DEUTSCHE
BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 08/2022

Cybersecurity and financial stability

Kartik Anand
(Deutsche Bundesbank)

Chanelle Duley
(University of Auckland)

Prasanna Gai
(University of Auckland)

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–876–8
ISSN 2749–2958

Non-technical summary

Research Question

From mobile phone-based solutions for customers to virtual data rooms automating due diligence processes, modern banks bear little resemblance to the brick and mortar institutions of the turn of the century. But as the digital transformation in banking has gathered pace, so have cyber risks to financial stability. Yet, despite growing concern in cyber risk there is, as yet, no formal model of cyber attacks on financial stability and the implications for the regulation of cybersecurity. Our paper fills this gap.

Contribution

Our analysis builds on the premise that banks use shared digital services provided by third-party vendors. But while cost-saving, shared services create cybersecurity dependencies – one bank's access can become the 'back door' through which attackers can deploy malicious code in the shared services to impact others. Investing in cybersecurity to monitor for unauthorised intrusions allows a bank to protect both itself and others. Cybersecurity thus bears the hallmarks of a weakest-link public good. A successful cyber attacks, however, creates outages in the shared services, which impair banks' businesses and can precipitate self-fulfilling bank runs.

Results

Banks' investments in cybersecurity are influenced by two countervailing effects. First, the desire to ex ante free-ride on the public good contributions of other banks towards collective security means that banks underinvest in cybersecurity. And second, the prospect of failing due to runs following successful cyber attacks encourages banks to invest more in cybersecurity. And as the proclivity for runs increases, so too do the incentives to invest more in cybersecurity. Regulatory and supervisory tools that account for how the ex ante free-riding incentives interact with ex post run risk for banks may be used to implement socially optimal cybersecurity investments.

Nichttechnische Zusammenfassung

Fragestellung

Von mobiltelefonbasierten Kundenlösungen bis hin zu virtuellen Datenräumen für die Automatisierung von Due-Diligence-Prozessen – moderne Banken haben kaum mehr Ähnlichkeit zu der analogen Bankenwelt der Jahrtausendwende. Mit der digitalen Transformation des Bankwesens sind allerdings auch sogenannte Cyberrisiken entstanden. Ungeachtet der wachsenden Besorgnis über diese Risiken ist bislang kein formales theoretisches Modell verfügbar, das die Auswirkungen von Cyberattacken für die Finanzstabilität und die aufsichtliche Regulierung der Cybersicherheit erfasst. Das vorliegende Forschungspapier schließt diese Lücke.

Beitrag

Unsere Analyse beruht auf der Annahme, dass Banken gemeinsame digitale Dienste nutzen, die von Drittanbietern zur Verfügung gestellt werden. Dadurch werden einerseits Kosten gesenkt, andererseits aber auch Abhängigkeiten im Hinblick auf die Cybersicherheit der Banken geschaffen, denn die Nutzung solcher Dienste kann zum Einfallstor für Schadcodes geraten, mit denen Dritte angegriffen werden können. Mit Hilfe von Investitionen in ihre IT-Sicherheit, die den Schutz vor unberechtigten Eindringlingen verstärkt, schützt eine Bank sowohl sich selbst als auch andere. Cybersicherheit weist mithin die Merkmale eines öffentlichen Gutes auf, bei der die Höhe der Sicherheit vom schwächsten Glieds in der Technologiekette abhängt (weakest link problem). Erfolgreiche Cyberangriffe führen zum Ausfall gemeinsamer Dienste, was wiederum die Geschäftstätigkeit einzelner Banken empfindlich stören und dadurch sich selbsterfüllende Bank-Runs auslösen kann.

Ergebnisse

Die Investitionen der Banken in die Cybersicherheit unterliegen dem Einfluss zweier gegenläufiger Effekte: Zum einen kann die einzelne Bank kostenlos von den Investitionen anderer Banken zur kollektiven IT-Sicherheit profitieren, woraus ein Anreiz zu Unterinvestition in die eigene Cybersicherheit entsteht. Zum anderen

sieht sich die Bank einer drohenden Schieflage gegenüber, wenn es nach erfolgreichen Cyberangriffen zu einem Bank-Run und einem verstärkten Mittelabfluss kommt. Dieses Risiko veranlasst Banken wiederum, höhere Investitionen in ihre IT-Sicherheit zu tätigen. Das gesellschaftlich optimale Niveau an Investitionen in die Cybersicherheit kann mithilfe von regulatorischen und aufsichtlichen Maßnahmen erreicht werden, welche der Wechselwirkung zwischen Ex-ante-Anreizen zur Ausnutzung von Mitnahmeeffekten und dem Ex-post-Risiko eines Bank-runs Rechnung tragen.

Deutsche Bundesbank Discussion Paper No 08/2022

Cybersecurity and financial stability*

Kartik Anand, Chanelle Duley and Prasanna Gai

Abstract

Cyber attacks can impair banks operations and precipitate bank runs. When digital infrastructure is shared, banks defend themselves by investing in cybersecurity but can free-ride on the security measures of others. Ex ante free-riding by banks interacts with the ex post coordination frictions underpinning bank runs. While the temptation to free-ride induces underinvestment in cybersecurity, the prospect of a run encourages greater investment. System-wide cybersecurity is suboptimal and sensitive to rollover risk and bank heterogeneity. Regulatory measures, including negligence rules, liquidity regulation and cyber hygiene notices, facilitate constrained efficient cybersecurity investment. We suggest testable hypotheses to inform empirical work in this area.

Keywords: Cyber attacks, bank runs, global games, weaker-link public goods.

JEL classifications: G01, G21, G28, H41.

*Disclaimer: *The views expressed in this paper are those of the authors and do not necessarily represent those of the Deutsche Bundesbank or the Eurosystem.* Anand: Deutsche Bundesbank, Research Department, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt, Germany. Email: kartik.anand@bundesbank.de. Duley and Gai: University of Auckland, 12 Grafton Rd, Auckland 1010, New Zealand. Email: chanelle.duley@auckland.ac.nz and p.gai@auckland.ac.nz. We thank Mei Dong, Thomas Eisenbach, Charlie Kahn, Philipp J. König, Silvio Petriconi, Xavier Vives and seminar participants at the Deutsche Bundesbank, European Banking Theory Brown Bag Webinar, 2021 IFABS Conference, Oxford, 2021 Ridge Virtual Forum Workshop on Financial Stability, Montevideo, and the 2021 European Winter Meeting of the Econometric Society, Barcelona, for helpful comments. All remaining errors are our own.

"Cyber security is a public good...the social benefit conveyed by a well functioning and resilient financial system...requires a higher level of investment in cyber security than what individual firms would like to do on their own. In addition, many individual firms rely on shared services....an individual firm may rely on others in the shared network to make investments to increase the security of the network, but if every firm thinks this way, there will be underinvestment in security."

— Loretta J. Mester, Reserve Bank of Cleveland, 21 November 2019

1 Introduction

From mobile phone-based solutions for customers to virtual data rooms automating due diligence processes, modern banks bear little resemblance to the brick and mortar institutions of the turn of the century. But as the digital transformation in banking has gathered pace, so have cyber risks to financial stability.[1] Duffie and Younger (2019) describe how cyber attacks on banks can trigger runs involving wholesale depositors. Operational deposits, which account for 60% of wholesale unsecured funding for US banks, are typically the most vulnerable. These deposits are associated with high-frequency transactions such as payroll and payment processing. Eisenbach, Kovner, and Lee (2022) estimate that a cyber attack on any of the five most active US banks in wholesale funding markets would have significant spillovers to other banks. Despite the growing interest in cyber risk there is, as yet, no formal model of cyber attacks on financial stability and the implications for regulation. Our paper fills this gap.

Our analysis builds on the premise that many banks use shared digital services provided by third-party vendors (Kashyap and Wetherilt, 2019). For example, commercially available cloud-based solutions allow banks to avoid the costs of owning and maintaining in-house services (Boot, Hoffmann, Laeven, and Ratnovski, 2021). But they also create cybersecurity dependencies one bank's access can become the 'back door' through which an attacker can impact others. By gaining access to a bank's systems, they can deploy malicious code to exploit vulnerabilities – which are often unknown to the vendor (Perlroth, 2021) – in the shared service and cause outages.[2] Investment in cybersecurity, such as the hiring of IT

[1] Data from the Carnegie Endowment for International Peace indicates that the number of cyber attacks on financial institutions is increasing four-fold, year-on-year (Mauer and Nelson, 2020).

[2] Vulnerabilities refer to coding and hardware flaws or weaknesses that can be exploited to: (i) gain privileged access to victims' devices and applications; (ii) access data on the affected device or application and (iii) disrupt access for legitimate users. Vulnerabilities are ubiquitous in major platforms, from Amazon Web Services to VMware, and other large and small digital

specialists to continually monitor for intrusions, allows a bank to protect both itself and others using the shared service. Cybersecurity thus has the hallmarks of a *weakest-link public good* (Hirshleifer, 1983; Cornes, 1993).

Banks' ex ante investments in cybersecurity are shaped by the coordination failures that fuel ex post bank runs and threaten financial stability. But since banks do not internalise the cybersecurity investment decisions of others, they have incentives to free-ride on these investments. While free riding discourages investment in cybersecurity, the prospect of a bank run encourages more investment. And as the proclivity for runs, i.e., *rollover risk*, increases, so too do the incentives to invest more in cybersecurity. At the system level, cybersecurity is under-provided when rollover risk is low. But when the opportunity cost to withdraw is low and rollover risk is high, there can be over-provision of cybersecurity. Although cast as a model of bank runs, our results have broader applicability to other settings where the ex ante provision of a public good takes place in the shadow of coordination failure.

Our framework incorporates cyber attacks and cybersecurity into a model of bank runs (Rochet and Vives, 2004; Vives, 2014). Cyber attacks have three distinct elements. First, the greater the *intensity* of an attack, the more likely it is that an intrusion will be successful in penetrating cybersecurity defences. Ex ante uncertainty over the intensity of an attack reflects uncertainty about the identity of the attacker. This 'attribution problem' is a distinguishing feature of cyber attacks (Hayden, 2011). Second, following a successful intrusion and the deployment of malicious code, the shared services may suffer *temporary outages* that disrupt operations for all banks. In 2021Q2, for example, businesses experienced, on average, 23 days of downtime following ransomware attacks (Coveware, 2021). And third, even after the attack has been repelled there maybe *longer-lasting damage*. These deadweight losses can stem from the theft or corruption of data, or even the destruction of physical systems. Bouveret (2018) estimates that the annual average loss to banks from cyber attacks amounts to some US$100 billion, or 9% of banks' net income globally.

Banks invest in safe and liquid projects by issuing demandable debt claims to depositors.[3] To manage these projects, banks contract digital services from a third-party vendor, or 'platform'. Each bank then chooses how much of its working

services. Examples include the 'Meltdown' and 'Spectre' hardware vulnerabilities present in every Intel processor since 1995, but only uncovered in 2018 (Lipp, Schwarz, Gruss, Prescher, Haas, Fogh, Horn, Mangard, Kocher, Genkin, Yarom, and Hamburg, 2018; Kocher, Horn, Fogh, Genkin, Gruss, Haas, Hamburg, Lipp, Mangard, Prescher, et al., 2019).

[3]Consistent with much evidence, uninsured bank debt is assumed to be demandable. Demandability can arise endogenously due to liquidity needs (Diamond and Dybvig, 1983) or as a commitment device to overcome an agency conflict (Calomiris and Kahn, 1991; Diamond and Rajan, 2001).

capital it must allocate towards cybersecurity to bolster the defences surrounding the platform against cyber attacks (a public good), and towards operational resilience and business continuity provisions in the event of a cyber attack (a private good). A cyber attack that causes a temporary outage to the shared service impairs banks' ability to manage their projects. This, in turn, can lead to deposits being withdrawn early and precipitate a bank run. And, even after services resume, banks may suffer permanent losses that further compromise their ability to repay depositors. We pin down a unique equilibrium using global game methods (Morris and Shin, 2003); a run occurs if the outage exceeds a certain threshold which depends on the bank's investment decisions.

The more that each bank invests in cybersecurity, the greater the vigilance and fortification of the platform. Alternatively, banks can devote resources towards operational resilience in the event of an outage. This ensures that banks are better able to cope with outages by migrating project management operations to back-up systems, and thereby service greater withdrawals without failing. Investing more in cybersecurity comes at the expense of investing in operational resilience. So by investing more in cybersecurity, banks reduce the likelihood of cyber attacks causing outages in shared services and disrupting their operations. But conditional on the cyber attack being successful, banks have lower operational resilience, which increases their fragility, i.e., the risk of failing due to runs.

Banks' investments in cybersecurity are strategic complements, i.e. the more a bank invests in cybersecurity, the greater are the incentives for others to do likewise. Specifically, the cybersecurity investment game between banks is supermodular, with expected equity values displaying increasing differences (Van Zandt and Vives, 2007). As such, the ex ante investments are characterised by multiple equilibria. In one equilibrium, all banks find it optimal to invest nothing in cybersecurity and to focus on shoring up their own operational resilience. And in the other, there is a positive level of investment in cybersecurity that is to the benefit of all banks. The distribution of working capital across banks shapes system-level cybersecurity. Greater heterogeneity of working capital across banks can lead to lower system-wide cybersecurity, depending on the support of the distribution relative to a critical threshold that shapes the trade-off between cybersecurity and operational resilience.

Cybersecurity investments are constrained inefficient. The desire to free-ride on the public good contributions of other banks towards collective security means that a bank underinvests in cybersecurity. The bank equates its individual marginal rate of substitution between cybersecurity and operational resilience with its marginal rate of transformation, in contrast to the Samuelson (1954) rule for optimal public good provision. In this benchmark, each bank's contribution is set by equating the sum of the marginal rate of substitution over all banks to its marginal rate of transformation. Rollover risk, on the other hand, provides a countervailing force.

Banks have stronger incentives to stave off cyber attacks entirely by investing more in cybersecurity, which increases their marginal rates of substitution. In particular, if rollover risk is sufficiently large, banks over-invest in cybersecurity relative to the Samuelson benchmark.

The suboptimal system-wide investment in cybersecurity can be corrected by regulatory measures. A regulator can achieve the constrained efficient outcome by establishing a minimum standard of due care, with a negligence rule penalising banks that exert insufficient care. Penalties of the kind imposed by the US Securities and Exchange Commission (SEC) on financial firms that have deficient cybersecurity practices can be viewed in this light. Regulators can also attain the constrained efficient outcome by requiring banks to invest at the level necessary for optimal provision of the public good. Such minimum guidelines for heightened cyber protection form the basis of cyber hygiene notices and stress tests, such as those recently undertaken by the Monetary Authority of Singapore (MAS). Regulatory measures that impact the stability of banks' funding structures can also be used to influence banks' investments. In particular, by stipulating that each bank's funding structure must conform to a given level of rollover risk, a regulator can also implement the constrained efficient outcome.

Our model generates several testable implications that may inform future empirical work on cybersecurity investment and financial stability. An increase in the average intensity of cyber attacks invites greater reliance on operational resilience vis--vis cybersecurity. Greater rollover risk, such as the share of uninsured and unsecured wholesale debt, increases banks' incentives to invest in cybersecurity. As the losses following a cyber attack grow, banks also prefer to guard against intrusion by investing more in cybersecurity rather than operational resilience. But as banks increasingly fund their investments with equity, they have more to lose from a cyber attack and this increases the returns to investing more in operational resilience over cyber safeguards.

