FEDERAL RESERVE BANK OF NEW YORK

ECONOMIC POLICY REVIEW

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

Thomas M. Eisenbach, Anna Kovner, and Michael Junho Lee

Volume 31, Number 1
January 2025

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

Thomas M. Eisenbach, Anna Kovner, and Michael Junho Lee

OVERVIEW

- The authors examine whether the impact of a cyberattack varies based on financial conditions, finding that systemic consequences are heightened when markets are more volatile and when financial intermediaries' balance sheets are strained.

- The onset of the COVID-19 pandemic was one such period: in February-March 2020, the estimated impact of a cyberattack on one of the five largest banks was about 50 percent greater than it would have been in the rest of 2020.

- Both the conditional likelihood of a cyber attack and its financial amplifications may rise with adverse financial conditions.

- Official-sector interventions to stabilize financial markets likely have a mitigating, if unintended, effect on cyber vulnerability: a decline in hypothetical network impact in the second week of March 2020 coincided with large liquidity injections by the Federal Reserve.

Much attention has been paid to the way in which cyber risk may be amplified by the financial system (see, for example, 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. If cyber risk and its amplifications were unrelated to financial risks, it would be reasonable to place cyber risk with other forms of operational risk that do not vary with financial stability vulnerabilities, such as severe weather. However, cyberattacks can be orchestrated by strategic actors, and consequently, the target and timing of a cyber incident are not random (Eisenbach, Kovner, and Lee 2022). This strategic element introduces systemic considerations. An attacker's motive may be to inflict maximum disruption or damage on the system. Attackers motivated by financial gain—ransomware attackers, for

Thomas M. Eisenbach is a financial research advisor and Michael Junho Lee a financial research economist at the Federal Reserve Bank of New York. Anna Kovner, currently director of research at the Federal Reserve Bank of Richmond, was director of Financial Stability Policy Research at the New York Fed at the time this article was written. Emails: thomas.eisenbach@ny.frb.org; anna.kovner@rich.frb.org; michael.j.lee@ny.frb.org.

The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Banks of New York or Richmond, the Federal Reserve Board, or the Federal Reserve System. To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/epr/2025/epr_2025_cyber-vulnerability_eisenbach

https://doi.org/10.59576/epr.31.1.1-24

example—may time attacks to occur during periods of high financial volatility, expecting to increase both the likelihood and the amount of a payout. Furthermore, if cyber risk and its systemic spillovers do, in fact, co-move with financial risks, then policymakers and institutions must be prepared for attacks to be launched precisely under circumstances when both cyber and financial vulnerabilities are acute. This raises a natural question: How does systemic cyber risk interact with other financial vulnerabilities? When it rains and negative shocks lead to financial market dislocations, does it also pour by increasing the systemic vulnerability to a cyberattack?

In this article, we examine how the consequences of a cyberattack evolve over the financial cycle, in particular, during periods of financial stress. We begin with cyber risk, the risk of loss stemming from an attack on or breach of computer systems and digital technologies (Brando et al. 2022; Curti et al. 2023). This loss could arise from any number of attack methods, including denial of service (DoS) attacks, whereby access to a firm's website is impaired; ransomware attacks, whereby an attacker prevents access to data and/or systems; and email attacks, whereby a firm's email system is compromised, enabling the attackers to submit fraudulent requests. All of these methods were identified by the Financial Services Information Sharing and Analysis Center (FS-ISAC) as potential threats.[1] Another method by which a cyberattack could affect the financial system is by targeting components of the financial supply chain, such as key service providers—as seen, for example, in the 2023 attack on ION Markets, which affected some customers' ability to book and process derivatives trades.

To study how cyber vulnerability evolves over the financial cycle, we analyze the market turmoil at the onset of the COVID-19 pandemic. March 2020 was marked by sudden, severe stress across asset classes and global financial markets (see, for example, Federal Reserve Board 2020; Haddad, Moreira, and Muir 2021), and market uncertainty, proxied by the VIX, spiked to its highest level since the global financial crisis (GFC). This time period offers a unique opportunity to study how cyber vulnerability is affected by financial market volatility because the shock is not directly related to financial institutions' business models. This period is also unique because it marks the first economic downturn since the GFC and the first episode of extreme market turmoil in an ample reserves regime.[2] The increased market volatility brought about 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, since many financial institutions shifted to remote working, thus potentially increasing the points at which cyber vulnerabilities can be exploited.

We start by documenting three broad stylized facts from early 2020 that lend support to the potential for heightened system-level vulnerability to cyber risk. First, wholesale payment activity strongly co-moves with market uncertainty, reaching a peak correlation of 0.87 in the period following the pandemic shock. Second, the wholesale payment system becomes more concentrated, with the top five banks increasing their payment share by 3 percentage points. In other words, times of high uncertainty give rise to even greater dependence on the resiliency of key large financial institutions in an already concentrated system. Third, banks exhibit signs of both intra- and interday liquidity stress, as evidenced by increases in payment-related liquidity needs, reserve balance volatility, and strategic settlement behavior. The 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 consistent feature of the past two decades.

