Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Cyberattacks and Financial Stability: Evidence from a Natural Experiment

Antonis Kotidis and Stacey L. Schreft

2022-025

Please cite this paper as:
Kotidis, Antonis, and Stacey L. Schreft (2022). "Cyberattacks and Financial Stability: Evidence from a Natural Experiment," Finance and Economics Discussion Series 2022-025. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2022.025.

Cyberattacks and Financial Stability: Evidence from a Natural Experiment

by

Antonis Kotidis and Stacey L. Schreft*

March 23, 2022

Abstract

This paper studies the effects of a unique multi-day cyberattack on a technology service provider (TSP). Using several confidential daily datasets, we identify and quantify first- and second-round effects of the event. For banks using relevant services of the TSP, the attack impaired their ability to send payments over Fedwire, even though the Federal Reserve extended the time they had to submit payments. This impairment (first-round effect) caused other banks to receive fewer payments (second-round effect), leaving them at risk of having too few reserves to send their own payments (a potential third-round effect). These innocent-bystander banks responded differently depending on their size and reserve holdings. Those with sufficient reserves drew down their reserves. Of the others, smaller banks borrowed from the discount window, while larger banks borrowed in the federal funds market. These significant adjustments to operations and funding prevented the second-round effect from spilling over into third-round effect and broader financial instability. These findings highlight the important role for bank contingency planning, liquidity buffers, and the Federal Reserve in supporting the financial system's recovery from a cyberattack.

JEL Classification Codes: E42, E50, E58, G21, G28

Keywords: cyberattack, cybersecurity, payments, discount window, federal funds market

* Authors' contacts: Antonis Kotidis, Board of Governors of the Federal Reserve System, email: antonis.kotidis@frb.gov; Stacey Schreft, Board of Governors of the Federal Reserve System and U.S. Treasury Office of Financial Research, email: stacey.schreft@frb.gov. This paper was produced while Schreft was on assignment to the Federal Reserve Board. We are deeply grateful to Andreas Lehnert, Beth Klee and Jill Cetina for their support and comments at various stages of this project. We also thank Gara Afonso, Adam Copeland, Jeff Gerlach, Anna Kovner, Michael Lee, Kris Natoli and Margaux MacDonald for helpful discussions. Tristan D'Orsaneo provided outstanding research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Board of Governors of the Federal Reserve System, the Office of Financial Research, or the U.S. Treasury Department.

1. Introduction

Cyber risk is a major concern for financial system regulators. In interviews, the Chair of the Federal Reserve, Jerome Powell, has repeatedly cited cybersecurity as the risk of greatest concern to him (Powell, 2019; Powell, 2021). European Central Bank President Christine Lagarde has stated that a cyberattack could trigger a financial crisis (Lagarde, 2020). In the same spirit, economists have begun to highlight the need for cyber monitoring and macroprudential regulation (Kashyap and Wetherilt, 2019) and emphasize the importance of understanding cyberattacks as a financial stability risk by studying hypothetical cyber events (Duffie and Younger, 2019; Eisenbach, Kovner, and Lee, 2021). Yet, to date, no paper has traced through and quantified the impact of an actual cyberattack that potentially threatened financial stability, which is the main contribution of our paper.

We analyze a unique event as a natural experiment to understand how a multi-day cyberattack that disrupts the operations of a technology service provider (TSP) can affect payment flows and bank liquidity.[1] When the TSP discovered the cyberattack, it took its computer systems offline to limit the damage done.[2] In doing so, some bank customers of the TSP ("users," which make up the treatment group) lost the ability to send payments over Fedwire using their usual processes.[3] The banks affected had backup processes available for accessing Fedwire, but those processes were more time consuming to use, and the banks generally did not switch over to them quickly.[4] In contrast, other banks that were not affected by the TSP's recovery efforts ("non-users," which make up the control group) remained able to access and send payments over Fedwire. These non-

[1] Throughout this paper, we refer to the firm impacted by the cyberattack as "the TSP" and maintain the confidentiality of the event by not referring to the firm by name, lines of business, or location; the specific dates or days of the week when the cyberattack occurred; the event's duration; the cybersecurity vulnerability exploited; and the exact number of institutions affected. Cyberattacks often do not become public knowledge unless they involve the release of personally identifiable information.

[2] Effective cybersecurity requires not just defending against cyberattacks and identifying those that occur, but also responding to and recovering from them. In this paper, we refer to the cyberattack as starting with its discovery and ending when the recovery is complete. Thus, we use the term "cyberattack" to refer to the active intrusion and the disruption to the TSP's systems. This captures the reality that the response phase, which in this case involved taking systems offline, can be more disruptive to a target's operations than the period when the attacker is active in the target's computer systems.

[3] The number of bank users was greater than 50, which is important in our empirical analysis. Not all of the users were banks; some were credit unions or other financial institutions, but all routinely sent payments over Fedwire. For simplicity, we refer to the users as "banks" throughout the paper.

[4] See FedLine® Solutions for a description of the various ways financial institutions can send Fedwire payments.

user banks serve as a counterfactual for what would have happened in the absence of the cyberattack.

Our focus on this particular cyberattack and its effect on payment flows through the impairment of some banks' access to Fedwire is motivated by the criticality of the payment system for financial stability and broader economy. The U.S. financial system is a complex network of financial markets, with the Fedwire interbank payments network at its heart. As a result, any disruption to banks' ability to send payments through Fedwire is a potential risk to financial stability (e.g., Denbee et al., 2021; Duffie and Younger, 2019; Armantier, Arnold, and McAndrews, 2008). The payment system's relevance for inflation, securities prices, and the transmission of monetary policy (Piazzesi and Schneider, 2021) means that such disruptions can affect the macroeconomy more generally. The pivotal role of the payment system in the economy makes it crucial to understand the impact a cyberattack could have on payment flows and the policies that can mitigate them.

Our empirical identification strategy exploits several matched confidential datasets, which include: (i) interbank payments from the Fedwire Funds Service, (ii) interbank federal funds loans, (iii) bank-level loans from the discount window, and (iv) bank-level data on reserve balances at the Federal Reserve. Importantly, the frequency of all these datasets is daily, which allows us to sharpen the identification and carefully quantify the impact of the cyberattack. We rely as well on unique confidential data on which banks had their usual access to Fedwire disrupted by the cyberattack at the TSP, matched with balance sheet information from Call Reports.

Our findings can be summarized as follows.[5] During the event and relative to the non-users, users sent 16.4% fewer payments on average across the days of the event. As a result, the value of payments sent by users dropped 26.5%, which resulted in the accumulation of reserves in their accounts at the Federal Reserve. However, the impact was not evenly distributed across days. On the first day of the cyberattack,[6] when the impact was most severe, users sent 36% fewer payments,

[5] To briefly describe our terminology at this point, a sender-bank sends payments through Fedwire to a receiver-bank. A sender-bank can be either a user of the TSP (treatment group) or a non-user (control group). We refer to these banks as "users" and "non-users," respectively. On the receiving end, all receiver-banks are non-users of the affected operations of the TSP (innocent bystanders). We refer to them as "receiver-banks" throughout the paper.
[6] The first day of the attack in this paper refers to the first day the TSP reacted to stop the attack by taking its servers offline. The attack may have occurred earlier and not immediately have been discovered, or it may have been discovered but thought not to require such action. We carefully check for this possibility in our empirical analysis.

which resulted in a 50.7% drop in the value of payments. This drop would have been nearly 61% had the Federal Reserve not extended the trading day beyond its normal close.[7] We find that users took advantage of the trading day extension by sending larger, but not more, payments, which is in line with the Federal Reserve's recommendations to financial institutions in contingency situations and findings from previous research (e.g. Klee, 2010). On the subsequent days, as the TSP gradually restored its service, the magnitude of the disruption in payment flows became gradually smaller.

The impact on users (the *first-round effect*) was not the only effect of the TSP's service disruption. When some users had difficulty sending payments, some non-user banks (i.e., innocent bystanders to the event) might have received fewer payments than they otherwise would have (the *second-round effect*). Because banks rely heavily on incoming payments to provide the funds to send their own payments, a disruption of incoming payments received could leave a bank with a shortage of liquidity to send their own payments. In this case, the impact of a cyberattack could propagate through the financial system (creating a *third-round effect*) unless banks find alternative sources of funding to make their own payments. We find that although receiver-banks had fewer incoming payments during the cyberattack, and in particular on the first day, they were able to make their outgoing payments, either because the Federal Reserve provided liquidity to them or because they found other sources of liquidity. In particular, relatively smaller banks were more likely to tap the discount window for funding; this was especially true for those with relative fewer reserves as a share of their assets.[8] In contrast, relatively larger banks borrowed from the fed funds market, with the exception of a few very large banks that had a high share of reserves relative to total assets. We find that those very large banks were more likely to draw down their own reserves on the first and subsequent days of the cyberattack.

All in all, we find that interventions by the Federal Reserve—both as a lender of last resort and an operator of Fedwire—and actions by user and non-user banks prevented the cyberattack from becoming a financial stability event. First, users requested—and the Federal Reserve granted--extensions of the Fedwire day. That, along with user banks' switch to alternative methods to send

[7] The Fedwire business day runs weekdays, except holidays, from 9 p.m. eastern time on the preceding day until 7 p.m. The deadline for a bank to initiate a payment for a third party is 6 p.m. See Fedwire Funds Services.
[8] Ashcraft, McAndrews, and Skeie (2011) shows that small banks are less likely to access the fed funds market for funding and, as such, more likely to borrow from the discount window to address a funding shortfall. Ennis and Klee (2021) shows that borrowing from the discount window is more pronounced for those banks with relatively fewer reserves.

payments over Fedwire, mitigated the drop in payments sent, but was not enough to prevent the impact on users from spilling over to reduce non-users' funding. Non-user banks with sufficiently high reserves drew down those reserves. Of the remaining non-users, smaller banks requested, and the Federal Reserve Banks granted, discount window loans, while larger banks borrowed in the interbank market. All of these actions together prevented a third-round effect of the cyberattack from spillover of the associated disruption of payment flows to other banks and parts of the financial system. While the drop in the share of total payments sent was not large, we present estimates that it could have been significantly larger if the cyberattack had hit a more dominant TSP or more or larger banks.

Our paper relates to three strands of the literature. First, we contribute to the literature on the effects of a cyberattack on the financial system. Several papers have described the transmission channels through which cybersecurity events could impair financial stability (for example, Boer and Vasquez, 2017; ESRB, 2020; Healey et al., 2018; Kopp, Kaffenberger, Wilson, 2017; Office of Financial Research, 2017; Ros, 2020; Warren, Kaivanto, and Prince, 2018). Others have estimated the cost of cyberattacks (Aldasoro et al., 2020; Bouveret, 2019; Council of Economic Advisors, 2018; Gogolin, Lim, and Vallascas, 2021; Wellburn and Strong, 2019) and whether stock prices reflect the potential losses (Florackis et al., 2020; Kamiya et al., 2021).

Some papers aim to quantify the impact of cyberattacks, whether actual or, because of a lack of data, hypothetical. Crosignani, Macchiavelli, and Silva (2021) quantifies the impact of the NotPetya attack, and while NotPetya did not impact the financial sector, the paper finds that access to bank credit lines and the use of trade credit mitigated the damage done. Duffie and Younger (2019) estimates that twelve systemically important U.S. financial institutions would have enough liquid assets to withstand wholesale funding runs in a hypothetical severe cyber event. However, the payment system might not be able to process payments sufficiently quickly during a cyber event to avoid economic harm. Eisenbach, Kovner, and Lee (2021) considers a hypothetical one-day cyberattack that stops any one of the five largest banks from sending payments. The impact is sizable; for example, almost 40% of bank assets are affected, excluding those of the affected bank. The impact could be 38% larger if the attack was timed to hit on the day with the highest number of payments.

Our paper contributes to this strand of the literature, and stands in contrast to it, by being the first paper to explore the financial stability implications of an actual cyberattack on the banking

system, as well as the central bank's actions to mitigate the impact.[9] The data allow us to conduct a unique natural experiment because we can precisely identify the banks that used the TSP's affected services and had to send Fedwire payments using contingency methods. Likewise, we can identify banks that were not users of affected services and offer a counterfactual regarding what would have happened in the absence of the cyberattack, thereby allowing us to carefully quantify the impact of the cyberattack on payment flows.

Second, our paper connects to the literature that studies the payment system as a part of the financial system's critical infrastructure. Fedwire has been found to have a complex network (Soramäki et al., 2007) through which stress easily can spread between banks. This network has been disrupted in the past by events that affected access to Fedwire and required Federal Reserve intervention. One example is the September 11, 2001, terrorist attacks (McAndrews and Potter, 2002; Lacker, 2003). Another is the software failure Bank of New York (BoNY) experienced in 1985 (Ennis and Price, 2015).[10] In a study of hypothetical cyberattacks (Eisenbach, Kovner, and Lee, 2021), an attack that impaired the ability of an average of 24 small banks' (those with less than $10 billion in assets) to send Fedwire payments could reduce liquidity at least one of the five largest financial institutions. Fedwire disruptions, whatever their source, can interrupt the normal dynamics of liquidity and risk sparking liquidity spirals in payment systems (McAndrews and Rajan, 2000; Klee, 2010; Afonso and Shin, 2011).