Related literature

Our paper contributes to the nascent and largely empirical literature on cybersecurity and financial stability. Duffie and Younger (2019) describe cyber-runs and conduct a stress test to understand the resilience of systemically important US banks to wholesale depositor withdrawals following a cyber attack.[4] Eisen-

[4] Jamilov, Rey, and Tahoun (2021) contribute to the empirical literature on cyber risk and contagion. They construct a text-based measure of cyber risk, which they use to test the link between balance sheet and income statement information from publicly-listed companies and their exposure to cyber risk. They also explore whether cyber risk exposure influences asset pricing, and whether the effects of this exposure propagate to unaffected peer firms. They demonstrate that cyber risk exposure has a negative and significant effect on stock returns of

bach et al. (2022) examine how cyber attacks impair a bank's ability to repay withdrawing creditors and discuss how this influences creditors' incentives to run. They suggest that, since cyber attacks impair the ability of the bank to repay early withdrawals, the first-mover advantage of creditors is weakened. The sequential service constraint means that creditors who withdraw face a lower probability of being repaid in full. Our analysis complements this argument. In our model, whenever rollover risk is low, bank failure following a cyber attack is not driven by coordination failure but instead by the deadweight losses. In this case, since the impairment to banks' ability to repay is permanent, there is no scope for an inefficient run. But when rollover risk is large, the impairment suffered by a bank following a cyber attack is more transient, and inefficient runs are a source of bank failure.

Our paper also informs the growing policy literature in this area (Kashyap and Wetherilt, 2019; Adelmann, Ergen, Gaidosch, Jenkinson, Morozova, Schwarz, and Wilson, 2020; Aldasoro, Gambacorta, Giudici, and Leach, 2020; Aldasoro, Frost, Gambacorta, and Whyte, 2021). Many of the policy arguments for cyber stress tests and other measures are informal – our positive analysis formalises the concerns expressed by Mester (2019) in the epigraph, while our normative analysis makes a rigorous case for why regulatory intervention may be justified.

There are also points of contact with the economics of security. Gordon and Loeb (2002) analyse the factors that influence how much a firm is willing to protect itself. Varian (2004) and Grossklags, Christin, and Chuang (2008) extend Gordon and Loeb (2002) to the case of multiple firms with security externalities. We contribute to this literature in two ways. First, in modelling the incentives of banks, we clarify the role of limited liability in shaping protection decisions. And second, we demonstrate how the ex ante incentive to invest in security is shaped by ex post coordination failure.

Finally, we add to the large literature on bank runs and global games (Morris and Shin, 2003; Goldstein and Pauzner, 2005). We specifically build on Rochet and Vives (2004), where unsecured debt holders delegate their rollover decisions to professional managers, so the decisions to rollover are global strategic complements. Our contribution shows how cyber risk management interacts with run risk in such a setting.

---

affected firms and find evidence that cyber risk is a source of systematic risk in financial markets due to contagion effects or firm-to-firm networks. See also Kamiya, Kang, Kim, Milidonis, and Stulz (2021), Woods, Moore, and Simpson (2021) and Florakis, Louca, Michaely, and Weber (2020).

2 Model

A single-good economy comprising N risk-neutral banks, indexed by $b = \{1, \ldots, N\}$, extends over three dates, $t = 0, 1, 2$. Banks differ in their endowments of working capital, $W_b < 1$, and are protected by limited liability. At the initial date, each bank operates a risk-free legacy project of unit size that is liquid and yields return, $R > 1$, at $t = 2$. Projects are funded by uninsured demandable debt, D, and outside equity, $E = 1 - D$. The debt is issued to creditors who delegate roll over decisions to professional fund managers. We assume that the face value of debt, F, is exogenous and independent of the withdrawal date. In Appendix A, we show that the key trade-offs remain qualitatively unchanged with an endogenous face value.

2.1 Digital service platforms.

Banks require digital services to manage their projects. For example, credit risk monitoring may be digitised using software from third parties, with the data generated stored on secured sites (Deutsche Bundesbank, 2021). At $t = 0$, banks outsource these operations to perfectly competitive digital service platforms that can simultaneously provide services to any number of banks with negligible operating costs. We suppose that there is a representative platform and normalize the fee it charges to zero.

Digital service platforms have vulnerabilities that are unknown to the platform vendor. Would-be attackers become aware of the vulnerabilities and seek to exploit them.[5] They do this by gaining access to local systems within banks that connect to the platform. By causing outages to digital services, cyber attacks impair banks' operations and can potentially precipitate bank runs.

Since the platform vendor is unaware of the vulnerabilities, it cannot require banks to undertake mitigating actions.[6] Banks must, therefore, invest in cybersecurity to monitor and repel unauthorised intrusions into their systems. The more they invest, the more capable they become in avoiding cyber attacks. Banks can also expend working capital on business continuity to shore up their operational resilience to outages. We suppose that bank b invests its working capital into cybersecurity, S_b, and operational resilience, O_b, such that $S_b + O_b \leq W_b$.

[5] There is an underground market for "zero day" vulnerabilities where hackers sell software and hardware vulnerabilities that they discover to interested buyers, such as government agencies, other hackers, and information security firms (Perlroth, 2021). The Stuxnet malicious code that spread via Microsoft Windows and targeted industrial control systems is an example of a code that exploited several zero day vulnerabilities (McDonald, Murchu, Doherty, and Chien, 2013).

[6] Due to this Knightian uncertainty, insurance markets against cyber attacks are unavailable (LeRoy and Singell, 1987).

2.2 Cybersecurity as a public good.

The security of banks' digital operations on the platform is a shared responsibility, which we model as a "weaker-link" public good (Cornes, 1993). Specifically, the level of cybersecurity available to all banks is

$$\chi = \left[\prod_{b=1}^{N} S_b\right]^{1/N}. \tag{1}$$

A higher level of cybersecurity entails better detection and response to intrusions of banks' computer systems. In particular, we can picture the "security blanket" over the platform as a circular region with N banks situated along the perimeter. Each bank is responsible for maintaining security along its portion of the perimeter. But an attacker who breaches the section of the perimeter guarded by bank b can disrupt the platform and adversely impact all banks. Importantly, our analysis abstracts from the security provided by the platform vendor – our focus is on cybersecurity investments by banks *over and above* that provided by the platform.

The weaker-link formulation implies that investment in cybersecurity generates positive externalities for all banks since $\partial\chi/\partial S_b > 0$. The marginal product of investment in cybersecurity is also higher for those banks with low levels of investment. Equation (1), moreover, exhibits constant elasticity of substitution and is homogeneous to degree one. So scaling up each bank's investment by a constant factor increases the level of cybersecurity for all banks by the same factor. Were we to instead model cybersecurity as a weakest-link public good (Hirshleifer, 1983), then investments across banks would be perfect complements. In this case, the marginal product of investment to a bank – from investing at a level greater than the lowest level of investment across all banks – would be zero and so all banks would choose the same lowest level of investment. The weaker-link formulation, thus, allows us to better explore the impact of bank heterogeneity on cybersecurity investments.

2.3 Cyber attacks.

At $t = 1$, banks sharing the platform are subject to a cyber attack. The intensity of the attack, $\lambda \in [0, \bar{\lambda}]$, is a uniformly distributed random variable that reflects the capability of the attacker. For example, state-sponsored attackers will have considerable resources to launch a high-intensity attack (high λ). By contrast, typical cyber-criminals will likely have fewer resources so their attacks are less intense (low λ). The uncertainty over attack intensity thus reflects uncertainty over the identity of the attacker, i.e., the so-called "attribution problem" (Hayden, 2011).

A cyber attack successfully breaches the security blanket whenever $\lambda > \chi$. The

attack disrupts the platform causing temporary outages.[7] The outage, $\alpha \in [0,1]$, is a uniformly distributed random variable that impairs banks' abilities to manage their projects and access key functions. But investments in business continuity and operational resilience can mitigate the outage. By investing O_b in backup systems, the effective accessibility shock experienced by bank b is $\alpha[1-h(O_b)] < \alpha$, where $h(O_b)$ is an increasing and concave function. Unlike cybersecurity, which is a public good, operational resilience is a *private good*, the benefits of which are not shared with other banks.[8]

2.4 Bank failure conditions.

Platform outages can precipitate runs on banks due to funding illiquidity (Duffie and Younger, 2019). If a fraction $\ell_b \in [0,1]$ of debt is withdrawn at $t=1$, the the bank fails due to *illiquidity* whenever

$$R\big[1-\alpha(1-h(O_b))\big] - \ell_b FD < 0\,, \tag{2}$$

i.e., the value of available assets is insufficient to service withdrawals. So the bank fails whenever $\alpha > \alpha_b^{IL}(\ell_b) \equiv \frac{R-\ell_b FD}{R\left(1-h(O_b)\right)}$.

Cyber attacks can also have longer-lasting repercussions. These include the loss of secret information pivotal to the bank's role as a financial intermediary (Dang, Gorton, Holmström, and Ordonez, 2017), losses incurred from paying ransom demands, and even physical damage to critical systems. The credit downgrading in 2019 of the Maltese bank, Valletta PLC, following a cyber attack and concerns over the bank's operational risk management highlights the risks to bank insolvency (S&P Global Market Intelligence, 2019).[9]

We capture such possibilities by assuming that bank b is subject to deadweight losses proportional to the effective outage experienced. Bank b's project yields $R - R\delta\alpha(1-h(O_b))$, where $\delta < 1$ reflects the deadweight loss incurred through the cyber attack. The greater the deadweight loss, the larger are bank losses and the lower the return from the project. Bank b fails due to *insolvency* at $t=2$

[7] For example, the recent distributed denial of service (DDoS) attack on the New Zealand Stock Exchange prevented the posting of market announcements and led to trading suspensions over several days (Tarabay, 2021). And an attack on Nicehash – a crypto-currency platform – resulted in a freeze of customer funds for 24 hours (Daniel, 2020).

[8] An alternate specification for the private good, which we considered in an early working paper, is for banks to invest directly in projects and thereby lower their leverage. With a lower leverage, the consequences of outage shocks are also less pronounced and, hence, the scope for banks to fail following a cyber attack is reduced.

[9] The Dark Seoul malware attack on South Korean banks induced losses of some US$ 378 million (Kopp, Kaffenberger, and Wilson, 2017). See also Bouveret (2018).

whenever

$$R\big[1-\delta\alpha(1-h(O_b))\big]-\ell_b FD<(1-\ell_b)FD\,, \tag{3}$$

i.e., the gross returns from the project are insufficient to repay the total debt claims against the bank. So the bank fails whenever $\alpha>\alpha_b^{IN}\equiv\frac{R-FD}{R\delta(1-h(O_b))}$, which is independent of the fraction of withdrawals at the interim date.

2.5 Rollover decisions.

Creditors delegate rollover decisions to professional fund managers who are rewarded for making the right decision – if the bank does not fail, a fund manager's payoff difference between withdrawing and rolling over is $-c<0$; if the bank fails, the differential payoff is $b-c>0$.[10] The *conservatism ratio*, $\gamma\equiv\frac{b-c}{b}$, summarises these payoffs. More conservative managers (higher γ) are less inclined to roll over since the cost of withdrawing is low. When $\gamma>0$, fund managers' actions are strategic complements and the bank is subject to rollover risk.

We suppose there is a continuum of fund managers, $k\in[0,D]$, for creditors in each bank b, who base their roll over decisions on private and noisy signals about the outage shock, namely

$$x_{bk}=\alpha+\epsilon_{bk}\,, \tag{4}$$

where the noise term ϵ_{bk} is independent of the outage shock, α, and is independently and identically distributed across fund managers according to a normal distribution with mean zero and standard deviation σ. There is no overlap between the sets of fund managers across the different banks.

Table 1 illustrates the timing of events in the model.

$t=0$	$t=1$	$t=2$
1. Cybersecurity and operational resilience investments	1. Cyber attack causes outage	1. Projects mature and deadweight losses accrue
	2. Private signals on outage disruption	2. Outstanding debts repaid
	3. Fund managers roll over or withdraw debt	

Table 1: Timeline of events.

[10] As an example, assume the cost of withdrawing is c; the benefit from withdrawing when the bank fails is $b-c>0$; the payoff from rolling over when the bank fails is zero.

3 Analysis

The symmetric pure strategy perfect Bayesian equilibrium for each bank comprises a critical shock, α_b^*, a private threshold signal, x_b^*, and investments in cybersecurity, S_b^*, and operational resilience, O_b^*, such that

1. at $t = 1$, fund managers' rollover decisions, x_b^*, are optimal and the run threshold induces bank failure whenever $\alpha > \alpha_b^*$, given S_b^* and O_b^*;

2. at $t = 0$, the bank's investments, S_b^* and O_b^* maximise expected equity value, given the run threshold (x_b^*, α_b^*), and subject to the resource constraint $S_b^* + O_b^* \leq W_b$.

We construct the equilibrium backwards, solving for the optimal rollover decision of fund managers before establishing the optimal investment in cybersecurity and operational resilience.

3.1 Rollover risk

The levels of investment in cybersecurity and operational resilience, together with the outage shock, shape the dynamics of rollover risk. For a given mass of early withdrawals, ℓ_b, bank b does not fail provided the outage shock, α, is sufficiently small. But the criteria determining the largest shock that bank b can withstand depend on whether failure is driven by illiquidity or insolvency.

Lemma 1. *There exists a unique threshold*

$$\widehat{\gamma} = \frac{1}{\delta} - \frac{R}{FD}\left[\frac{1}{\delta} - 1\right] \tag{5}$$

such that bank failure is driven by illiquidity if and only if the mass of withdrawals is sufficiently large, $\ell_b > \widehat{\gamma}$.

Figure 1 illustrates Lemma 1 by plotting the illiquidity and insolvency conditions together with their "envelope" – the red line encapsulating the region where the bank does not fail. In region I, where the fraction of withdrawals is small, $\ell_b < \widehat{\gamma}$, the bank is able to service them following the outage. But in doing so, due to deadweight losses, the bank has too few resources to repay claims that are rolled over. So although the bank can meet interim liquidity needs, the cyber attack renders it insolvent at $t = 2$.

In region II, the fraction of withdrawals is so large, given the outage shock, that the bank is unable to service them at $t = 1$. This is despite the bank having sufficient resources at $t = 2$, even allowing for deadweight losses, to satisfy all its claims. So the bank fails due to illiquidity even though it is solvent.