Motivated by these empirical regularities, we estimate the hypothetical impact of a cyberattack by adapting the methodology of Eisenbach, Kovner, and Lee (2022), which involves a scenario-based approach to evaluating the financial stability risks of cyberattacks amplified through the Fedwire Funds wholesale payment network. To date, there has not been a cyber event with truly systemic consequences for the U.S. financial system. In the absence of actual examples, the wholesale payment network is a natural setting in which 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 Eisenbach, Kovner, and Lee (2022) methodology in other countries (see, for example, Kosse and Lu 2022) and confirmed that the key dynamics implied by the methodology played out in an actual cyberattack on a U.S. financial service provider (Kotidis and Schreft 2022).

We find that the systemic consequences of a successful cyberattack are higher when markets are more volatile and when financial intermediaries' balance sheets are strained. Specifically, cyber vulnerability, defined as the likely amplification of a cyberattack through the interruption of payment flows, was elevated in late February and early March 2020, with the impact of a cyberattack on one of the five largest banks averaging about 50 percent greater than it 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 complementarities.[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]

These results can be generalized to other periods of heightened uncertainty. We document that the systemic consequences of a cyberattack rise during times of elevated market volatility as a result of increased exposure to transaction processing by financial intermediaries and financial market utilities. In the full year of 2020, wholesale payment activity showed a remarkable correlation with the VIX, a correlation of 0.72 at a 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 of about $70 billion per day.

Increasing digitization is accompanied by increasing cyber risk, which *unconditionally raises* financial stability vulnerabilities. While there is a substantial technical literature on cyber defenses and documented attacks, this article does not consider the intensity or the probability of an attack. That said, it is worth noting that cyberattacks in pursuit of geopolitical goals may coincide with financial volatility, such 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 cyberattack and increasing financial market volatility, offering a window for a cyberattack 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 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 cyberattack on a counterparty. This conclusion, however, may underestimate the impact on markets should a cyberattack impair the trading books and records of a bank and delay settlement 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 affect financial market functioning.

Our findings on the relation between cyber risk and the financial cycle expand our understanding of the systemic implications of cyber risk along two key dimensions. First, at the level of an individual financial institution, our article provides an explicit framework for understanding and quantifying the implications of an attack over the financial cycle, as well as assessing 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.

Our article 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, Macchiavelli, and Silva 2021) or through critical service providers, utilities, and technology infrastructure (Welburn and Strong 2022). Our article examines macroeconomic risks in the context of the financial system, using the complete payments network, which offers a unique and holistic 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 becoming 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 industry experts, policymakers, and academics generally recognize cyber risk as significant (Brando et al. 2022), no systemic event has been triggered by a cyberattack to 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 (see, for example, Aldasoro et al. 2020a). Researchers have studied the frequency of cyberattacks and how they may be mitigated by bank lending (see, for example, Aldasoro et al. 2020b; Crosignani, Macchiavelli, and Silva 2021). To overcome this problem, studies have estimated cyber risk by examining scenarios (see, for example, Duffie and Younger 2019; Eisenbach, Kovner, and Lee 2022). A notable exception is Aldasoro et al. (2020b), who examine how operational costs, including cyber-related costs, changed for financial institutions around the time of the GFC.

From a microprudential viewpoint, they show that operational risk vulnerabilities build up in a boom but are realized in a downturn. Our results complement this finding by providing a macroprudential perspective: Cyber vulnerabilities emerge at the same time that market volatility spikes. Our article 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, Duley, and Gai (2022) and demonstrate the feedback loop between cyber risk and financial conditions that becomes particularly acute during times of high market stress.

The article proceeds as follows. Section 1 shows the effects of market stress on payment activity in early 2020 and the past two decades more broadly. Section 2 applies cyber scenarios to understand vulnerabilities during adverse market conditions. Section 3 discusses the mitigating effects of policy responses and Section 4 concludes.

1. 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 the Fedwire Funds Service (Fedwire), the U.S. wholesale payment system operated by the Federal Reserve that provides detailed information on the accounts of and flows between a diverse set of financial institutions.

1.1 Fedwire Payments Network

In 2020, the Fedwire payments network processed more than 700,000 transfers per day, with a total daily transfer value of over $3 trillion.[5] The Fedwire payments network has a core-periphery structure (Soramäki et al. 2007) in which payments are highly concentrated among a small set of large institutions and smaller institutions do not have many links between one another. The left panel of Chart 1 shows that the five most active banks in Fedwire account for roughly 50 percent of total payments, and the top ten account for more than 60 percent—a level of concentration that has been quite stable in the post-GFC period. The right panel shows how reliant the network is on the largest participants at the core; the value-weighted average of missing links among an institution's neighbors—that is, the percentage of potential links between the institution's neighbors that are not active in the data—is roughly 80 percent, while the equal-weighted average is less than 40 percent (see Eisenbach, Kovner, and Lee [2022] for additional statistics).

1.2 Patterns in Early 2020

We first illustrate each pattern at the beginning of 2020 and then use regressions to show that the patterns hold over a longer sample period using data from 1997 to 2020.

