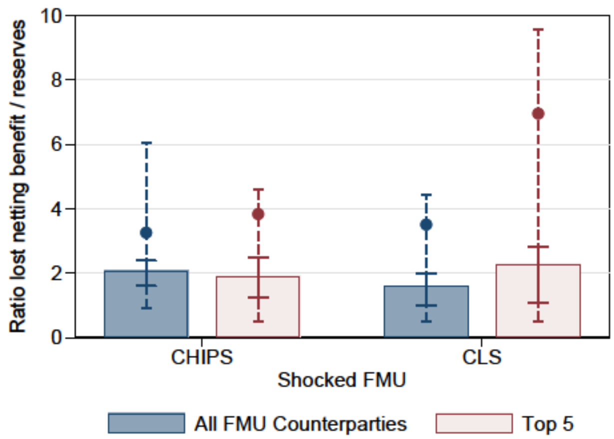

Figure 10: Lost netting benefits in DFMU scenario. The figure shows the distribution of the estimated value of failed payments on CHIPS and CLS, scaled by daily reserve balances of each bank. Bars represent the average ratio; solid whiskers represent the p25/p75 range; dashed whiskers represent the p1/p99 range; dots show the maximum impact days in March.

relevant payments to Fedwire. However, both DFMUs offer specific benefits that Fedwire does not. In particular, member institutions would no longer be able to realize the liquidity and capital savings associated with netting (CHIPS and CLS) and counterparty risk protections (CLS).

Although we cannot analyze directly the transactions between institutions on DFMUs, both CHIPS and CLS depend on Fedwire to settle participants' net payment obligations. Using the observed flows between banks and the DFMUs on Fedwire, and the gross value of payments processed within CHIPS and CLS from public data, we can approximate the netting benefits of a DFMU by taking the ratio of the (gross) activity on the DFMU to the (net) flows to the DFMU on Fedwire. To approximate an individual bank's daily netting benefits, we scale the daily aggregate gross-to-net ratio by the bank's daily payments to the DFMU.

Figure 10 summarizes the lost netting benefits, normalized by banks' reserves for CHIPS and CLS. The additional payment value that would need to be executed in Fedwire is significant, about two times banks' reserve balances on average. The dots correspond to the largest impact days in March 2020, which are in the right tail of the distribution for both CHIPS and CLS.

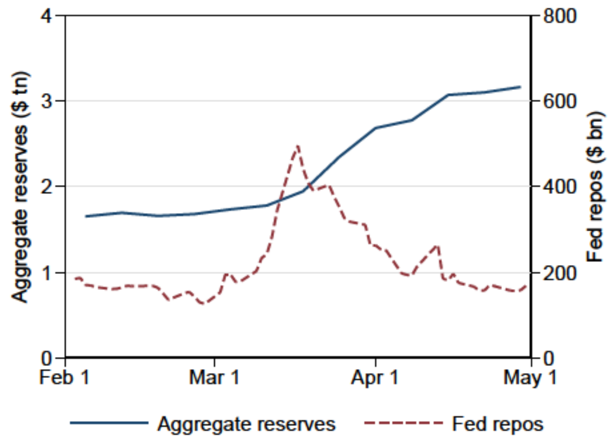

Figure 11: Aggregate reserves and Fed repo operations. The figure shows the level of aggregate reserves and the outstanding amount of Federal Reserve repo operations (overnight and term).

4 Mitigating factors and policy response considerations

The previous section showed that system-level cyber vulnerability was generally elevated during the market turmoil in early 2020. These results may be an underestimate as they do not incorporate any feedback effect on financial markets trading in response to a cyber attack. Perhaps surprisingly, the largest impact as measured by the weighted share of institutions impaired was in late February and early March and then declined dramatically through the end of March (Figure 6). Policy interventions intended to stabilize financial markets therefore had the unintended benefit of also mitigating systemic cyber vulnerability. In particular, the steep decline in network impact that starts in the second week of March coincides with the large increase in aggregate reserves resulting from Federal Reserve asset purchases (Figure 11). With the steep increase in reserves across the system, banks' liquidity positions became more resilient to disruptions in payment flows.

