

Ros, Greg

Research Report

The making of a cyber crash: A conceptual model for systemic risk in the financial sector

ESRB Occasional Paper Series, No. 16

Provided in Cooperation with:

European Systemic Risk Board (ESRB), European System of Financial Supervision

Suggested Citation: Ros, Greg (2020) : The making of a cyber crash: A conceptual model for systemic risk in the financial sector, ESRB Occasional Paper Series, No. 16, ISBN 978-92-9472-132-7, European Systemic Risk Board (ESRB), European System of Financial Supervision, Frankfurt a. M., <https://doi.org/10.2849/915512>

This Version is available at:

<https://hdl.handle.net/10419/238320>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.


Occasional Paper Series

No 16 / May 2020

The making of a cyber crash: a conceptual model for systemic risk in the financial sector

by
Greg Ros

with contributions from Gabriella Biro, Tom Keating,
Wiebe Ruttenberg, Elise Vik Sætre, Dinant Veenstra,
Benedetto Andrea de Vendictis, Maarten Willemen and
Paul Williams



ESRB
European Systemic Risk Board
European System of Financial Supervision

Contents

Executive summary	3
1 Context	6
1.1 Cyber risk	6
1.2 Cyber threat	7
1.3 Vulnerability	10
1.4 Assets	13
1.5 Countermeasures	15
1.6 Starting point	16
2 Shock	17
2.1 Consequences	17
2.2 Technical impacts	18
2.3 Business impacts	19
2.4 Impact measurement	20
3 Amplification	22
3.1 Complex adaptive systems	22
3.2 Vulnerability in a system context	23
3.3 Amplifiers	25
3.4 System amplifiers	26
3.5 Cyber-specific amplifiers	41
3.6 Alignment of amplifiers	47
3.7 Contagion channels	49
4 Systemic event	53
4.1 Impact thresholds	53
4.2 Absorptive capacity	54
4.3 Systemic mitigants	57
5 Using the model	61



5.1	Modelling scenarios	61
5.2	Scenario examples	63
5.3	Observations from scenario analysis	66
	References	68
	Abbreviations	73
	Imprint and acknowledgements	74



Executive summary

In October 2017, the European Systemic Risk Board (ESRB) set up a group whose objective was to examine cyber security vulnerabilities within the financial sector, and their potential impact on financial stability and the real economy. In its first year, the European Systemic Cyber Group (ESCG) sought to develop a shared understanding of Common Individual Vulnerabilities (CIVs) across ESRB members, and to identify the unique characteristics of cyber risk that could contribute to a systemic event. Building on this work, this paper describes a **conceptual model for systemic cyber risk**, and aims to:

- provide a structured approach that can be used to describe cyber incidents, from genesis through to a potential systemic event;
- demonstrate the link between the crystallisation of cyber risk in a firm-specific context (portraying microprudential concerns), and the possible ramifications for the financial system (applying a macroprudential focus);
- identify system-wide vulnerabilities and the unique characteristics of cyber incidents which can act as amplifiers, thereby propagating shocks through the financial system;
- support the use of historical or theoretical scenario-based analysis to demonstrate the viability of the model;
- suggest system-wide interventions that could act as systemic mitigants.

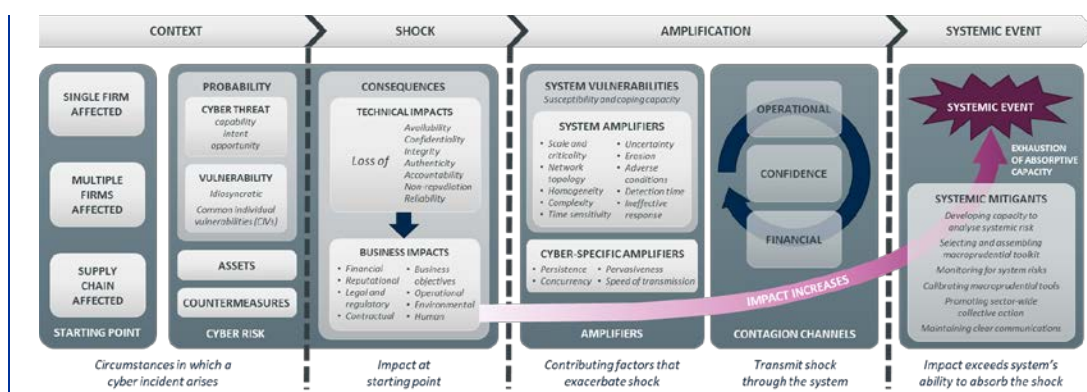
Although the model is geared towards disruption arising from cyber incidents, it can also be used for any source of operational disruption (although some elements of the model may be less relevant).

Model overview

In order to deconstruct and describe the macro-financial implications of operational and cyber risks, the systemic cyber risk model is split into four distinct phases (as illustrated overleaf):

- **Context** – the circumstances in which a cyber incident arises, in the form of a crystallised cyber risk;
- **Shock** – a description of the technical and business impacts experienced at the moment the cyber incident originates;
- **Amplification** – the systemic amplifiers and contagion channels which exacerbate the shock through a system, increasing the magnitude of the impact experienced;
- **Systemic event** – the point at which a system is no longer able to absorb the shock.





Each phase is described in further detail in Sections 1-4 of this paper, while Section 5 provides a guide to using the model for historical-event analyses or scenario-based analyses.

Conclusions

During the development and testing of our approach, a number of observations were noted.

(a) With regard to the model:

- The Context phase is useful for scenario design, but is not essential for assessing systemic vulnerabilities or relevant mitigants. It is possible to adopt a **cause-agnostic** approach which ignores the circumstances of disruption and focuses solely on impact.
- From a microprudential perspective, it is important to maintain a **dual focus** on both idiosyncratic individual vulnerabilities and CIVs. The latter are prioritised locally, but informed by a broader, collective view.
- Measuring impact is challenging and remains primarily a **judgement-based, qualitative approach**. Although some quantitative indicators exist, they should be used to complement and inform impact assessments.

(b) With regard to policy considerations arising from the model:

- A systemic event arising from a cyber incident is **conceivable**. Cyber incidents resulting in near-systemic consequences have occurred, in circumstances that can be described as **"severe, but plausible"**. However, a truly systemic event would require an alignment of amplifiers and a lack of effective systemic mitigants that would be **"extreme, but existential"** in nature.
- A cyber incident which causes only operational-to-operational contagion may have system-wide impacts. However, the current base of evidence suggests that a systemic event **requires** the confidence and/or financial contagion channels to be triggered.

Building on this conceptual model, future work could be undertaken to: (i) study the efficacy of individual systemic mitigants; (ii) use quantitative or data-driven methods to more accurately



express each phase of amplification; or (iii) further study the interaction and measurement of impact at institutional and at aggregate-system level.

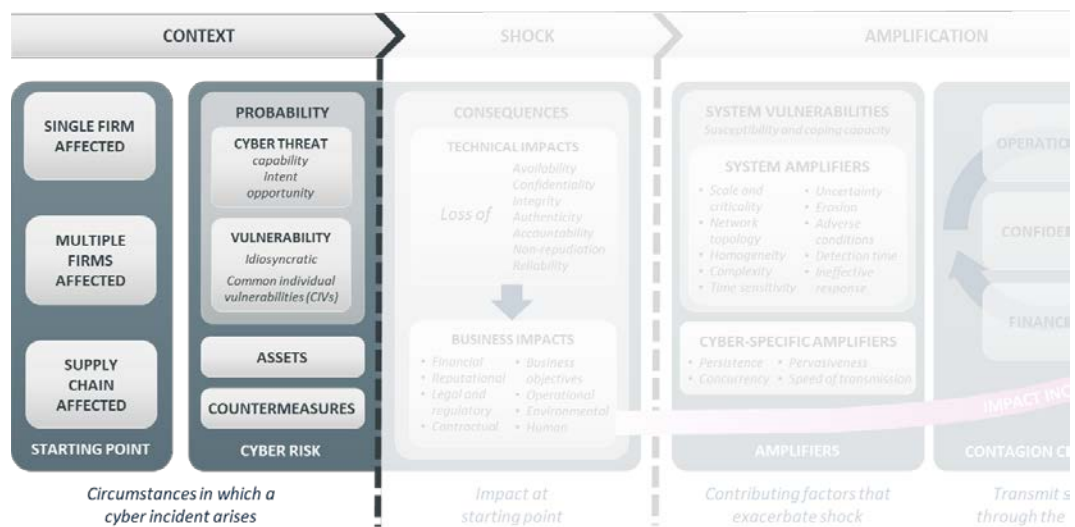
JEL codes: E17, G01, G20, K24, L86, M15, O33



1 Context

The circumstances in which a cyber incident arises

As an entry point into the systemic cyber risk model, the **Context** phase conceptualises the starting point of a cyber incident, in the form of a crystallised cyber risk. In order to convey a degree of realism when performing a scenario-based analysis using the model, it is helpful to describe the individual elements that constitute cyber risk.