Chart 1
Network Concentration and Missing Links

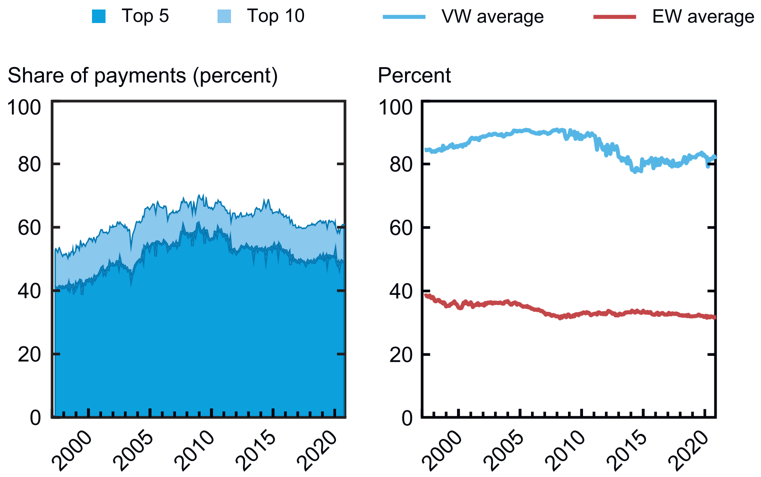

Sources: Authors' calculations based on Fedwire Funds Service data.

Notes: The left panel shows the share of payments sent by the top five and top ten institutions by payments activity. The right panel shows value-weighted (VW) and equal-weighted (EW) averages across participants and the percentage of missing links among an institution's neighbors.

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 Chart 2, 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 a daily frequency. From February to April 2020, the period when market stress was most 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 the top-5banks accounting for roughly 50 percent of payment value (Eisenbach, Kovner, and Lee 2022). 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 network's dependence on core institutions 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. Chart 3 plots the trailing five-day average share of payment value of the top five banks. The top five banks' share of daily payment value rises by about 3 percentage points at its peak on March 18

Chart 2

Wholesale Payment Activity and Market Uncertainty

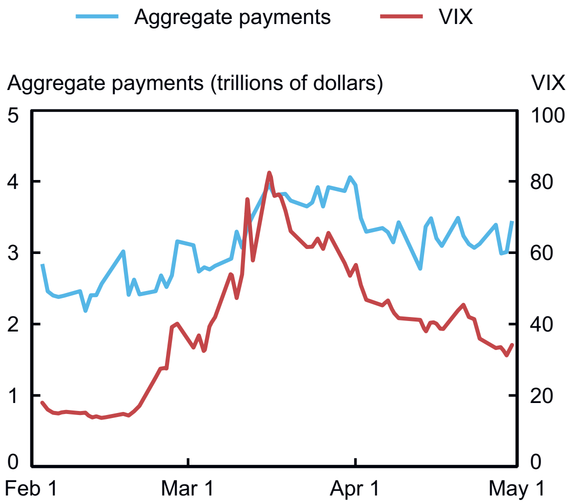

Sources: Authors' calculations based on Fedwire Funds Service data and CBOE Volatility Index.

Note: The chart shows aggregate payment activity in Fedwire Funds and the CBOE Volatility Index.

Chart 3

Concentration of Payment Activity

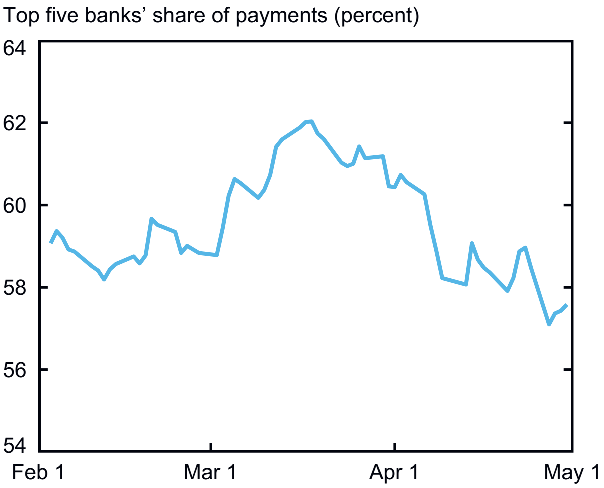

Sources: Authors' calculations based on Fedwire Funds Service data.

Note: The chart shows the trailing five-day average share of payment value of the top five banks.

(roughly two times the share's standard deviation in 2020), before falling and stabilizing at levels comparable to those at 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 positions, since 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 represent a form of coordination failure.

There were several indications that the wholesale payment system became more susceptible to coordination failure in the market turmoil spurred by COVID-19. 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 five banks, the ratio of daily payments over reserves increased by almost 50 percent, rising 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-five banks (Chart 4, left panel).

Banks typically manage their reserve balance to maintain a desirable level of liquidity. In theory, payments volume could increase but in a predictable manner, and banks might still be able to manage their reserve balances. However, if greater and more volatile payments had made reserve management more difficult, individual banks would have faced greater liquidity risks to processing payments. The increase in payment-related liquidity needs appears to have affected banks' abilities to manage a stable reserve balance. The right panel of Chart 4 plots the trailing thirty-day standard deviation in reserve balances for the most active banks in Fedwire. From mid-February to mid-April, the standard deviation in daily reserves for the top five 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 noticeably diminished post-crisis owing 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 the day. Chart 5 shows the timing of the 90th percentile of intraday payment value. Delays begin in late February, around the point when 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.