Our paper contributes to this second strand by analyzing an actual cyberattack that disrupts banks' ability to send payments through Fedwire. Those banks were not extremely large, yet their inability to send payments caused banks with no direct connection to the cyberattack to find themselves short reserves to send their own payments. This finding offers support for the Eisenbach, Kovner, and Lee (2021) result that the impact of a cyberattack at smaller institutions can be amplified through its effect on larger institutions. In our study, the Federal Reserve used the same tools to support the affected banks that it used in the September 11 and BoNY disruptions.

Third, our paper connects to the literature on banks' reserves management. Banks manage their reserves in response to liquidity or payments shocks (e.g. Poole, 1968; Frost, 1971; Bhattacharya

[9] Closely related to our paper is Klee (2010), which analyzes the impact of hypothetical operational outages on the flow of payments. Klee finds that discount window borrowing picks up on the day of the outage and that extensions to trading hours are more likely to be granted on days with possible outages.

[10] The software failure damaged the database BoNY used to identify the receivers of the Fedwire payments it sent. The Federal Reserve extended the deadline for BoNY to send funds until 2:15 a.m. the next morning, and then made an unusually large discount window loan to BoNY (Ennis and Price, 2015).

and Gale, 1987) by borrowing from the fed funds market (private liquidity) or the discount window (public liquidity). The fed funds market is an over-the-counter market where depository institutions borrow reserves that banks hold at the Federal Reserve. The micro-decisions of market participants in the fed funds market, where they negotiate loans bilaterally, have been modeled by Ho and Saunders (1985) and, more recently, by Afonso and Lagos (2015) and Bianchi and Bigio (2021). On the empirical side, the literature has studied the dynamics of lending in the fed funds market, providing evidence of illiquidity after the Global Financial Crisis (e.g. Ashcraft and Duffie, 2007; Afonso, Kovner, and Schoar, 2011). In contrast, the Federal Reserve's discount window facility provides short-term liquidity insurance to eligible depository institutions by injecting reserves into the system. For example, Clouse et al. (2000) argues that the facility is an important backup source of liquidity in crisis times, while Fischer (2016) promotes the idea that discount window lending can be useful at all times. However, although discount window lending is meant to deal with illiquidity concerns (e.g. Bagehot, 1873), Schwartz (1992) argues that large amounts of discount window lending have gone to banks experiencing insolvency problems. Goodfriend and King (1988) makes a similar critique of the usefulness of discount window lending. Ennis and Klee (2021) presents empirical evidence that many banks use of the discount window in normal (non-stress) times and consistently adjust their behavior to influence their need to borrow. They conclude that banks, in principle, could adjust to not having the discount window facility open at all times, as some critics might prefer.

We contribute to this third strand of the literature by documenting the usefulness of both the fed funds market and the discount window facility for reserves management during a cyberattack. In particular, we document the heterogeneous responses of banks in their reserve management based on their size and initial reserve holdings. We show that large banks, except for a few very large ones with ample reserves they can draw down, borrow in the fed funds market. This is consistent with the analysis in Duffie and Younger (2019), which shows systemically important banks' resilience to hypothetical cyberattacks because of their holdings of high-quality liquid assets. In contrast, small banks turn to the discount window facility for funding. with the effect

being more pronounced among small banks that hold fewer reserves as a proportion of their assets, which echoes the findings in Ennis and Klee (2021).[11]

In focusing on the role the fed funds market and discount window can play in mitigating a cyberattack, we raise some new arguments for the value of these funding sources. Our results suggest that closing the discount window in normal times likely would be a mistake. Cyberattacks are the new normal, and discount window funding can stop the spread of an attack's disruption to payment flows and possible impairment of financial stability. In addition, banks, whether small or large, may not be able to borrow in the fed funds market to cope with a cyberattack. To borrow, they need to have borrowing relationships established in advance; to have sufficient capacity to borrow, given the borrowing limits their counterparties have set; and to have counterparties willing to lend to them. If knowledge of a cyberattack becomes public, counterparties' willingness to lend to affected financial institutions could quickly evaporate. The Federal Reserve can offset this coordination failure in the fed funds market and restore payment flows by lending through its discount window facility. Our results thus reinforce the call in McAndrews and Potter (2002) for the Federal Reserve to serve as a coordinator of last resort (COLR) for payments. While the risk of moral hazard from the Federal Reserve serving as a COLR cannot be ruled out, discount window borrowing is costly for banks, which should somewhat mitigate the risk.

2. Description of the Cyber Event

Financial firms rely on a wide range of third-party TSPs to operate digitally. As evidenced by the cyberattacks that exploited vulnerabilities at SolarWinds (McMillan, 2021) and Microsoft's Exchange Servers (Volz and McMillan, 2021), a cyber event at a TSP, whether an intentional attack or not, can affect the TSP's customers, either through a service outage or by compromising their own IT systems.

The cyber event studied in this paper effectively started when the TSP discovered evidence of an attack on its computer network. To contain the attack, the TSP disconnected affected servers from the internet, and in doing so, some banks lost the ability to access Fedwire in their usual and

[11] We also contribute to the literature on central banks as lenders of last resort (e.g. Rochet and Vives, 2004), which documented both the benefits (van Bekkum, Gabarro and Irani, 2018; Carpinelli and Crosignani, 2021, Jasova, Mendicino and Supera; 2021) and costs (Drechsler et al., 2016; Crosignani, Faria-e-Castro and Fonseca, 2020; Jasova et al., 2021) of such interventions.

preferred way.[12] Figure 1 presents the size distribution of users, as measured by total assets, compared with non-users.[13] Users were not extremely large banks but were generally larger than non-users. U.S. global systemically important banks (G-SIBs), which were non-users, are excluded from this chart—because they account for a very large share of payments sent over Fedwire (Eisenbach, Kovner, and Lee, 2021).[14] Our empirical analysis of the cyberattack's effect on user banks' ability to send payments is conducted with and without the G-SIBs included, and the findings are unchanged when they are included. However, their inclusion in Figure 1 would skew the graph of non-users to the right. We do include the G-SIBs as non-users in our analysis of spillover effects from users' difficulty sending payments.

The service disruption resulted in many payments that users would have sent over Fedwire being delayed. The number and value of payments sent by users over Fedwire fell significantly during the cyberattack (Figures 2 and 3, respectively).[15] Especially on the first day of the cyberattack, when the impact was most severe, users sent 81% fewer payments compared with the same day the week before, which resulted in a drop of 72% in the value of payments. It is noteworthy that the drop is neither 100% nor 0%. We cannot observe whether users shifted to alternative methods of sending payments through Fedwire, but had they not, the drop would have been 100%; had they shifted quickly and completely, there would have been no drop. Hence, so we infer that users did shift gradually and imperfectly to other methods. The drop on the first day is very similar quantitatively whether we compare it with the number and value of payments sent on the same day two weeks, one month, and even one year before the actual cyberattack. Many factors can drive changes in payments sent, so our empirical analysis identifies the fraction that can be attributed to the cyberattack.

[12] As stated above, the cyberattack may have occurred earlier and not immediately have been discovered, or not been realized to require such action. The time at which the TSP took its servers offline, creating the service outage, is confidential.

[13] To protect the confidentiality of the cyberattack, as well as the identities of the banks that were impacted by it, we remove the units of measurement from all charts in the paper, unless we state otherwise.

[14] We also winsorize the upper 99th percentile of transactions in Fedwire to avoid having a few abnormally large payments shape our conclusions. None of the conclusions of this paper changes if we do not winsorize the payment flows in Fedwire.

[15] To protect the confidentiality of the event, we normalize the number and value of payments to 1. We also anonymize the duration of the event by referring only to a first and last day and the mid-period. This implies that the event lasted at least three business days. In these figures, as well as Appendix Figures 1 and 2, we average the values for the middle days and plot that average as the value over a single middle day. Thus, whether the mid-period lasted one day or 12, it appears in these figures as one day.

An alternative way to visualize the magnitude of the disruption is presented in Figure 4. The top chart plots the distribution of the number of payments by the users on the first day (line in red, $t = 0$) compared with the same day in previous and following weeks (lines in neutral colors). We observe that the distribution of the number of payments shifts to the left on the first day of the cyberattack, which is not the case either before or after the cyberattack. This suggests that the set of user banks sent fewer payments on the first day compared with the same day in previous and following weeks. The bottom chart plots the distribution of the value of payments, which also shifted left, showing that on the first day of the cyberattack (line in red, $t = 0$) user banks sent fewer high value payments.

Because Figure 4 captures all payments during the day, it understates the impact of the disruption by omitting the effect of the Federal Reserve's granting requests from users to extend Fedwire's trading hours. Figure 5 displays the distribution of the number and value of payments sent by hour. The normal deadline for a bank to initiate a payment is 6:00 p.m., with settlement occurring by 6:30 p.m. Some users that could not meet the deadline to send payments asked the Federal Reserve to extend the trading day, resulting in more payments settling after 6:30 p.m. Compared with payments sent on the same day in previous and following weeks (lines in neutral colors), the number of payments sent during the first day of the cyberattack was substantially lower overall (top chart). However, during that first day, as users started to switch to alternative ways to access Fedwire, the number of payments settling slightly increased during normal hours (i.e., before 6:30 p.m.). After 6:30 p.m., because extensions were granted, the number of payments sent was positive and gradually declined.[16]

The bottom chart in Figure 5 plots the hourly distribution of the log value of payments. Compared with the same day in previous and following weeks, the value of payments sent was lower overall, but increased during the day, consistent with users' switch to alternative ways to access Fedwire. There is a spike in the value of payments settling after 6:30 p.m., which, combined with the top panel, suggests that users sent larger, but not more, payments during the extension. This is consistent with the Federal Reserve's recommendations that banks prioritize critical

[16] See Armantier, Arnold and McAndrews (2008) for an analysis of the normal timing distribution of Fedwire payments.

payments during contingency situations.[17] A similar situation arose on September 11, 2001, following the terrorist attacks in New York: banks sent larger, but not more, payments after the normal close of the Fedwire trading day (McAndrews and Potter, 2002).

The result of this pattern in payments on the first day of the cyberattack was an increase in the average value of payments, which is important from a financial stability point of view. Figure 6 shows that the average value of payments sent was increasing during that first day before it spiked after 6:30 p.m. We corroborate these observations by plotting the share of payments gained from extending trading hours on the first and subsequent days in Figure 7. On the first day of the event, users sent 9.7% of their payments after 6:30 p.m., or approximately 13% more payments in terms of value. Extensions were also requested and granted on subsequent days of the event, but they accounted for a smaller share of payments sent, reflecting the gradual restoration of services by the TSP and users' adapting to alternative methods of sending Fedwire payments.

3. An Illustrative Example

Figure 8 illustrates hypothetical Fedwire payment flows before and after a cyberattack and provides guidance for our empirical exercise. Before the attack (top panel), Bank 1 sends payments of $100 each to Bank 2, a non-user of the TSP, and Bank 3, a user of the TSP. Both Bank 2 and Bank 3 receive those payments and need to send $100 to Bank 4, which also is a non-user. Bank 4, in turn, needs to send $200 to Bank 1. The TSP's response to the cyberattack disrupts Bank 3's ability to send payments, but not Bank 2's. As a result of the disruption (bottom panel), Bank 2 sends funds to Bank 4, but Bank 3 cannot. As a result, Bank 3's reserves grow from $50 to $150, while Bank 4 finds itself with only $150 in reserves ($100 received from Bank 2 and the $50 it originally held in reserves) to send to Bank 1. Bank 4, though a non-user and not directly affected by the TSP's response, faces a liquidity shortage of $50.

In practice, there are several ways Banks 3 and 4 can react. Bank 3 could switch to an alternative TSP if it had business arrangements already in place with them, use alternative methods for sending payments over Fedwire, ask the Federal Reserve to extend the Fedwire trading day to

[17] According to these recommendations, banks should "*be prepared to prioritize their offline transactions to those that the institution has identified as the most critical transactions, particularly later in the business day*." See Fedwire Offline Services as a Limited Contingency Tool (frbservices.org).

provide more time to send payments, or delay sending the payment until a later date.[18] If Bank 3 is delayed in sending its payment, Bank 4 might choose to borrow funds from the fed funds market or at the discount window to avoid incurring an overnight overdraft on its reserves account at the Federal Reserve. We explore these alternative responses in the following sections of the paper.

4. Data

Our paper brings together several confidential datasets, which we describe below.

List of Users of the TSP: We obtain confidential data on the list of user banks. These data are not publicly available and cannot be obtained by other proprietary or commercial datasets. Because of our access to the data, we are able to trace payment flows of the users (the treatment group) relative to the non-users (the control group).

Fedwire Funds Service: Fedwire is a real-time gross settlement (RTGS) system, where requests to send payments are processed and settled by the Federal Reserve after they are initiated by a bank. The dataset provides detailed information on the daily payment flows between a diverse set of financial institutions. Although banks are identified by their 9-digit ABA (American Bankers Association assigned) routing number, our level of analysis is the depository institution level as identified by its RSSD number. This is because regulations, including reserve requirements, are at the depository institution level.[19] We exclude settlement institutions, such as CHIPS, from our analysis.

