

SOME PRINCIPLES FOR REGULATING CYBER RISK*

By Anil K Kashyap and Anne Wetherilt*

*Kashyap: University of Chicago Booth School of Business, National Bureau of Economic Research, Centre for Economic Policy Research and Bank of England (anil.kashyap@chicagobooth.edu); Wetherilt: Bank of England (anne.wetherilt@bankofengland.co.uk). We thank Christopher Dawson for helpful assistance; Kevin Murphy for extremely valuable suggestions; and colleagues at the Bank of England for helpful feedback. This research has been supported by a grant from the Alfred P. Sloan Foundation to the Macro Financial Modeling (MFM) project at the University of Chicago. All views here are our own and do not necessarily reflect the views of the Bank of England, the Financial Policy Committee or any of the other organizations with which we are affiliated

As part of the international framework for capital standards, banks are required to fund themselves with loss-absorbing capital to guard against risk of loss resulting from inadequate or failed internal processes, people and systems or from external events. One estimate is that between 2012 and 2017, major banks lost nearly \$200 billion from operational risk events.¹ Cyber risk is commonly cited as one of the highest operational risk concerns.² In this paper, we argue that cyber risk creates new microprudential and macroprudential challenges, and develop six regulatory

principles that capture the unique risks posed by cyber threats.

Superficially, cyber and some other operational risks look similar. Both can involve the failure of some process or technology that could cripple a firm and potentially have broader consequences. We argue that upon closer inspection cyber is special in two ways:

- a) the way a shock occurs; and
- b) the impact of the shock after it occurs.

Although the transmission to the broader economy operates through familiar channels, the unique nature of the shock and the subsequent impact mean that the appropriate regulatory response is likely to differ.³ We explain along the way why the private sector left to its own will not be able to take adequate steps to address cyber risk. Hence, cyber risk requires some special regulatory responses.

¹ ORX (2018) reports losses by major global banks, estimated at Euro 170 billion for the 2012-2017 period (p. 6).

² See e.g. Risk Magazine, Top 10 operational risks for 2018, 22 February 2018.

³ Healey et al (2018), Kopp et al (2017) and Kaivanto and Warren (2018) also discuss the unique nature of cyber risk and what it means for financial stability.

I. What's special about a cyber shock?

In January 2018, Ciaran Martin, head of the UK's National Cyber Security Centre (NCSC) warned that a major cyber attack was a matter of 'when' rather than 'if.'⁴ In its annual report on cyber threats to UK business, the NCSC notes that 'the race between hackers' and defenders' capabilities will increase in pace and intensity.'⁵ In the US, the Ponemon Institute has estimated that the cost of cyber crime rose by 23% between 2016 and 2017.⁶

Recent research estimates that global corporate spending on cyber security will be as high as \$124 billion in 2019.⁷ Other estimates suggest that financial services firms spend about 12% of their IT budgets on cyber security.⁸

A cyber attack can come in more than one form. Some attacks cause disruption to computer systems, slowing down or totally halting critical processes. Others affect the data supporting these processes, either by gaining unauthorised access or by corrupting data. Both types of cyber shocks have common characteristics, which distinguish them from other operational shocks.

First, the intent. Disruptive attacks are conducted with malicious intent, and designed to inflict maximum damage, perhaps by combining attacks on multiple systems, or by selecting a critical date. Second, the probability. As noted earlier, it is widely accepted amongst experts that probability of success is now much higher, and a high-impact event is a matter of 'when,' rather than 'if.' Third, the timing. The attack might involve a hidden phase, where malicious code is inserted and data is compromised and manipulated to create problems. Once the attack becomes known, it can be difficult to appreciate the extent of the damage and to identify effective solutions. As an example, experts believe that the 2017 NotPetya virus had been present for several-weeks in targeted hardware.⁹

And fourth adaptability. New tools and techniques available to cyber attackers reduce the cost of attacks and heighten their impact, whilst at the same time increasing the cost of defence.¹⁰ They also enable attackers to exploit previously untapped vulnerabilities.

Some operational shocks share some of these features. For example, terrorist activity is malicious and adaptive. But as we will argue

⁴ Major cyber attack on UK a matter of 'when, not if' – security chief, Guardian newspaper, 23 January 2018.

⁵ NCSC (2018), p. 6.

⁶ See Richards et al. (2017), p. 2.

⁷ See www.Gartner.com, Press Release, 15 August 2018.

⁸ Hiscox (2018), p. 13.

⁹ See Greenberg (2018).

¹⁰ See e.g. Lewis (2018), p. 5.

next, when all four characteristics are present, managing the risk – i.e. preventing and recovering from cyber attacks – becomes prohibitively expensive.