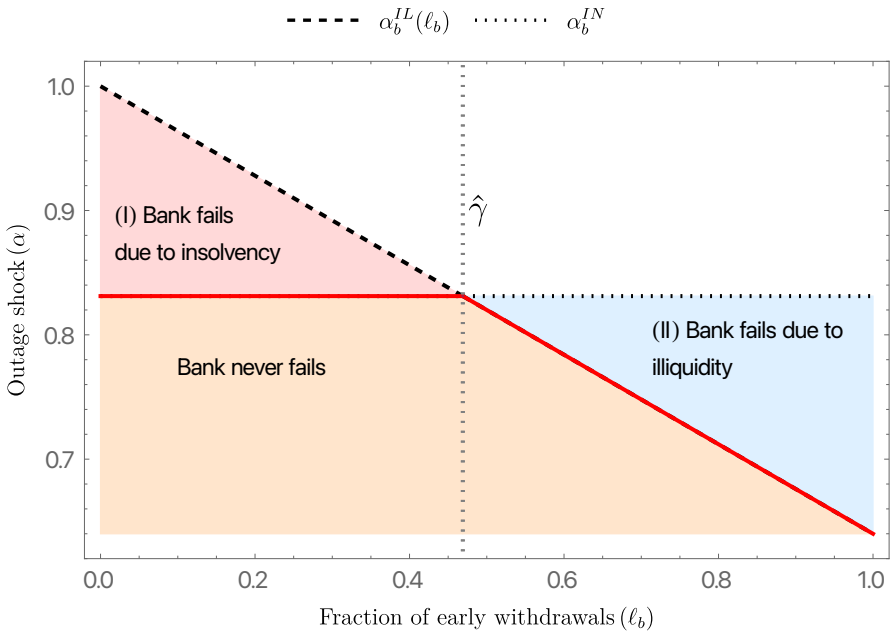

Figure 1: Failure conditions for a bank following a successful cyber attack.

Assumption 1. $R - FD < R\delta[1 - h(W_b)]; \quad R > FD.$

Assumption 1 ensures a tripartite classification over values of the outage shock (Morris and Shin, 1998). The first part of the assumption implies that if the outage completely disrupts the platform, $\alpha = 1$, and all debt claims are rolled over, $\ell_b = 0$, the deadweight loss is so large that the bank cannot service its debts and fails due to insolvency. The second part of the assumption ensures that if there is no outage, $\alpha = 0$, and all fund managers withdraw early, $\ell_b = 1$, then the bank has sufficient resources to repay creditors in full at $t = 1$.

In the limit of vanishing private noise, $\sigma \to 0$, we pin down a unique run threshold, x_b^*, for each bank b following a cyber attack.

Proposition 1. *There exists a unique signal, x_b^*, such that fund manager k rolls over debt at $t = 1$ if and only if $x_b \leq x_b^*$. The threshold signal, x_b^*, corresponds to a unique outage shock threshold*

$$\alpha_b^*(O_b) = \begin{cases} \alpha_b^{IN} \equiv \frac{R-FD}{R\delta[1-h(O_b)]} & \text{if } \gamma < \widehat{\gamma} \\ \alpha_b^{IL}(\gamma) \equiv \frac{R-\gamma FD}{R[1-h(O_b)]} & \text{if } \gamma \geq \widehat{\gamma}. \end{cases} \tag{6}$$

Bank b fails whenever $\alpha > \alpha_b^*(O_b)$.

Bank failure is driven by illiquidity only when rollover risk is sufficiently large, i.e. $\gamma \geq \widehat{\gamma}$. Following an outage shock, the bank is only able to service debt claims

if few fund managers withdraw. But if a sufficient proportion of fund managers withdraw, it is in each fund manager's best interest to also withdraw because the bank will fail as a result of the run. By contrast, when $\gamma < \widehat{\gamma}$, bank failure is driven by insolvency concerns and rollover risk plays no role. Following a large enough outage, $\alpha > \alpha_b^{IN} \equiv \frac{R-FD}{R\delta[1-h(O_b)]}$, each fund manager has a strictly dominant strategy to withdraw since the bank is sure to fail. In what follows, we define *bank fragility* as the risk of a bank failing due to rollover risk and measure this by the threshold $\alpha_b^*(O_b) = \alpha_b^{IL}(\gamma) \equiv \frac{R-\gamma FD}{R[1-h(O_b)]}$, for each bank $b = 1, \ldots, N$.

Lemma 2. *Under a binding budget constraint, $S_b + O_b = W_b$, greater investment in cybersecurity increases bank fragility, i.e. $\frac{\partial \alpha_b^*}{\partial S_b} < 0$, for all banks $b = 1, ..., N$.*

With a binding budget constraint, investing more in cybersecurity implies a reduction in operational resilience. While this ensures that the bank is better able to thwart a cyber attack, in the event of a high-enough intensity attack that penetrates the security blanket, the consequences for the bank are worse. By investing too little in business continuity and back-up systems, the bank is more exposed to the outage shock, rendering it more fragile.We next consider banks' ex ante investments and argue that the budget constraints always bind in equilibrium.

3.2 Optimal cybersecurity investment

At $t = 0$, each bank must decide how much to invest in cybersecurity and operational resilience, taking as given the levels of investment by other banks and the failure threshold, $\alpha_b^* \equiv \alpha_b^*(O_b)$. The ex ante probability that a cyber attack is thwarted, given a level of cybersecurity, χ, is $\Pr[\lambda < \chi] = \chi/\overline{\lambda}$.

The expected profit for bank b is the sum of its profits in the event that the cyber attack fails, and its residual profit following a cyber attack that is not severe enough to trigger bank failure, i.e.

$$\pi_b(S_b, O_b) = \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}(R - FD) + \left[1 - \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}\right] \left(\int_0^{\alpha_b^*} EV_2(\alpha, O_b) d\alpha\right). \quad (7)$$

Bank b's equity value following an outage is $EV_2(\alpha, O_b) = R[1 - \alpha\delta(1 - h(O_b))] - FD$, and $\vec{S}_{-b}$ is a vector representing the cybersecurity investments by the other $N-1$ banks. The bank chooses S_b and O_b to maximise equation (7) subject to the constraint $O_b + S_b \leq W_b$. But since expected profits are increasing in both S_b and O_b, the constraint always binds in equilibrium (Bergstrom, Blume, and Varian, 1986).

The bank balances the marginal benefit and cost of cybersecurity when selecting its investment. A higher level of cybersecurity investment improves the security blanket around the platform, lowering the probability of a successful attack. But

this comes at the cost of reduced investment in operational resilience, which makes the bank more fragile in the event of a successful attack. As a result, the bank has an incentive to free-ride on the cybersecurity investments of other banks.

To clarify the interaction between ex ante free riding in cybersecurity by banks and ex post coordination failure, we examine two benchmark cases. In the first, we consider the allocation problem of a social planner who chooses how much each bank should invest in cybersecurity, taking into account how this influences other banks' investments in the absence of any rollover risk. In the second, we introduce coordination failure and examine how rollover risk impacts the social planner's choice. As a final step, we characterise the privately optimal investment for individual banks in the face of both the free-riding problem and rollover risk.

3.3 Benchmark 1

Absent any free riding and coordination failure, $\gamma < \widehat{\gamma}$ and bank b is fundamentally insolvent whenever $\alpha > \alpha_b^* = \alpha_b^{IN}$ and runs are efficient. The allocation problem for a social planner is

$$\begin{aligned}
\max_{\{S_b, O_b\}} \quad & \Pi = \sum_{b=1}^{N} \left[(\chi/\overline{\lambda})\,(R - FD) + \left[1 - (\chi/\overline{\lambda})\right] \left(\int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b) d\alpha \right) \right] \\
\text{subject to} \quad & \sum_{b=1}^{N} W_b = \sum_{b=1}^{N} \Big(S_b + O_b \Big) \\
& \alpha_b^{IN} = \frac{R - FD}{R\delta[1 - h(O_b)]} \qquad \forall\, b = 1, ..., N \\
& \chi = \left(\prod_{b=1}^{N} S_b \right)^{\frac{1}{N}} .
\end{aligned}$$

In other words, the social planner sets banks' investments in cybersecurity and operational resilience to maximise the aggregate expected profits of all banks, taking into account the aggregate resource constraint, their failure thresholds, and how cybersecurity investment influences the level of public good provision. The

solution to the planner's problem is given by the following system of equations:

$$\sum_{b=1}^{N} \frac{\overbrace{(R - F\,D) - \int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b)d\alpha}^{\equiv \frac{\partial \pi_b}{\partial \chi}}}{\underbrace{(\bar{\lambda} - \chi)\int_0^{\alpha_b^{IN}} \Big(\partial EV_2/\partial O_b\Big)d\alpha}_{\equiv \frac{\partial \pi_b}{\partial O_b}}} = \frac{1}{\partial \chi / \partial S_{b'}}, \qquad \forall\ b' = 1, \ldots, N\,, \quad (8)$$

where $\sum_{b=1}^{N} W_b = \sum_{b=1}^{N}(S_b + O_b)$. Equation (8) is a version of the "Samuelson rule" (Samuelson, 1954). The planner chooses an allocation to equate the marginal rate of transformation of working capital into cybersecurity for each bank (the right-hand side) with the sum of the marginal rates of substitution between investing in cybersecurity and operational resilience for all N banks (the left-hand side). The numerator on the left-hand side of equation (8) reflects the effect of a marginal increase in cybersecurity, χ, on each bank's equity value. With a higher level of cybersecurity, it is more likely that cyber attacks fail and, hence, the bank earns a higher equity value. The denominator represents the marginal effect for each bank from investing in operational resilience. A higher O_b implies that bank b can reduce the deadweight losses it would face in the event of a successful attack and thereby obtain a higher equity value.

3.4 Benchmark 2

Next consider the case where banks must contend with rollover risk, i.e. $\gamma > \widehat{\gamma}$. In the event of an outage, $\alpha > \alpha_b^* = \frac{R - \gamma\, F\, D}{R(1 - h(O_b))}$, the impairment of bank b's project is so severe that it cannot service withdrawals and fails. But had creditors rolled over their claims, the bank could have repaid them in full at $t = 2$ despite deadweight losses. So the bank fails due to illiquidity at $t = 1$ even though its equity value is positive, i.e. $EV_2(\alpha_b^*) = (R - FD) - \delta(R - \gamma FD) > 0$. The Samuelson rule thus

becomes

$$\sum_{b=1}^{N} \frac{\overbrace{(R - F\,D) - \int_0^{\alpha_b^*} EV_2(\alpha, O_b)d\alpha}^{\equiv \frac{\partial \pi_b}{\partial \chi}}}{\underbrace{(\bar{\lambda} - \chi)\left[EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial O_b} + \int_0^{\alpha_b^*}\Big(\partial EV_2/\partial O_b\Big)d\alpha\right]}_{\equiv \frac{\partial \pi_b}{\partial O_b}}} = \frac{1}{\partial \chi/\partial S_{b'}}\,, \tag{9}$$

for all $b' = 1, \ldots, N$, and where the aggregate resource constraint is satisfied. The influence of rollover risk is twofold. First, the range of outage shocks where banks obtain positive equity value is smaller, since $\alpha_b^* = \frac{R - \gamma\, F\, D}{R(1-h(O_b))} < \alpha_b^{IN}$. So the marginal benefit from greater cybersecurity – in thwarting attacks and earning higher equity value – is greater. This effect increases banks' marginal rates of substitution. Second, the marginal benefit of investing in operational resilience also decreases as banks care more about protecting their equity value $EV_2(\alpha_b^*) > 0$ by thwarting cyber attacks in the first place. Summing up, banks' marginal rates of substitution are higher with rollover risk. Relative to Benchmark 1, there is an over-provision of cybersecurity.

3.5 Privately optimal solution

When banks choose their investments, taking as given the investments of all other banks, they equate their marginal rate of transformation with their own marginal rate of substitution. When both free-riding and rollover risk are present, the privately optimal level of cybersecurity investment is constrained inefficient. To characterise the solution, we make the following assumption:

Assumption 2. *Fund managers' conservatism satisfies $\gamma > \widehat{\gamma}$ and bank failures are driven by illiquidity.*

Proposition 2. *The privately optimal levels of investment (S_b^*, O_b^*) for bank b are given by the solution to*

$$\frac{\partial \pi_b/\partial \chi}{\partial \pi_b/\partial O_b} = \frac{1}{\partial \chi/\partial S_b}\,, \tag{10}$$

subject to the constraint $W_b = S_b^* + O_b^*$ and where

$$\frac{\partial \pi_b}{\partial \chi} = R - FD - \int_0^{\alpha_b^*} EV_2(\alpha, O_b)d\alpha \tag{11}$$

is the marginal benefit to bank b from greater cybersecurity,

$$\frac{\partial \pi_b}{\partial O_b} = \left(\overline{\lambda} - \chi\right) \left[EV_2(\alpha_b^*) \frac{\partial \alpha_b^*}{\partial O_b} + \int_0^{\alpha_b^*} \frac{\partial EV_2}{\partial O_b} d\alpha \right] \tag{12}$$

is the marginal benefit from investing in operational resilience, and α_b^* *is given by equation (6).*

3.6 Working capital

A feature of the weaker-link formulation is that it allows us to examine how heterogeneity across banks' working capital influences cybersecurity investment. Before exploring the consequences at the system level, we first clarify how the endowment of working capital shapes cybersecurity investment for an individual bank.

Proposition 3. *There exists a critical threshold* $\widehat{W}$*, such that bank b's investment in cybersecurity is increasing in working capital,* $\partial S_b^*/\partial W_b > 0$*, if and only if* $W_b \leq \widehat{W}$*.*

An increase in working capital induces two opposing forces on a bank's cybersecurity investment. On one hand, for a given level of investment, S_b, an increase in working capital mechanically increases investment in operational resilience. So the difference in equity value when a cyber attack succeeds and when it fails is smaller. This reduces the incentive for the bank to invest more in cybersecurity. On the other hand, the marginal benefit from greater investment in operational resilience is subject to diminishing returns. When the endowment of working capital is small, the marginal rate of substitution is increasing with the endowment and the bank invests more in cybersecurity. But when working capital exceeds $\widehat{W}$, the marginal rate of substitution is decreasing and there is less cybersecurity investment. The critical threshold, $\widehat{W}$, is larger when there is rollover risk since the marginal rate of substitution is generally higher.

3.7 Other comparative static results

Propositions 4, 5 and 6 report how changes in (i) cyber risk; (ii) bank equity; and (iii) conservatism of fund managers influence bank behavior.

Proposition 4. *Cybersecurity investment is decreasing in the maximum intensity of cyber attacks,* $\partial S_b^*/\partial \bar{\lambda} < 0$*, and increasing in the deadweight loss,* $\partial S_b^*/\partial \delta > 0$*. Conversely, fragility,* α_b^**, is decreasing in* $\bar{\lambda}$ *and increasing in* δ*.*

An increase in $\bar{\lambda}$ implies an increase in the average intensity of cyber attacks. As the intensity of attacks increases, so too does the likelihood that it will be successful

and cause outages on the common platform. Banks are, therefore, incentivised to invest in operational resilience in order to mitigate the outages shocks. This, in turn, leads to a decrease in investment in both cybersecurity and fragility.

An increase in the deadweight loss, δ, induces two effects. First, the difference in the expected equity value earned by the bank, between when the cyber attack is unsuccessful and successful is larger. This is because, for any given outage size, the losses suffered are larger. This effect tends to increase the marginal rate of substitution. Second, the benefits from investing in operational resilience are lower. In particular, the bank is better off trying to avoid outages that lead to deadweight losses than attempting to mitigate them by investing in operational resilience. In sum, the two effects ensure that the bank invests more in cybersecurity and, as a consequence, fragility also increases.

Proposition 5. *A marginal increase in bank equity leads to a decrease in cybersecurity investments,* $\frac{\partial S_b^*}{\partial E} < 0$*, and an decrease in fragility,* $\frac{d\alpha_b^*}{dE} > 0$*.*

Given the fixed balance sheet assumption, an increase in bank equity mechanically leads to a decrease in debt issuance and therefore bank leverage. But since the bank now has more 'skin in the game', there is more to lose in the event that the bank fails. Since bank failure is driven by outages, following a successful cyber attack, the marginal returns to investing more in operational resilience (and thereby reducing fragility) are greater. So the bank reduces its investment in the public good of cybersecurity.