Federal Funds: The federal funds market is an over-the-counter market where depository institutions negotiate inter-bank loans and their terms directly with each other. These loans are essentially uncollateralized loans of reserve balances held at the Federal Reserve. We identify fed funds transactions between lenders and borrowers using the Furfine algorithm (Furfine, 1999) and restrict our attention to loans extended by Federal Home Loan Banks (FHLBs), which are the main suppliers of fed funds to eligible depository institutions.[20] We cross-check the validity of our fed

[18] We do not observe whether users switched to alternative TSPs for sending Fedwire payments. However, if they did to a meaningful extent, then we would not have identified a statistically significant difference in their payment flows relative to the non-users.

[19] A similar aggregation at the depository-institution level can be found in Eisenbach, Kovner, and Lee (2021) and Copeland, Duffie, and Yang (2021).

[20] In a thorough analysis of the increased role of FHLBs in funding markets, Gissler and Narajabad (2017) discusses the importance of FHLBs for the provision of fed funds to depository institutions. On some days, FHLBs account for almost the entire supply of federal funds.

funds data with two sources: the universe of fed funds transactions as reported in confidential FR2420 forms,[21] as well as the 10K filings of FHLBs. By focusing on fed funds loans extended by FHLBs, we overcome the known type I and type II errors of the Furfine algorithm identified by Armantier and Copeland (2015).

Discount Window: Eligible depository institutions can post collateral and borrow funds from the Federal Reserve's discount window. These depository institutions can borrow from the discount window as long as they are illiquid but solvent and have taken the necessary steps in advance of setting up the systems and collateral-pledging processes to access the window. Our dataset provides information on depository institutions' daily borrowing from the discount window in the period around the cyberattack.

Other Datasets: We also exploit confidential Federal Reserve accounting records, which include end-of-day reserve balances banks hold in their accounts with the Federal Reserve. We match reserve balances from these records with balance sheet data obtained from bank Call Reports.

5. First-round Effect

5.1. Disruption of payments sent by users during the cyberattack

Because all Fedwire payments are initiated by the sender, the first-round effect of the cyberattack would be on a user bank's ability to send payments. We thus start our empirical analysis using a difference-in-differences model to study the impact of the TSP's service disruption on user banks' ability to send payments over Fedwire. We define a sender-bank as a bank that sends a payment over Fedwire, and a receiver-bank as a bank that receives a payment. Our variables of interest are the change in the number and value of payments sent by a sender-bank s to a receiver-bank r on a specific day t compared with the same day a week before in order to account for seasonality in payment flows (e.g., Treasury settlement days, which occur Thursdays, mid-month, and end of month). Since we are interested in the number of payments between a sender-bank and a receiver-bank, we aggregate Fedwire's transaction-level data at the sender-bank–receiver-bank–day level and count the number of transactions for each pair of banks on each

[21] These data are used by the Federal Reserve Bank of New York (FRBNY) to publish the daily volume in the fed funds market.

day. For the value of payments, we aggregate by taking the sum of the value of all transactions for each pair on each day. Our empirical models are:

$$\Delta Log(Number\ of Payments)_{srt} = \beta_1 \times Users_{\ s} \times Cyberattack_t + FE + \varepsilon_{srt},$$

$$\Delta Log(Value\ of Payments)_{srt} = \beta_1 \times Users_{\ s} \times Cyberattack_t + FE + \varepsilon_{srt},$$

where *Users* is a dummy variable that takes value one if a sender-bank was a user and is zero otherwise. *Cyberattack* is a dummy variable that takes value one during the period of the cyberattack and is zero otherwise. Summary statistics for these variables are presented in Appendix Table 2. Our model includes a set of fixed effects that we add progressively to isolate the impact of the cyberattack on payment flows and describe in detail below. The model allows for rather conservative standard errors that are two-way clustered at the sender-bank and day level.[22]

Table 1 reports the effect of the cyberattack on payments sent. The first three columns consider the number of payments sent. The coefficient in column 1 suggests that relative to non-users, users sent 14.1% fewer payments during the cyberattack. We obtain this estimate after controlling for unobserved differences across sender-banks (sender-bank fixed effects) and time-varying shocks that are common to all sender-banks (day fixed effects). In column 2, we expand the model by including a set of receiver-bank*day fixed effects to compare payments sent by users and non-users to the same receiver-bank on the same day, thereby tightening the identification of the impact of the cyberattack. Our estimates suggest that users sent 16.4% fewer payments relative to non-users. Finally, in column 3 we control for the interaction of a sender-bank's size, as measured by the log of assets, with the cyberattack dummy. This interaction term accounts for the fact that relatively large banks may have arrangements to switch more quickly to alternative methods to send payments during a cyberattack, so failing to control for the size channel may lead to biased estimates. Our estimates in column 3 suggest that size does not matter for a user bank's ability to send payments during the cyberattack.

Columns 4-6 consider the effect on the value of payments sent. The estimate from our preferred specification (column 5) suggests that the value of payments sent by users relative to non-users

[22] Our standard errors are not biased downward because there are at least 50 clusters in both the sender-bank and time dimensions (e.g., Bertrand, Duflo, and Mullainathan, 2004). Our results remain robust when we triple-cluster the standard errors at the sender-bank, receiver-bank, and day level.

dropped by 26.5% during the cyberattack. The magnitude of the effect is very similar to the one obtained when we allow for a less strict specification (column 4) and when we control for the size channel (column 6).

Taken together, our results show that users sent fewer and smaller payments as a result of the cyberattack. With payment flows disrupted, reserves would have accumulated in the Federal Reserve accounts of users (as illustrated in Figure 8). To explore this effect, we regress the change in the reserve balances of sender-banks on day *t* compared with the same day a week before on the *Users* dummy variable that we introduced earlier. The results, shown in Table 2, suggest that reserves of users were about 18% higher during the cyberattack (column 1), even after accounting for the size channel of the sender-banks (column 2).

5.2. Robustness tests

In this section, we perform a number of tests to check the robustness of our findings that, as a result of the cyberattack, users sent fewer and smaller payment and saw their reserves increase.

First, we check for pre- and post-trends in the number and value of payments sent by users and non-users. Our previous findings relied on the assumption that before the cyberattack and after the recovery from it, the number and value of payments sent by the two groups followed similar trends (a parallel trends assumption). However, as we stated above, the cyberattack may have occurred earlier and not immediately have been discovered, or not been realized to require such action. Although Figures 2 and 3 provide visual support for the parallel trends assumption, they are not a formal test for existing pre- and post-trends. We construct a *Pre-Cyberattack* dummy variable, which takes value one before the cyberattack and zero otherwise. We construct this variable to account for policy changes before the cyberattack that may have affected the number and value of payments sent over Fedwire. However, we cannot refer to the exact policy changes or dates to protect the confidentiality of the event. Similarly, we construct a *Post-Cyberattack* dummy variable that takes value one for all days after the resolution of the cyberattack and zero otherwise. Our results in Appendix Table 3 suggest that there are no trends in the number and value of payments sent by users and non-users either before or after the resolution of the cyberattack.[23] This supports our conclusions that the observed differences during the cyberattack can be attributed to the disruption of the TSP's service.

[23] These results do not depend on whether we control for the size channel.

Second, we consider a placebo cyberattack one year before the actual cyberattack.[24] For example, if the attack started on the second Wednesday of June, then the placebo event is assumed to start on the second Wednesday of June one year earlier. The purpose of this exercise is to ensure that the day the attack started is not special in any other aspect that could explain our results. Our estimates are presented in columns 1 and 2 of Appendix Table 4. We do not find a statistically significant difference in the number and value of payments sent by users and non-users a year before the actual cyberattack, which suggests that our baseline results indeed reflect the impact of the disruption of the TSP's services.

Third, we consider an alternative definition of our dependent variables by comparing changes in the number and value of payments sent on the same day one month before the actual cyberattack began (columns 3 and 4 of Appendix Table 4). The purpose of this exercise is to correct for potential seasonality in payments on a monthly basis. Our estimates suggest that the effect of the disruption is very similar to the one obtained in Table 1.

Finally, we rerun our empirical model by including the U.S. G-SIBs as sender-banks in our analysis. We present results on the number and value of payments sent in Appendix Table 5 and the accumulation of reserves in Appendix Table 6. The effect of the disruption with the U.S. G-SIBs included as sender-banks is very similar to the effect when they are excluded. Hence, our results are not sensitive to whether we include these very large financial institutions.

5.3. Analysis by day

An important feature of cyberattacks is the steps taken by affected firms to restore disrupted services and systems. In the unfortunate situation where a cyberattack disrupts operations for more than one day, the impact likely will be more severe on the first day and diminish thereafter as firms adapt to the event.

We explore the impact of the cyberattack on the flow of payments by day by introducing three new variables in our analysis. *First Day of Cyberattack* is a dummy variable that takes a value one on the first day of the event and zero otherwise. *Mid-Period of Cyberattack* is a dummy variable that takes a value one between the first and the last day of the cyberattack and zero otherwise. Similarly, *Last Day of Cyberattack* is a dummy variable that takes a value one on the last day of

[24] Appendix Figures 1 and 2 present graphical evidence of our empirical exercise, with the event's duration anonymized as done in Figures 2 and 3

the cyberattack and zero otherwise.[25] Each of these variables is then interacted with the dummy variable *Users*, the interaction of which captures the impact the cyberattack had on the number and value of payments sent by users on each day during this multi-day event.

Table 3 reports the day-by-day effect on the number and value of payments sent by users relative to non-users. We find that users sent 36% fewer payments on the first day of the attack (column 1), which resulted in a drop by 50.7% in the value of payments (column 2). These magnitudes are significantly larger compared with the 16.4% and 26.5% declines, respectively, that we documented in Table 1 for the entire period of the event. As expected, as the TSP gradually restored service and banks gained experience using alternative means of sending payments over Fedwire over subsequent days, the impact on the number and value of payments gradually decreased. These results suggest that the first day of the cyberattack was indeed the most severe, and the actions taken by the TSP and the affected user banks to restore normal operations also helped mitigate the cyberattack's impact on the payment system and financial stability.

As we discussed earlier, an important step taken by the Federal Reserve was to grant requests to extend the Fedwire trading day so banks could initiate payments after 6 p.m.[26] To evaluate the effect of the extensions on the number and value of payments, we repeat our analysis by excluding the payments settled after 6:30 p.m. We report these results in columns 3 and 4. On the first day of the cyberattack, the extension of the Fedwire trading day had little impact on the drop in the number of payments sent by the users. In contrast, the drop in the value of payments would have been 10 percentage points larger (close to 61%) had the Federal Reserve not extended the trading day. However, that effect did not persist: the extensions had little impact on the value of payments sent in the mid-period or on the last day of the event.

[25] The confidentiality of the cyberattack does not allow us to reveal the duration of the event. As such, we construct these variables in a way so that they do not reveal this information. To anonymize the length of the event, the first and last days of the cyberattack are simply referred to as such. To get the "mid-period," we define for each middle day our dependent variables as the log difference in payments (number or value) between a given day during the cyberattack and its values the week before the cyberattack. We then anonymize those values by taking the average over all the days of the mid-period of the attack. In other words, we take the average of the log difference. It is important to note that we do not take the difference of the log average as this would not properly adjust for the seasonality in payment flows. A similar point has also been made in the context of bilateral trade flows in the international trade literature (e.g. Baldwin and Taglioni, 2006). The conclusions of this paper do not depend on the anonymization of the event. All results hold when we analyze the event without anonymizing the duration.

[26] The latest a bank can ask for an extension is 15 minutes before the normal close of Fedwire.

In line with the graphical evidence in Figure 7, our results suggest that users sent larger value payments but not a greater number of payments after 6:30 p.m. We corroborate these findings with additional evidence in Table 4. We construct two dependent variables: one measures the ratio of the number of payments sent after 6:30 p.m. to the total number of payments sent during the day, and the other is similar, but replaces the number of payments with the value of payments. We regress these measures on the interaction of the *Users* dummy with the dummy variables that represent each day separately. We find that users sent 9.7% of their payments and 11.2% of the value of their payments after 6:30 p.m. on the first day of the attack. The effect of the extensions gradually declined over subsequent days.

5.4. Economic magnitude of the first-round effect

In this section, we estimate the economic magnitude of the first-round effect on payment flows on the first day of the cyberattack, when the disruption was most severe. Our analysis finds that relative to non-users, users sent 50.7% fewer payments on the first day of the event. We multiply our estimate by user banks' share of the total value of payments sent over Fedwire on the same day one week before the cyberattack (Appendix Table 1), which is 0.6%. This implies that the total value of payments dropped by 0.3% (= 0.507*0.006) on the first day, which would have been 0.4% had the Federal Reserve not extended the Fedwire trading day.

While the first-day drop in payments was small, as noted earlier, the G-SIBs were not users of the service, and Figure 1, which omits the G-SIBs, shows that there were some large non-G-SIB non-user banks compared with user banks. It is easy to imagine a similar cyberattack on a more dominant TSP or on more or larger banks directly that would have a larger impact. To get a sense of how large the impact might be, we consider the hypothetical scenario in which one of the top-5 most active U.S. G-SIBs in terms of trading through Fedwire is affected by the cyberattack. This scenario is in the spirit of Eisenbach, Kovner, and Lee (2021), but we use our estimates on the first-round drop in payments sent on the first day of the attack, along with the G-SIBs' market shares, to obtain our estimate of the average impact. We estimate that if the cyberattack had affected a single U.S. G-SIB, it would have resulted on average in 7% fewer payments sent—an impact 17.5 times larger than what we obtain for the actual banks involved. While the G-SIBs might be expected to have better contingency practices than other banks, they also are responsible

for dramatically more payments in number and volume than the user banks in this paper's case study.