1.1 Cyber risk

According to Bernoulli's **Expected Utility Theory (EUT)**, risk is the product of the probability of a certain outcome and its consequence. This concept still endures, as can be seen in the Financial Stability Board's (2018) definition of **cyber risk**:

"The combination of the probability of cyber incidents occurring and their impact."¹

Within the field of EUT, there are many methods used for risk assessment – these can broadly be classified as "single-asset" versus "systems of assets". A single-asset risk assessment involves a single target such as a building, airport or computer, while a system risk assessment involves many assets connected together to form a system. In the field of information and communications technology (ICT), well-established methodologies for ICT risk assessment from the leading standard setters (e.g. the International Organization for Standardization (ISO)², the National

¹ Financial Stability Board (2018), *FSB Cyber Lexicon*.

² International Organization for Standardization (2018), *Information technology – Security techniques – Information security risk management*.



Institute of Standards and Technology (NIST)³, the Information System Audit and Control Association (ISACA)⁴) are widely used across the financial sector.

For the Context phase, the model employs the former method of risk assessment, using a modified form of the **Probabilistic Risk Assessment** devised by Norman Rasmussen in 1975:

$$\text{cyber risk} = \frac{\text{cyber threat} \times \text{vulnerability} \times \text{assets} \times \text{consequences}}{\text{countermeasures}}$$

This expression captures the probability of a cyber incident occurring (a combination of cyber threat and vulnerability), the perceived value of assets affected, and the technical or business impacts resulting from the cyber incident (consequences). To complete the picture, countermeasures are included as a form of mitigation, to express a net risk position.

The modified expression for cyber risk provides a simplification that is sufficient for use in this conceptual model. However, the following caveats should be considered when using this approach.

1. Although the expression uses a mathematical notation, this format is used to convey a relationship between constituent parts – it is not a strictly numerical approach.
2. The model assumes that the threat and vulnerability components are independent – i.e. it assumes that a threat actor is not influenced by a target's vulnerability. In reality, there is a relationship between threat and vulnerability and this may be more accurately modelled using **game theory**. In this case, threat actors may seek out weaker targets so they can achieve their objectives at minimal cost.
3. Decision-making is motivated by emotion, but rational decision-making based on EUT ignores the emotional dimension entirely. The concept of **Prospect Theory (PT)** argues that humans cannot evaluate expected utility risk on a consistent basis. With subconscious risk aversion or risk-seeking behaviours undermining rationality, risk exposure predictions are likely to be under or overestimated.⁵

1.2 Cyber threat

In an operational context, the study of threats has historically focused on forces of nature, such as hurricanes or earthquakes. As technology has become more pervasive, its scope has been extended to address incidents arising from accidental human actions including software errors or coding mistakes. Similarly, deliberate actions, driven by malevolent motivation and enabled by this technological revolution, are now at the forefront of discussions of threats to the financial sector and the wider economy. This evolution is reflected in the World Economic Forum's annual Global Risks

³ National Institute of Standards and Technology (2012), *Guide for Conducting Risk Assessments*.

⁴ ISACA (2009), *The Risk IT Framework*.

⁵ Kahneman, D. and Tversky, A (1979), "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, Vol. 47, No 2, pp. 263-291.



Report, in which malicious cyber incidents rank in the top ten of risks in terms of both likelihood and impact.⁶

Based on the classification set out in ISACA⁷, threats can generally be described as:


- natural events;
- accidental events;
- deliberate acts.

However, the approach proposed by Lucas et al. offers a more complete methodology for threat classification⁸. This methodology divides the threat space into sub-spaces (see Figure 1.1) on the basis of the following three orthogonal dimensions:

- **agent** – the threat source or threat actor that imposes the threat on one or more assets. Aside from the “human” classification, technological threats relate to physical or chemical processes acting on materials, while force majeure agents are environmental.
- **motivation** – a binary classification that distinguishes between deliberate or accidental acts.
- **localisation** – a binary classification of origin related to the perimeter of the affected entity.

Figure 1.1

Three-dimensional categorisation model for threats



Source: Modified from Lucas et al.

⁶ World Economic Forum (2020), *The Global Risks Report 2020, 15th Edition*.

⁷ Gelbstein, E. (2013), “Quantifying Information Risk and Security”, *ISACA Journal*, Vol. 4.

⁸ Ruf, L., Thorn, A., Christen, T., Gruber, B., Portmann, R. and Luzern, H. (2008), “Threat Modeling in Security Architecture – The Nature of Threats”, *Information Security Society Switzerland (ISSS) – Working Group on Security Architecture*.

This classification scheme provides a structured approach which can be used by those wishing to describe threats as a part of scenario modelling (see Section 5). When considered in the context of cyber risk within the financial sector, the scope of **cyber threat** is more narrowly defined as:

“A circumstance with the potential to exploit one or more vulnerabilities that adversely affects cyber security.”⁹

All of the threat dimensions previously referred to are still relevant in the light of this definition, although the resulting shock negatively affects cyber security (see Section 2.1).

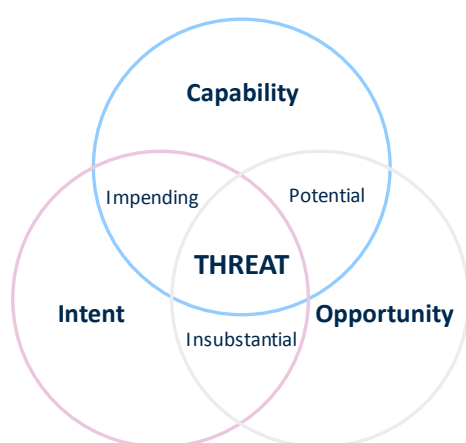
Beyond a description and classification of cyber threats, the model's evaluation of cyber risk requires at least taking a qualitative view of a likelihood under assessment. For the purposes of this paper, where the nature of a threat is either non-human or accidental, likelihood is based on the frequency and severity of historical events, in order to determine a “probability of occurrence”.

For deliberate acts where a human actor is involved, the categorisation of cyber threats should be enhanced to account for three further factors (see Figure 1.2):

- **capability**- the ability of threat actors to successfully achieve their intended goal and leverage an opportunity, through a combination of resources and knowledge;
- **intent** – the degree to which a threat actor has demonstrated its role, aims or purposes in causing harm;
- **opportunity** – the threat actor's timing and knowledge of the target space, including its vulnerabilities.

Figure 1.2

Probability characteristics for cyber threats arising from deliberate acts



Source: Modified from the Idaho National Library.¹⁰

⁹ Financial Stability Board (2018), *FSB Cyber Lexicon*.

¹⁰ Rodriguez, J. and Gasper, P.D. (2015), *Understanding the Value of a Computer Emergency Response Capability for Nuclear Security*, Idaho National Laboratory.