Chart 4
Payment-Related Liquidity Needs and Reserve Balance Volatility

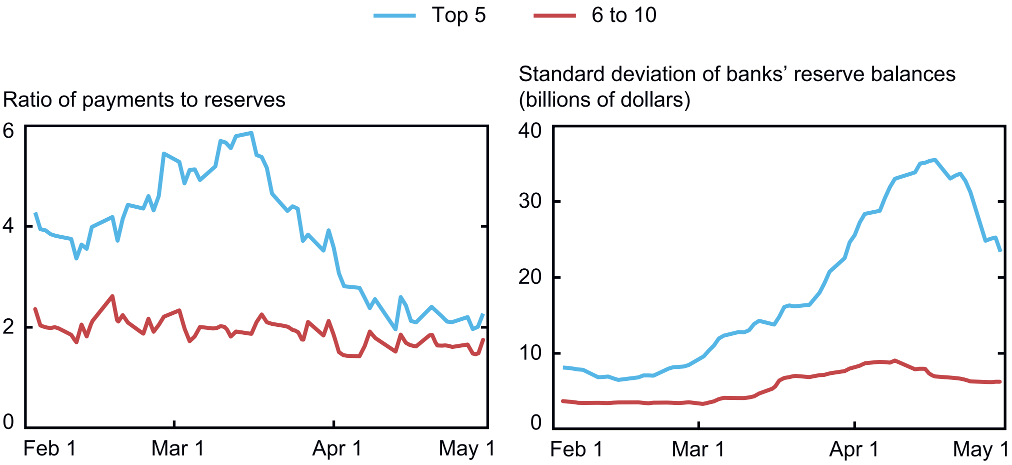

Sources: Authors' calculations based on Fedwire Funds Service data.

Notes: The left panel shows the ratio of average payment value to average reserves for the top five banks and for the banks ranked six to ten. The right panel shows the trailing thirty-day standard deviation of banks' reserve balances, averaged across the top five banks and for the banks ranked six to ten.

Chart 5
End-of-Day Payment Delay

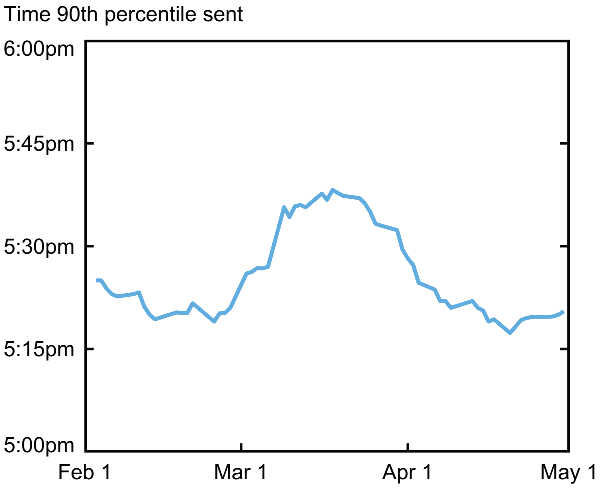

Sources: Authors' calculations based on Fedwire Funds Service data.

Note: The chart shows the five-day moving average of the time of day by which 90 percent of intraday payment value has been sent.

1.3 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 regressions of the key variables characterizing payment activity on the VIX over a sample period of more than 20 years (1997-2020). The regressions are at a monthly frequency with year fixed effects and control for the level of aggregate reserves, which strongly correlates with payment activity after 2008 (Eisenbach, Frye, and Hall 2019).

Column (1) shows that the relationship 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 five 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 by which 90 percent 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

Table 1
Wholesale Payment Activity and Market Uncertainty

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

Sources: Authors' calculations based on Fedwire Funds Service data and the CBOE Volatility Index.

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

implications for the amplification of a cyberattack through the financial system. In the context of the wholesale payment system, a cyberattack could be timed to occur during periods when 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 making a greater impact on 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 payment volumes are high and volatile and banks may have incentives to conserve reserves or think strategically about the timing of payments. The relationship between financial market uncertainty and cyber vulnerabilities in the payment system indicates that for at least the past two decades, when it rains, it pours.

2. Cyber Vulnerability during Adverse Market Conditions

We adopt the cyber scenario approach used in Eisenbach, Kovner, and Lee (2022) in a form modified for the analysis of adverse market conditions in 2020 in order to gain more insights into the impact of financial uncertainty on vulnerability to a cyberattack. All of our scenarios assume that a cyberattack 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 cyberattack may impair the availability of relevant data or the communication systems of an institution, or it may compromise the integrity of the institution's operations by either manipulating or corrupting the data. In both instances, the attacked institution's ability or willingness to make 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 owing to the institutional features of Fedwire, whereby 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 because of 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 that was observed following the attacks on September 11, 2001 (Lacker, 2004). Our scenarios therefore calculate counterfactual end-of-day reserve balances for all institutions, that is, 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 thirty-day average reserve balance of bank *i* on day *t*, and σ_{ref}^i and $\bar{r}_{\text{ref}}^i$ are the trailing thirty-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 the 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.[6]

2.1 Baseline Scenario

The baseline scenario examines the impact of an attack on a single top five bank, assuming no reaction by other banks and focusing only on the first day. We focus on an attack on a top five bank because of the highly concentrated network structure of Fedwire payments (as described in Section 1.1) where an attack on one of the central institutions can have an outsize impact for two related reasons. First, because these institutions account for a disproportionate share of payments value (Chart 1, left panel), an attack on one of them traps a disproportionate share of liquidity. Second, because many of the central institutions' smaller neighbors are not linked (Chart 1, right panel), most payment flows between smaller institutions effectively go through the central institutions and would therefore be disrupted by an attack on the central institution. In Eisenbach, Kovner, and Lee (2022), we also consider a "reverse scenario" where we calculate how many small institutions would have to be attacked at the same time in order for one of the top five to become liquidity impaired.