5.5. Lessons for policy

An important lesson from the previous analysis is that the Federal Reserve's extension of the trading day allows banks to send larger, and presumably more critical, payments, which is important from a financial stability perspective. The pattern of payments observed in the cyberattack after 6:30 p.m. is similar to the pattern seen on September 11, 2001 in the aftermath of the terrorist attacks in New York (McAndrews and Potter, 2002). The Federal Reserve's ability to extend the trading day injects processing time into the payment system to relieve stress from operational disruptions, similar to the way its ability to inject liquidity into the banking system can relieve funding stress. It is akin to an automatic stabilizer in the way unemployment insurance is: just as an unemployed worker can apply for benefits, and the state's unemployment insurance office may or may not grant benefits, so too a bank can ask for an extension of the trading day via an established process, and the Federal Reserve may or may not grant the request.

6. Second-round Effect

Because banks rely heavily on incoming payments to provide the funds to send their own payments, a disruption in incoming payments, like the one we documented (the *first-round effect*), can leave banks on the receiving-end with a shortage of liquidity and unable to send their own payments *(*the *second-round effect)*. Unless receiver-banks have sufficient reserves or are able to tap funding from alternative sources, the effect of the cyberattack could propagate to yet other banks and parts of the financial system that were not directly exposed to the cyberattack (*third-round effect*).[27] We explore these issues in this section.

6.1. Incoming payments of exposed receiver-banks

Banks on the receiving end of payments from user sender-banks can either be users or non-users. A receiver-bank that is a user would still receive payments but could not send payments, akin to the first-round effect described above. In this section, we consider how the cyberattack can

[27] Schreft and Zhang (2018) illustrates the scope for an operational disruption, including one from a cyberattack, to spread through the complex networks that make up the financial system, generating multiple rounds of disruption.

spill over to receiver-banks who were innocent bystanders. As such, we focus on receiver-banks who were non-users themselves and analyze the drop in their incoming payments as a result of their indirect exposure to the cyberattack through sender-banks in Fedwire. To this end, we construct a new variable that we call Exposed Receiver-bank and define it as the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. For example, if a receiver-bank was receiving a total of $100 over the days before the cyberattack, of which $20 was from user sender-banks and $80 was from non-user sender-banks, the exposure to the shock of the receiver-bank would be 20% (= 0.2*1 + 0.8*0).[28] Intuitively, higher values of this measure indicate that a receiver-bank is at greater exposure to the cyberattack because it had more incoming payments from users in the period before the attack. We then regress the log change of incoming payments of a receiver-bank *r* on day *t* compared with the same day a week earlier on this measure of indirect exposure, interacted with dummies for each day of the attack separately. Our empirical model is:

$$\Delta \log (Payments)_{rt} = \beta_1 \times exposed\ receiver - bank_r \times Day\ Dummies_t + FE + \varepsilon_{rt}.$$

Our model accounts for all time-invariant observed and unobserved heterogeneity among receiver-banks (receiver-bank fixed effects) and time-varying shocks that are common to all receiver-banks (day fixed effects). Finally, we allow for rather conservative standard errors, which are two-way clustered at the receiver-bank and day level.[29]

We present the results of this regression in Table 5. Our estimates in column 1 imply that a one standard deviation increase in the exposure of a receiver-bank to the cyberattack (0.182) is associated with a reduction in the bank's incoming payments by approximately 13% on the first day of the event. This number drops to 10% in the mid-period, which is in line with the gradual restoration of service by the TSP. On the last day of the disruption, the estimated effect, although

[28] We have look-ahead and look-back windows we use in our analysis. To ensure anonymity of the dates and length of the event, we cannot state the lengths of those windows, but these windows are of sufficient length for our empirical analysis.

[29] As we discussed above, our standard errors are not biased downwards because the number of clusters is above 50 in both dimensions (for receiver-bank and day).

negative, is not statistically significant. We expand our model in column 2 by controlling for the size channel of the receiver-banks, but this does not materially affect our estimates.

6.2. Alternative sources of funding

Our results suggest that the incoming payments of receiver-banks fell during the period of the cyberattack, with the drop being more pronounced on the first day than on subsequent days. We next explore to what extent receiver-banks responded by obtaining funds from alternative sources. Typically, larger banks are more likely to access the fed funds market (e.g., Ho and Saunders, 1985), while smaller banks are more likely to access the discount window. This may also depend on the ex-ante share of reserves a bank holds in its account at the Federal Reserve (e.g., Ennis and Klee, 2021). For instance, Duffie and Younger (2019) shows that twelve systemically important U.S. banks have enough liquid assets to withstand the wholesale funding runs that might occur in a severe cyber event.

Table 6 shows our findings regarding whether banks accessed the discount window. We construct a dummy variable that takes value one if a receiver-bank borrowed from the discount window at time *t*, conditional on no past use at time *t-1*.[30] We then regress this dummy variable on the interaction of our exposure measure with the dummy variables that represent each day of the event separately. In column 1, we consider all banks in our sample and find no evidence of borrowing from the discount window. However, when we split our sample between large (above median) and small (below median) banks, we find that larger banks were less likely to borrow from the discount window, while smaller banks were more likely to do so. Economically, a one standard deviation (0.206) increase in the exposure of small receiver-banks is associated with an increase by 0.47% in their chance to tap the discount window for funding on the first day of the cyberattack (column 3). If we further restrict our attention to those small entities that did not access the fed funds market that day, we find that the probability of their borrowing from the discount

[30] In contrast to the construction of the dependent variables above, the construction of a dummy variable in Table 6 is trickier when we anonymize the length of the mid-period of the event. As we already discussed, all variables are constructed before that anonymization process (e.g., log differences). Then, we take the average of those dependent variables during the mid-period of the cyberattack (i.e., the average of the log differences). However, in the case of a dummy variable that may take values of zero, one, and zero during the three days of a hypothetical three-day "mid-period" event, the average will not be zero or one anymore, and the dummy variable ceases to be a dummy. In this case, we assign the dummy to be equal to one if a bank accessed the discount window in any one of the days during the mid-period of the cyberattack. Our conclusions do not depend on whether we anonymize the event.

window that first day increased by 0.51% (column 4). In the mid-period of the attack, the probability of obtaining funding from the discount window remained positive but dropped in terms of economic significance.

We study the role of liquidity buffers (columns 6 and 7) by analyzing a receiver-bank's ex-ante share of reserves to its total assets. We split the set of small banks that did not access the fed funds market into those with higher and lower reserves relative to assets based on the median bank in our distribution. As expected, we find that small banks with fewer reserves were much more likely to tap the discount window for funding. A one standard deviation (0.179) increase in a receiver-bank's exposure is associated with a 0.76% increase in the probability of discount window borrowing on the first day of the event, when the disruption was most severe. The cyber event had a positive and statistically significant impact on discount window borrowing by small banks with relatively high reserves as well, but the magnitude of the effect is relatively smaller.[31]

Having documented how small banks reacted in response to a drop in their incoming payments, we now turn our attention to large banks. It is generally true that larger banks access the fed funds market, where they predominately borrow from FHLBs. In column 1 of Table 7, we regress the log of fed funds borrowing on our exposure measure for all large banks interacted with the day dummies separately. The coefficients for the first day and mid-period of the cyberattack are positive but not statistically significant. We split this set of banks into two groups, those that are relatively smaller (below median size) and relatively larger (above median size). Our estimates in columns 2 and 3 suggest that smaller-than-median large banks borrowed more from the fed funds market on the first day and mid-period of the event. In contrast, larger-than-median large banks borrowed less from the fed funds market on the first and subsequent days of the event. To better understand the reaction of those banks, we further split the larger-than-median large banks into two groups: those with high and those with low reserves as a share of their total assets.[32] We find that on the first day of the cyberattack, relatively larger banks with relatively higher (lower) reserves borrowed less (more) from the fed funds market (columns 4 and 5, respectively). Table 8 shows that the large banks with relatively high reserves drew down their reserves on the first and

[31] Banks could have accessed the discount window for reasons not necessarily related to the cyberattack. We check for this and find no evidence of pre-trends in discount window borrowing. Results are available upon request.

[32] For the relatively large banks, we use the average and not the median bank of our distribution to split the sample into more and less liquid banks because using the median left no meaningful variation.

subsequent days of the disruption, both compared to the previous week and day (columns 1 and 2, respectively). This result is consistent with the finding of Duffie and Younger (2019), which documents the important role of reserves during a hypothetical cyber event affecting twelve systematically important financial institutions. These results support the importance of liquidity buffers in mitigating the impact of a cyberattack on payment flows.

7. Third-round effect from spillovers from exposed receiver-banks

The results of the previous section suggest that receiver-banks across the size spectrum (from small and less liquid banks to large and highly liquid banks) reacted differently in response to the drop in payments received because of the cyberattack. We next consider whether their reactions were sufficient to address their liquidity shortfall and avoid a decline in their own ability to send payments, with the risk of accompanying spillover effects.

We explore this issue in Table 9, where we regress the log change of receiver-banks' outgoing payments on the interaction of their exposure measure with dummies for each phase of the event separately. Starting from the baseline specification in column 1, our results suggest that there is no meaningful third-round impact on payments sent, especially on the first day of the cyberattack, when the disruption was most severe. In other words, receiver-banks sufficiently replenished the shortfall in their incoming payments with alternative sources of funding, or by drawing on their own reserves, and were able to send their own payments. This result holds when we control for the impact of receiver-bank size (column 2), as well as when we test for pre- and post-trends in the outgoing payments of receiver-banks (columns 3 and 4). Taken together, we find no third-round effect from the cyberattack. The Federal Reserve's support and the actions taken by sender and receiver banks were sufficient to prevent further cascading of the event to yet other banks or other parts of the financial system.

8. Conclusion

We conduct a unique natural experiment using confidential data to understand how an actual cyberattack on a TSP affected payment flows and bank liquidity. We quantify the many dimensions of the attack, including the spread of the event from users of the TSP to non-users. This quantification provides insight into how much worse the impact could have been had a more dominant TSP or more or even larger banks been affected.

We document that the first-round effect of the cyberattack was on users. Those users found their ability to send payments impaired, even though they switched to other methods of sending payments and the Federal Reserve extended Fedwire operating hours for them. Their inability to send payments spilled over to banks that did not use the TSP. Those innocent bystanders received fewer incoming payments from the user banks, leaving them with fewer reserves with which to send payments themselves. We quantify the extent of this contagion—the second-round effect—and the heterogeneous reactions of those bystanders. Those with more ample reserves drew down their reserves. Of the banks with fewer reserves, the smaller ones borrowed from the discount window, while the larger ones borrowed from the fed funds market. With these arguably aggressive actions, the bystanders avoided the attack's impairing their own ability to send payments, which could have resulted in a third-round effect from spillovers to yet other bystanders. These actions were not costless. The bystander banks either incurred costs from holding more reserves in advance or costs of borrowing funds during the event.

Three policy lessons stand out. First, contingency planning matters. Banks had, and likely used, alternative processes to access Fedwire; otherwise, the impact of the cyberattack would have been more pronounced. However, they did not switch to them quickly enough to avoid a material drop in the number and volume of payments sent during the event. Second, liquidity buffers matter. Banks that had sufficient reserves drew down those reserves to send payments themselves. Third, Federal Reserve support matters. The Federal Reserve extended both processing time for Fedwire payments and liquidity through discount window loans to mitigate the effect of the cyberattack. The fact that the Federal Reserve extended liquidity supports the conclusion that a cyber event can also be a liquidity event. The ability to extend processing time for Fedwire payments helps prevent the first-round effect from generating second-round effects. The Federal Reserve's ability to extend discount window liquidity helps prevent the second-round effect from spilling over into a third-round effect and broader financial instability. Cyberattacks are the new normal, and these findings highlight the important role for bank contingency planning and for the Federal Reserve in supporting the financial system's recovery from them.

References

Afonso, Gara, Anna Kovner, and Antoinette Schoar (2011). "Stressed, not frozen: The federal funds market in the financial crisis," *Journal of Finance*, vol. 66, no. 4 (August), pp. 1109-1139.

Afonso, Gara and Ricardo Lagos (2015). "Trade dynamics in the market for federal funds," *Econometrica*, vol. 83, no. 1 (January), pp. 263–313.

Afonso, Gara and Hyun Shin (2011). "Precautionary demand and liquidity in payment systems," *Journal of Money, Credit, and Banking*, vol. 43, no. 2 (October), pp. 589-618.

Aldasoro, Iñaki, Leonardo Gambacorta, Paolo Giudici, and Thomas Leach (2020). "The drivers of cyber risk," BIS Working Paper 865. Basel: Bank for International Settlements, May.

Armantier, Oliver, Jeffrey Arnold, and James McAndrews (2008). "Changes in the timing distribution of Fedwire funds transfers," Federal Reserve Bank of New York, *Economic Policy Review*, vol. 14, no. 2 (September), pp. 83-112.

Armantier, Oliver, and Adam Copeland (2015). "Challenges in identifying interbank loans," Federal Reserve Bank of New York, *Economic Policy Review*, vol. 21, no, 1 (October), pp. 1-17.