The treatment of cyber threats in this paper is not intended to be exhaustive – instead it aims to provide a degree of initial consistency to support comparability in scenario modelling. Further extensions such as describing the threat actors/sources and the vectors used, and exploiting sequencing, can be explored using frameworks such as MITRE ATT&CK¹¹, although these are of lesser importance as we step beyond the Context phase of the model.

It is also important to note that threats are contextual in respect of individual assets, services, entities or systems that may be affected. In a European context, publications such as the annual ENISA Threat Landscape Report provide a detailed overview of the entire cyber threat “ecosystem” which the financial sector is partly exposed to¹².

1.3 Vulnerability

The term **vulnerability** is often used in a more technical context by the cyber security community to describe specific occurrences and their associated characteristics, e.g. as in the Common Vulnerabilities and Exposures (CVE) database¹³. However, the conceptual model in this paper begins with the broader Financial Stability Board (FSB) definition to support the use of the term **vulnerability** in a thematic context.

“A weakness, susceptibility or flaw of an asset or control that can be exploited by one or more threats.”¹⁴

In this case, the target of the vulnerability is either an asset (see Section 1.4) or a control measure (a type of countermeasure as discussed in Section 1.5). When viewed through the lens of “operational risk”, as stated by the Basel Committee on Banking Supervision (BCBS), this base definition may benefit from two further extensions:

*“The risk of loss resulting from inadequate or failed internal **processes**, people and systems or from external events.”¹⁵*

First, the types of vulnerabilities considered could be extended to account for gaps (either incomplete or entirely missing), or to complement weakness (*inadequate quality*), susceptibility (*can be affected by something else*) and flaws (*defects or imperfections*). Second, the inclusion of defective processes allows the scope of vulnerability to include circumstances that are **causal** and that lead to the spreading of further vulnerabilities. For example:

- a deficient controls testing **process (flaw)** may result in an undetected **weakness** in **control measures** that could be exploited by a threat;

¹¹ The MITRE Corporation (2018), *MITRE ATT&CK: Design and Philosophy*.

¹² European Union Agency For Network and Information Security (2018), Threat Landscape Report.

¹³ The MITRE Corporation (2019), **Common Vulnerabilities and Exposures (CVE)**.

¹⁴ Financial Stability Board (2018), *FSB Cyber Lexicon*.

¹⁵ Basel Committee on Banking Supervision (2011), *Principles for the Sound Management of Operational Risk*.



- the lack of an internal cyber security awareness campaign (**process gap**) may lead to a lack of knowledge in individuals (**asset**) and a higher **susceptibility** to compromise.

With these changes, the Context phase of the model is able to support two categories of vulnerability:

- **idiosyncratic** – individual vulnerabilities with unique characteristics for each individual entity;
- **common** – the prevalent occurrence of the same or similar vulnerabilities with shared characteristics across a system. This category is also referred to as **thematic vulnerabilities or CIVs**.

A non-localised concept of vulnerability, i.e. using a system-wide context, is covered in Section 3.2.

Individual entities have to deal with a vast array of idiosyncratic vulnerabilities, any of which, if exploited, could give rise to a cyber incident. However, the presence of CIVs raises an additional concern, with simultaneous and independent risk crystallisation potentially occurring in different parts of a system.

In addition, vulnerabilities found within assets or control measures tend to be more localised and more bound to specific affected assets than those found in processes. Process-based CIVs have far greater ramifications, due to their prevalence and potential for generating new vulnerabilities within a system. Process-based CIVs that can also be exploited directly are of particular concern.

It is therefore essential to maintain a dual focus on:

1. the most significant idiosyncratic vulnerabilities at individual institutions;
2. the most prevalent, directly exploitable and “infectious” CIVs within a system.

Idiosyncratic vulnerability prioritisation is performed using a variety of techniques (e.g. OWASP¹⁶), and is based on factors, which may include but are not restricted to:

- **severity** – the possible consequences if a threat exploits the vulnerability;
- **prevalence** – observed occurrences within a system;
- **discoverability** – the degree of ease with which the vulnerability can be discovered;
- **exploitability** – the degree of ease with which the vulnerability can be exploited;
- **persistence** – the measure of longevity of the vulnerability (how long it has been in existence);
- **urgency** – the degree to which the vulnerability is being actively exploited;
- **detection** – the likelihood of exploitation detection;

¹⁶ Open Web Application Security Project (OWASP), **OWASP Risk Rating Methodology**.

- **effort** – the time/resources required to eliminate or minimise the potential for exploitation.

In 2018, 14 ESRB member authorities collectively identified the 13 CIVs listed in Figure 1.3. Notably, the top three CIVs identified are process-based, judged to be capable of direct causation, and highly prevalent (i.e. present in all/most of the 14 reporting jurisdictions). Prioritisation is based on individual jurisdictions' supervisory judgements and is not based on the potential systemic impact of the vulnerabilities. This information helps to inform microprudential authorities by:

- informing judgement for local prioritisation;
- providing a challenge process where a CIV is not locally identified. In other words, is the CIV not present/prevalent, or are the methods for detection deficient?

However, there is no evidence that any given CIV is more likely to lead to a systemic event. Prioritisation approaches which imply such a relationship may mislead authorities, leading them to focus on a subset of issues, and promoting a false sense of security. There may be a correlation between CIVs and a higher frequency of cyber incidents at individual institutions, but it is not possible to infer a systemic link from this correlation alone.

Mitigating idiosyncratic vulnerabilities and CIVs should reduce the likelihood of cyber incidents, although it will not eliminate them altogether. Given the premise that disruption will occur (i.e. cyber incidents are a matter of “when” not “if”), it is equally important to focus on system vulnerabilities in order to minimise systemic consequences.

Figure 1.3
Common Individual Vulnerabilities (CIVs) identified by the ESCG in 2018

Rank	Common individual vulnerability (CIV)	Category	Causation	Prevalence
1	Insufficient industry oversight of third party suppliers and supply-chain	Weakness in process	Direct	1
2	Inadequate cyber hygiene	Weakness in process	Direct	2
3	Ineffective testing of people, processes and technology	Flaw in process	Direct	5
4	Insufficient cyber strategic planning and board level influence on cyber resilience	Weakness in process	Indirect	3
5	Lack of investment in cyber threat intelligence	Gap in process	Indirect	4
6	Presence of end of life systems	Susceptibility/flaw in asset	Direct	6
7	Technology centric focus underestimating responsibility of people and processes	Weakness in process	Indirect	7
8	Organisational culture change needed for secure cybersecurity behaviours	Gap in process	Indirect	8
9	Cyber undermines existing operational resilience arrangements	Weakness in control measures	Direct	9
10	High risk internet use requires better controls	Weakness in control measures	Direct	12
11	Firm scale and resources impact effective cyber risk management	Susceptibility in process	Indirect	10
12	Insufficient credible third line of defence challenge in firms	Weakness in process	Indirect	11
13	Cyber maturity targets not defined	Gap in process	Indirect	13

Source: Vulnerability information sourced from a survey of ESRB members performed in Q1 2018.

Notes: CIVs are sorted by aggregate priorities across the 14 ESCG member authorities, then categorised using the terminology in Section 1.3, with the final column representing the ranking order of CIV presence across jurisdictions.



1.4 Assets

From a narrative perspective, the term **asset** has been included in the Context phase to describe the target, i.e. where the risk crystallises. However, assets do not modify the probability dimension of the cyber risk expression but have, instead, a stronger connection to consequences based on the asset value at risk. The model uses the FSB definition of an asset:

“Something of either tangible or intangible value that is worth protecting, including people, information, infrastructure, finances and reputation.”¹⁷

This definition lends itself to a financial accounting view of assets as shown in Figure 1.4¹⁸. The classification starts with a distinction between **tangible** and **intangible** assets, based on whether these take a physical form or carry finite monetary value. Within tangible assets, there is a further split between assets that can readily be converted into cash without disturbing normal operations (**current assets**), and those purchased for continued and long-term use (**fixed assets**). Intangible assets can also be divided into those that can be separated from other assets (**identifiable**) and those which cannot (**non-identifiable**).