Proposition 6. *A marginal increase in fund managers' conservatism leads to an increase in cybersecurity investment,* $\frac{\partial S_b^*}{\partial \gamma} > 0$*, and fragility,* $\frac{d\alpha_b^*}{d\gamma} < 0$*.*

Following an increase in γ, each fund manager, k, is less likely to roll over claims for any given signal, x_{bk}. Ex ante, this means that bank b is more likely to fail for any realisation of outage shock, α_b. Therefore, the marginal benefit to allocating an additional unit of endowment to cybersecurity is high, relative to the marginal opportunity cost of investing more in operational resilience. This is because the bank is more likely to fail even following a relatively low intensity cyber attack as conservatism increases. So the bank is better off shoring up its cybersecurity defences to ward off cyber attacks entirely, even though this leads to an increase in fragility in the event of a successful attack.

4 System-wide cybersecurity

To establish the joint equilibrium for all bank investment decisions, we first show how each bank, b, responds to an increase in cybersecurity investment by another bank.

Lemma 3. *An increase in cybersecurity investment by any bank $b' \neq b$ elicits bank b to increase in cybersecurity investment by bank $\frac{\partial S_b^*}{\partial S_{b'}} > 0$, for all $b = 1, \ldots, N$.*

The security blanket is jeopardised by the banks that invest the least in cybersecurity. As each bank increases its investment, the benefits to greater cybersecurity investment also rise for all other banks. Increases in cybersecurity investment by bank b', from a low base, are met by substantial increases in bank b's contribution. As $S_{b'}$ becomes sufficiently large, the net benefit to bank b from additional cybersecurity investment is still positive but smaller in magnitude. This reflects the increased opportunity cost from investing additional units of working capital in cybersecurity.

A consequence of Lemma 3 is that the cybersecurity investment game between banks is supermodular with increasing differences in banks' expected profits (Van Zandt and Vives, 2007). This ensures that banks' best response correspondences for the level of cybersecurity investment, given the investments of other banks, intersect to yield greatest and least Nash equilibria.

4.1 Joint equilibrium

The system-wide consequences for cybersecurity are summarised in the following proposition:

Proposition 7. *There exist two Nash equilibria. In the first, all banks invest nothing in cybersecurity, $S_b^* = 0$ for all $b = 1, ...N$. In the second, all banks split their endowments between cybersecurity and operational resilience, $S_b^* \in (0, W_b)$ for all $b = 1, ...N$, and the equilibrium level of cybersecurity is*

$$\chi^* = \left(\prod_{b=1}^{N} S_b^* \right)^{\frac{1}{N}}.$$

Figure 2 illustrates Proposition 7 for the two-bank case. If bank b anticipates that others will invest more of their working capital into operational resilience, then it expects the level of cybersecurity to be low. So cyber attacks, including those of low intensity, are likely to be successful and disruptive. It is therefore a best response for bank b to also devote more of its working capital on operational resilience. Since all banks reason and behave this way, there is no investment in this "bad" equilibrium.

In the other equilibrium, each bank invests $S_b^* \in (0, W_b)$ on cybersecurity. If bank b expects others to invest a small amount in cybersecurity, then it is in bank b's interests to do so as well. But once the level of investment by all other banks is sufficiently high, system-wide security attains a relatively high level. Beyond this

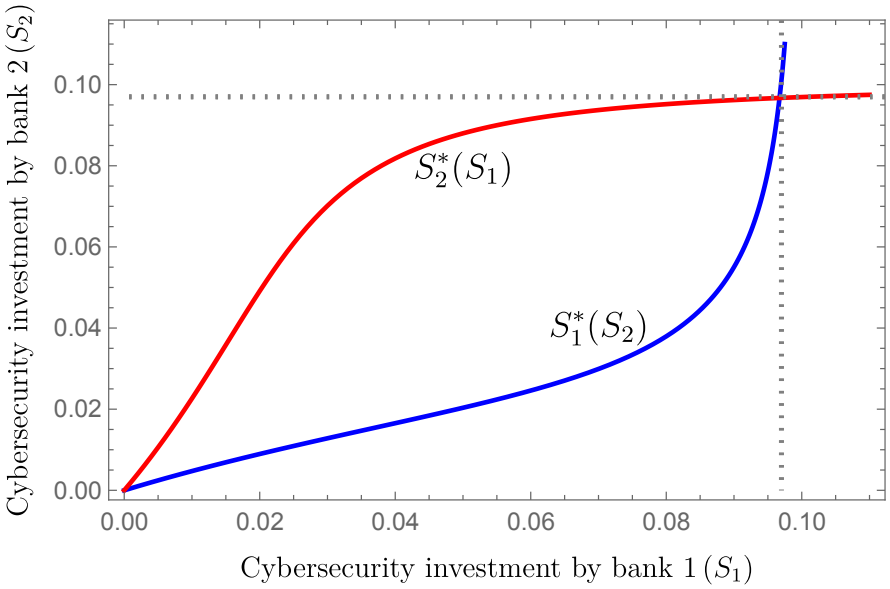

Figure 2: The best response of each bank to the level of investment in cybersecurity by the other bank. The best response curves overlap at a least and greatest Nash equilibrium in which each bank invests nothing, or a positive amount, into cybersecurity respectively.

point, there are negative returns to bank b from further investing in cybersecurity and it refrains from doing so. The "good" equilibrium corresponds to the point of diminishing returns for all banks.

4.2 Distribution of working capital

At the system level, we can order all banks according to their endowments of working capital, $W_1 < W_2 < ... < W_N$, without loss of generality. We also make the following assumption:

Assumption 3. *Initial endowments of working capital satisfy* $W_1 < \widehat{W}$ *and* $W_N > \widehat{W}$.

Assumption 3 implies that, following a marginal increase in its endowment, bank 1 increases its cybersecurity investment, while bank N decreases this investment in favour of shoring up operational resilience.

Proposition 8. *A mean-preserving spread increase in banks' endowments of working capital reduces equilibrium cybersecurity,* χ^*.

Following a mean-preserving spread increase in working capital at the system level, bank 1 and bank N both choose to decrease their cybersecurity investment

in favour of operational resilience. Anticipating this, all other banks decrease their investments as well since decisions are strategic complements. If the shift in endowments is not symmetric around $\widehat{W}$, the net effect is generally ambiguous. For example, if the endowments of only 'poor' banks would increase, i.e. those with $W_b < \widehat{W}$, then the net effect is an increase in cybersecurity. Not only would poor banks increase investments, but so too would 'rich' banks.

But if the endowments of only rich banks increases (banks for whom $W_b \geq \widehat{W}$), the opposite is true. Rich banks would want to invest more in operational resilience to lower fragility, which influences poor banks to do the same. The provision of system-level cybersecurity thus depends crucially on the distribution of working capital.

5 Regulatory implications & testable hypotheses

The laissez-faire equilibrium in Proposition 2 can be fruitfully compared with Benchmark 1, which is akin to the problem confronting a 'regulator'. Recall in the setting of Benchmark 1, the regulator operates in the absence of rollover risk, and internalises the impact that banks' cyber investments have on each other, when choosing allocations for the system as a whole. We denote by S^R the level of investment in cybersecurity set by the regulator, which is common for all banks.

Proposition 9. *For each bank, $b = 1, \ldots, N$, there exists a critical threshold, γ_b^c, such that there is underinvestment in cybersecurity, $S_b^* < S_b^R$, if and only if $\gamma < \gamma_b^c$.*

When rollover risk is low, then so too is the risk of bank failure due to a run. Ex ante, this does not present a strong incentive for a bank to invest in cybersecurity. Compared with Benchmark 1, free-riding exerts a stronger influence on banks' incentives and implies under-provision of the public good. As rollover risk increases, the risk of bank failure due to runs also rises. In this situation, banks have greater incentives to stave off cyber attacks and, accordingly, invest in cybersecurity. Rollover risk considerations dominate free-riding and lead to over-provision of cybersecurity and under-investment by banks in their own operational resilience.

Importantly, the point at which this coordination failure concerns begin to dominate those are free-riding depend on the bank's initial endowment of working capital. Thus, at the system-level the normative implications are more ambiguous. However, for $\gamma < \min\{\gamma_1^c, \ldots, \gamma_N^c\}$, it is clear that free-riding motives dominate for all banks and so there is a collective underinvestment in cybersecurity. While, for $\gamma > \max\{\gamma_1^c, \ldots, \gamma_N^c\}$ the influence of rollover risk is dominating for all banks, thereby leading to an overinvestment in cybersecurity at the system level.

5.1 Liquidity regulation

Our analysis offers a novel perspective on regulation aimed at ensuring stable bank funding. The parameter γ can be broadly interpreted as a measure of the fragility of banks' funding structures. A bank that is exposed to large rollover risk may be characterised, for example, by a small net stable funding ratio (NSFR). Our model suggests that by relaxing such liquidation regulations, banks would become more exposed to rollover risk and thereby invest more in cybersecurity. Moreover, if the regulations could be tailored for individual banks, such that for each bank, b, we have that $\gamma_b = \gamma_b^c$, the regulator can elicit the constrained efficient outcome that accounts for the interaction between ex ante free riding and ex post run risks. Blanket liquidity regulations, however, will always lead to sub-optimal outcomes given the heterogeneity across banks. But once broader risks to financial stability are taken into account, policymakers may opt for stricter liquidity regulation that reduces banks' rollover risk, albeit at the expense of greater cyber risk.

5.2 Duty of care

The regulator can achieve the Benchmark 1 outcome by establishing a negligence rule. To see this, consider the case $\gamma < \min\{\gamma_1^c, \ldots, \gamma_N^c\}$, where all banks under-invest in cybersecurity and free-riding dominates rollover risk concerns. We then have the following result.

Proposition 10. *The regulator can achieve the constrained efficient outcome by introducing a negligence rule at $t = 2$ with a penalty κ_b, in the event of a successful attack, that is proportional to the bank's investment, S_b^*, when it is less than S_b^R.*

The negligence rule establishes the regulator's solution as a minimum level of *due care* for each bank (Brown, 1973; Shavell, 2009). If a cyber attack is successful and banks suffer losses, then there is no further liability so long as all banks exercise the due care standard, i.e. so long as all banks exercise the due care standard, i.e. so long as each bank invests at least S_b^R. Otherwise, banks that exert insufficient care are penalised proportionally to the level of their under-investment with a penalty of κ_b.[11]

The penalty, κ_b, reflects the distance between the regulator's optimum and what the bank chooses in the absence of any intervention. Although the bank does not directly internalise how its investment impacts other banks, the penalty ensures that the constrained efficient allocation provides the best private outcome

[11]The ability of the US Securities and Exchange Commission (SEC) to sanction and fine financial firms for deficient cybersecurity procedures is an example of such a negligence rule. In 2021, the SEC fined eight firms \$200,000 – \$300,000 each for poor cybersecurity that resulted in the disclosure of customer information.

for the bank. And since the incentives to invest in cybersecurity are increasing in rollover risk, γ, we have that $\frac{\partial \kappa_b}{\partial \gamma} < 0$. The penalty required to enforce the negligence rule is smaller when rollover risk concerns prompt the bank to invest more in cybersecurity.

Imposing a conditional penalty increases the marginal returns to each additional unit invested in cybersecurity because the penalty erodes residual profits in the event of a cyber attack. So banks are better off substituting operational resilience for cybersecurity. Banks continue to substitute away from operational resilience until the constrained efficient outcome is reached. When constrained efficient, banks pay no penalty since their investment meets the standard of due care.

5.3 Cyber hygiene notices and stress tests

The regulator can also achieve the constrained efficient outcome by imposing that each bank invests S^R at $t = 0$. Clearly, constraining banks to invest at least the level set by the regulator will ensure that there is optimal provision of the public good. But the choices are sub-optimal for banks since they would choose to invest at lower levels. An example is the approach taken by the Monetary Authority of Singapore (MAS). The MAS sets minimum regulatory guidelines in the form of a Cyber Hygiene Notice that obliges banks to implement a set of cybersecurity measures, including network perimeter defenses, malware protection, and baseline configuration standards. Compliance with these regulatory requirements and expectations are verified and enforced by the MAS (Goh, Kang, Koh, Lim, Ng, Sher, and Yao, 2020).

The use of cyber stress tests, such as those implemented by the Bank of England, is another form of such a regulatory approach. The Financial Policy Committee of the Bank tests the resilience of the UK financial system to cyber attacks by requiring financial firms to meet a system-wide tolerance threshold set by the regulator (Kashyap and Wetherilt, 2019).

5.4 Testable implications

Our model suggests a rich set of testable hypotheses that may inform future empirical work. Data on cybersecurity investment can be obtained, for example, from the Network and Information Systems (NIS) survey conducted by the European Union Agency for Cybersecurity (ENISA). This annual survey provides information on firms, IT budgets and information security spending.

prediction 1. *Banks invest more on operational resilience following an uptick in cyber attacks.*

Our analysis suggests that, following an increase in the intensity of cyber attacks, banks will be incentivised to mitigate the consequences by setting up recovery procedures that trigger in the event of a breach. In the US, for example, a potential proxy for such mitigating action is the extent to which banks have obtained certification from an industry standard such as Sheltered Harbor for their resilience planning. Using state-level data, an empirical specification might regress the fraction of banks obtaining Sheltered Harbor certification against measures of how large and widespread a cyber attack is. Our model predicts that the coefficient should be positive.

prediction 2. *Greater rollover risk increases cybersecurity investment.*

Heightened rollover risk increases the marginal benefits of preventative cybersecurity relative to the marginal costs of mitigating actions on business continuity. In an empirical specification that regresses cybersecurity investment against banks' shares of uninsured and unsecured wholesale debt, the coefficient should be positive.

prediction 3. *Banks with lower leverage invest less in cybersecurity.*

Banks that predominantly fund investments with equity have more to lose from a cyber attack and will tend to invest in backup procedures and mitigating actions to lower failure risk. Our model suggests that there should be a positive coefficient when cybersecurity investment is regressed against measures of bank leverage.

prediction 4. *Banks invest more in cybersecurity following an increase in losses from cyber attacks.*

Banks that experience more costly attacks should, according to our analysis, invest more in cybersecurity than those experiencing less disruption since their incentives to protect equity value are greater. Information on bank losses following a cyber attack is available from data sources such as Advisen (Aldasoro et al., 2020).

6 Conclusion

We provide an analytical framework to show how cyber attacks might morph into bank runs, and which takes seriously the notion that cybersecurity is a public good (Mester, 2019). In our model, banks trade off strengthening their operational resilience to ward off private run risk against taking measures that benefit the security of the system as a whole. System-wide investment in cybersecurity is suboptimal as a result. We show how cybersecurity at the system level depends on the distribution of banks' working capital and derive several comparative static

results that lend themselves to empirical testing. We also discuss how negligence rules, cyber hygiene notices, and regulatory guidance on bank funding structures can facilitate constrained efficient outcomes.