Chart 6 summarizes 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 attacked institution itself).[7] 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 the attacked institution. Our focus is on the maximal impact during the month of March 2020, as represented by the dots in the chart. As anticipated, the worst impact at the time of market turmoil is close to or even above the 99th percentile across all of 2020, in terms of both raw share and weighted share.

Chart 6

Impact of an Attack on a Top Five Bank

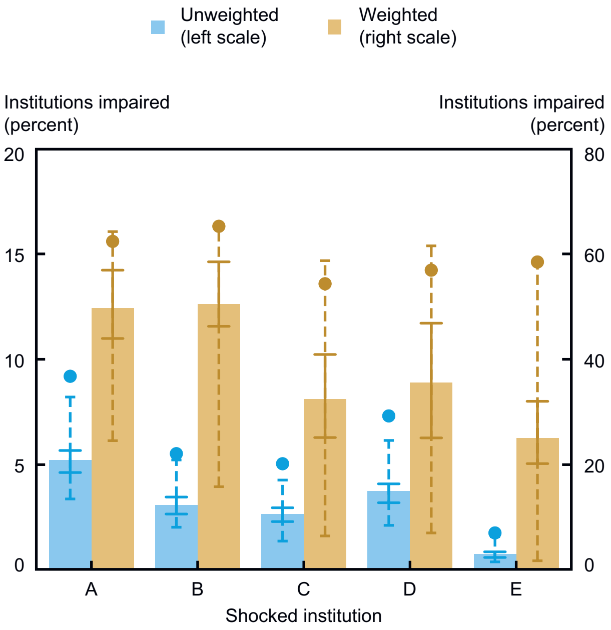

Sources: Authors’ calculations based on Fedwire Funds Service data.

Notes: The chart shows the distribution across days of the unweighted share of institutions impaired by a shock to each of the top five institutions and of the share weighted by payments (excluding the attacked institution). Bars represent the average impact; solid whiskers represent the p25/p75 range and dashed whiskers the p1/p99 range; dots show the maximum impact days in March.

Chart 7 shows the time series of the impact of an attack, averaged across the top five 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 percent of institutions by payment share would have been impaired in an attack on a top five bank. Compared to around an average of 40 percent across all of 2020, vulnerability is therefore about 50 percent higher. While the raw share of impaired banks increased throughout March and peaked on March 30, the weighted share decreased dramatically through the end of March, which we explore in greater detail when discussing the mitigating effects of policy interventions in Section 3.

2.2 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 nonreaction of banks assumes

Chart 7

Average Impact of an Attack on a Top Five Bank

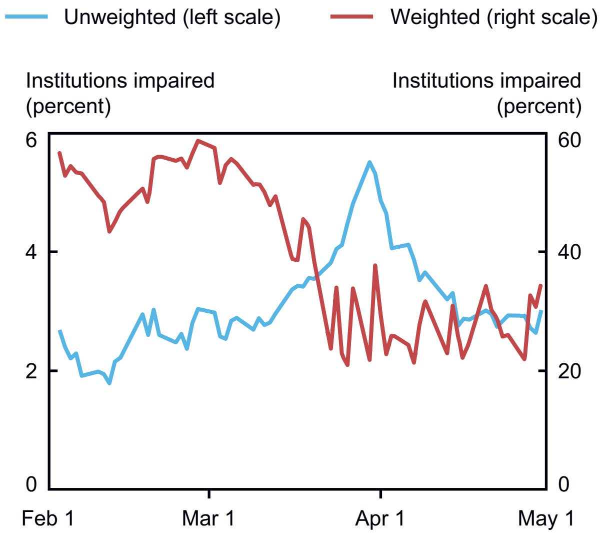

Sources: Authors' calculations based on Fedwire Funds Service data.

Note: The chart shows the daily time series of the impact of an attack, averaged across the top five banks.

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 1 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 cyberattack 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, which may be reluctant to disclose to counterparties and clients the exact state of their internal systems or data. For example, in the context of the GFC, Afonso, Kovner, and Schoar (2011) find evidence that banks respond to uncertainty by reducing lending in the interbank market.

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 net payment deficit of the institution in the entire year of 2020.[8]

A priori, it is not clear whether the impact should be greater in the cascade scenario, since 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 prompted to hoard liquidity in the cascade scenario is endogenous to the payment network structure. Chart 8 compares the impact in the simple and