In the context of cyber incidents, the asset scope in which risk can directly crystallise is limited to information or information systems (highlighted in **orange**), with the latter defined by the FSB as a:

“...set of applications, services, information technology assets or other information-handling components, which includes the operating environment.”¹⁹

Information systems therefore include all items that come under the categories of technology and software. Although people, machinery or reputation may be the targets of a malicious cyber incident, it is the directly affected information or information systems related to these assets which have a subsequent effect.

¹⁷ Financial Stability Board (2018), *FSB Cyber Lexicon*.

¹⁸ International Financial Reporting Standards (2018), *Conceptual Framework for Financial Reporting*.

¹⁹ Financial Stability Board (2018), *FSB Cyber Lexicon*.



Figure 1.4

A non-exhaustive breakdown of asset types

Classification		Type	Examples
tangible	fixed	property	buildings, equipment, machinery, vehicles, land, office space, office equipment, furnishings
		technology	information and communication technology (ICT) hardware storage equipment, servers, mainframes, back-up facilities, desktop equipment, network equipment, communications, voice services
			operational technology (OT) hardware Building management controls systems, SCADA systems, Industrial Control Systems (ICS), Distributed Control Systems, Intrusion Detection Systems, Physical Access Control Systems, Emergency Management Systems
	current	financial	cash (or cash equivalent), short-term investments (equity instrument), contractual claim (e.g. bank deposit, bond, stock)
Intangible	identifiable	people (human capital)	skills, talents, abilities
		software	operating systems (incl. virtual), applications (internal), applications (third party), middleware components, web components
		Information	datastores (RDBMS, key/value stores, document stores), file-based data (electronic or physical), code (in-house), third party libraries (purchased / open source), archived information
		Trade	knowledge, intellectual property, designs
		marketing	trademarks / copyright, customer lists, contract, distribution channels
	non-identifiable	goodwill / brand value	reputation, loyalty, market share

Source: Drawn from the IFRS' Conceptual Framework.

Notes: Asset types highlighted in orange reflect assets where risk arising from cyber incidents can crystallise.

An important characteristic of assets is their ability to act as a **store of value**, i.e. as a form of **capital** which can be saved, retrieved or exchanged at a later time. To that end, the three generally accepted functions of money are as a store of value, a medium of exchange and a unit of account. In an organisational context, value creation is a complex interplay between an entity, the capital that it holds, and its external environment. In 2013, the International Integrated Reporting Council (IIRC) published the model for value creation shown in Figure 1.5, which describes six different forms of capital: financial, manufactured, intellectual, human, social and natural²⁰. These pools of capital are constantly changing and are intrinsically tied to underlying assets that act as vessels for value.

With the onset of digitisation, money and other financial assets have increasingly converted from a physical to a binary form, and information assets have been transformed into stores of value. In the context of a cyber incident, an adverse impact on information or information systems may be reflected in a change to their inherent **value at risk**. The concept of value at risk (VaR) is commonly used, from a financial perspective, as a risk measure for losses on investments. With the aim of bringing a similar level of quantification to cyber risk, industry participants are actively developing approaches that mimic VaR, including the following notable examples: the World

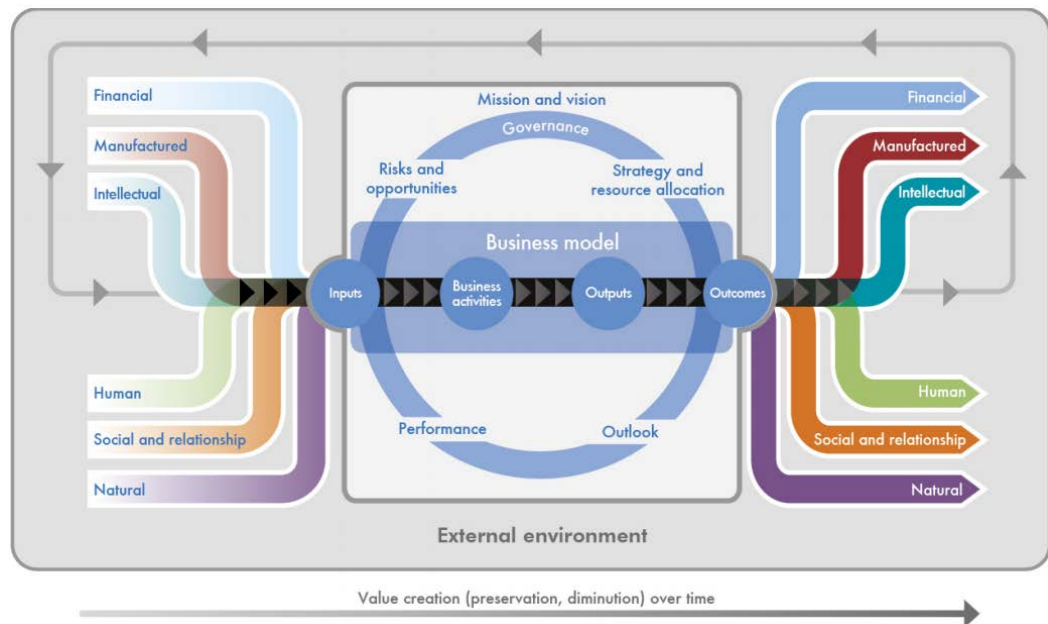
²⁰ International Integrated Reporting Council (2013), *The International <IR> Framework*.



Economic Forum's CyberVaR²¹, the International Monetary Fund's Framework for Quantitative Assessment²² and the Factor Analysis of Information Risk (FAIR) methodology²³.

Figure 1.5

An organisational model for value creation



Source: The IIRC.

1.5 Countermeasures

The model has so far described elements whose presence can increase cyber risk.

Countermeasures have the opposite effect, and NIST defines these in the context of cyber risk as:

“Actions, devices, procedures, techniques, or other measures that reduce the vulnerability of an information system. Synonymous with security controls and safeguards.”²⁴

Countermeasures can take many forms, using administrative, managerial, technical or legal methods to modify or manage cyber risk. These can be further broken down into the following common control measure types:

- **preventative** – designed to discourage errors or irregularities from occurring, e.g. the segregation of duties, access controls, authorisation;

²¹ World Economic Forum (2015), *Partnering for Cyber Resilience – Towards the Quantification of Cyber Threats*.

²² International Monetary Fund (2018), *Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment*.

²³ The Open Group (2013), *Technical Standard – Risk Taxonomy (O-RT), Version 2.0*.

²⁴ National Institute of Standards and Technology (2018), *Risk Management Framework for Information Systems and Organizations*.



- **detective** – designed to find errors or irregularities after these have occurred, e.g. exception reports, reconciliations, control totals, error reports;
- **directive** – designed to encourage a desirable event, e.g. accounting manuals, documented procedures, training and supervision;
- **corrective** – designed to help mitigate damage once a risk has materialised, e.g. complaint handling, virus isolation;
- **deterrent** – designed to reduce the likelihood of a vulnerability being exploited, without actually reducing the exposure, e.g. notices, cameras, fences;
- **compensatory** – alternative arrangements to use other control measures when the original control measures have failed or cannot be used;
- **recovery** – applied in more significant disruptions to restore services to within expected levels, e.g. DR/BC mechanisms, backup systems.

The role of countermeasures in the model is flexible, and Section 5 shows the two forms the model may take. Initially, a scenario-based analysis may disregard the modifying properties of countermeasures as if they were non-existent or ineffective, to present a worst-case scenario. Subsequent runs may re-introduce countermeasures to show a more realistic approximation.