Ashcraft, Adam, and Darrell Duffie (2007). "Systemic illiquidity in the federal funds market," *AEA Papers and Proceedings*, vol. 97, no. 2 (May), pp. 221-25.

Ashcraft, Adam, James McAndrews, and David Skeie (2009). "Precautionary reserves and the interbank market," Staff Report Number 370. New York: Federal Reserve Bank of New York, May.

Bagehot, Walter (1873). *Lombard Street*. London: H.S. King.

Baldwin, Richard, and Daria Taglioni (2006). "Gravity for dummies and dummies for gravity equations," NBER Working Paper 12516. Cambridge, Mass: National Bureau of Economic Research, September.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004). "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics*, vol. 119, no. 1 (February), pp. 249-275.

Bhattacharya, Sudipto, and Douglas Gale (1987). "Preference shocks, liquidity and central bank policy." W. Barnett and K. Singleton K., eds., *New Approaches to Monetary Economics*, chapter 4, pp. 69–88. Cambridge: Cambridge University Press.

Bianchi, Javier, and Saki Bigio (2021). "Banks, liquidity management and monetary policy," *Econometrica*, forthcoming.

Board of Governors of the Federal Reserve System (2021). "Fedwire funds services," webpage, accessed September 1.

Boer, M., and J. Vasquez (2017). "Cyber security and financial stability: How cyber-attacks could materially impact the global financial system," Institute of International Finance, online publication, September.

Bouveret, Antoine (2019). "Estimation of losses due to cyber risk for financial institutions," *Journal of Operational Risk*, vol. 14 no. 2 (June), pp. 1-20.

Carpinelli, Luisa, and Matteo Crosigani (2021). "The design and transmission of central bank liquidity provisions," *Journal of Financial Economics*, vol. 141, no. 1 (July), pp. 27-47.

CBS News (2019). "Full 60 Minutes interview with Fed Chair Jerome Powell," *60 Minutes*, online publication, March 10.

CBS News (2021). "Fed Chair Jerome Powell tells 60 Minutes America is going back to work," *60 Minutes*, online publication, April 12.

Clouse, James A, Dale Henderson, Athanasios Orphanides, David Small, and Peter Tinsley (2000). "Monetary policy when the nominal short-term interest rate is zero," Finance and Economics Discussion Serices 2000-51. Washington, DC: Board of Governors of the Federal Reserve System, November.

Copeland, Adam, Darrell Duffie, and Yilin Yang (2021). "Reserves were not so ample after all," NBER Working Paper 29090. Cambridge, Mass: National Bureau of Economic Research, July.

Council of Economic Advisors (2018). "The cost of malicious cyber activity to the U.S. economy," Executive Office of the President, online publication, February.

Crosignani, Matteo, Marco Macchiavelli, Andre F. Silva (2021). "Pirates without borders: The propagation of cyberattacks through firms' supply chains," Staff Report No. 937. New York: Federal Reserve Bank of New York, May.

Crosignani, Matteo, Miguel Faria-e-Castro, and Luís Fonseca, (2020). "The (unintended?) consequences of the largest liquidity injection ever," *Journal of Monetary Economics*, vol. 112 (June), pp. 97-112.

Denbee, Edward, Christian Julliard, Ye Li, and Kathy Yuan (2021). "Network risk and key players: A structural analysis of interbank liquidity," *Journal of Financial Economics*, vol. 141 (September), pp. 831-859.

Drechsler, Itamar, Thomas Drechsel, David Marquez-Ibanez, and Philipp Schnabl (2016). "Who borrows from the lender of last resort?" *The Journal of Finance* vol. 71 (October), pp. 1933–1974.

Duffie, Darrell, and Joshua Younger (2019). "Cyber runs," Hutchins Center on Fiscal & Monetary Policy Working Paper #51. Washington, DC: Brookings Institution, June.

Eisenbach, Thomas, Anna Kovner, and Michael Junho Lee (2021). "Cyber risk and the U.S. financial system: A pre-mortem analysis," *Journal of Financial Economics*, forthcoming.

Ennis, Huberto M., and David A. Price (2015). "Discount window lending: Policy trade-offs and the 1985 BoNY computer failure," Economic Brief 15-05. Richmond, Vir.: Federal Reserve Bank of Richmond, May.

Ennis, Huberto, and Elizabeth Klee (2021). "The Fed's discount window in 'normal' times," Finance and Economics Discussion Series 2021-016. Washington: Board of Governors of the Federal Reserve System, March.

European Systemic Risk Board (ESRB) (2020). "Systemic cyber risk," online publication, February.

FRBServices.org (2021). "FedLine® solutions," webpage, accessed September 1, 2021.

FRBServices.org (2021). "Fedwire® offline services as a limited contingency tool," webpage, accessed September 1, 2021.

Fischer, Stanley (2016). "The lender of last resort function in the United States," Speech at the Committee on Capital Markets Regulation, February 10.

Florackis, Chris, Christodoulos Louca, Roni Michaely, and Michael Weber (2020). "Cybersecurity risk," Becker Friedman Institute for Economics Working Paper 2020-178. Chicago: University of Chicago, December.

Frost, Peter (1971). "Banks' demand for excess reserves," *Journal of Political Economy*, vol. 79 (July-August), pp. 805–825.

Goodfriend, Marvin and Robert G. King (1988). "Financial deregulation, monetary policy, and central banking." Federal Reserve Bank of Richmond, *Economic Review*, vol. 88, no 1 (May/June), pp. 3–22.

Gogolin, Fabian, Ivan Lim, and Francesco Vallascas (2021). "Cyberattacks on small banks and the impact on local banking markets," online publication, April 8.

Healey, Jason, Patricia Mosser, Katheryn Rosen, and Alexander Wortman (2018). "The ties that bind: A framework to assess the linkage between cyber risks and financial stability,"

Columbia School of International and Public Affairs, Project on Cyber Risk to Financial Stability Working Paper. New York: Columbia University, December.

Ho, Thomas, and Anthony Saunders (1985). "A micro model of the federal funds market," *Journal of Finance*, vol. 40, no. 3 (July), pp. 977-988.

Jasova, Martina, Caterina Mendicino, and Dominik Supera (2021). "Policy uncertainty, lender of last resort and the real economy," *Journal of Monetary Economics* forthcoming.

Jasova, Martina, Luc Laeven, Caterina Mendicino, José-Luis Peydró and Dominik Supera, (2021). "Systemic risk and monetary policy: The haircut gap channel of the lender of last resort," online publication, June 1.

Kamiya, Sinichi, Jun-Koo Kang, Jungmin Kim, Andreas Milidonis, and René M. Stulz (2021). "Risk management, firm reputation, and the impact of successful cyberattacks on target firms," *Journal of Financial Economics*, vol. 139 no. 3 (March), pp. 719-749.

Kashyap, Anil K., and Anne Wetherilt (2019). "Some principles for regulating cyber risk," *AEA Papers and Proceedings*, vol. 109, no. 3 (May), pp. 482-487.

Klee, Elizabeth (2010). "Operational outages and aggregate uncertainty in the federal funds market," *Journal of Banking & Finance*, vol. 34, no. 10 (October), pp. 2386–2402.

Kopp, Emanuel, Lincoln Kaffenberger, and Christopher Wilson (2017). "Cyber risk, market failures, and financial stability," IMF Working Paper 17/185. Washington, DC: International Monetary Fund, August.

Lacker, Jeffrey M. (2003). "Payment system disruptions and the Federal Reserve following September 11, 2001," Working Paper 03-16. Richmond, Vir.: Federal Reserve Bank of Richmond, December 23.

Lagarde, Christine (2018). "Estimating cyber risk for the financial sector," Linkedin, online publication, accessed September 1, 2021.

McAndrews, James J., and Simon Potter (2002). "Liquidity effects of the events of September 11, 2001," Federal Reserve Bank of New York, *Economic Policy Review*, vol. 8 no. 2, November), pp. 59-79.

McAndrews, James and Samira Rajan (2000). "The timing and funding of Fedwire funds transfers," Federal Reserve Bank of New York, *Economic Policy Review*, vol. 6, no. 2 (July), pp. 17-32.

McMillan, Robert (2021). "Hackers lurked in SolarWinds email system for at least 9 months, CEO says," *The Wall Street Journal*, February 2.

Office of Financial Research (2017). "Cybersecurity and financial stability: Risks and resilience," OFR Viewpoint 17-01. Washington, DC: Office of Financial Research, February 15.

Piazzesi, Monika, and Martin Schneider, (2021). "Payments, credit and asset prices," online publication, May.

Poole, William. (1968). "Commercial bank reserve management in a stochastic model: Implications for monetary policy," *Journal of Finance*, vol. 23, no. 5 (December), pp. 769–791.

Rochet, Jean-Charles, and Xavier Vives (2004). "Coordination failures and the lender of last resort: Was Bagehot right after all?" *Journal of the European Economic Association*, vol. 2, no. 6 (December), pp. 1116–1147.

Ros, Greg (2020). "The making of a cyber crash: A conceptual model for systemic risk in the financial sector," Occasional Paper No. 16. Frankfurt: European Systemic Risk Board, May.

Schreft, Stacey, and Simpson Zhang (2018). "Network analysis: Defending financial stability by design," OFR Brief Series 18-02. Washington, DC: Office of Financial Research, August 1.

Soramäki, Kimmo, Morten L. Bech, Jeffrey Arnold, Robert J. Glass, and Walter E. Beyeler (2007). "The topology of interbank payment flows," *Physica A: Statistical Mechanics and its Applications*, vol. 379, no. 1 (June), pp. 589-619.

van Bekkum, Sjoerd, Marc Gabarro, and Rustom M. Irani (2018). "Does a larger menu increase appetite? Collateral eligibility and bank risk-taking," *The Review of Financial Studies*, vol. 31, no. 3 (October), pp. 943-979.

Volz, Dustin, and Robert McMillan (2021). "Massive hacks linked to Russia, China exploited U.S. internet security gap," *The Wall Street Journal*, March 10.

Warren, Phil, Kim Kaivanto, and Dan Prince (2018). "Could a cyber attack cause a systemic impact in the financial sector?" Quarterly Bulletin, 2018 Q4. London: Bank of England, Fourth Quarter.

Wellburn, Jonathan William, and Aaron Strong (2019). "Systemic cyber risk and aggregate impacts," Institute for Civil Justice WR-1311. Santa Monica: RAND Institute, September.

Figure 1: Size of Users and Non-Users of the TSP

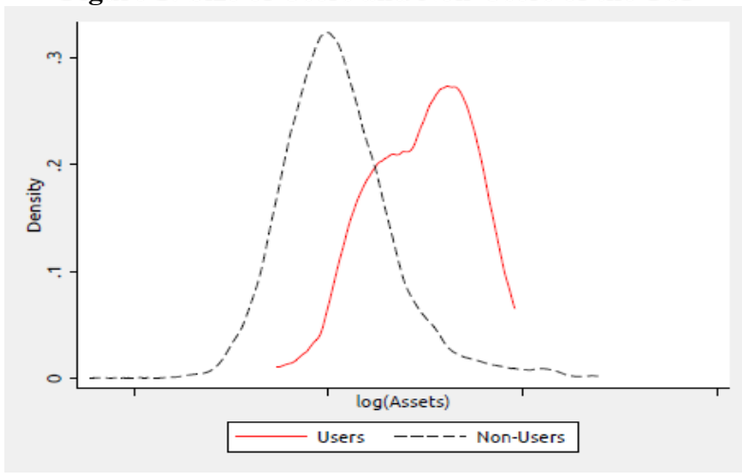

Note: The chart plots the size distribution measured by assets of users and non-users of the TSP that send payments over Fedwire. The U.S. G-SIBs were non-users and are excluded in the chart.

Figure 2: Number of Payments

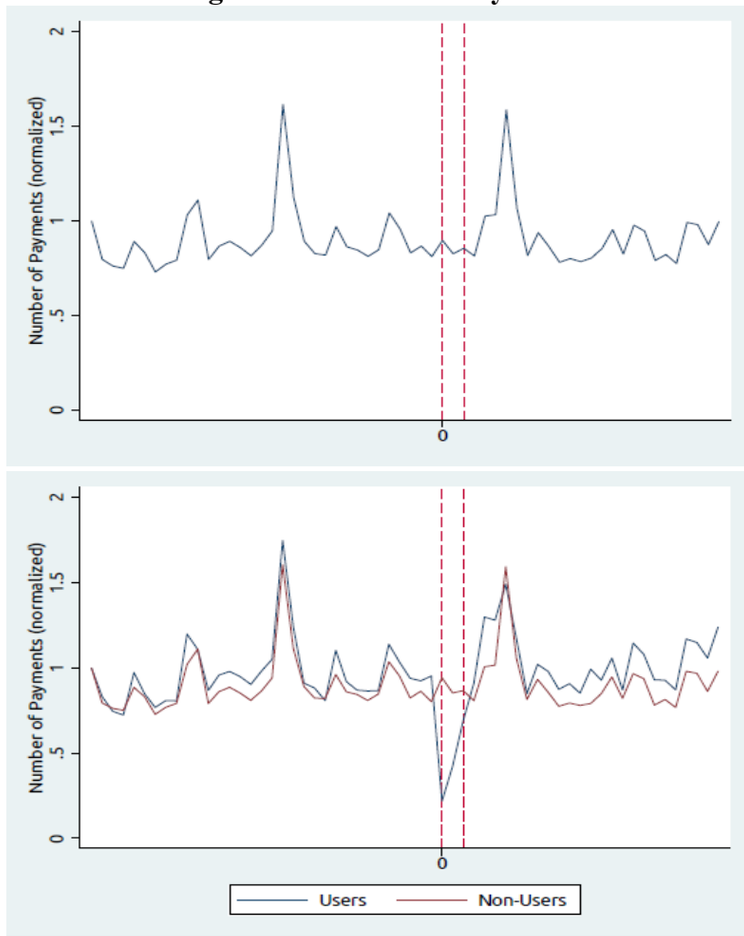

Note: The charts plot the number of payments in total (top) and by type of bank (bottom) before and after the cyberattack. The red vertical dashed lines correspond to the first day of the cyberattack and the last day of the cyberattack. We anonymize the mid-period by averaging the values for the middle days and plot that average as the value over a single middle day.