Financial regulators are increasingly focussed on cyber risks to financial stability. For example, the European Central Bank has introduced a Threat Intelligence-Based Ethical Red Team (TIBER-EU) framework for EU-based financial entities (Panetta, 2020). And the Monetary Authority of Singapore has examined how banks' capital and liquidity buffers might cope in the face of a 24-hour system outage triggered by a cyber event (Goh et al., 2020). Our work provides a formal basis for such regulatory emphasis. In highlighting the important role of shared IT services in generating cyber risk dependencies across banks, our results highlight the importance of data on the network of linkages between banks and digital platforms. Such data might usefully inform "top-down" macroprudential cyber stress testing in much the same way as stress tests on interbank networks (Gai, Haldane, and Kapadia, 2011; Glasserman and Young, 2016).

Future work on the interaction between cyber risk management and rollover risk might usefully explore the role of cyber insurance markets in shaping the trade-offs identified in this paper. Deeper analysis of the drivers of the deadweight losses from cyber attacks is also warranted. Arguably, cyber attacks compromise the ability of a bank to both make and keep secret information and it is the loss of such information that is, ultimately, most devastating for the integrity of the financial system.

Bibliography

Adelmann, F., I. Ergen, T. Gaidosch, N. Jenkinson, A. Morozova, N. Schwarz, and C. Wilson (2020). Cyber risk and financial stability: It's a small world after all. Staff Discussion Notes (007), International Monetary Fund, Washington, DC.

Aldasoro, I., J. Frost, L. Gambacorta, and D. Whyte (2021). Covid-19 and cyber risk in the financial sector. BIS Bulletin No. 37.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach (2020). The drivers of cyber risk. BIS Working Paper No. 865.

Bergstrom, T., L. Blume, and H. Varian (1986). On the private provision of public goods. Journal of Public Economics 29(1), 25–49.

Boot, A., P. Hoffmann, L. Laeven, and L. Ratnovski (2021). Fintech: what's old, what's new? Journal of Financial Stability 53, 100836.

Bouveret, A. (2018). Cyber risk for the financial sector: A framework for quantitative assessment. International Monetary Fund, Working Paper No. 18/143, Washington DC.

Brown, J. P. (1973). Toward an economic theory of liability. The Journal of Legal Studies 2(2), 323–349.

Calomiris, C. and C. Kahn (1991). The role of demandable debt in structuring optimal banking arrangements. American Economic Review 81(3), 497–513.

Cornes, R. (1993). Dyke maintenance and other stories: Some neglected types of public goods. The Quarterly Journal of Economics 108(1), 259–271.

Coveware (2021). Q2 ransom payment amounts decline as ransomware becomes a national security priority. Available at https://www.coveware.com/blog/2021/7/23/q2-ransom-payment-amounts-decline-as-ransomware-becomes-a-national-security-priority.

Dang, T. V., G. Gorton, B. Holmström, and G. Ordonez (2017). Banks as secret keepers. American Economic Review 107(4), 1005–29.

Daniel, E. (2020). GoDaddy: Hackers gain control of crypto domains in social engineering scam. Verdict.

Deutsche Bundesbank (2021). Digital risks in the banking sector. Monthly Report.

Diamond, D. and P. Dybvig (1983). Bank runs, deposit insurance and liquidity. Journal of Political Economy 91, 401–419.

Diamond, D. and R. Rajan (2001). Liquidity risk, liquidity creation, and financial fragility: A theory of banking. Journal of Political Economy 109(2), 287–327.

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

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

Florakis, C., C. Louca, R. Michaely, and M. Weber (2020). Cybersecurity risk. NBER Working Paper 28196.

Frankel, D., S. Morris, and A. Pauzner (2003). Equilibrium selection in global games with strategic complementarities. Journal of Economic Theory 108(1), 1–44.

Gai, P., A. Haldane, and S. Kapadia (2011). Complexity, concentration and contagion. Journal of Monetary Economics 58(5), 453–470.

Glasserman, P. and H. P. Young (2016). Contagion in financial networks. Journal of Economic Literature 54(3), 779–831.

Goh, J., H. Kang, Z. X. Koh, J. W. Lim, C. W. Ng, G. Sher, and C. Yao (2020). Cyber risk surveillance: A case study of Singapore. MAS Staff Paper No. 57.

Goldstein, I. and A. Pauzner (2005). Demand deposit contracts and the probability of bank runs. Journal of Finance 60(3), 1293–1327.

Gordon, L. and M. Loeb (2002). The economics of information security investment. ACM Transactions on Information and System Security 5(4), 438–457.

Grossklags, J., N. Christin, and J. Chuang (2008). Secure or insure? A game-theoretic analysis of information security games. In WWW '08: Proceedings of the 17th international conference on World Wide Web, pp. 209–218.

Hayden, M. (2011). Statement for the Record, House Permanent Select Committee on Intelligence, The Cyber Threat. National Security Agency. Available at https://www.hsdl.org/?view&did=689629.

Hirshleifer, J. (1983). From weakest-link to best-shot: The voluntary provision of public goods. Public Choice 41(3), 371–386.

Jamilov, R., H. Rey, and A. Tahoun (2021). The anatomy of cyber risk. NBER Working Paper No. 28906.

Kamiya, S., J.-K. Kang, J. Kim, A. Milidonis, and R. M. Stulz (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. Journal of Financial Economics 139(3), 719–749.

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

Kocher, P., J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, et al. (2019). Spectre attacks: Exploiting speculative execution. In 2019 IEEE Symposium on Security and Privacy (SP), pp. 1–19. IEEE.

Kopp, E., L. Kaffenberger, and C. Wilson (2017). Cyber risk, market failures, and financial stability. International Monetary Fund, Working Paper No. 17/185.

LeRoy, S. F. and L. D. Singell (1987). Knight on risk and uncertainty. Journal of Political Economy 95(2), 394–406.

Lipp, M., M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg (2018). Meltdown: Reading kernel memory from user space. In 27th USENIX Security Symposium (USENIX Security 18), Baltimore, MD, pp. 973–990. USENIX Association. August.

Mauer, T. and A. Nelson (2020). International strategy to better protect the financial system against cyber threats. Technical report, Carnegie Endowment for International Peace.

McDonald, G., L. Murchu, S. Doherty, and E. Chien (2013). Stuxnet 0.5: The missing link. Symantec security response, Symantec.

Mester, L. J. (2019). Cybersecurity and financial stability. Speech at the Federal Reserve Bank of Cleveland, Cleveland, Ohio. 21 November.

Morris, S. and H. Shin (2003). Global games: Theory and applications. In M. Dewatripont, L. Hansen, and S. Turnovsky (Eds.), Advances in Economics and Econometrics (Proceedings of the 8th World Congress of the Econometric Society). Cambridge University Press.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. American Economic Review 88(3), 587–597.

Panetta, F. (2020). Keeping cyber risk at bay: our individual and joint responsibility. Introductory remarks at the fifth meeting of the Euro Cyber Resilience Board for pan-European Financial Infrastructures. Frankfurt, 16 December.

Perlroth, N. (2021). This is how they tell me the world ends: The cyberweapons arms race. Bloomsbury Publishing.

Rochet, J.-C. and X. Vives (2004). Coordination failures and the lender of last resort: was Bagehot right after all? Journal of the European Economic Association 2(6), 1116–47.

Samuelson, P. (1954). The pure theory of public expenditure. The Review of Economics and Statistics 36(4), 387–389.

Shavell, S. (2009). Economic analysis of accident law. Harvard University Press.

S&P Global Market Intelligence (2019). S&P downgrades Malta-based Bank of Valletta. https://www.spglobal.com/marketintelligence/en/news-insights/trending/5mvfiykwlxliliri78qd-q2.

Tarabay, J. (2021). How a dated cyber-attack brought a stock exchange to its knees. Bloomberg Businessweek.

Van Zandt, T. and X. Vives (2007). Monotone equilibria in Bayesian games of strategic complementarities. Journal of Economic Theory 134(1), 339–360.

Varian, H. (2004). System reliability and free riding. In L. J. Camp and S. Lewis (Eds.), Economics of information security, pp. 1–15. Springer.

Vives, X. (2014). Strategic complementarity, fragility, and regulation. Review of Financial Studies 27(12), 3547–92.

Woods, D. W., T. Moore, and A. C. Simpson (2021). The county fair cyber loss distribution: Drawing inferences from insurance prices. Digital Threats: Research and Practice 2(2), 1–21.

Mathematical Appendix

A1 Proof of Lemma 1

Bank b fails due to insolvency whenever

$$\alpha > \alpha_b^{IN} \equiv \frac{R - FD}{R\delta[1 - h(O_b)]}, \tag{A1}$$

while it fails due to illiquidity whenever

$$\alpha > \alpha_b^{IL}(\ell_b) \equiv \frac{R - \ell_b FD}{R[1 - h(O_b)]}. \tag{A2}$$

While α_b^{IN} is invariant to the proportion of withdrawals, the threshold $\alpha_b^{IL}(\ell_b)$ is decreasing in ℓ_b. The proportion of withdrawals, $\hat{\gamma}$, for which the two failure conditions intersect is given by $\alpha_b^{IL}(\hat{\gamma}) = \alpha_b^{IN}$, i.e.,

$$\frac{R - FD}{R\delta[1 - h(O_b)]} = \frac{R - \hat{\gamma}FD}{R[1 - h(O_b)]}, \tag{A3}$$

which is independent of the amount invested in operational resilience.

A2 Proof of Proposition 1

The proof is in three steps. First, we show that the dominance regions at $t = 1$ are well defined. If all fund managers withdraw early, $\ell_b = 1$, then the illiquidity failure threshold is given by $\alpha_b^{IL}(1) = \frac{R-FD}{R[1-h(O_b)]}$, where $\alpha_b^{IL}(1) < \alpha_b^{IN}$. If, however, no fund manager withdraws at $t = 1$, then $\ell_b = 0$. In this case, the bank never fails due to illiquidity since $\alpha_b^{IL}(0) = \frac{1}{1-h(O_b)} > 1$. But the bank can, nevertheless, fail at $t = 2$ due to insolvency whenever $\alpha > \alpha_b^{IN}$, since $\alpha_b^{IN} < 1$.

Under Assumption 1, it follows that $\underline{\alpha_b} \equiv \alpha_b^{IL}(1) > 0$ is the largest accessibility shock under which bank b's survival is independent of the number of fund managers withdrawing early. When $\alpha \in [0, \underline{\alpha_b}]$, fund managers have a dominant strategy to roll over their claims. Let $\overline{\alpha_b} \equiv \alpha_b^{\overline{IN}} < 1$ denote the upper dominance bound, beyond which bank b fails regardless of the number of fund managers who withdraw early. When $\alpha \in [\overline{\alpha}_b, 1]$, fund managers have a dominant strategy to withdraw early. Finally, for $\alpha \in (\underline{\alpha}_b, \overline{\alpha}_b)$, If the outage shock were common knowledge, the run dynamics of fund managers would be characterized by multiple, self-fulfilling equilibria, as shown in Figure A1.

Second, we define a threshold strategy which, for well defined dominance

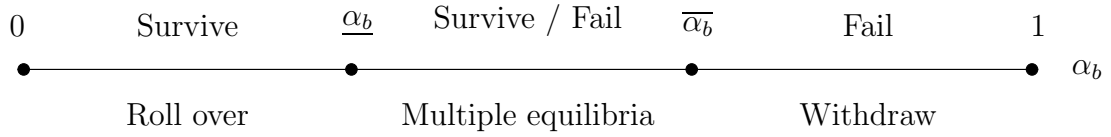

Figure A1: Tripartite classification of the accessibility shock.

bounds and sufficiently precise private information, survives iterated deletion of strictly dominated strategies (Morris and Shin, 2003; Frankel, Morris, and Pauzner, 2003). Denote this threshold point x_b^*, with corresponding strategy

$$s(x_{bk}) = \begin{cases} \text{withdraw} & \text{if } x_{bk} > x_b^*, \\ \text{roll over} & \text{if } x_{bk} \leq x_b^*. \end{cases} \tag{A4}$$

Finally, we characterise this equilibrium. With switching point x_b^*, the proportion of fund managers who withdraw at $t = 1$, given some realisation of the outage shock α is

$$\ell_b(\alpha, x_b^*) = \Pr(x_{bk} > x_b^* | \alpha) = 1 - \Phi\left(\frac{x_b^* - \alpha}{\sigma}\right), \tag{A5}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. At some critical mass of withdrawals, $\ell_b^*(\alpha_b^*, x_b^*)$, the bank just reaches its survival threshold, α_b^*:

$$[R(1 - \alpha_b^*[1 - h(O_b)]) - \ell_b^*(\alpha_b^*, x_b^*)FD]_+ + R\alpha_b^*[1-h(O_b)](1-\delta) = [1-\ell_b^*(\alpha_b^*, x_b^*)]FD, \tag{A6}$$

whenever $\ell_b^*(\alpha_b^*, x_b^*) \leq \widehat{\gamma}$, and

$$R(1 - \alpha_b^*[1 - h(O_b)]) = \ell_b^*(\alpha_b^*, x_b^*)FD, \tag{A7}$$

whenever $\ell_b^*(\alpha_b^*, x_b^*) > \widehat{\gamma}$.

To determine $\Pr(\alpha > \alpha_b^* | x_{ik})$, fund managers use Bayes' rule. At the threshold signal x_b^*, fund managers should be indifferent between withdrawing and rolling over, so that

$$\gamma = \Pr(\alpha \leq \alpha_b^* | x_{bk} = x_b^*). \tag{A8}$$

Let $G(\alpha | x_{bk} = x, \sigma)$ denote the posterior distribution over the accessibility shock, conditional on observing x_{bk}.

From equation (A5), we can write the critical accessibility shock as follows

$$\alpha_b^* = x_b^* - \sigma\Phi^{-1}\left(1 - \ell_b^*\right). \tag{A9}$$

Inserting this into (A8), we have

$$\gamma = G\left(x_b^* - \sigma\Phi^{-1}\left(1-\ell_b^*\right)|x_{bk} = x_b^*, \sigma\right). \tag{A10}$$

As private signals become infinitely precise, i.e., as $\sigma \to 0$, the right-hand side of condition (A10) simplifies to

$$G(x_b^* - \sigma\Phi^{-1}\left(1-\ell_b^*\right)|x_{bk}, \sigma) \to \Phi\left(-\Phi^{-1}(1-\ell_b^*)\right). \tag{A11}$$

The anticipated critical mass of withdrawals by a fund manager who observes $x_{bk} = x_b^*$ thus approaches $\ell_b^* = \gamma$. Substituting for ℓ_b^* in (A6) and (A7) gives us a unique outage shock

$$\alpha_b^* = \begin{cases} \alpha_b^{IN} \equiv \frac{R-FD}{R\delta[1-h(O_b)]} & \text{if } \gamma < \widehat{\gamma}, \\ \alpha_b^{IL}(\gamma) \equiv \frac{R-\gamma FD}{R[1-h(O_b)]} & \text{if } \gamma \geq \widehat{\gamma}, \end{cases} \tag{A12}$$

and the run threshold in Proposition 1 follows.

A3 Proof of Lemma 2

Under a binding budget constraint, we have $O_b = W_b - S_b$. First, suppose that $\gamma \leq \widehat{\gamma}$ so that banks fail due to insolvency. Taking the partial derivative of α_b^{IN} with respect to S_b, we have

$$\frac{\partial \alpha_b^{IN}}{\partial S_b} = \frac{\alpha_b^{IN}(S_b)}{[1-h(W_b-S_b)]}\frac{\partial h}{\partial S_b} < 0\,. \tag{A13}$$

Since operational resilience is lower when the bank invests more in cybersecurity, i.e., $\partial h/\partial S_b < 0$.