1.6 Starting point

In addition to the elements described using the cyber risk expression, the concept of **starting point** is also captured as part of the Context phase. The starting point represents one of three generalised entry points into the model, from which a cyber incident could originate:

- **single firm affected** – impairs one component of the financial system, e.g. a systemically important firm or financial market infrastructure;
- **multiple firms affected** – impairs multiple components simultaneously and independently;
- **supply chain affected** – affects a dependency upon which the financial sector relies, e.g. a technology provider, which in turn becomes the source of disruption or introduces an infection vector through which disruption can propagate.

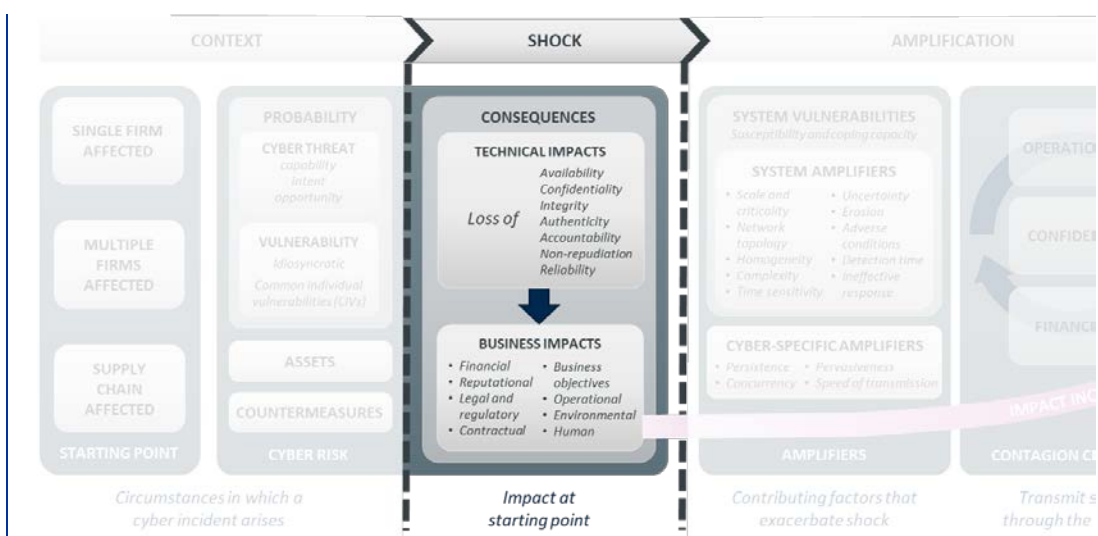


2 Shock

The initial outcome of cyber risk crystallisation at the starting point

In economics, a shock is loosely described as an unexpected or unpredictable event that produces significant change within an economy, either positively or negatively.

The **Shock** phase of the model describes the localised effects of a cyber incident which could subsequently affect a system more broadly.



2.1 Consequences

The final element of the cyber risk expression (see Section 1.1) describes the **consequences** of one or more cyber threats successfully exploiting a vulnerability. Consequences arising from disruption are typically expressed in the form of **impacts**, defined by the ISO as:

“...the evaluated consequence of a particular outcome...”²⁵

In the model, impacts are presented in two stages:

- **technical impacts** which describe the immediate effects of disruption on the asset(s) affected (see Section 2.2);
- the subsequent **business impacts** which describe organisational repercussions (see Section 2.3).

²⁵ International Organization for Standardization (2018), *Security and resilience – Vocabulary*.



The measurement of impact involves studying lagging indicators that can only be collected after an incident occurs. Impacts can be expressed either quantitatively or qualitatively, and may not be immediately observable. In Section 2.4 the model explores the implications of impact measurement in further detail.

2.2 Technical impacts

Immediate technical impacts occur at the moment cyber risk crystallises and results in a **cyber incident**, which the FSB defines as:

“...a cyber event...that jeopardizes the cyber security of an information system or the information the system processes, stores or transmits...”²⁶

In turn, **cyber security** aims to preserve the specific properties of information and information system assets:

“Preservation of confidentiality, integrity and availability of information and/or information systems through the cyber medium. In addition, other properties, such as authenticity, accountability, non-repudiation and reliability can also be involved.”²⁷

Each cyber security property is defined in Figure 2.1, accompanied by examples of disruptive events that could lead to the loss of each property. Commonly, the **CIA triad** (confidentiality, integrity, availability) is the typical focus. However, a broader approach, which includes all seven properties, is required where non-CIA properties are a contributing factor to cyber incidents, e.g. where reputation is targeted.

In many cases, a combination of properties can be affected across related assets. The scenario case studies referred to in Section 5 provide hypothetical and real-world examples of such circumstances. The technical impacts associated with the loss of cyber security often result in business impacts for related entities and the services they provide.

²⁶ Financial Stability Board (2018), *FSB Cyber Lexicon*.

²⁷ Idem.



Figure 2.1

Asset properties which cyber security aims to preserve, and examples of disruptive events related to the loss of each property

Property	Definition	Examples of related disruptive event(s)
Availability	property of being accessible and usable on demand by an authorised entity	• three types of availability loss: total, partial, or intermittent
Integrity	property of accuracy and completeness	• data manipulation (creation, addition, duplication, modification, re-sequencing, deletion) • data corruption (unreadable, but recoverable or can be reconstituted) • data destruction (irrevocable)
Confidentiality	property that information is neither made available nor disclosed to unauthorised individuals, entities, processes or systems	• unintended / unauthorised disclosure (e.g. data leakage) • unauthorised acquisition (e.g. exfiltration, interception)
Authenticity	property that an entity is what it claims to be	• impersonation / cloned identity
Accountability	property that ensures that the actions of an entity may be traced uniquely to that entity	• injection (e.g. man-in-the-middle)
Non-repudiation	ability to prove the occurrence of a claimed event or action and its originating entities	• rumour / speculation • disinformation
Reliability	property of consistent intended behaviour and results	• (technological) degradation of operations

Source: Property definitions taken from FSB Cyber Lexicon.

2.3 Business impacts

The study of business impacts has been a cornerstone of business continuity planning since it originated in the 1950s. This concept has since matured, with **business impact analysis (BIA)** featuring as a core activity within modern resilience management. For the model, a generalised set of impact categories are proposed based on the ISO technical specification for BIAs²⁸. The standard describes a core set of impact categories, with examples as follows:

- **Financial** – financial losses due to fines, penalties, lost profits or diminished market share;
- **Reputational** – negative opinion or brand damage;
- **Legal and regulatory** – litigation liability and withdrawal of licence of trade;
- **Contractual** – breach of contracts or obligations between organisations;
- **Business objectives** – failure to deliver on objectives or take advantage of opportunities.

For the model, three impact categories are added to capture other consequences of disruption.

- **Operational** – discontinued or reduced service levels, workflow disruptions, or supply chain disruptions;

²⁸ International Organization for Standardization (2015), *Societal security – Business continuity management systems – Guidelines for business impact analysis (BIA)*.

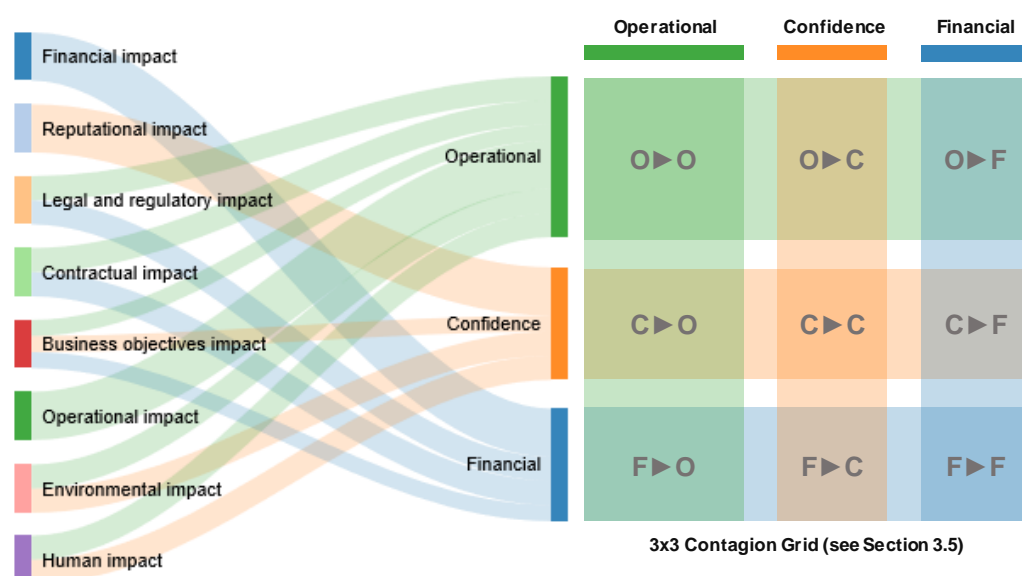


- **Environmental** – harmful effects on the biophysical environment;
- **Human** – loss of life, injury, impact to community, short and long-term emotional impact.