Figure 3: Value of Payments

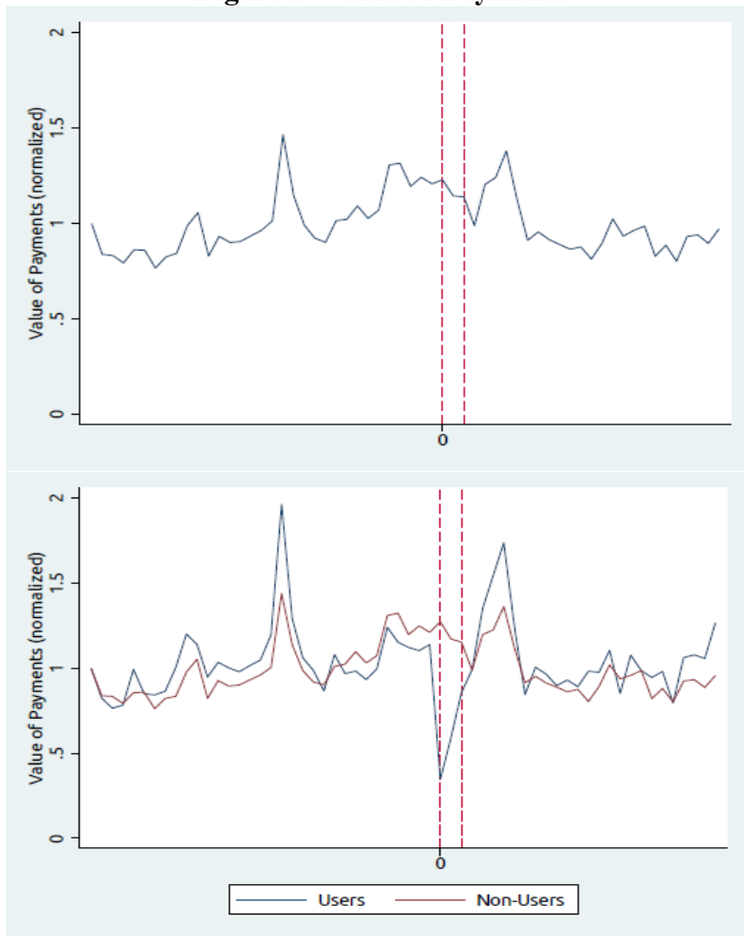

Note: The charts plot the value of payments in total (top) and by type of bank (bottom) before and after the cyberattack. The red vertical dashed lines correspond to the first day of the cyberattack and the last day of the cyberattack. We anonymize the mid-period by averaging the values for the middle days and plot that average as the value over a single middle day.

Figure 4: Distribution of Payments by Day

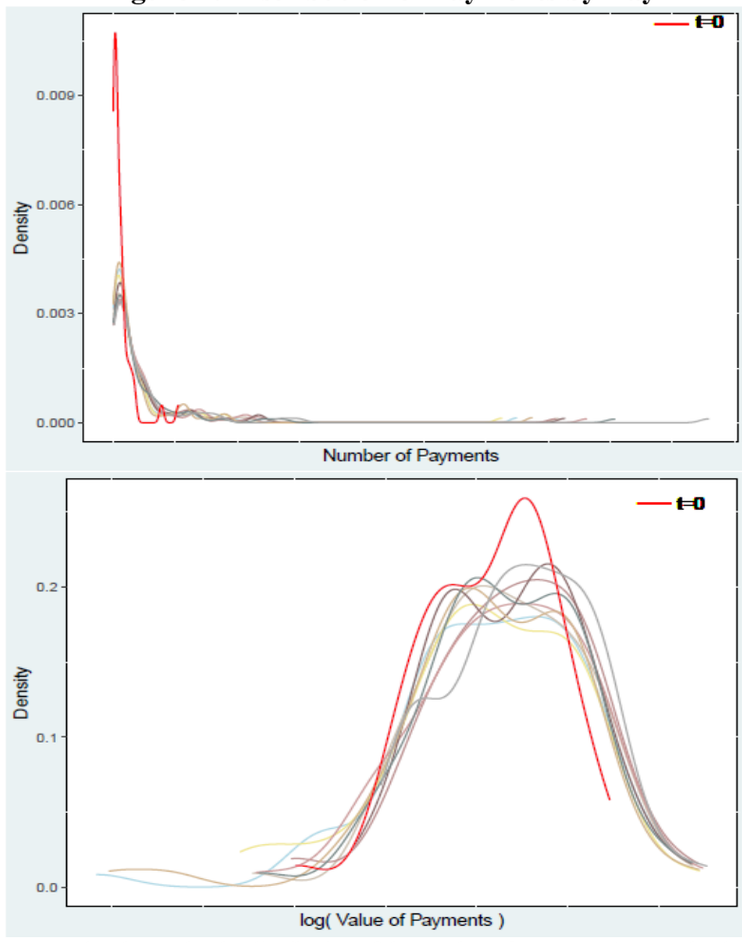

Note: The charts plot the distribution of the number of payments (top) and value of payments (bottom) by users of the TSP the first day of the cyberattack (t=0) and the same day in the previous and following weeks.

Figure 5: Distribution of Payments by Half-Hour Interval

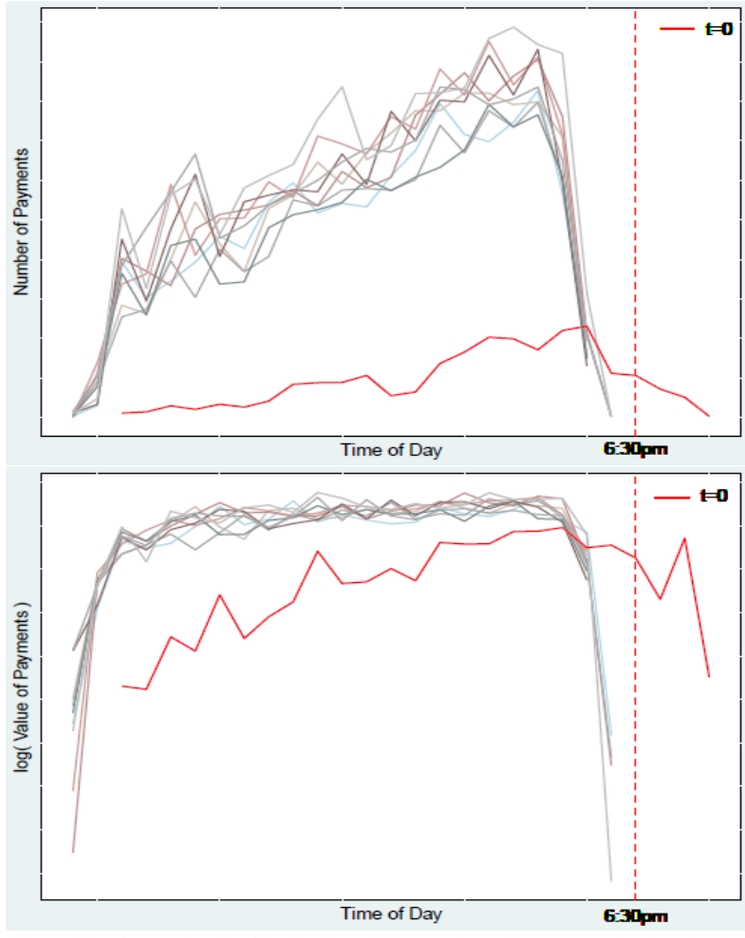

Note: The charts plot the distribution of the number of payments (top) and value of payments (bottom) by users of the TSP by half-hour interval on the first day of the cyberattack (t=0) and the same day in the previous and following weeks.

Figure 6: Average Value of Payments by Half-Hour Interval

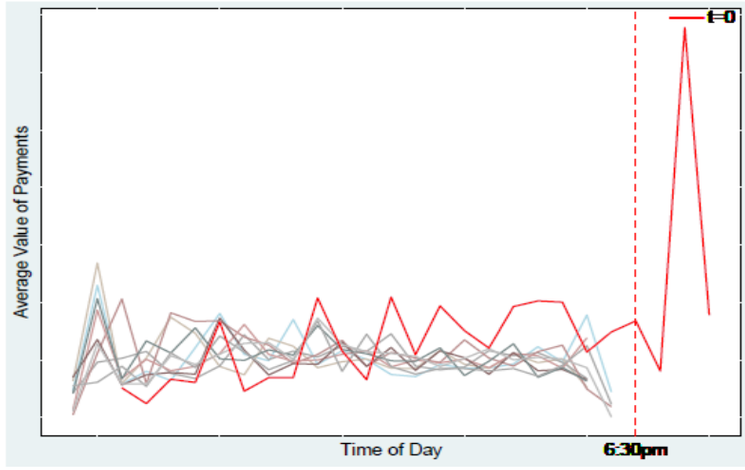

Note: The chart plots the distribution of the average value of payments by users of the TSP by half-hour interval on the first day of the cyberattack (t=0) and the same day in the previous and following weeks.

Figure 7: Share of Payments Gained from Extending Trading Hours

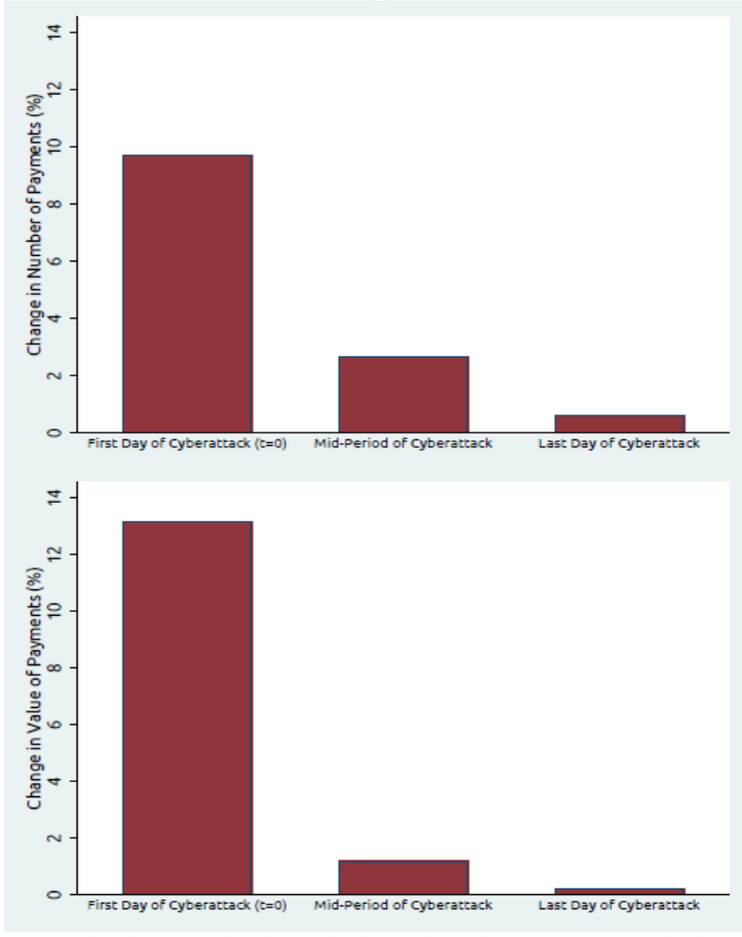

Note: The charts plot the change in the number of payments (top) and value of payments (bottom) gained from extending Fedwire trading hours for users of the TSP on the first, mid-period and last day of the cyberattack.