Next, suppose that $\gamma \geq \widehat{\gamma}$ so that banks fail due to illiquidity. Similarly, we have

$$\frac{\partial \alpha_b^{IL}}{\partial S_b} = \frac{\alpha_b^{IL}(S_b)}{[1-h(W_b-S_b)]}\frac{\partial h}{\partial S_b} < 0\,, \tag{A14}$$

since, as above, $\partial h/\partial S_b < 0$. Because the bank fails whenever the outage shock is above the threshold, fragility increases as the threshold falls. Therefore, fragility is increasing in cybersecurity investments.

A4 Samuelson benchmark conditions

We next set out the optimisation problem for the planner and derive the Samuelson (1954) rule for each of Benchmarks 1 and 2.

The planner optimises the combined surplus from each bank $b = 1, ..., N$, given by expected equity value

$$\Pi = \sum_{b=1}^{N} \left[(\chi/\overline{\lambda})(R - F\,D) + [1 - (\chi/\overline{\lambda})] \left(\int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b) d\alpha \right) \right], \qquad \text{(A15)}$$

where, for $\gamma < \widehat{\gamma}$, bank b's failure condition is α_b^{IN}, and subject to the level of cybersecurity given by

$$\chi = \left(\prod_{b=1}^{N} S_b \right)^{\frac{1}{N}},$$

and the aggregate resource constraint $\sum_{b=1}^{N} W_b = \sum_{b=1}^{N} (S_b + O_b)$. The planner's Lagrangian is, thus,

$$\begin{aligned} \mathcal{L} = \sum_{b=1}^{N} & \left[(\chi/\overline{\lambda})(R - FD) + [1 - \chi/\overline{\lambda}] \left(\int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b) d\alpha \right) \right] \\ & + \phi_1 \left(\sum_{b=1}^{N} W_b - \sum_{b=1}^{N} O_b - \sum_{b=1}^{N} S_b \right) + \phi_2 \left(\left(\prod_{b=1}^{N} S_b \right)^{\frac{1}{N}} - \chi \right). \end{aligned} \qquad \text{(A16)}$$

The necessary and sufficient Kuhn-Tucker conditions are

$$(1) \quad \frac{\partial \mathcal{L}}{\partial O_b} = 0 : (\overline{\lambda} - \chi) \left[\int_0^{\alpha_b^{IN}} \frac{\partial EV_2}{\partial O_b} d\alpha \right] = \overline{\lambda} \phi_1 \quad \forall b = 1, ..., N;$$

$$(2) \quad \frac{\partial \mathcal{L}}{\partial \chi} = 0 : \sum_{b=1}^{N} \left[R - FD - \int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b) d\alpha \right] = \overline{\lambda} \phi_2;$$

$$(3) \quad \frac{\partial \mathcal{L}}{\partial S_b} = 0 : \phi_2 \frac{\left(\prod_{b=1}^{N} S_b \right)^{\frac{1}{N}}}{N S_b} = \phi_1 \quad \forall b = 1, ..., N.$$

From conditions (2) and (3), it is clear that

$$\frac{(1/\overline{\lambda}) \sum_{b=1}^{N} \left[R - FD - \int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b) d\alpha \right]}{\phi_1} = \frac{1}{\left(\prod_{b=1}^{N} S_b \right)^{\frac{1}{N}} / N S_b}.$$

Substituting for ϕ_1 from condition (1), we have

$$\sum_{b=1}^{N} \frac{\left[R - FD - \int_0^{\alpha_b^{IN}} EV_2(\alpha, O_b)d\alpha\right]}{(\overline{\lambda} - \chi)\left[\int_0^{\alpha_b^{IN}} \frac{\partial EV_2}{\partial O_b} d\alpha\right]} = \frac{1}{\partial\chi/\partial S_b}\,. \tag{A17}$$

In the case with rollover risk, $\gamma > \widehat{\gamma}$, the critical outage shock is $\alpha_b^* = \frac{R - \gamma\, F\, D}{R[1-h(O_b)]}$. We repeat the same exercise to obtain Benchmark 2. In the main text, we argue that Benchmark 2 elicits an over-investment in cybersecurity relative to Benchmark 1. To see this, consider the marginal rate of substitution for bank b under Benchmark 2, which is given by

$$\frac{(R - F\,D) - \int_0^{\alpha_b^*} EV_2(\alpha, O_b)d\alpha}{(\bar{\lambda} - \chi)\left[EV_2\big(\alpha_b^*\big)\frac{\partial \alpha_b^*}{\partial O_b} + \int_0^{\alpha_b^*}\Big(\partial EV_2/\partial O_b\Big)d\alpha\right]}\,. \tag{A18}$$

An increase in rollover risk, γ causes the numerator to increase, since $\partial\alpha_b^*/\partial\gamma < 0$ for all banks. But it also causes the denominator to decrease, since

$$(\overline{\lambda} - \chi)\left[\frac{\partial EV_2(\alpha_b^*)}{\partial\gamma}\frac{\partial\alpha_b^*}{\partial O_b} + EV_2(\alpha_b^*)\frac{\partial^2\alpha_b^*}{\partial O_b\partial\gamma} + R\delta\alpha_b^*\frac{\partial h}{\partial O_b}\frac{\partial\alpha_b^*}{\partial\gamma}\right] < 0\,.$$

The first and last terms cancel out, leaving the middle term which is negative, since

$$\frac{\partial^2\alpha_b^*}{\partial O_b\partial\gamma} = \frac{\partial h/\partial O_b}{1 - h(O_b)}\frac{\partial\alpha_b^*}{\partial\gamma} < 0\,.$$

Therefore, with rollover risk, the marginal rates of substitution for each bank is higher, leading to greater investment in cybersecurity relative to Benchmark 1.

A5 Proof of Proposition 2

Each bank, $b = 1, \ldots, N$, chooses cybersecurity investments, S_b, and operational resilience, O_b, to maximise its expected equity value. Repeating an analogous exercise to that in the previous section – but for a single bank – produces the

condition in Proposition 2. Substituting for $O_b = W_b - S_b$, we have that

$$\frac{\partial \pi_b}{\partial S_b} = \frac{\partial \chi}{\partial S_b}\left(R - FD - \int_0^{\alpha_b^*} EV_2(\alpha)d\alpha\right) + (\overline{\lambda} - \chi)\left(EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial S_b} + \int_0^{\alpha_b^*} \frac{\partial EV_2}{\partial S_b}d\alpha\right). \tag{A19}$$

The optimal S_b^* is thus given by the first-order condition $\frac{\partial \pi_b}{\partial S_b}\Big|_{S_b = S_b^*} = 0$. Since $\lim_{S_b \to 0} \partial\chi/\partial S_b = +\infty$, it follows that $\lim_{S_b \to 0} \frac{\partial \pi_b}{\partial S_b} > 0$. Thus, it is always optimal for the bank to invest a positive level in cybersecurity, given that other banks also invest at a positive level as well. Moreover, for $S_b \to W_b$, as long as $h(0) > 1 - \frac{2(\alpha_b^*(0)-1)(R-FD)}{R\delta(\alpha_b^*(0))^2}$, where $\alpha_b^*(0)$ is the bank's illiquidity threshold when the bank invests nothing in operational resilience, $O_b = 0$, the first term in Equation (A19) is negative. And since the second term is strictly negative, it follows that $\lim_{S_b \to W_b} \frac{\partial \pi_b}{\partial S_b} < 0$. Therefore it is never optimal for the bank to invest everything in cybersecurity. Thus, by the intermediate value theorem, an optimum for the investment in cybersecurity exists within the interval $(0, W_b)$ for each bank $b = 1, \ldots, N$.

Finally, we need to show that the optimum is a maximum. For this, the second partial derivative of bank b's expected equity value with respect to cybersecurity investment is given by

$$\begin{aligned}\frac{\partial^2 \pi_b}{\partial S_b^2} = &\left(\frac{\partial^2 \chi}{\partial S_b^2}\Big/\overline{\lambda}\right)\left[R - FD - \int_0^{\alpha_b^*} EV_2(\alpha)d\alpha\right] \\ &- \left(2\frac{\partial \chi}{\partial S_b}\Big/\overline{\lambda}\right)\left[EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial S_b} + R\delta\frac{\partial h}{\partial S_b}\frac{(\alpha_b^*)^2}{2}\right] \\ &+ (1 - \chi/\overline{\lambda})\left[EV_2(\alpha_b^*)\frac{\partial^2 \alpha_b^*}{\partial S_b^2} + R\delta\frac{\partial^2 h}{\partial S_b^2}\frac{(\alpha_b^*)^2}{2}\right].\end{aligned}$$

Evaluating the above expression at $S_b = S_b^*$, we get

$$\begin{aligned}\frac{\partial^2 \pi_b}{\partial S_b^2}\Big|_{S_b = S_b^*} = &\left(R - FD - \int_0^{\alpha_b^*} EV_2(\alpha)d\alpha\right)\left[\frac{\partial^2 \chi}{\partial S_b^2} + \frac{2\left(\frac{\partial \chi}{\partial S_b}\right)^2}{1 - \chi}\right] \\ &+(1 - \chi)\left\{EV_2(\alpha_b^*)\frac{\partial^2 \alpha_b^*}{\partial S_b^2} + R\delta(\alpha_b^*)^2\left(\frac{1}{2}\frac{\partial^2 h}{\partial S_b^2} + \frac{(\partial h/\partial S_b)^2}{1 - h(W_b - S_b^*)}\right)\right\}.\end{aligned} \tag{A20}$$

For ease of notation, we only sparingly highlight some terms that depend on the equilibrium level of investment in cybersecurity. Denoting by $\overline{W} = \max\{W_1, \ldots, W_N\}$, then as long as $N > \left\lceil \frac{1+W_N}{1-W_N} \right\rceil$, the term in the squared brackets in Equation (A20)

is negative. Next, note that

$$\frac{\partial^2 \alpha_b^*}{\partial S_b^2} = \frac{2\alpha_b^*}{1-h(W_b - S_b^*)}\left[\frac{1}{2}\frac{\partial^2 h}{\partial S_b^2} + \frac{(\partial h/\partial S_b)^2}{1-h(W_b-S_b)}\right].$$

Rearranging the term inside the curly braces on the second line in Equation (A20) gives us

$$\left(EV_2(\alpha_b^*)\frac{2\alpha_b^*}{1-h(W_b-S_b^*)} + R\delta(\alpha_b^*)^2\right)\left[\frac{1}{2}\frac{\partial^2 h}{\partial S_b^2} + \frac{(\partial h/\partial S_b)^2}{1-h(W_b-S_b^*)}\right]. \tag{A21}$$

Thus a sufficient condition for for $\partial^2\pi_b/\partial S_b^2|_{S_b^*} < 0$ is

$$\frac{1}{2}\frac{\partial^2 h}{\partial S_b^2} + \frac{(\partial h/\partial S_b)^2}{1-h(W_b-S_b^*)} < 0\,. \tag{A22}$$

An example of a function that satisfies this condition is $h(O_b) = \zeta + O_b^{\psi}$, where $\zeta > 0$ and $\psi < 1$. In particular, whenever $\zeta < 1 - \frac{\overline{W}^{\psi}(1+\psi)}{(1-\psi)}$, we have that the condition in Equation (A22) is satisfied and that $\partial^2\pi_b/\partial S_b^2|_{S_b=S_b^*} < 0$, and so the optimum is indeed a maximum.

A6 Proof of Proposition 3

By the implicit function theorem, the equilibrium investment by bank b changes in response to a change in endowment, W_b, (for a fixed level of investment by all other contracting banks) as follows

$$\frac{\partial S_b^*}{\partial W_b} = \frac{-1}{\frac{\partial^2\pi_b}{\partial S_b^2}}\left[\frac{\partial\chi}{\partial S_b}\left(-EV_2(\alpha_b^*)\frac{\partial\alpha_b^*}{\partial W_b} - R\delta\frac{(\alpha_b^*)^2}{2}\frac{\partial h}{\partial W_b}\right) + [1-\chi(S_b)]\left\{EV_b(\alpha_b^*)\frac{\partial^2\alpha_b^*}{\partial S_b\partial W_b}\right.\right.$$
$$\left.\left. + R\delta(\alpha_b^*)^2\left(\frac{1}{2}\frac{\partial^2 h}{\partial S_b\partial W_b} + \frac{1}{1-h(W_b-S_b^*)}\frac{\partial h}{\partial S_b}\frac{\partial h}{\partial W_b}\right)\right\}\right]. \tag{A23}$$

Rearranging the cross derivative, we have

$$\alpha_b^*\left\{\frac{2EV_2(\alpha_b^*)}{1-h(W_b-S_b)} + R\delta\alpha_b^*\right\}\left[\frac{\partial h/\partial S}{1-h(W_b-S_b^*)} + \frac{\partial^2 h/\partial S_b\partial W_b}{2\partial h/\partial W_b} - \frac{(1/\overline{\lambda})\partial\chi/\partial S_b}{2(1-\chi/\overline{\lambda})}\right].$$

Both terms inside braces are positive. Therefore, a necessary and sufficient

condition for S_b^* to be increasing in W_b is

$$\frac{\partial^2 h/\partial S_b \partial W_b}{\partial h/\partial W_b} + \frac{2\partial h/\partial S_b}{1-h(W_b-S_b^*)} > \frac{\frac{\partial \chi}{\partial S_b}\Big/\overline{\lambda}}{1-\chi/\overline{\lambda}}\,. \tag{A24}$$

Let $\widehat{W}$ be the level of endowment such that $\partial S_b^*/\partial W_b = 0$. Using the example above, where $h(O_b) = \zeta + O_b^{\psi}$, the threshold, $\widehat{W}$, is defined implicitly by

$$\left\{\psi + 1 - \frac{2\psi(1-\zeta)}{1-\zeta-(\widehat{W}-\widehat{S})^{\psi}}\right\} = (\widehat{W}-\widehat{S})\left(\frac{\frac{\partial \chi}{\partial S_b}\Big/\overline{\lambda}}{1-\chi/\overline{\lambda}}\right), \tag{A25}$$

where $\widehat{S}$ is the equilibrium investment in cybersecurity elicited by the endowment, $\widehat{W}$. To establish that $\widehat{W}$ is well defined, first consider an endowment, W_b, that is marginally larger than $\widehat{W}$. Given $\widehat{S}$, this would imply that $\partial S_b^*/\partial W_b < 0$. And so the new investment in cybersecurity is lower than $\widehat{S}$. Consequently, $W_b - S_b^* > \widehat{W} - \widehat{S}$ for $W > \widehat{W}$. This, in turn implies that the violation of the condition in Equation (A24) grows in equilibrium. Next consider a W_b that is marginally smaller than $\widehat{W}$. As before, given $\widehat{S}$ this implies that at W_b, we have that $\partial S_b^*/\partial W_b > 0$. Consequently, $S_b > \widehat{S}$ and therefore $W_b - S_b^* < \widehat{W} - \widehat{S}$, which implies that the condition in Equation (A24) continues to hold in equilibrium.