The model creates a relationship between the business impacts where the cyber incident originated, and the contagion channels (see Section 3.7) that may propagate the shock through a system (as illustrated in Figure 2.2).

Figure 2.2

Business impact “inputs” into contagion channels



Source: Impact categories modified from ISO/TS 23317:2015.

2.4 Impact measurement

The measurement of business impacts within individual institutions is typically composed of two complementary approaches:

- **qualitative** (judgement-based) – the use of descriptive statements to define levels of increasing severity for each impact category. Over the course of an incident, appraisals are regularly performed against these qualitative scales to approximate impact and to drive appropriate organisational responses. As it is easy to implement, this approach is by far the most common, although it relies on the consistent judgement of individuals who may introduce bias or subjectivity.
- **quantitative** (metric-based) – the use of data-driven indicators to individually or collectively determine a degree of impact. Quantitative approaches generally involve more challenges in initially defining and sourcing accurate and timely data to use as part of an incident response.

Measuring impact can require significantly different **time horizons** following an incident – these may be short-term (days), medium-term (months) or long-term (years) in nature.

Figure 2.3 describes potential impact indicators alongside the time horizon for their measurement. Horizons are constrained by the sourcing of data for each indicator, some of which are more immediate and “point-in-time”, compared with others which can only be measured over time, or after a prolonged period of time. Further study of these indicators to establish their relative efficacy and accuracy would be beneficial as the subject of future work.

Figure 2.3
Potential impact indicators

Category	Measurement subject	Potential impact indicators	Horizon
Financial	Incident related expenditure	Costs incurred to handle the incident (e.g. technical investigation, return to normal operations, public relations)	Short
	Customer detriment	Costs incurred for incident notification, customer protection, financial reimbursement	Short to Medium
	Deposit stability	Measuring deposit rates and flows over time, e.g. Liquidity Coverage Ratio (LCR)	Short to Medium
	Market value	Stock valuation using Cumulative Abnormal Returns (CAR)	Medium to Long
	Profitability ratios	Gross/net profit margin, EBITDA, Return on Equity (ROE)	Medium to Long
	Perceived business risk	Costs for raising debt relative to credit rating, change in insurance premiums	Medium to Long
	Loss of competitive advantage	Value of lost contract revenue or loss of intellectual property	Long
	Regulatory costs	Fines or fees levied for non-compliance	Long
	Capital charges	Pillar II capital add-on	Long
	Litigation	Penalties for breach of contract, settlement costs	Long
Confidence	Media coverage	Negative news signals, volume of attention, adverse social media activity, reputation index, number of press enquiries	Short to Medium
	Investor sentiment	Market-based measures (e.g. CBOE Volatility Index) and technicals (e.g. Relative Strength Index, Money Flow Index)	Short to Medium
	Brand valuation	Customer surveys for top-of-mind awareness, familiarity, advocacy	Medium to Long
	Customer metrics	Measuring customer satisfaction, loyalty, acquisition, retention, problem incidence	Medium to Long
Operational	Business service disruption	Trading volumes and values, service availability, performance indicators	Short to Medium
	Downstream impact	Number of dependent services disrupted, number of third parties / customers affected	Short to Medium

Source: Developed by the ESRB's ESCG.

Notes: Further research on the following indicators is referenced as follows: the Deloitte study on business impacts arising from cyber incidents²⁹; a stock valuation using CAR³⁰; investor sentiment³¹ and reputation management³².

²⁹ Deloitte (2016), *Beneath the surface of a cyberattack – A deeper look at business impacts*.

³⁰ Yayla, A.A. and Hu, Q. (2010), *The impact of information security events on the stock value of firms: the effect of contingency factors*.

³¹ Baker, M. and Wurgler, J. (2007), “Investor Sentiment in the Stock Market”, *Journal of Economic Perspectives*, Vol. 21, No 2, pp. 129-151.

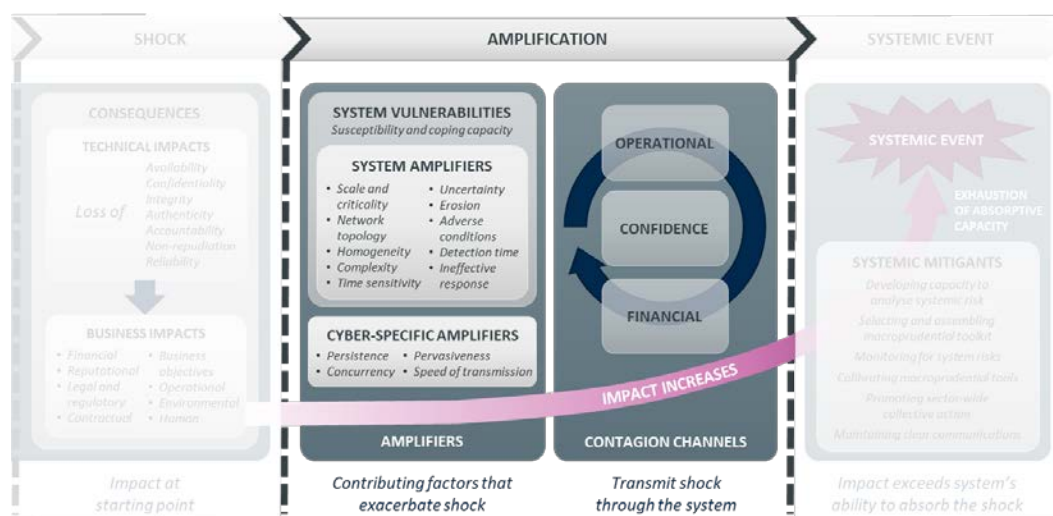
³² Eisenegger, M. and Imhof, K. (2007), *The True, the Good and the Beautiful: Reputation Management in the Media Society*.



3 Amplification

The propagation of shocks inside the financial system

As the initial shock emanates from its starting point, the **Amplification** phase of the model describes how these impacts can intensify and spread across the financial system. Contagion channels trigger in an escalating manner, with the presence and alignment of one or more systemic amplifiers fuelling the resulting cascading effect.



3.1 Complex adaptive systems

Before exploring the effects of shocks on the financial system, it is necessary to establish how the term **system** and its associated characteristics are used within the model. Although the notion of systems has existed for centuries, Ludwig von Bertalanffy (1968) is credited with pioneering the field of **general systems theory**, which considers a system as an organised entity made up entirely of interrelated and interdependent parts³³.

Systems are delineated by **boundaries** that define and distinguish them from other systems in the environment. The financial system as a whole may be considered not as a single entity, but as a **system of systems**, bounded in the following ways:

- **global** – the entire financial system ecosystem, and its intersection with the real economy;
- **jurisdiction** – to the extent it is covered by supra-national or national laws, rules or regulations for a given financial sector;

³³ von Bertalanffy, L. (1968), *General Systems Theory: Foundations, Development, Applications*.

- **functional** – the entities and services which are necessary for the delivery of a specific economic function (which may span jurisdictions).