Cyber shocks differ from other shocks in a second way, namely their widespread impact on organisations and the wider financial system. In part, this comes because the interconnectedness of the financial system makes wide-scale disruption possible. Indeed, malicious software may be introduced directly into firms, or indirectly via their counterparties or third parties, thus creating a vast network for attackers to exploit. Through supply chain attacks, attackers can also gain access to confidential data from a wide range of sources.¹¹ These dependencies also arise from the use of common software. The 2017 Wannacry incident exploited a common vulnerability in Windows systems across multiple organisations and sectors. Disruption to critical processes was widespread, affecting over 300,000 computers in 150 countries.¹²

The second unusual feature arises from the repercussions of a quiet, hidden phase of an attack. A terrorist might spend a long time planning an attack, but the damage of the attack would be instantly visible. In contrast, the

impact of a cyber attack may remain unknown for a long period. The resulting uncertainty over the extent of damage can cause special problems for recovery, in particular when it is not known whether and when the integrity of data has been compromised.¹³ This is why cyber attacks causing data damage or theft are typically more expensive to organisations.¹⁴ As an example, in 2017, NotPetya malware caused significant damage to several global companies, as data were permanently corrupted and hence unrecoverable.¹⁵

A cornerstone of most contingency planning is a commitment to rapidly restoring services via backup systems. Uncertainty about the integrity of backup may put this plan at risk.

Together, the scale and timing uncertainty of cyber shocks imply that not only risk management, but also incident management carry very significant costs for individual organisations.

II. Why is regulation needed at all?

To argue for regulation, we should ask whether firms will adequately invest in both preventive and recovery capabilities, especially since they have clear and strong commercial interests in doing so. For example, Wannacry

¹¹ See e.g. NCSC (2018), p. 13.

¹² See e.g. NCSC (2018), p. 8.

¹³ See NCSC (2018), p. 15.

¹⁴ Richards et al. (2017), p. 28-9.

¹⁵ See e.g. NCSC (2018), p. 15; Greenberg (2018).

did not affect firms who had applied the most recent patches to their Windows systems.

So why would the social and private interests in guarding against cyber risks diverge? Despite self-interest, there are four reasons why social and private incentives for addressing cyber risk can differ. First, firms may have adequate incentives to prepare for idiosyncratic risk scenarios, but they may not fully account for system-wide effects of a successful attack. For example, a cyber attack that knocks out one firm or piece of infrastructure could undermine confidence in unaffected firms. Individual firms have fewer incentives to internalize concerns about how an incident at their firm might affect overall confidence in the financial system (or potentially the overall functioning of the system if they provide a critical service).

Second, firms' exposure to common risks may not be fully priced. Shared services or software create common vulnerabilities. In making their purchases, firms may not internalize the associated risks of having many parties that have similar openness to an attack.

Third, regarding recovery, the management at any firm considering cyber risk typically rely on a combination of internal defences and recourse to outside experts (e.g. specialist

suppliers, consultants or government cyber experts). If multiple firms are simultaneously attacked, each individual firm's assumptions about the availability of external resources may prove incorrect. Management might believe that they should not be expected to prepare for a scenario where they cannot access specialists to help. Alternatively, management may choose to withdraw from the provision of services, rather than keeping them running partially or through manual workarounds.¹⁶

Fourth, individually, firms may face information constraints. Society might want firms to share information with each other following an attack, as this may deepen their understanding of common vulnerabilities. But firms may be reluctant to do so, to safeguard their reputation.¹⁷ Likewise, regulators may benefit from seeing firms' own cyber resilience assessment reports. Firms that report weaknesses uncovered through their own tests should not necessarily be penalized or receive greater attention from the regulatory bodies.

Together, these four factors may explain why regulators might reach different judgments about risk tolerances than firms, thus creating a role for regulation. Indeed, we doubt that firms will choose to protect themselves to the degree that society might want and to the specific

¹⁶ We thank Patricia Mosser for suggesting this last point.

¹⁷ See e.g. Kopp et al (2017).

shocks that might prove most damaging. Regulation can attempt to remedy this problem, without necessarily being overly prescriptive. And faced with the knowledge that a future cyber attacks could cause severe damage to the finance sector, a regulatory response based purely on prevention is going to be inadequate. In the next Section, we set out some general principles for microprudential regulation, and in the following one, we discuss the role of macroprudential regulation.

III. Regulatory principles and microprudential policy

Supervisory authorities have an interest in ensuring that the firms and Financial Market Infrastructures (FMIs) they supervise are run in a sustainable manner. The starting point for managing systemic risks is a robust microprudential policy framework.¹⁸ We propose three principles related to cyber risk that regulators can adopt to help deliver this outcome.