A7 Proof of Proposition 4

By the implicit function theorem, the response of S_b^* to a marginal increase in $\overline{\lambda}$ is given by

$$\partial S_b^*/\partial \overline{\lambda} = \frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}}\left\{EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial S_b} + \int_0^{\alpha_b^*}\frac{\partial EV_2(\alpha)}{\partial S_b}d\alpha\right\} < 0\,. \tag{A26}$$

The terms inside the braces are both negative, since

$$\frac{\partial \alpha_b^*}{\partial S_b} = \frac{\alpha_b^*}{1-h[W_b-S_b]}\frac{\partial h}{\partial S_b} < 0\,,$$

and

$$\int_0^{\alpha_b^*}\frac{\partial EV_2(\alpha, S_b)}{\partial S_b}d\alpha = R\delta\frac{(\alpha_b^*)^2}{2}\frac{\partial h}{\partial S_b} < 0\,.$$

Together, implies that $\partial S_b^*/\partial \overline{\lambda} < 0$.

To see that fragility is decreasing in $\overline{\lambda}$, note that fragility is affected only

indirectly by the change in S_b^*. Therefore,

$$\frac{d\alpha_b^*}{d\overline{\lambda}} = \frac{\partial \alpha_b^*}{\partial S_b}(\partial S_b^*/\partial \overline{\lambda}) > 0\,. \tag{A27}$$

A marginal increase in the deadweight loss, δ, on the other hand, increases equilibrium cybersecurity investment:

$$\frac{\partial S_b^*}{\partial \delta} = \frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}} \Bigg\{ \frac{\partial \chi}{\partial S_b} R[1 - h(W_b - S_b^*)] \frac{(\alpha_b^*)^2}{2} + (\overline{\lambda} - \chi) \left[-R\alpha_b^*[1 - h(W_b - S_b^*)] \frac{\partial \alpha_b^*}{\partial S_b} + R \frac{(\alpha_b^*)^2}{2} \frac{\partial h}{\partial S_b} \right] \Bigg\},$$

which can be simplified to

$$\frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}} \left\{ \frac{\partial \chi}{\partial S_b} R[1 - h(W_b - S_b^*)] \frac{(\alpha_b^*)^2}{2} \right\} > 0\,. \tag{A28}$$

Since α_b^* is independent of δ, fragility is affected only indirectly by the deadweight loss:

$$\frac{d\alpha_b^*}{d\delta} = \frac{\partial \alpha_b^*}{\partial S_b} \frac{\partial S_b^*}{\partial \delta} < 0\,. \tag{A29}$$

A8 Proof of Proposition 5

Bank equity, $E = 1 - D$, is used to fund the portion of the project that is not funded by debt. Banks' equilibrium investments respond as follows to a marginal increase in D as follows

$$\frac{\partial S_b^*}{\partial D} = \frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}} \Bigg\{ \frac{\partial \chi}{\partial S_b} \left(-F(1 - \alpha_b^*) - EV_2(\alpha_b^*) \frac{\partial \alpha_b^*}{\partial D} \right) + (\overline{\lambda} - \chi) \left[\frac{\partial \alpha_b^*}{\partial S_b} \frac{\partial EV_2}{\partial D} + EV_2(\alpha_b^*) \frac{\partial h / \partial S_b}{1 - h(W_b - S_b^*)} \frac{\partial \alpha_b^*}{\partial S_b} + R\delta\alpha_b^* \frac{\partial h}{\partial S_b} \frac{\partial \alpha_b^*}{\partial D} \right] \Bigg\}. \tag{A30}$$

All the terms in the square brackets are positive. Therefore, a sufficient condition for the expression in (A30) to be increasing is

$$\gamma > \frac{-Rh(W_b - S_b^*)}{EV_2(\alpha_b^*) - F\,D}\,.$$

From Lemma 1 we have,

$$\widehat{\gamma} = \frac{1}{\delta}\left[1 - \frac{R(1-\delta)}{F\,D}\right],$$

and since $EV_2(\alpha_b^*) = 0$ at $\widehat{\gamma}$, the condition is satisfied by Assumption 1. Substituting for $D = 1 - E$, we have $\partial S_b^*/\partial E < 0$.

The effect on bank fragility is two-fold.

$$\frac{d\alpha_b^*}{dE} = \frac{\partial \alpha_b^*}{\partial E} + \frac{\partial \alpha_b^*}{\partial S_b}\frac{\partial S_b^*}{\partial E} > 0\,. \tag{A31}$$

An increase in equity directly reduces fragility, since a lower portion of the bank's project is funded by debt that is subject to runs. Moreover, an increase in equity causes banks to allocate more of their resources towards operational resilience, causing a reduction in cybersecurity investment. This also lowers fragility by bolstering operational resilience.

A9 Proof of Proposition 6

As rollover risk rises, the response of each bank's equilibrium investment in cybersecurity is given by

$$\begin{aligned}\frac{\partial S_b^*}{\partial \gamma} = \frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}}\Bigg\{ &- \frac{\partial \chi}{\partial S_b} EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial \gamma} + (\overline{\lambda} - \chi)\Big[- R\delta[1 - h(W_b - S_b^*)]\frac{\partial \alpha_b^*}{\partial S_b}\frac{\partial \alpha_b^*}{\partial \gamma} \\ &+ EV_2(\alpha_b^*)\frac{\partial h/\partial S_b}{1 - h(W_b - S_b^*)}\frac{\partial \alpha_b^*}{\partial \gamma}\Big] + R\delta\frac{\partial h}{\partial S_b}\alpha_b^*\frac{\partial \alpha_b^*}{\partial \gamma}\Bigg\},\end{aligned} \tag{A32}$$

which simplifies to

$$\frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}}\left\{-\frac{\partial \chi}{\partial S_b} EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial \gamma} + (\overline{\lambda} - \chi)EV_2(\alpha_b^*)\frac{\partial h/\partial S_b}{1 - h(W_b - S_b^*)}\frac{\partial \alpha_b^*}{\partial \gamma}\right\} > 0\,. \tag{A33}$$

The first term in the parenthesis is positive since $\partial \alpha_b^*/\partial \gamma = \frac{-FD}{R[1-h(W_b - S_b)]} < 0$, and the second term is also positive since $\partial h/\partial S_b < 0$. Therefore, $\partial S_b^*/\partial \gamma > 0$.

An increase in rollover risk introduces a direct and indirect effect on bank fragility.

$$\frac{d\alpha_b^*}{d\gamma} = \frac{\partial \alpha_b^*}{\partial \gamma} + \frac{\partial \alpha_b^*}{\partial S_b}\frac{\partial S_b^*}{\partial \gamma} < 0\,. \tag{A34}$$

For each outage shock, the likelihood of bank failure increases in γ, which rep-

resents the equilibrium proportion of fund manager withdrawals in Proposition 1, and so $\partial \alpha_b^*/\partial \gamma < 0$. Furthermore, banks reallocate resources towards cybersecurity in an effort to avoid an outage entirely which, by Lemma 2, increases bank fragility.

A10 Proof of Lemma 3

By the implicit function theorem, the marginal response of bank b to a change in investment by one other bank, b', holding all other bank investments fixed, is given by

$$\frac{\partial S_b^*}{\partial S_{b'}} = \frac{-1}{\frac{\partial^2 \pi_b}{\partial S_b^2}} \left[\frac{\partial^2 \chi}{\partial S_b S_{b'}} \left(R - FD - \int_0^{\alpha_b^*} EV_2(\alpha) d\alpha \right) - \frac{\partial \chi}{\partial S_{b'}} \left(EV_2(\alpha_b^*) \frac{\partial \alpha_b^*}{\partial S_b} + R\delta \frac{(\alpha_b^*)^2}{2} \frac{\partial h}{\partial S_b} \right) \right].$$

By concavity of the profit function, $\partial^2 \pi_b / \partial S_b^2 < 0$, bank b's response is increasing if and only if

$$\left[\frac{\partial^2 \chi}{\partial S_b S_{b'}} \left(R - FD - \int_0^{\alpha_b^*} EV_2(\alpha) d\alpha \right) - \frac{\partial \chi}{\partial S_{b'}} \left(EV_2(\alpha_b^*) \frac{\partial \alpha_b^*}{\partial S_b} + R\delta \frac{(\alpha_b^*)^2}{2} \frac{\partial h}{\partial S_b} \right) \right] \geq 0.$$

Using the first order condition, we can rewrite the above inequality as follows

$$\left(EV_2(\alpha_b^*) \frac{\partial \alpha_b^*}{\partial S_b} + R\delta \frac{(\alpha_b^*)^2}{2} \frac{\partial h}{\partial S_b} \right) \left[\frac{-(\overline{\lambda} - \chi(S_b, \vec{S}_{-b})) \frac{\partial^2 \chi}{\partial S_b S_{b'}}}{\partial \chi / \partial S_b} - \frac{\partial \chi}{\partial S_{b'}} \right].$$

Since both $\partial \chi / \partial S_b > 0$ and $\partial \chi / \partial S_{b'} > 0$, and since

$$\frac{\partial^2 \chi}{\partial S_b S_{b'}} = \frac{\chi(S_b, \vec{S}_{-b})}{N^2 S_b S_{b'}} > 0\,,$$

the term inside square brackets is negative. With $\partial \alpha_b^* / \partial S_b < 0$ and $\partial h / \partial S_b < 0$, the inequality is satisfied and $\partial S_b / \partial S_{b'} > 0$.

A11 Proof of Proposition 7

We use the results of Van Zandt and Vives (2007) to show that monotone supermodularity of the ex ante investment decision is sufficient to establish the existence of a greatest and least Nash equilibrium.

Define bank b's best response correspondence, $BR_b(\vec{S}_{-b}) : \vec{S}_{-b} \to [0, W_b]$, as follows:

$$BR_b(\vec{S}_{-b}) \equiv \underset{S_b \in [0,W_b]}{\arg\max}\, \pi_b(S_b|\vec{S}_{-b}). \tag{A35}$$

To prove that the investment decision is supermodular, it is sufficient to establish:

1. The profit function, $\pi_b(S_b|\vec{S}_{-b})$, is supermodular:

$$\pi_b(S_b|\vec{S}_{-b}) + \pi_b(S'_b|\vec{S}_{-b}) \leq \pi_b(S_b \wedge S'_b|\vec{S}_{-b}) + \pi_b(S_b \vee S'_b|\vec{S}_{-b}),$$

   where $S_b, S'_b \neq S_b \in [0, W_b]$ and $\wedge$ and $\vee$ denote the meet and join of investments S_b and S'_b respectively. Since the investment space is a lattice, it follows that the meet and join of S_b and S'_b are in the investment space. Supermodularity of the profit function follows from the concavity of the profit function.

2. The action profile has increasing first differences:

$$\frac{\partial^2 \pi_b}{\partial S_b S_{b'}} \geq 0,$$

   for all $b, b' \neq b \in \mathcal{N}$, where $\mathcal{N}$ denotes the set of banks $b = 1, ..., N$ that contract on the platform. This holds from the Proof of Lemma 3.

Together, this is sufficient to establish that the investment decision is monotone supermodular. By the results of Van Zandt and Vives (2007), the best response mapping $BR_b(\vec{S}_{-b})$ must, therefore, contain a well-defined greatest ($\overline{BR}_b(\vec{S}_{-b})$) and least element ($\underline{BR}_b(\vec{S}_{-b})$) and the set of all greatest (least) best responses for each $b = 1, \ldots, N$ form a greatest and a least Nash equilibrium.

Next, we characterise these equilibria. It is straightforward to show that the zero-investment outcome is a (least) Nash equilibrium. In the two-bank case, if the other bank invests nothing in cybersecurity, bank b's marginal payoff is negative for any non-zero investment:

$$\left.\frac{\partial \pi_b}{\partial S_b}\right|_{S_{b'}=0} = EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial S_b} + R\delta\frac{(\alpha_b^*)^2}{2}\frac{\partial h}{\partial S_b} < 0 \quad \forall S_b \in [0, W_b].$$

When the other bank sets its investment in cybersecurity to 0, bank b's profits are decreasing in S_b for all $S_b \in [0, W_b]$ and so it maximises its profits by setting $S_b = 0$. This is a Nash equilibrium since

$$\pi_b(0|0) \geq \pi_b(S'_b|0) \quad S'_b \neq 0.$$

That is, given that the other bank has invested 0, the profits accruing to b for investing any amount other than 0 are lower than those received by investing 0.

We can also show that the greatest Nash equilibrium is different from zero (i.e., $\underline{BR}_b(\vec{S}_{-b}) \neq \overline{BR}_b(\vec{S}_{-b})$) by a contradiction. Suppose that $S^*_{b'} > 0$ and that $BR_b(\vec{S}_{-b})|_{S_{b'}>0} = 0$. By the definition of $BR_b(\vec{S}_{-b})$, profits are maximised whenever S_b forms a best response. By concavity of the profit function, there is an interior solution whenever $W_b > \hat{W}_b$. If the derivative of a strictly concave function is zero at some point, then that point is a global maximum. But we also show that $\partial\pi_b/\partial S_b|_{S_b\to 0} > 0$. This implies that $BR_b(\vec{S}_{-b})|_{S_{b'}>0} > 0$, a contradiction. The two-bank case can be easily generalised into a game with $N > 2$ banks. This is sufficient to prove that $\overline{BR}_b(\vec{S}_{-b}) \neq \underline{BR}_b(\vec{S}_{-b}) = 0$ and the cybersecurity investment game has multiple Nash equilibria.

A12 Proof of Proposition 8

From the proof of Proposition 3, we have $\partial S^*_1/\partial W_1 > 0$ since $W_1 < \widehat{W}$. Therefore, bank 1 decreases its equilibrium investment following a decrease in its endowment. Since $W_N > \widehat{W}$, Proposition 3 also implies that $\partial S^*_N/\partial W_N < 0$. Both banks, thus, reduce their equilibrium investments in response to a mean-preserving spread shift in endowments.

Moreover, by Lemma 3, the equilibrium investments by all *other* banks on the platform also decrease in response to the reduction by banks 1 and N. Together, these effects result in a decrease in equilibrium cybersecurity, $\chi^*(\vec{S})$.

A13 Proof of Proposition 9

As previously argued, in the absence of rollover risk, banks under-invest in cybersecurity due to their free riding incentives. But, as rollover risk increases, banks individual choose to invest more in cybersecurity, while Benchmark 1 remains unchanged. And, as γ increases, the gap between what banks individual choose to invest and that from Benchmark 1 decreases. Thus, for a critical level of rollover risk, γ^c_j, bank j's investment in cybersecurity are equal to that given under Benchmark 1. For each bank, $j = 1, \ldots, N$, the critical γ^c_j is given by the solution to

$$\frac{\partial\pi_j(\alpha^*_j, \gamma^c_j)/\partial\chi}{\partial\pi_j(\alpha^*_j, \gamma^c_j)/\partial O_j} = \sum_{b=1}^{N} \frac{\partial\pi_b(\alpha^{IN}_b)/\partial\chi}{\partial\pi_b(\alpha^{IN}_b)/\partial O_b}. \tag{A36}$$

Thus, for $\gamma < \min\{\gamma^c_1, \ldots, \gamma^c_N\}$, there is underinvestment by all banks relative to Benchmark 1, while for all $\gamma > \max\{\gamma^c_1, \ldots, \gamma^c_N\}$, there is overinvestment by

all banks at the system level. For intermediate values, however, the results are ambiguous.