Chart 8

Comparison of Simple and Cascade Scenarios for February-April 2020

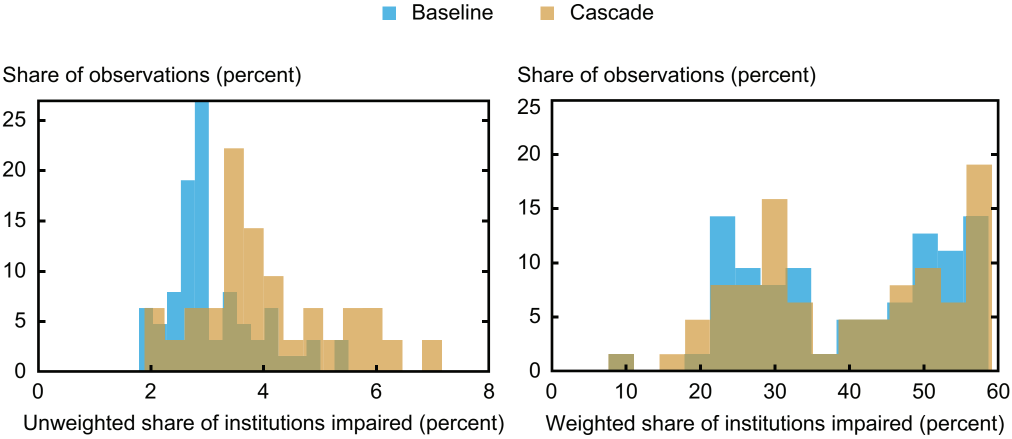

Sources: Authors' calculations based on Fedwire Funds Service data.

Notes: The chart shows the distribution across days of the impact of the baseline scenario and the cascade scenario for February to April 2020, averaged across the top five 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).

cascade scenario involving the top five 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 comes 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. Chart 9 shows the average daily forgone payment value in the cascade scenario, both the payments forgone by the attacked institution and the payments forgone by other institutions owing to the cascade. Both increase considerably in March, but the forgone payments from the cascade notably more so, reaching close to the level of the forgone payments of the attacked institution.

2.3 The Consequences of Delayed Recovery

So far, the analysis has focused on the single-day impact of a cyberattack. However, an extended cyber incident, owing to delays in operational recovery, 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 a timely recovery to normal operations to reduce the impact from a cybersecurity incident. The extent to which an attacked

Chart 9
Forgone Payment Value in Cascade Scenario

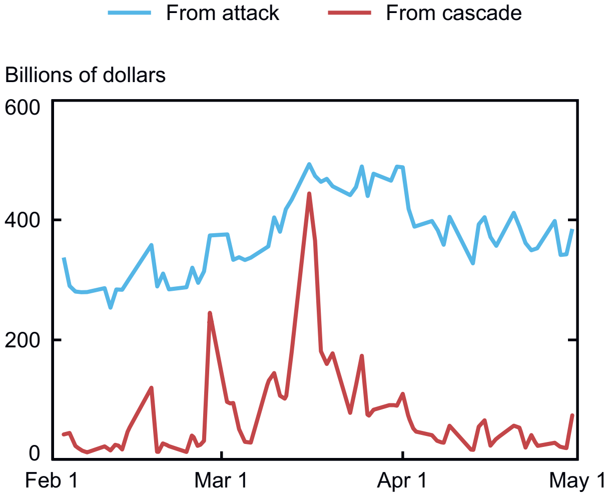

Sources: Authors' calculations based on Fedwire Funds Service data.

Note: The chart shows the payments forgone by the attacked institution and the payments forgone by other institutions due to the cascade, averaged across the top five institutions.

institution is able to restore the capabilities and services that were impaired depends on the strengths of its recovery function.

We examine the impact of a delayed recovery from a cyberattack by extending the baseline one-day scenario to consider a multiday scenario. The multiday 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 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 Chart 10. The impact of an attack averaged across the top five banks substantially increases with a delayed recovery in late February to mid-March, with the net weighted share increasing from about 45 percent on day one to more than 70 percent by day five. In other words, by the fifth day, the vast majority of the network (by payment share) is put in a compromised liquidity position as a result of the cyberattack.

However, the increase in the share of impaired institutions provides only a partial picture of the severity of a prolonged disruption to a top five bank, especially during a period of adverse market conditions. In particular, delayed recovery increases the severity of liquidity dislocations, since 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 Chart 10 shows the aggregate liquidity shortfall, defined as the gap between institutions' impairment thresholds and their counterfactual reserve balance, aggregated across all

Chart 10
Multiday Scenario

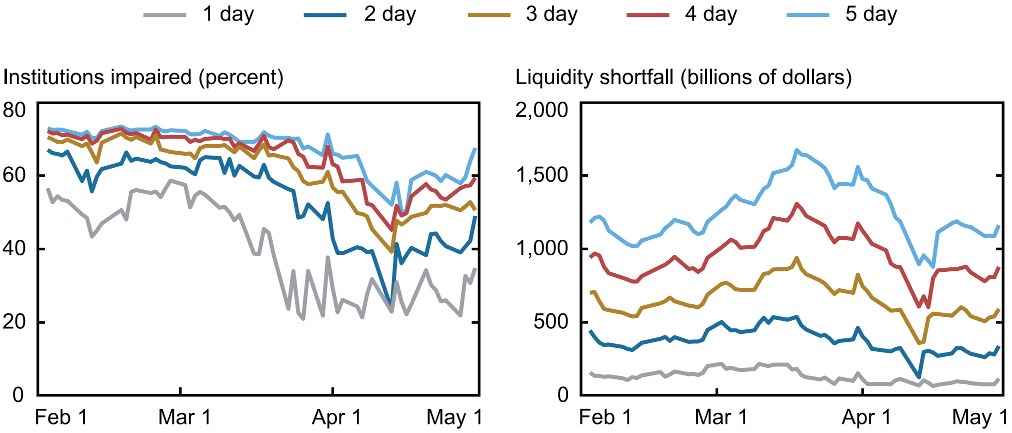

Sources: Authors' calculations based on Fedwire Funds Service data.