Principle 1: Insist that firms operate with the presumption that a successful attack is inevitable.

Principle 1 is a foundational principle of the UK approach to cyber resilience. In a recent Discussion Paper, the UK financial authorities note that firms should assume that disruption to their systems and processes supporting will occur. Furthermore, the UK authorities expect firms to set a tolerance for disruption to their most important business services. This in turn requires firms to identify those services, and the systems and processes that are critical for their delivery.¹⁹ Principle 1 is also captured by recent G7 guidance.²⁰

Principle 2: Insist that firms plan for prolonged and system-wide disruption, with particular attention to resourcing for response and recovery.

Principle 2 acknowledges that resources may be constrained if multiple entities are compromised simultaneously, and/or there is widespread data corruption. The principle encourages firms to plan for a wide range of scenarios and go beyond their pure idiosyncratic concerns. It also suggests that there are strong incentives for collective industry action, allowing firms to share knowledge and pool resources.²¹

¹⁸ See e.g. Woods (2018).

¹⁹ Bank of England (2018b), p. 16; Woods (2018), p. 5.

²⁰ See G7 (2017).

²¹ For example, in 2016, the largest US banks created the Financial Systemic Analysis and Resilience Center (FSARC) to combine their cyber-related capabilities. In the UK, the Cross Market Operational Resilience Group (CMORG) promotes cooperation in the finance

Principle 3: Aim for two-way dialogue between firms and supervisors about appropriate recovery times.

Principle 3 encourages firms and supervisors to discuss the externalities that may arise from insufficient investment in cyber security. It also acknowledges that firms face difficult choices, given the scale and the uncertain timing of a cyber attack on the one hand, and the cost of cyber protection and incident management on the other hand. Building on Principles 1 and 2, this third principle suggests that firms and supervisors have a supervisory conversation that internalizes social concerns, whilst also recognizing private constraints.

IV. A role for macroprudential policy

Our first three principles help correct some of the imbalances between firms' incentives for managing their idiosyncratic risks and society's risk tolerance. They do not fully address the concerns that society might have for the stability of the overall financial system. The financial crisis taught us that regulators need to think about the viability of the whole financial system and not just individual firms.

So we now propose three further principles that are macroprudential in nature.

Principle 4: Conduct cyber stress tests that explore common vulnerabilities that may amplify the impact of a cyber shock.

Principle 4 characterizes the approach taken by the UK Financial Policy Committee (FPC). In 2018, the FPC announced that it would test the resilience of the UK finance system by asking firms whether they could meet a system-wide tolerance set by the FPC for the delivery of critical economic functions.²² The FPC will ask firms at the core of the finance system to consider a common stress scenario that assumes severe disruption and/or data corruption. Firms will need to demonstrate that they have plans in place to resume operations within the FPC's tolerance. Stress testing may reveal weaknesses, such as reliance on common infrastructure or software with limited substitutability. They can also identify the extent to which firms' plans for recovery are jointly realistic.

Principle 5: Plan for system-wide disruption by setting appropriate recovery expectations for the delivery of critical economic functions.

sector. CMORG also oversees a regular programme of exercises to test sector-wide capabilities.

²² Bank of England (2018), pp. 40-1.

Principle 5 complements Principle 2 and highlights the importance of planning for system-wide disruption. Principle 5 aims to align planning assumptions and resourcing decisions made by individual firms with system-wide recovery objectives. This principle is a key objective for the UK authorities, as they develop their approach to cyber resilience.²³

Principle 5 explicitly links financial stability to the ability of the finance system to provide critical economic functions.²⁴ A severe cyber attack could undermine this in two ways.

First, disruption at a single firm could have a systemic impact, for example if this firm is a sufficiently large provider of a function, or a dominant market participant. In this case, regulators need to be assured that the firm's recovery planning is robust enough to be able to deliver enough of its critical functions to support the overall system, without making it prohibitively expensive to run its business.

Second, a cyber incident may cause disruption at multiple firms. In that case, the principle implies that collectively, the remaining firms must be able to support critical functions.²⁵

In either case, firms may need to demonstrate that they are able to conduct business (for example, by relying on alternative providers), and the regulators will need to assess how well this can be done when setting recovery expectations. Overall, macroprudential regulators ought to prepare for both of these cases.

Principle 6: Encourage firms to avoid common vulnerabilities and to make more diverse infrastructure or software choices

Finally, Principle 6 recognizes that some of the finance sector vulnerabilities stem from investment choices made by firms, which determine their exposure to common risks. Generally, regulators can try to approach this in two ways.