As pointed out earlier, the weighted impact of an attack on a top-5 bank begins to drop even earlier, in the first week of March while the unweighted impact continues to increase through the end of March. A potential explanation is the liquidity injected through the Fed's repo operations that started increasing at the beginning of March and were further expanded on March 9 and March 16.[8] Dealers' repo activity during this time has been linked to trades with affiliated large banks (Carlson, Saravay, and Tian, 2021). If these operations increased mainly the liquidity positions of the largest banks, they can explain the disconnect between the weighted and unweighted impacts over the course of March.

[8] See statements https://www.newyorkfed.org/markets/opolicy/operating_policy_200309 and https://www.newyorkfed.org/markets/opolicy/operating_policy_200316, respectively

Indeed, in an environment with abundant reserves, the potential for a cyber attack to have broader systemic impact is dramatically reduced by some measures. For one, the significant increase in aggregate reserves contributed to a lower average impact in the post-March period. From an ex-ante standpoint, operating under an abundant reserves regime can improve the resiliency of the system to illiquidity episodes caused by a cyber event. Another potential benefit of an abundant reserve environment is lowering the propensity for banks to strategically hoard liquidity in response to abnormal payment activity resulting from a cyber event. From an ex-post standpoint, offering easy access to emergency liquidity to banks experiencing short-term shortages reduce the risk of coordination failure and of transmission to other counterparties and markets.

The provision of liquidity is an effective, if blunt solution to improving resiliency to systemic cyber risk. However, as shown in the multi-day scenario analysis, a cyber event that goes unresolved for an extended period of time can require extraordinary levels of emergency liquidity injections (Figure 9). Although the discount window could, in principle, facilitate short-term access to liquidity, the levels required could quickly exceed permissible amounts based on impaired banks' unencumbered collateral. Furthermore, the multi-day scenario does not account for run-like behavior in other financial markets. The failure to remedy the operational issues sprouting from a cyber event could trigger financial instability across markets.

Another potential policy response involves directly addressing payment disruptions by using an emergency payment processor that can make payments on behalf of a bank directly impaired by an attack.[9] This form of response, which targets the root of the operational issue, has the advantage of containing the impact to those directly affected by a cyber event, and can reduce the set of counterparties with whom regulators must coordinate to maintain normal functioning. In addition, it has a stabilizing effect on the wholesale payment system by negating potential spillovers to other banks, thereby reducing the scope for coordination failures among other banks.

Implementation could involve a combination of an emergency payment processing system and a latent data back-up system for key institutions of the network, e.g., in the spirit of Sheltered Harbor.[10] When activated, clients of the impaired institution could be granted access to submit payment requests directly to the payment processing system. The data back-up system could be used to identify clients and assist the impaired institution in

[9] Although Fedwire can facilitate emergency payments for banks experiencing operational issues, only a set of prioritized payments can be processed in a timely manner.

[10] Sheltered Harbor is a not-for-profit industry-led initiative to have institutions regularly back up critical customer account data in a standardized format in case of an operational outage (https://www.shelteredharbor.org/)

authorizing payments. A related proposal put forth by Duffie and Younger (2019) recommends a standby narrow payment-bank utility that provides emergency payment processing services to critical non-bank financial institutions during operational emergencies. At heart, the goal would be to develop operational redundancies for the broader financial system that would be activated only in emergency situations.

The two forms of policy responses, the emergency provision of liquidity and of operational support, are complements. An abundance in aggregate reserves and accessibility to emergency liquidity can increase general resiliency to short-term cyber disruptions. For severe cyber incidents involving longer durations of recovery and for those involving key institutions of the network, an emergency payment system could be more efficient and effective at ensuring that markets function as usual, in parallel with the process of recovering an affected institution's operations.

To the extent that the payments proxy for financial transactions and that consumers and businesses would be reluctant to engage in those transactions with a bank impaired by a cyber attack, reserves and payment solutions may not have the same ameliorative effect. For example, if an impaired bank is a lender and cannot access books and records to authorize funds, payments may not be able to be made. While liquidity would not be the key source of amplification, it is possible that financial and real transactions would be hampered even if the payments issues were solved.