Notes: The left panel shows the weighted share of impaired institutions for each multiday scenario, averaged across the top five institutions. The right panel shows the total liquidity shortfall of impaired institutions for each multiday scenario, averaged across the top five institutions.

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.

2.4 Attack on Designated Financial Market Utilities

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 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 directly analyze 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.

Chart 11 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.

3. 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, since they do not incorporate any feedback effect on financial markets trading in response to a cyberattack. 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 (Chart 7). 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 (Chart 12). 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 five 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.[9] Dealers' repo activity during this time has been linked to trades with affiliated large banks (Carlson, Saravay, and Tian 2021). If these operations mainly increased the liquidity positions of the largest banks, that could explain the disconnect between the weighted and unweighted impacts over the course of March.

Indeed, in an environment with abundant reserves, the potential for a cyberattack to have a 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 system's resiliency 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 reduces 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 multiday scenario analysis, a cyber event that goes unresolved for an extended period of time can require extraordinary levels of emergency liquidity injections (Chart 10). Although the discount window could, in principle, facilitate

Chart 11

Lost Netting Benefits in DFMU Scenario

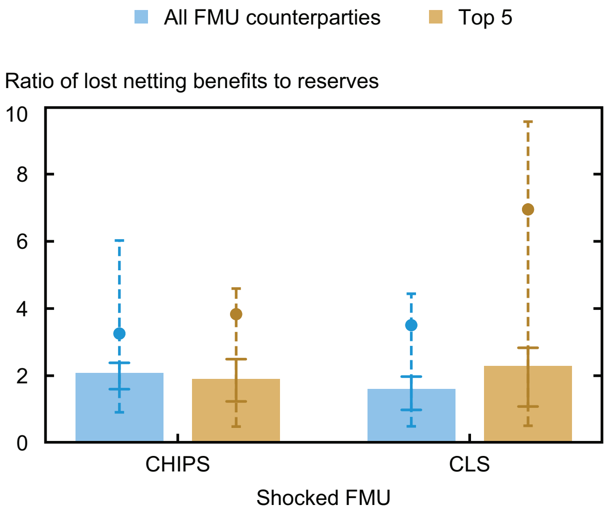

Sources: Authors' calculations based on Fedwire Funds Service data.

Notes: The chart 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 and dashed whiskers represent the p1/p99 range; dots show the maximum impact days in March.

Chart 12

Aggregate Reserves and Fed Repo Operations

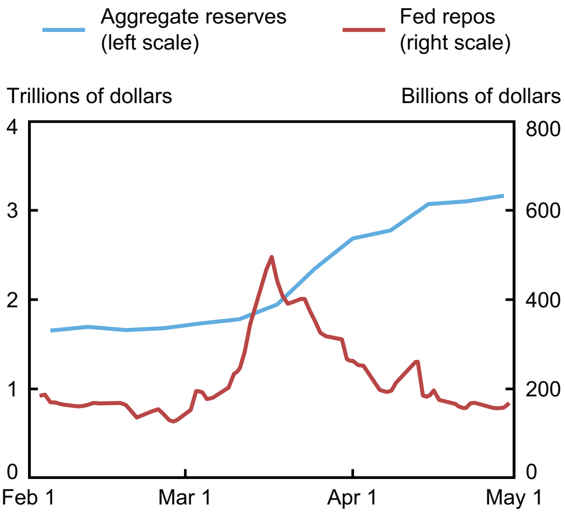

Sources: Authors' calculations based on Fedwire Funds Service data.

Note: The chart shows the level of aggregate reserves and the outstanding amount of Federal Reserve repo operations (overnight and term).

short-term access to liquidity, the levels required could quickly exceed permissible amounts based on impaired banks' unencumbered collateral. Furthermore, the multiday 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.[10] 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, for example, in the spirit of Sheltered Harbor.[11] When such a system is 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 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 nonbank financial institutions during operational emergencies. In essence, 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 operational support, are complements. An abundance of aggregate reserves and access 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 can be used as a proxy for financial transactions and that consumers and businesses would be reluctant to engage in those transactions with a bank impaired by a cyberattack, 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, then it may not be able to make payments. 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 resolved.

4. 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 affecting financial markets and raising the threat level of cyber risk. In this article, we show that the financial system is particularly vulnerable to cyberattacks 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, cyberattacks, in contrast to other forms of operational risk, may be strategically timed to coincide with a period of financial stress. In such a scenario, financial amplifications *conditional* on a successful attack *and* the *conditional* likelihood of a cyberattack may both rise with adverse financial conditions.

Our study points to several avenues of future research. First, we show evidence for a significant interaction between cyber and financial conditions that could magnify systemic risk during times of financial stress. This poses concerns regarding the possibility of gaps between the social and private value of safeguarding the technical infrastructure of the financial system. In light of the importance of financial market utilities and banks, future research could explore the cyber resiliency of the financial system and provide insight into how the (changing) financial environment affects technological resilience as well as document the extent of single-source providers of key financial services and data. This would naturally lend itself to a prescriptive exercise on the optimal supervisory and regulatory framework to prevent the likelihood of a systemic cyber event. Second, our analysis largely focuses on propagation through financial and technological linkages. In the midst of a cyberattack, however, there is significant scope for panic and misinformation. A promising and underexplored area is how incomplete information can exacerbate the original shock and even be manipulated to magnify the impact. Finally, we show that liquidity injections by the central bank were largely effective at mitigating amplification channels in March 2020.