Figure 8: Illustration of How a Cyberattack Could Disrupt Payment Flows

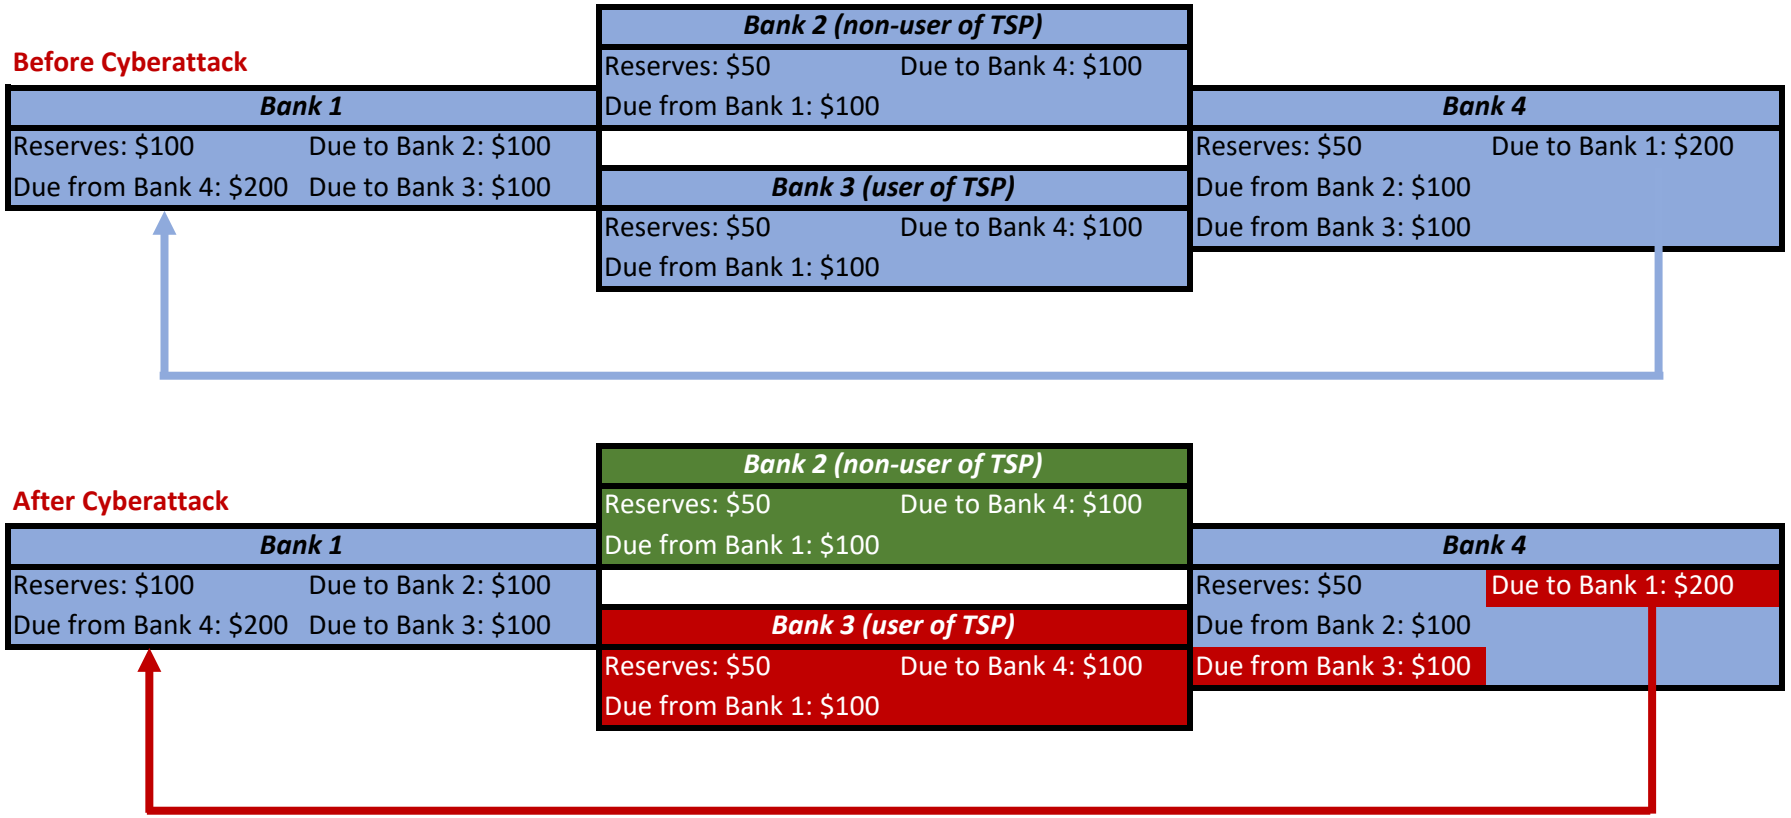

Note: The chart describes how a cyberattack can disrupt the payment system and provides a graphical representation of our empirical exercise.

Table 1: Effect of Cyberattack on Number and Value of Payments

	Δlog(Number of Payments)			Δlog(Value of Payments)		
	1	*2*	*3*	*4*	*5*	*6*
Users * Cyberattack	-0.141**	-0.164***	-0.165***	-0.233***	-0.265***	-0.265***
	(0.055)	(0.057)	(0.058)	(0.085)	(0.095)	(0.095)
Size * Cyberattack			-0.006			-0.001
			(0.005)			(0.006)
Sender-Bank FE	yes	yes	yes	yes	yes	yes
Day FE	yes	no	no	yes	no	no
Receiver-Bank x Day FE	no	yes	yes	no	yes	yes
Observations	562961	562961	562961	562961	562961	562961
R^2	0.022	0.104	0.104	0.008	0.103	0.103

Note: Δlog(Number of Payments) is the log change in the number of Fedwire payments compared with the previous week. Δlog(Value of Payments) is the log change in the value of Fedwire payments compared with the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Size is the log of assets of sender-banks. Cyberattack is a dummy variable that takes value one during the period of the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 2: Accumulation of Reserves by User Sender-Banks

	Δlog(Reserves)	
	1	*2*
Users * Cyberattack	0.177*	0.183*
	(0.102)	(0.103)
Size * Cyberattack		-0.004
		(0.008)
Sender-Bank FE	yes	yes
Day FE	yes	yes
Observations	72725	72725
R^2	0.043	0.043

Note: Δlog(Reserves) is the log change in the reserves of sender-banks compared with the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Size is the log of assets of sender-banks. Cyberattack is a dummy variable that takes value one during the period of the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 3: Analysis of Cyberattack by Day

	Including extensions in trading hours		*Excluding extensions in trading hours*	
	Δlog(Number of Payments)	Δlog(Value of Payments)	Δlog(Number of Payments)	Δlog(Value of Payments)
	1	*2*	*3*	*4*
Users * First Day of Cyberattack	-0.360***	-0.507***	-0.358***	-0.606***
	(0.025)	(0.101)	(0.022)	(0.097)
Users * Mid-Period of Cyberattack	-0.126***	-0.220***	-0.128***	-0.225***
	(0.039)	(0.079)	(0.037)	(0.074)
Users * Last Day of Cyberattack	-0.104***	-0.190***	-0.103***	-0.186***
	(0.038)	(0.069)	(0.037)	(0.065)
Sender-Bank FE	yes	yes	yes	yes
Receiver-Bank x Day FE	yes	yes	yes	yes
Observations	562961	562961	562671	562671
R^2	0.105	0.103	0.105	0.103

Note: Δlog(Number of Payments) is the log change in the number of Fedwire payments compared with the previous week. Δlog(Value of Payments) is the log change in the value of Fedwire payments compared with the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack and zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack and zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack and zero otherwise. Columns 1 and 2 include extensions in trading hours, while columns 3 and 4 exclude extensions in trading hours. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 4: Share of Payments Gained from Extending Trading Hours

	%Number of Payments gained during extension	%Value of Payments gained during extension
	1	*2*
Users * First Day of Cyberattack	0.097***	0.112***
	(0.012)	(0.014)
Users * Mid-Period of Cyberattack	0.027***	0.045***
	(0.010)	(0.012)
Users * Last Day of Cyberattack	0.007***	0.022***
	(0.002)	(0.003)
Sender-Bank FE	yes	yes
Day FE	yes	yes
Observations	100131	100131
R^2	0.075	0.106

Note: %Number of Payments gained during extension is the ratio of the number of payments after 6:31 p.m. to the number of payments during the day. %Value of Payments gained during extension is the ratio of the value of payments after 6:31 p.m. to the value of payments during the day. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack and zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack and zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 5: Incoming Payments of Exposed Receiver-Banks

	Δlog(Value of Payments)	
	1	*2*
Exposed Receiver-Bank* First Day of Cyberattack	-0.718***	-0.707***
	(0.122)	(0.124)
Exposed Receiver-Bank * Mid-Period of Cyberattack	-0.530***	-0.517***
	(0.092)	(0.094)
Exposed Receiver-Bank * Last Day of Cyberattack	-0.047	-0.035
	(0.113)	(0.116)
Size * Cyberattack		0.016
		(0.014)
Receiver-Bank FE	yes	yes
Day FE	yes	yes
Observations	58505	58505
R^2	0.030	0.030

Note: Δlog(Value of Payments) is the log change in the value of Fedwire payments compared to the previous week. Exposed Receiver-Bank is the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack, zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack, zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack, zero otherwise. Standard errors are two-way clustered at the receiver-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 6: Borrowing from the Discount Window

	$P(DW_t > 0 \mid DW_{t-1} = 0)$					
	All banks	*Large banks*	*Small banks*			
			accessed FF market	*did not access FF market*	*did not access FF market*	
					High Reserves/Assets	*Low Reserves/Assets*
	1	*2*	*3*	*4*	*5*	*6*
Exposed Receiver-Bank* First Day of Cyberattack	0.003	-0.023***	0.023***	0.025***	0.013***	0.042***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.008)
Exposed Receiver-Bank * Mid-Period of Cyberattack	0.002	-0.007	0.010	0.014**	0.015*	0.012
	(0.005)	(0.007)	(0.006)	(0.007)	(0.008)	(0.011)
Exposed Receiver-Bank * Last Day of Cyberattack	-0.019***	-0.039***	-0.004	-0.005	-0.012***	0.006
	(0.003)	(0.006)	(0.003)	(0.003)	(0.004)	(0.005)
Receiver-Bank FE	yes	yes	yes	yes	yes	yes
District x Day FE	yes	yes	yes	yes	yes	yes
Observations	58505	29217	29288	29031	16978	12053
R^2	0.067	0.071	0.093	0.094	0.077	0.141

Note: $P(DW_t > 0 \mid DW_{t-1} = 0)$ is the probability of discount window borrowing by a receiver-bank at time *t* conditional on no past use at time *t-1* . Exposed Receiver-Bank is the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. Large is a dummy variable that takes value one for banks above the median bank in terms of size. Small is a dummy variable that takes value one for banks below the median bank in terms of size. Accessed (did not access) FF market denotes a bank that with positive (zero) fed funds borrowing on the relevant days. High Reserves/Assets is a dummy variable that takes value one for banks above the median bank in terms of their reserves to assets ratio before the cyberattack. Low Reserves/Assets is a dummy variable that takes value one for banks below the median bank in terms of their reserves to assets ratio before the cyberattack. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack and zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack and zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack and zero otherwise. Standard errors are two-way clustered at the receiver-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 7: Borrowing from the Fed Funds Market

	log(Fed Funds)				
	All large banks	*Relatively smaller banks*	*Relatively larger banks*		
				High Reserves/Assets	*Low Reserves/Assets*
	1	*2*	*3*	*4*	*5*
Exposed Receiver-Bank* First Day of Cyberattacks	2.578	3.620**	-11.189*	-23.683***	16.473**
	(2.572)	(1.345)	(5.461)	(4.415)	(4.196)
Exposed Receiver-Bank * Mid-Period of Cyberattacks	1.833	3.238*	-2.638	-3.036	3.699
	(2.555)	(1.568)	(4.508)	(5.569)	(3.741)
Exposed Receiver-Bank * Last Day of Cyberattacks	-1.174	2.044	-5.755*	-6.452	-1.808
	(2.377)	(1.214)	(2.836)	(5.290)	(2.850)
Receiver-Bank FE	yes	yes	yes	yes	yes
Day FE	yes	yes	yes	yes	yes
Observations	463	261	201	82	102
R^2	0.876	0.924	0.752	0.775	0.807

Note: Log(Fed Funds) is the log fed funds borrowing. Exposed Receiver-Bank is the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. All large banks are banks that accessed the fed funds market. Relatively smaller banks is a dummy variable that takes value one for banks below the median bank within the set of banks that accessed the fed funds market. Relatively larger banks is a dummy variable that takes value one for banks above the median bank within the set of banks that accessed the fed funds market. High Reserves/Assets is a dummy variable that takes value one for banks above the average bank in terms of their reserves to assets ratio before the cyberattack. Low Reserves/Assets is a dummy variable that takes value one for banks below the average bank in terms of their reserves to assets ratio before the cyberattack. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack and zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack and zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack and zero otherwise. Standard errors are two-way clustered at the receiver-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 8: Large Bank Reliance on Their Reserves

	Δlog(Reserves)	log(Reserves)
	1	*2*
Exposed Receiver-Bank* First Day of Cyberattack	-17.549***	-17.657***
	(1.780)	(3.626)
Exposed Receiver-Bank * Mid-Period of Cyberattack	-3.503*	-7.404
	(1.714)	(5.435)
Exposed Receiver-Bank * Last Day of Cyberattack	-8.308***	-10.109
	(1.593)	(5.427)
Receiver-Bank FE	yes	yes
Day FE	yes	yes
Observations	82	82
R^2	0.694	0.875

Note: Δlog(Reserves) is the log change in the reserves of receiver-banks compared to the previous week. Log(Reserves) is the log reserves of receiver-banks. Exposed Receiver-Bank is the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack and zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack and zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack and zero otherwise. Standard errors are two-way clustered at the receiver-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Table 9: Outgoing Payments of Exposed Receiver-Banks