5 Conclusion

Recent events demonstrate that cyber and financial stress may also be driven by a common third factor. Notably, geopolitical conflict can increase financial market volatility, and simultaneously increase the likelihood of cyber warfare. The Russian invasion of Ukraine at the end of February 2022 is a case in point, both negatively impacting financial markets and raising the threat level of cyber risk. In this paper, we show that the financial system is particularly vulnerable to cyber attacks when uncertainty is high and prices are changing rapidly. This increase in vulnerability arises from the increase in payments volumes that accompany increased trading, as well as through the concentration of high dollar value payments among a relatively small set of systemically important banks. However, cyber attacks, in contrast to other forms of operational risk, may involve a strategic actor who times attacks to coincide with a period of financial stress. In such scenario, both financial amplifications *conditional* on a successful attack and the *conditional* likelihood of a cyber attack may rise with adverse financial conditions.

When considering policy responses, the optimal response to a malicious cyber attack may be very different, however, as the system must be resilient to the potentially higher amount of liquidity required in a situation with financial market volatility. Measures such as asset purchases may result in the increased reserves which can help to buffer these shocks. We note, however, that standard financial stability tools could be complemented by surgical policy tools specifically aimed at resolving cyber-related disruptions. The design and implementation of possible tools is an important question to be addressed by future work.

References

Afonso, G., D. Duffie, L. Rigon, and H. S. Shin (2022). Interbank payment timing is still closely coupled. Technical report, Working paper, Federal Reserve Bank of New York.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach (2020a). The drivers of cyber risk. Working paper 865, BIS.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach (2020b). Operational and cyber risks in the financial sector. Working paper 840, BIS.

Anand, K., C. Duley, and P. Gai (2022). Cybersecurity and financial stability.

Bech, M. L. and R. Garratt (2003). The intraday liquidity management game. *Journal of Economic Theory 109*(2), 198–219.

Bech, M. L., A. Martin, and J. McAndrews (2012). Settlement liquidity and monetary policy implementation—lessons from the financial crisis. *Economic Policy Review 18*(1).

Brando, D., A. Kotidis, A. Kovner, M. Lee, and S. L. Schreft (2022). Implications of cyber risk for financial stability. *FEDS Notes*.

Carlson, M., Z. Saravay, and M. Tian (2021). Use of the federal reserve's repo operations and changes in dealer balance sheets. *FEDS Notes* (August 06, 2021).

Chen, Q., C. Koch, P. Sharma, and G. Richardson (2020). Payments crises and consequences. Technical report, National Bureau of Economic Research.

Crosignani, M., M. Macchiavelli, and A. Silva (2021). Pirates without borders: The propagation of cyberattacks through firms' supply chains. Working Paper.

Curti, F., J. Gerlach, S. Kazinnik, M. Lee, and A. Mihov (2023). Cyber risk definition and classification for financial risk management. *Journal of Operational Risk 18*(2).

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

Eisenbach, T. M., K. Frye, and H. Hall (2019). Since the financial crisis, aggregate payments have co-moved with aggregate reserves. Why? *Federal Reserve Bank of New York Liberty Street Economics* (November 4, 2019).

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

Erol, S. and R. Vohra (2020). Network formation and systemic risk. *Working paper, Carnegie Mellon University*.

Federal Reserve Board (2020). Financial stability report. November.

Haddad, V., A. Moreira, and T. Muir (2021, 01). When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response. *Review of Financial Studies 34*(11), 5309–5351.

Kashyap, A. K. and A. Wetherilt (2019). Some principles for regulating cyber risk. *AEA Papers and Proceedings 109*, 482–87.

Kosse, A. and Z. Lu (2022). Transmission of cyber risk through the Canadian wholesale payment system. Staff working paper 2022-23, Bank of Canada.

Kotidis, A. and S. L. Schreft (2022). Cyberattacks and financial stability: Evidence from a natural experiment. FEDS working paper.

Lacker, J. M. (2004). Payment system disruptions and the Federal Reserve following September 11, 2001. *Journal of Monetary Economics 51*(5), 935–965.

Logan, L. K. (2020). A return to operating with abundant reserves. Remarks before the Money Marketeers of New York University, December 1, 2020. Available at https://www.newyorkfed.org/newsevents/speeches/2020/log201201.

McAndrews, J. and S. Rajan (2000). The timing and funding of Fedwire Funds transfers. *Economic Policy Review 6*(2).

Welburn, J. W. and A. M. Strong (2022). Systemic cyber risk and aggregate impacts. *Risk Analysis 42*(8), 1606–1622.