A14 Proof of Proposition 10

With a negligence rule in place, expected profits include a penalty, $\kappa_b(O_b)$, that is implemented conditional on a successful cyber attack

$$\begin{aligned}\pi_b(S_b, O_b) &= (\chi(S_b, \vec{S}_{-b})/\bar{\lambda})\,(R - F\,D) \\ &+ \left(1 - \frac{\chi(S_b, \vec{S}_{-b})}{\bar{\lambda}}\right) \int_0^{\alpha_b^*(O_b)} \Big[R(1 - \alpha\,\delta(1 - h(O_b)) - F\,D - \kappa_b(O_b)\Big]\, d\alpha\,.\end{aligned}$$

The introduction of a negligence rule lowers expected profits in all events where a cyber attack is successful and the bank survives, increasing the relative benefits from investing more in cybersecurity. For the penalty to be effective in eliciting the social optimum, it must satisfy two conditions:

$$\frac{R - FD}{R\delta(1 - h(O_b))} - \frac{\kappa_b(O_b)}{R\delta(1 - h(O_b))} > \alpha_b^*(O_b)\,, \tag{A37}$$

i.e., the penalty is not so large that it leads the bank to fail for any successful attack, and

$$\pi_b(S_b^R, O_b^R) \geq \pi_b(S_b, O_b)\,,$$

for all $S_b \leq S_b^R$ and $O_b \geq O_b^R$. That is, the penalty should be large enough that profits subject to a negligence rule under the social optimum are preferable to those in the laissez faire optimum.

Suppose the penalty is structured in the following way

$$\kappa_b(O_b) = \begin{cases} \kappa_b \times O_b & \text{if } O_b > O_b^R, \\ 0 & \text{otherwise.} \end{cases} \tag{A38}$$

Then the optimal penalties, $\{\kappa_j^*\}_{j=1}^N$, that deliver the social optimum are each given by

$$\frac{R - F\,D - \int_0^{\alpha_j^*}\left[EV_2(\alpha) - \kappa_j^* O_j\right] d\alpha}{(\bar{\lambda} - \chi)\left[\left(EV_2(\alpha_j^*) - \kappa_j^* O_j\right)\frac{\partial \alpha_j^*}{\partial O_j} + \int_0^{\alpha_j^*}\left(\partial EV_2/\partial O_j - \kappa_j^*\right) d\alpha\right]} = \sum_{b=1}^{N} \frac{(R - F\,D) - \int_0^{\alpha_b^{IN}} EV_2(\alpha) d\alpha}{(\bar{\lambda} - \chi)\int_0^{\alpha_b^{IN}}\left(\partial EV_2/\partial O_b\right) d\alpha}. \tag{A39}$$

The left-hand side of (A39) is the marginal rate of substitution between the public good and operational resilience for bank j, and the right-hand side is the sum of the marginal rates of substitution over all banks in the absence of rollover risk. Taking the partial derivative of the left-hand side with respect to the penalty, we have

$$\frac{O_j \alpha_j^* \xi(O_j) + \left(R - F\,D - \int_0^{\alpha_j^*}\left[EV_2(\alpha) - \kappa_j^* O_j\right] d\alpha\right)(\bar{\lambda} - \chi)\left[O_j \frac{\partial \alpha_j^*}{\partial O_j} + \alpha_j^*\right]}{\left[\xi(O_j)\right]^2} > 0,$$

where $\xi(O_j) = (\bar{\lambda} - \chi)\left[\left(EV_2(\alpha_j^*) - \kappa_j^* O_j\right)\frac{\partial \alpha_j^*}{\partial O_j} + \int_0^{\alpha_j^*}\left(\partial EV_2/\partial O_j - \kappa_j^*\right) d\alpha\right]$. With the left-hand side increasing in κ_j^*, an optimal penalty exists by the intermediate value theorem.

The size of the optimal penalty is affected by the degree of rollover risk. As we have just shown, the marginal rate of substitution for bank j is increasing in κ_j^*. Further, as we establish in the derivation of the Samuelson benchmarks, each bank's marginal rate of substitution is also increasing in γ. Therefore, by the implicit function theorem, it must hold that $\partial \kappa_j^* / \partial \gamma < 0$

A Endogenising the face value of debt

In this section, we endogenise the face value and show that the main trade-offs and insights are unaffected in this more generalised model.

The value of a debt claim issued by bank b, which we denote by $V(F, \vec{S}^e)$, depends on creditors' expectations over how much banks invest in cybersecurity, $\vec{S}^e$. Under the simplifying assumption that creditors receive nothing in the event

of a bank failure, we obtain

$$V_b(F_b, \vec{S}^e) \equiv [\chi(S_b^e, \vec{S}_{-b}^e)/\overline{\lambda}]F_b + \left(1 - \chi(S_b^e, \vec{S}_{-b}^e)/\overline{\lambda}\right) \alpha_b^*(\underbrace{O_b^e}_{=W_b - S_b^e}, F_b)F_b\,.$$

So if the cyber attack is unsuccessful, creditors are repaid in full for sure. While, if the cyber attack succeeds and the platform suffers an outage, then creditors are only repaid if the bank does not fail due to a run. To keep notation succinct in what follows, we denote the fragility threshold as a function of investment in cybersecurity and face value of debt, so $\alpha_b^* \equiv \alpha_b^*(S_b, F_b)$.

If creditors have access to a safe outside investment option that yields $r > 0$, the face value of debt under perfect competition is given by the solution to

$$V(F_b, \vec{S}^e) = r\,. \tag{A40}$$

Lemma A1. *If bank equity is sufficiently high, $E > \bar{E}$, the equilibrium face value, $F_b^*(S_b^e)$, is increasing in cybersecurity investment, $\frac{\partial F_b^*}{\partial S_b^e} > 0$.*

Proof. The face value, $F_b^*(S_b)$ is the value of F_b that satisfies

$$V(F_b, S_b) \equiv \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}} F_b + \left(1 - \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}\right) \alpha_b^*(S_b, F_b)F_b = r. \tag{A41}$$

An increase in face value, increases the value function

$$\frac{\partial V}{\partial F_b} = \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}} + \left(\frac{1 - \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}}{R[1 - h(O_b)]}\right)[R - 2\gamma F_b D] > 0, \tag{A42}$$

as long as $E > \hat{E} \equiv \frac{2\gamma - 1}{2\gamma}$. Next, to see that the value function is decreasing in S_b, note that

$$\frac{\partial V}{\partial S_b} = F_b \left\{ \frac{\partial \chi / \partial S_b}{\overline{\lambda}} - \alpha_b^*(S_b, F_b) \left[\frac{\partial \chi / \partial S_b}{\overline{\lambda}} - \left(1 - \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}\right) \frac{\partial h / \partial S_b}{1 - h(W_b - S_b)} \right] \right\}. \tag{A43}$$

Since $\partial \alpha_b^* / \partial E < 0$, it is clear that there exists another threshold, $\tilde{E}$, such that for $E > \tilde{E}$ we have that

$$\partial \chi / \partial S_b < \alpha_b^*(S_b, F_b \,|\, \tilde{E}) \left[\frac{\partial \chi}{\partial S_b} - (\overline{\lambda} - \chi(S_b, \vec{S}_{-b})) \frac{\partial h / \partial S_b}{1 - h(W_b - S_b)} \right],$$

and so $\partial V / \partial S_b < 0$. Thus, for $E > \bar{E} \equiv \max\{\hat{E}, \tilde{E}\}$, we have by the implicit

function that

$$\frac{\partial F_b^*}{\partial S_b} = \frac{-1}{\partial V/\partial F_b}\left(\frac{\partial V}{\partial S_b}\right) > 0. \tag{A44}$$

□

Creditors demand compensation for additional cybersecurity investment to offset the heightened fragility to which the bank is exposed in the event of a breach. The marginal returns to operational resilience are large relative to cybersecurity. So each creditor lends to the bank at a rate that depends positively on expectations of the bank's investment in cybersecurity. The bank, for its part, takes face value as given when setting its optimal allocation between cybersecurity and operational resilience.

Proposition A1. *If $\gamma > \bar{\gamma}$, there is a unique equilibrium (S_b^*, F_b^*) such that creditors' participation constraints are satisfied and the bank optimally chooses its cybersecurity investment.*

Proof. To establish the equilibrium set (S_b^*, F_b^*), we first show that the level of investment in cybersecurity that the bank chooses to maximise equity value is increasing in F_b. By the implicit function theorem, the marginal response of S_b^* to a unit increase in F_b is given by

$$\frac{\partial S_b^*}{\partial F_b} = \frac{-1}{\partial^2 \pi_b/\partial S_b^2}\left[\frac{\partial^2 \pi_b}{\partial S_b^e \partial F_b}\right], \tag{A45}$$

where

$$\frac{\partial^2 \pi}{\partial S_b^e \partial F_b} = (\partial\chi/\partial S_b^e)(1/\overline{\lambda})\left[-D(1-\alpha_b^*(S_b,F_b)) - EV_2(\alpha_b^*)\frac{\partial \alpha_b^*}{\partial F_b}\right]$$
$$+\left(1-\chi(S_b,\vec{S}_{-b})/\overline{\lambda}\right)\left[EV_2(\alpha_b^*)\frac{\partial h/\partial S_b}{1-h(W_b-S_b)}\frac{\partial \alpha_b^*}{\partial F_b} + \frac{\partial EV_2(\alpha_b^*)}{\partial F_b}\frac{\partial \alpha_b^*}{\partial S_b}\right.$$
$$\left. +R\delta\alpha_b^*(O_b,F_b)\frac{\partial h}{\partial S_b}\frac{\partial \alpha_b^*}{\partial S_b}\right],$$

where the second term is positive. The first term can be simplified to

$$-(\partial\chi/\partial S_b^e)(1/\overline{\lambda})D\Big(1-(1-\gamma\delta)\alpha_b^*(S_b,F_b) - \gamma\delta\alpha_b^{IN}(S_b,F_b)\Big),$$

which is positive if and only if $E > \frac{(1-\delta)}{2-\gamma\delta}$. For $\gamma > \bar{\gamma} \equiv \frac{2+3\gamma-\sqrt{4-4\delta+9\delta^2}}{4\delta}$, it follows that this condition is always satisfied for $E > \bar{E}$. This establishes that $\partial S_b^*/\partial F_b > 0$. From Lemma A1, we have that $\partial F_b^*/\partial S_b^e > 0$. For a unique intersection, we

require that (i) the value $\tilde{F}_b$ such that $S_b^*(\tilde{F}_b) = W_b$ satisfies $\tilde{F}_b > F_b^*(W_b)$; and (ii) the value $\underset{\sim}{F}_b$ such that $S_b^*(\underset{\sim}{F}_b) = 0$ satisfies $\underset{\sim}{F}_b < r$. □

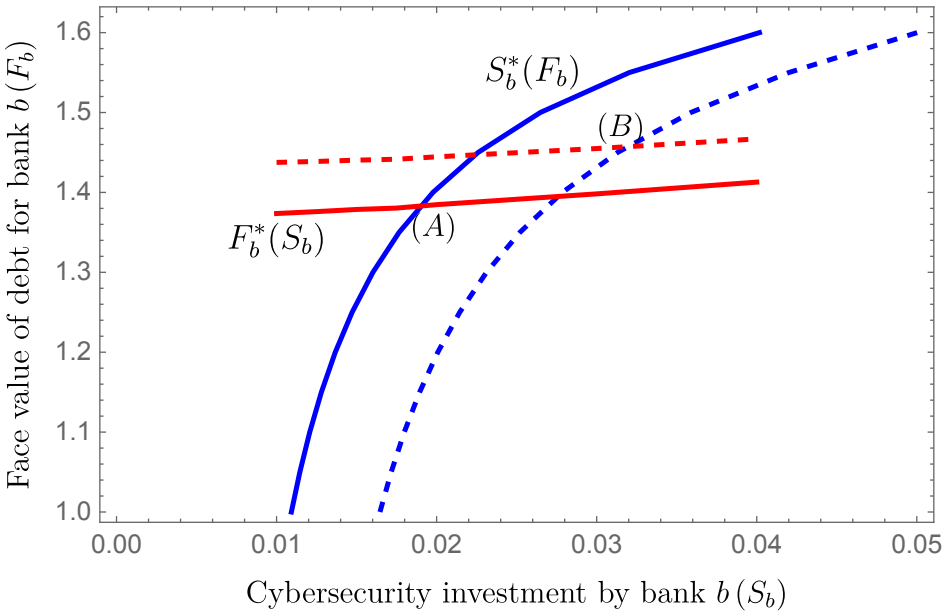

Figure A2: The face value of debt demanded by creditors, F_b^*, and level of cybersecurity investment initiated by each bank, S_b^*, are uniquely determined by the intersection (A) of bank b's reaction function (blue line) and creditors' reaction function (red line). Following an increase in rollover risk, the curves shift and the new equilibrium is given by intersection (B).

Figure A2 illustrates the equilibrium. The laissez-faire outcome that arises with an endogenous face value retains the link between rollover risk (ex post coordination failure) and cybersecurity investment (ex ante free riding on the public good). As the optimal level of cybersecurity increases, creditors demand higher face value to choose bank debt over their safe outside option. With a higher face value, bank b is made more fragile for every proportion of early withdrawals in the event that a cyber attack is successful, and so it is better off allocating its working capital towards cybersecurity. Since each bank only internalises its own marginal rate of substitution and face value, the free riding problem remains. The unique equilibrium (S_b^*, F_b^*) satisfies creditors' participation constraints and equates the bank's private marginal rate of substitution with its marginal rate of transformation.

Proposition A2. *Cybersecurity investment is increasing in rollover risk,* $\frac{\partial S_b^*}{\partial \gamma} > 0$.

Proof. We have already demonstrated that $\partial S_b^*/\partial \gamma > 0$. Next, consider the face value, $F_b^*(S_b^e)$. To see the direct effect on $F_b^*(S_b^e)$, by the implicit function theorem,

we have

$$\frac{\partial F_b^*}{\partial \gamma} = \frac{-1}{\partial V/\partial F_b}\left[\left(1 - \frac{\chi(S_b, \vec{S}_{-b})}{\overline{\lambda}}\right) F_b \frac{\partial \alpha_b^*}{\partial \gamma}\right] > 0. \tag{A46}$$

The expression is positive since $\partial \alpha_b^*/\partial \gamma < 0$. With an upward shift in both response functions, we have an increase in both S_b^* and F_b^* in equilibrium. □

Following an increase in γ, each fund manager k is less likely to rollover claims at $t = 1$ for any given signal x_{bk}. Ex ante, this means that bank b is more likely to fail for any realisation of the outage shock, α_b, and a given F_b. So the marginal benefit of an additional unit of cybersecurity is high relative to the marginal cost of further investment in operational resilience. The bank is therefore better off shoring up cyber defenses, even though this means an increase in fragility in the event of a successful attack. At the same time, with high rollover risk, creditors demand higher face value to choose bank debt over their safe outside option.

The effects on cybersecurity and face value are, thus, self-reinforcing. As such, the signs of the other comparative static results are unchanged. Both cybersecurity investment, $S_b^*(F_b)$, and face value, $F_b^*(S_b^e)$ are increasing in the deadweight loss of an attack, δ, and are decreasing in attack intensity, $\overline{\lambda}$, and bank equity, E.