The optimal policy response to a malicious cyberattack may be very different, however, since the system must be resilient enough to provide the potentially higher amount of liquidity required in a situation with financial market volatility. Measures such as asset purchases may result in increased reserves, which can help to buffer these shocks. We note, though, that standard financial stability tools could be complemented by instruments specifically designed to resolve cyber-related disruptions. The design and implementation of such tools are important topics for future research.

Notes

Acknowledgments: The authors are grateful for comments from Julian di Giovanni (editor), an anonymous referee, 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. They thank Ryan Craver for outstanding research assistance.

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

[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 (for example, Logan, 2020).

[3] For example, Afonso et al. (2022) find 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.

[5] See the publicly available statistics at https://www.frbservices.org/resources/financial-services/wires/volume-value-stats.

[6] The threshold differs in two ways from that used in Eisenbach, Kovner, and Lee (2022), 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 thirty-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 balances 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.

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

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

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

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

[11] 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/).

References

Afonso, G., D. Duffie, L. Rigon, and H. S. Shin. 2022. "How Abundant Are Reserves? Evidence from the Wholesale Payment System." Federal Reserve Bank of New York Staff Reports, 1040.

Afonso, G., A. Kovner, and A. Schoar. 2011. "Stressed, Not Frozen: The Federal Funds Market in the Financial Crisis." Journal of Finance 66, no. 4: 1109–39.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach. 2020a. "The Drivers of Cyber Risk." Bank for International Settlements Working Paper no. 865.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach. 2020b. "Operational and Cyber Risks in the Financial Sector." Bank for International Settlements Working Paper no. 840.

Anand, K., C. Duley, and P. Gai. 2022. "Cybersecurity and Financial Stability." Deutsche Bundesbank Discussion Paper no. 08/2022.

Bech, M. L., and R. Garratt. 2003. "The Intraday Liquidity Management Game." Journal of Economic Theory 109, no. 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, no. 1.

Brando, D., A. Kotidis, A. Kovner, M. Lee, and S. L. Schreft. 2022. "Implications of Cyber Risk for Financial Stability." FEDS Notes (May).

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).

Chen, Q., C. Koch, P. Sharma, and G. Richardson. 2020. "Payments Crises and Consequences." NBER Working Paper no. w27733.

Crosignani, M., M. Macchiavelli, and A. Silva. 2021. "Pirates without Borders: The Propagation of Cyberattacks through Firms' Supply Chains." Federal Reserve Bank of New York Staff Reports, no. 937.

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, no. 2.

Duffie, D., and J. Younger. 2019. "Cyber Runs." Brookings Institution Hutchins Center Working Paper no. 51.

References (Continued)

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).

Eisenbach, T. M., A. Kovner, and M. J. Lee. 2022. "Cyber risk and the U.S. Financial System: A Premortem Analysis." Journal of Financial Economics 145, no. 3: 802–26.

Erol, S., and R. Vohra. 2020. "Network Formation and Systemic Risk." Carnegie Mellon University working paper.

Federal Reserve Board. 2020. Financial Stability Report. (November).

Haddad, V., A. Moreira, and T. Muir. 2021. "When Selling Becomes Viral: Disruptions in Debt Markets in the COVID-19 Crisis and the Fed's Response." Review of Financial Studies 34, no. 11: 5309–51.

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." Bank of Canada Staff Working Paper 2022-23.

Kotidis, A., and S. L. Schreft. 2022. "Cyberattacks and Financial Stability: Evidence from a Natural Experiment." Finance and Economics Discussion Series Working Paper no. 2022-025.

Lacker, J. M. 2004. "Payment System Disruptions and the Federal Reserve Following September 11, 2001." Journal of Monetary Economics 51, no. 5: 935–65.

Logan, L. K. 2020. "A Return to Operating with Abundant Reserves." Remarks before the Money Marketeers of New York University (December 1).

McAndrews, J., and S. Rajan. 2000. "The Timing and Funding of Fedwire Funds Transfers." Economic Policy Review 6, no. 2.

Soramäki, K., M. L. Bech, J. Arnold, R. J. Glass, and W. E. Beyeler. 2007. "The Topology of Interbank Payment Flows." Physica A: Statistical Mechanics and Its Applications 379, no. 1: 317–33.

Welburn, J. W., and A. M. Strong. 2022. "Systemic Cyber Risk and Aggregate Impacts." Risk Analysis 42, no. 8: 1606–22.

FEDERAL RESERVE BANK OF NEW YORK
ECONOMIC POLICY REVIEW

The *Economic Policy Review* is published by the Research and Statistics Group of the Federal Reserve Bank of New York. The views expressed are those of the individual authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Series Editor: Julian di Giovanni. Editorial Staff: Robert Powell, Trevor Delaney, Maureen Egan, Anna Snider, and Peter Stevens. Design Staff: Theresa Izzillo, Alyona Botogoeva, and Laura Gharrity.

www.newyorkfed.org/research
Follow us on X (Twitter): *@NYFedResearch*