One approach is to reward firms that can continue to operate when a shared resource is compromised. For example, in assuring the delivery of electricity, the regulator can set prices so that suppliers that offer power when it is most needed are paid a premium for doing so. In the case of cyber risk, it is hard to think of mechanisms that reward individual firms for being able to deliver critical functions at times when their competitors cannot. For instance,

²³ Bank of England (2018b), p. 13.

²⁴ See Bank of England (2018), p. 40.

²⁵ A corollary to this is that once the critical mass is knocked out, the benefits to having others operating is probably small.

macroprudential regulators do not have control over contract design between private parties, so they cannot automatically adjust prices to reward a service provider for maintaining the viability of function during a period of distress.

The other approach is to tax behavior that might create shared risk. For example, market prices for software do not reflect the cost of the cyber risk for society that arises when many firms adopt the same package. Taxing usage to account for that cost would be the standard way to address this issue. Here, it is not obvious how to implement such a tax.

Stress testing can indirectly address this problem. The macroprudential regulator can devise stress scenarios that are tougher for commonly used resources. For instance, suppose there are two competing software options that firms could use. A stress test could assume that the dominant software option is compromised (while the alternative is not). That would implicitly penalize the firms that relied solely on the dominant option. Firms that made a different choice from the start (e.g. by having a robust fallback option) would not be required to undertake remedial action. The severity of the stress scenario could also be increased depending on the degree of concentration in firms' choices. While this is a blunt approach, it would provide incentives for

diversification and encourage innovation to develop alternative options.

V. Conclusion

While cyber risks are superficially similar to other operational risks, they differ importantly in the form they take and the impact they can have. Private incentives are unlikely to fully deliver the level of resilience that society is likely to prefer. The principles we have announced would help correct this gap.

Over the past two years, the G7 has issued high-level guidance to assist financial authorities and the sector in building greater cyber resilience. More detailed global guidance is available for supervisors of Financial Market Infrastructures.²⁶ Many authorities are currently in the process of developing more detailed cyber security expectations for their banking sector. We believe the six principles set out in this paper will help authorities as they review their microprudential and macroprudential frameworks.

More generally, by drawing on standard economic theory, we have highlighted the specific issues that make cyber problems special and need the attention of both microprudential and macroprudential

²⁶ See CPMI-IOSCO (2016).

authorities. Our principles are aimed at advancing the debate over what to do about cyber risk. The challenge in this area, as we see it, is to develop specific policies that respond to the unique nature of the shock, and encourage risk management solutions that acknowledge the unique impact of the shock.

REFERENCES

- Bank of England. 2018. "Financial Stability Report." Issue No. 43, June.
- Bank of England. 2018b. "Building the UK Financial Sector's Operational Resilience." Discussion Paper 1/18.
- CPMI-IOSCO. 2016. "Guidance on Cyber Resilience for Financial Market Infrastructures." BIS, June.
- G-7. 2017. "Fundamental Elements for Effective Assessment of Cybersecurity in the Financial Sector." October.
- G7. 2016. "Fundamental Elements of Cybersecurity for the Financial Sector." October.
- Greenberg, Andy. 2018. "The Untold Story of NotPetya, the Most Devastating Cyberattack in History." in Wired, August.
- Healey, Jason, Patricia Mosser, Katheryn Rosen and Adriana Tache. 2018. "The Future of Financial Stability and Cyber Risk." The Brookings Institution, October.
- Hiscox. 2018. "Cyber Readiness Report." February.
- Kaivanto, Kim and Philip Warren. 2018 "Could a Cyber Attack Cause a Systemic Impact in the Financial Sector??" Bank of England Quarterly Bulletin, forthcoming.
- Kopp, Emanuel, Lincoln Kaffenberger and Christopher Wilson. 2017. "Cyber Risk, Market Failure and Financial Stability." IMF Working Paper, 17/185.
- Lewis, James. 2018. "Economic Impact of Cyber Crime – No Slowing Down." McAfee report, February.
- National Cyber Security Centre. 2018. "The Cyber Threat to UK Business. 2017-2018 Report." April.
- Office of Financial Research. 2017. "Cyberscecurity and Financial Stability: Risks and Resilience", Viewpoint, February.
- ORX. 2018. "Annual Banking Loss Report. Operational Risk Loss Data for Banks Submitted between 2012 and 2017." June.
- Richards, Kevin. et al. 2017. "Cost of Cyber Crime Study. Insights on the Investments that Make a Difference." Ponemon Institute, October.
- Rosengren, Eric S. 2015. "Cyber Security and Financial Stability." Federal Reserve Board of Boston, speech, January.
- Woods, Sam. 2018. "Good cop, Bad Cop." Speech, Mansion House City Banquet, London, October.