A system is **complex** when it is hard to infer its dynamics from the behaviour and state of its individual components. This complexity relates to a variety of system properties³⁴, including:

- **emergence** – out of the interactions between the individual elements in a system, behaviour emerges at the level of a system as a whole;
- **non-linear dynamics** – systems may suddenly change behaviour or move to another regime, e.g. from a high degree of stability to very unstable behaviour;
- **direct and indirect feedback loops** – these can be both positive (enhancing, stimulating) or negative (detracting, inhibiting);
- **limited predictability** – the behaviour of complex systems cannot be predicted well. Small changes in initial conditions or history can lead to very different dynamics over time.

In a union of both systems and complexity theory, Haldane views the financial system as a **Complex Adaptive System (CAS)**, drawing parallels with other network disciplines such as ecology, epidemiology, biology and engineering³⁵. In particular, Haldane considers the highly interconnected financial system to be **robust-yet-fragile (RYF)**, and one whose feedback effects under stress add to these fragilities. Gai illustrates this RYF property using the example of the counterparty losses of a failing institution being widely dispersed to, and absorbed by, other entities (i.e. the system being robust)³⁶. If the failure of a single institution triggers contagious defaults, the high number of financial linkages also increases the potential for contagion to spread more widely (i.e. the system being more fragile). Market participants are subsequently more vulnerable to a second round default.

The model describes such fragilities as system vulnerabilities (see Section 3.2) that act as amplifiers (see Section 3.3) that spread impacts through multiple and often concurrent contagion channels (see Section 3.7).

3.2 Vulnerability in a system context

Section 1.3 introduces the notion of “vulnerability” in a localised context. However, it can also be conceptualised as a series of increased degrees of complexity and scale, as proposed by Birkmann³⁷.

As the concept is widened to the third sphere (see Figure 3.1), **system vulnerabilities** and their interactions with perturbations and stresses are considered as a function of:

³⁴ Cilliers, P. (1998), *Complexity and postmodernism - Understanding complex systems*.

³⁵ Haldane, A.G. (2009), *Rethinking the financial network*.

³⁶ Gai, P. (2013), *Systemic Risk: The Dynamics of Modern Financial Systems*.

³⁷ Birkmann, J. and Wisner, B. (2006), “Measuring the un-measurable: The challenge of vulnerability”, *United Nations University Institute for Environment and Human Security*.



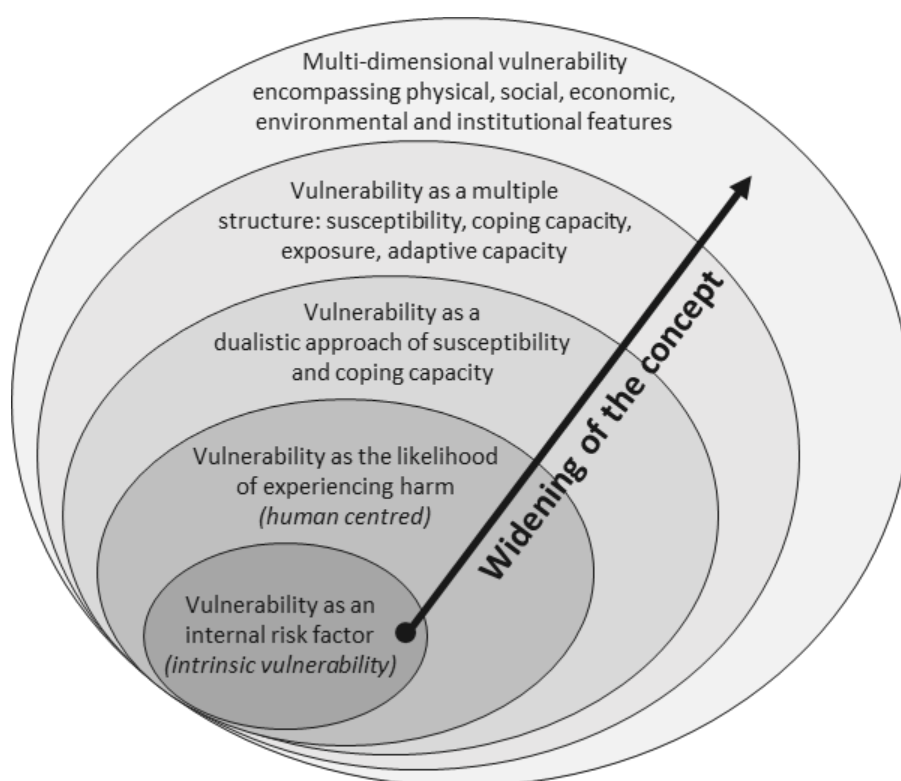
- **susceptibility** – system features, weaknesses, or states that increase the probability or consequences (or both) of adverse events³⁸ (this is further explored in Section 3.3);
- **coping capacity** – the extent to which a system is able to absorb shocks deriving from disruption (also referred to as absorptive capacity in Section 4.2).

The fourth sphere in Birkmann's model introduces two further parameters for consideration:

- **exposure** – the degree, duration and/or extension in which a system is in contact with, or subject to, perturbations;
- **adaptive capacity** – the extent to which a system can deal with, accommodate, recover, and learn from changing circumstances to facilitate desirable system outcomes.

Although these additional parameters may be relevant to understanding system behaviour during periods of disruption, the concepts fall outside the scope of the current model. Future efforts may seek to extend the model to account for these characteristics.

Figure 3.1
Different spheres of vulnerability



Source: Reproduced from Birkmann.

Notes: This paper explores the concepts of susceptibility and coping capacity in relation to system vulnerability found in the third sphere of this model.

³⁸ Hassel, H. (2007), "Risk and Vulnerability Analysis of Complex Systems: a basis for proactive emergency management", LUCRAM (Lund University Centre for Risk Analysis and Management).

3.3 Amplifiers

In its 2018 analysis, the ESCG identified the characteristics of the financial system that could make it susceptible to systemic disruption arising from a cyber incident:

- a high degree of **interdependence**;
- a lack of **transparency**;
- a reliance on **data**;
- a reliance on **confidence**.

These and other characteristics can represent system vulnerabilities with the potential to become **amplifiers** of a shock through a system. Figure 3.2 builds on the original list, differentiating between system amplifiers that may exacerbate any operational disruption (with the addition of **pre-existing weaknesses**), and cyber-specific amplifiers that relate to unique features of cyber incidents (e.g. intent, scale and speed). The paper will now explore each of these amplifiers in greater detail.

Figure 3.2

Amplifiers which, if present, may exacerbate impacts through a system

Characteristic	Amplifier	Description
System amplifiers (may feature in any operational disruption)		
High degree of interdependence	Scale and criticality	Business services and economic functions affected
	Network topology	Structure of a system and position of affected node(s) in that system
	Homogeneity	Common components or processes failing in the same way contributing to system fragility
Lack of transparency	Complexity	Unnecessary complication or insufficient understanding of the behavioural, cognitive, structural or contextual traits in a system
Reliance on data	Time sensitivity	Tipping points at which a system is more susceptible to disruption
Reliance on confidence	Uncertainty	Ambiguous or incomplete information which has a direct effect on confidence
	Erosion	System 'memory' where previous adverse events have a negative compound effect
	Adverse conditions	Environmental circumstances which exacerbate shocks
Pre-existing weakness	Detection time	The longer a disruption goes undetected, the greater the shock
	Ineffective response	Failure to coordinate a timely system-wide response
Cyber-specific amplifiers (characteristics which are unique to cyber incidents)		
Intent	Persistence	Prolonged and targeted efforts to relentlessly pursue motives
Scale	Concurrency	Compound shocks which can distract or overwhelm the target
	Pervasiveness	Ability to scale asymmetrically, with cross-border and cross-sectoral reach
Speed	Speed of transmission	Rapid failure propagation through hyper-connected financial and operational networks

Source: Developed by the ESRB's ESCG.