	Δlog(Value of Payments)			
	1	*2*	*3*	*4*
Exposed Receiver-Bank* First Day of Cyberattack	-0.153	-0.152	-0.161	-0.160
	(0.095)	(0.103)	(0.113)	(0.121)
Exposed Receiver-Bank * Mid-Period of Cyberattack	0.009	0.010	0.001	0.001
	(0.069)	(0.079)	(0.066)	(0.079)
Exposed Receiver-Bank * Last Day of Cyberattack	0.111	0.112	0.103	0.104
	(0.104)	(0.111)	(0.128)	(0.135)
Size * Cyberattack		0.000		0.000
		(0.010)		(0.010)
Exposed Receiver-Bank * Pre-Cyberattack			0.017	0.017
			(0.118)	(0.118)
Exposed Receiver-Bank * Post-Cyberattack			-0.025	-0.025
			(0.078)	(0.078)
Sender-Bank FE	yes	yes	yes	yes
Receiver-Bank x Day FE	yes	yes	yes	yes
Observations	320947	320947	320947	320947
R^2	0.120	0.120	0.120	0.120

Note: Δlog(Value of Payments) is the log change in the value of Fedwire payments compared to the previous week. Exposed Receiver-Bank is the weighted average of a receiver bank's incoming payments from sender-banks before the attack. The weights are the share of the receiver bank's total incoming payments sent by sender-banks, with user-senders' payments weighted by one and non-user-senders' payments weighted by zero. First Day of Cyberattack is a dummy variable that takes value one on the first day of the cyberattack, zero otherwise. Mid-Period of Cyberattack is a dummy variable that takes value one between the first and the last day of the cyberattack, zero otherwise. Last Day of Cyberattack is a dummy variable that takes value one on the last day of the cyberattack, zero otherwise. Pre-Cyberattack is a dummy variable that takes value one before the cyberattack, zero otherwise. Post-Cyberattack is a dummy variable that takes value one after the cyberattack, zero otherwise. Standard errors are two-way clustered at the receiver-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Appendix Figure 1: Number of Payments One Year Earlier

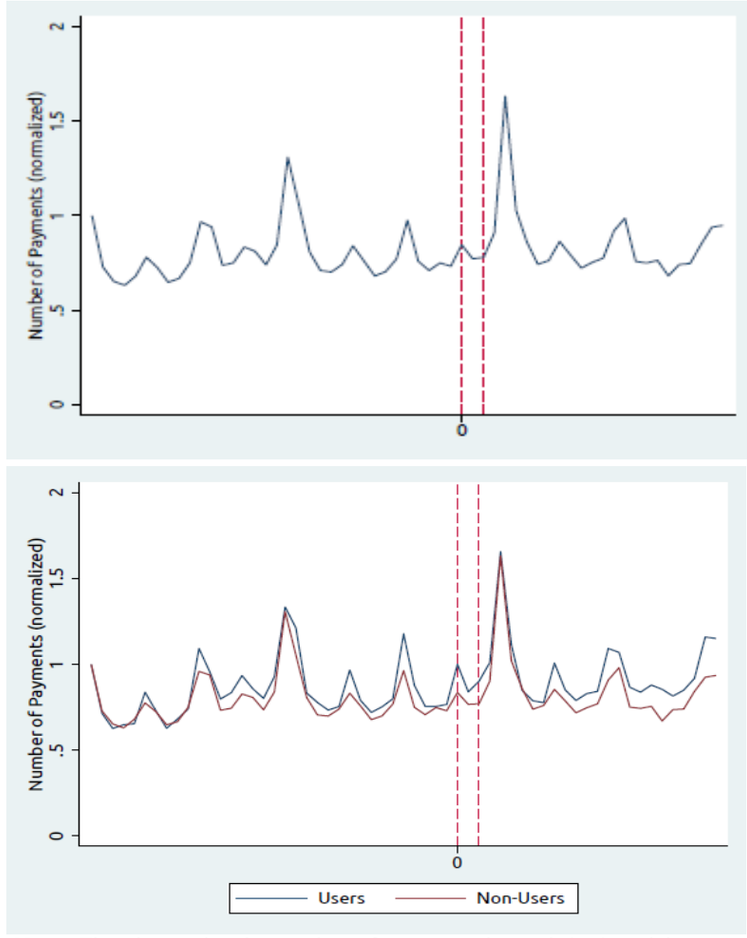

Note: The charts plot the number of payments in total (top) and by group (bottom) before and after a placebo cyberattack one year before the actual cyberattack. The red vertical dashed lines correspond to the first day of a placebo cyberattack and the last day of a placebo cyberattack. We anonymize the mid-period by averaging the values for the middle days and plot that average as the value over a single middle day.

Appendix Figure 2: Value of Payments One Year Earlier

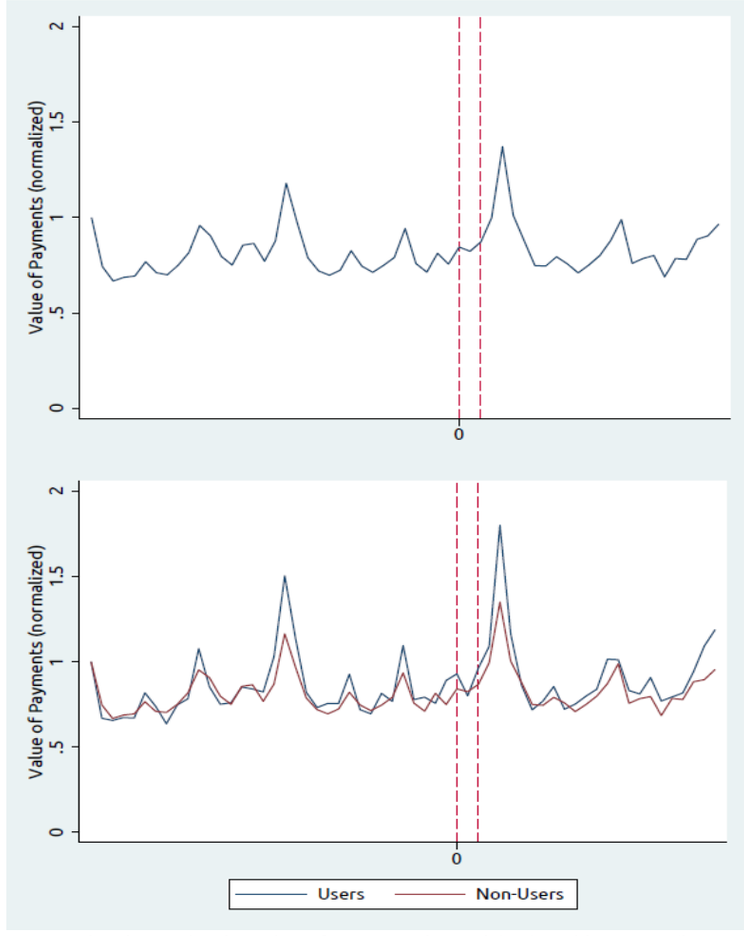

Note: The charts plot the value of payments in total (top) and by group (bottom) before and after a placebo cyberattack one year before the actual cyberattack. The red vertical dashed lines correspond to the first day of a placebo cyberattack and the last day of a placebo cyberattack. We anonymize the mid-period by averaging the values for the middle days and plot that average as the value over a single middle day.

Appendix Table 1: Payment Statistics on Users and Non-Users

	Full week before the cyberattack		Same day one week before the cyberattack	
Sender-Bank	*Share in Number of Payments*	*Share in Value of Payments*	*Share in Number of Payments*	*Share in Value of Payments*
Users	6.9%	2.6%	6.9%	2.1%
Non-Users	93.1%	97.4%	93.1%	97.9%

Note: The table presents payment statistics on users and non-users. GSIBs are excluded.

Appendix Table 2: Summary Statistics

Variables	*Unit*	*N*	*p(25)*	*mean*	*median*	*p(75)*	*sd*
A. Sender-Bank - Receiver-Bank level							
Δlog(Number of Payments)	%	562961	-0.287	0.006	0	0	0.630
Δlog(Value of Payments)	%	562961	-1.023	0.009	0	1.044	2.193
%Number of Payments gained during extensions	%	562961	0	0.000	0	0	0.017
%Value of Payments gained during extensions	%	562961	0	0.000	0	0	0.017
B. Receiver-Bank level							
Δlog(Value of Payments)	%	58505	-0.663	0.023	0.028	0.709	1.584
Access Discount Window	0/1	58505	0	0.007	0	0	0.086
log(Fed Funds)	-	463	19.133	20.172	20.292	21.488	1.424
Exposure (all banks)	%	58505	0.031	0.149	0.085	0.188	0.185
Exposure (small banks; accessed FF market)	%	29288	0.027	0.161	0.084	0.200	0.206
Exposure (small banks; did not access FF market)	%	29031	0.027	0.161	0.084	0.200	0.205
Exposure (small banks; did not access FF market; lower reserves to assets)	%	12053	0.025	0.132	0.072	0.161	0.179

Note: The table presents summary statistics of the main variables used in our analysis.

Appendix Table 3: Parallel Trends Analysis

	Δlog(Number of Payments)			Δlog(Value of Payments)		
	1	*2*	*3*	*4*	*5*	*6*
Users * Cyberattack	-0.171***	-0.164***	-0.176***	-0.273***	-0.263***	-0.281***
	(0.060)	(0.057)	(0.062)	(0.101)	(0.090)	(0.098)
Size * Cyberattack	-0.006	-0.006	-0.006	-0.001	-0.001	-0.001
	(0.005)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)
Users * Pre-Cyberattack	-0.019		-0.024	-0.028		-0.036
	(0.014)		(0.015)	(0.035)		(0.033)
Users * Post-Cyberattack		0.004	-0.009		0.005	-0.014
		(0.011)	(0.012)		(0.028)	(0.022)
Sender-Bank FE	yes	yes	yes	yes	yes	yes
Receiver-Bank x Day FE	yes	yes	yes	yes	yes	yes
Observations	562961	562961	562961	562961	562961	562961
R^2	0.104	0.104	0.105	0.103	0.103	0.103

Note: Δlog(Number of Payments) is the log change in the number of Fedwire payments on the day of the attack compared with the same weekday of the previous week. Δlog(Value of Payments) is the log change in the value of Fedwire payments compared with the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Size is the log assets of sender-banks. Cyberattack is a dummy variable that takes value one during the period of the cyberattack and zero otherwise. Pre-Cyberattack is a dummy variable that takes value one before the cyberattack and zero otherwise. Post-Cyberattack is a dummy variable that takes value one after the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Appendix Table 4: Additional Robustness Tests

	Placebo cyberattack one year before the actual cyberattack		Change in payments compared with the same day a month before	
	Δlog(Number of Payments)	Δlog(Value of Payments)	Δlog(Number of Payments)	Δlog(Value of Payments)
	1	*2*	*3*	*4*
Users * Cyberattack	0.005	-0.035	-0.126*	-0.222**
	(0.021)	(0.035)	(0.069)	(0.105)
Sender-Bank FE	yes	yes	yes	yes
Receiver-Bank*Day FE	yes	yes	yes	yes
Observations	423907	423907	432925	432925
R^2	0.109	0.095	0.100	0.105

Note: In columns 1 and 2, we consider a placebo cyberattack one year before the actual cyberattack. For example, if the attack started on the second Wednesday of June, then the placebo event is assumed to start on the second Wednesday of June one year earlier. In columns 3 and 4, the dependent variables denote changes compared with the same day a month before the cyberattack. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Cyberattack is a dummy variable that takes value one during the period of the (placebo) cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Appendix Table 5: Effects of Cyberattack with G-SIBs Included

	Δlog(Number of Payments)			Δlog(Value of Payments)		
	1	*2*	*3*	*4*	*5*	*6*
Users * Cyberattack	-0.139**	-0.159***	-0.160***	-0.235**	-0.243**	-0.231**
	(0.056)	(0.057)	(0.058)	(0.089)	(0.095)	(0.094)
Size * Cyberattack			-0.001			0.012*
			(0.004)			(0.007)
Sender-Bank FE	yes	yes	yes	yes	yes	yes
Day FE	yes	no	no	yes	no	no
Receiver-Bank x Day FE	no	yes	yes	no	yes	yes
Observations	760024	760024	760024	760024	760024	760024
R^2	0.023	0.121	0.121	0.007	0.120	0.120

Note: Δlog(Number of Payments) is the log change in the number of Fedwire payments compared with the previous week. Δlog(Value of Payments) is the log change in the value of Fedwire payments compared to the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Size is the log of assets of sender-banks. Cyberattack is a dummy variable that takes value one during the period of the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.

Appendix Table 6: Accumulation of Reserves with GSIBs Included

	Δlog(Reserves)	
	1	*2*
Users * Cyberattack	0.174*	0.182*
	(0.101)	(0.102)
Size * Cyberattack		-0.005
		(0.008)
Sender-Bank FE	yes	yes
Day FE	yes	yes
Observations	73336	73336
R^2	0.043	0.043

Note: Δlog(Reserves) is the log change in the reserves of sender-banks compared with the previous week. Users is a dummy variable that takes value one if a bank was a user of the TSP that was hit by the cyberattack and zero otherwise. Size is the log assets of sender-banks. Cyberattack is a dummy variable that takes value one during the period of the cyberattack and zero otherwise. Standard errors are two-way clustered at the sender-bank and day level. Statistical significance is denoted as *p<0.1, **p<0.05, ***p<0.01.