3.4 System amplifiers

Scale and criticality

The first systemic amplifier is not related to a system vulnerability, but instead to the attributes of the originating node in the context of a system (note that interconnectedness is tackled separately as a part of network topology). The node's systemic importance can be influenced by many factors, but will invariably be dictated by the extent to which it contributes to the delivery of overarching critical (economic) functions. This approach is consistent with the view of the BCBS which stresses that *“systemic importance should be measured in terms of the impact that a failure of a bank can have on the global financial system and wider economy rather than the risk that a failure can occur.”*³⁹

In its guidance for recovery and resolution planning, the FSB sets out an approach for determining criticality based on⁴⁰:

1. an analysis of the impact of the sudden discontinuance of the function;
2. an evaluation of the market for that function;
3. an assessment of the impact of the failure of a specific node that contributes to that function.

The scale of the initial shock, coupled with the criticality of the affected node, will directly affect the intensity of the shock that is transmitted through a system.

The FSB guidance also describes the identification of critical functions and critical shared services, including a non-exhaustive list of functions that could potentially exhibit some degree of criticality. The summary of this list in Figure 3.3 seeks to provide a starting point for evaluating whether an affected node within the system contributes to the provision of these critical functions which serve the real economy.

³⁹ Basel Committee on Banking Supervision (2011), *Global systemically important banks: Assessment methodology and additional loss absorbency requirement*.

⁴⁰ Financial Stability Board (2013), *Recovery and Resolution Planning for Systemically Important Financial Institutions: Guidance on Identification of Critical Functions and Critical Shared Services*.



Figure 3.3

FSB guidance on critical functions

Function	Description	Drivers of criticality
Deposit taking	Acceptance of deposits from non-financial intermediaries	<ul style="list-style-type: none"> • Ability to take new deposits where depositors are willing to use alternatives. • Deposits with little residual maturity carry liquidity risk exposure when experiencing sudden or extensive withdrawals. • Retail deposits are deemed more critical, as these tend to be less diversified. • Transaction (current) accounts with continuous access are deemed critical for the smooth settlement of day-to-day financial transactions. • Presence of effective deposit protection arrangements can significantly reduce criticality.
Lending and loan services	Provision of funds to non-financial counterparties, such as corporates or retail customers, and can extend through to loan servicing functions	<ul style="list-style-type: none"> • Lending can become critical if liquidity and funding strains occur for the borrowers before customers can find alternate sources of credit. • Particular lending products may be critical in a given jurisdiction. • Extent to which products, collateral terms, and underwriting are standardised. • Over-reliance on short-term lending to close temporary liquidity gaps. • Disruption to trade finance may have a spill-over effect on international flow of goods.
Payments, Clearing, Custody & Settlement	Multilateral systems among participating institutions, used for the purposes of clearing, settling, or recording payments, securities, derivatives, or other financial transactions; or provision of these services as an intermediary	<ul style="list-style-type: none"> • Payment systems that have the potential to trigger or transmit systemic disruptions, especially if they are the sole payment system in a country or the principal system in terms of the aggregate value of payments; systems that mainly handle time-critical, high-value payments; and systems that settle payments used to effect settlement in other systemically important FMI. • The presumption is that all CSDs, SSSs, CCPs, and TRs are systemically important, at least in the jurisdiction where they are located. • Presence of substitutability across payment channels and asset classes.
Wholesale Funding Markets	Lending and borrowing in wholesale markets to and from financial counterparties	<ul style="list-style-type: none"> • Systemic relevance of wholesale market and specific segments of that market within a jurisdiction. • Degree of interconnectedness on either the borrower or lender side of the market. • Individual market participants with high market share in wholesale activities. • Presence of excessive maturity transformation or leverage.
Capital Markets and Investments activities	Issuance and trading of securities, related advisory services, and related services such as prime brokerage, as well as investment of the firm's own capital in private equity or similar principal investments	<ul style="list-style-type: none"> • Significant presence of capital-market based financing or concentration in capital markets-related functions • Number of alternate operators in primary markets with sufficient distribution capacity and relevant expertise. • Transactional intensity in secondary markets during times of stress or disruption. • Degree of portability of client accounts.

Source: FSB⁴¹.

⁴¹ Financial Stability Board (2013), *Recovery and Resolution Planning for Systemically Important Financial Institutions: Guidance on Identification of Critical Functions and Critical Shared Services*.



Network topology

Many complex networks have been found, in practice, to exhibit the property of being “scale-free”, i.e. the characteristics of the network are independent of the size of the network. That is to say, they comprise a core set of nodes with a large number of connections and a large set of peripheral nodes with few connections (see Figure 3.4). As the network grows, the underlying structure remains the same. There is a **core-periphery**, or hub-and-spokes, network configuration.⁴²

These scale-free topologies have important, if subtle, implications for system resilience. For example, core-periphery models with higher degrees of concentration have been found to be very robust, at a systemic level, to random shocks. That is because these shocks are very likely to fall on peripheral nodes unconnected with, and hence unlikely to cascade through, a system as a whole. However, these systems are also vulnerable to targeted disruption on the core nodes – the “super-spreaders” – whose hyper-connectivity risks generating a systemic cascade.

Therefore, the position or **centrality** of the distressed or failed node in the network can have a significant amplification effect⁴³. Centrality may be measured by the number of links that terminate on a node (in degree), by the distance from other nodes (closeness), or by the existing connections to central nodes. A measure of centrality particularly suitable for financial networks is the **betweenness centrality** of a node, which is defined as the number of shortest paths that go through the node. However, the impact of network properties and the effectiveness of centrality measures depend on the behaviour of the nodes in each case. These properties have limitations in that they do not adequately capture all complex behaviour.

⁴² Haldane, A.G. (2015), *On microscopes and telescopes*.


⁴³ Borgatti, S.P. (2005), “Centrality and network flow”, *Social Networks*, No 27, pp. 55-71.




Figure 3.4

Illustration of random versus scale-free networks

a) Random network



b) Scale-free network



Source: Inspired by Barabási⁴⁴, and Dalziel and McManus⁴⁵.

Homogeneity

As in ecology, the long-term sustainability or vitality of a financial system is strongly associated with two structural attributes: efficiency and resilience⁴⁶. These attributes are related to the levels of **diversity** and **connectivity** found in a system, albeit in opposite directions (see Figure 3.5).

A highly interconnected and diverse system plays a positive role in resilience because additional options help a system to rebound from the loss or disruption of one or more pathways or nodes. Conversely, redundant pathways and excess diversity can hinder throughput efficiency, leading to stagnation that erodes vitality by dissipating weak throughput via various inefficient sectors. In short, resilience and efficiency are mutually competitive, because the streamlining that increases efficiency automatically reduces resilience. This inherent push-pull trade-off explains why, at a certain point, increasing a system's efficiency makes it more brittle, even as it becomes bigger and more directed.


⁴⁴ Barabási, A.-L. (2002), *Linked: The New Science of Networks*.

⁴⁵ Dalziel, E.P. and McManus, S. (2004), *Resilience, Vulnerability, and Adaptive Capacity: Implications for System Performance*.

⁴⁶ Goerner, S.J., Lietaer, B. and Ulanowicz, R. (2009), "Quantifying economic sustainability: Implications for free-enterprise theory, policy and practice", *Ecological Economics*, Vol. 69, No 1, pp 76-81.

The lack of diversity, or **homogeneity**, in a system can take form at different levels. At an asset level, it could be reflected in the use of identical components or products, while at a functional level there could be convergence on methodology or process. For example, the production and use of IT systems is characterised by significant economies of scale and, as a consequence, there is a limited variety of products in the market. If a vulnerability is identified, it has the potential to be simultaneously exploited in a multitude of IT systems, as shown by the WannaCry example overleaf. This homogeneity applies even to products that provide cyber security (e.g. antivirus software). Overall, a lack of diversity can contribute to the fragility or brittleness of the financial system, allowing small shocks to be magnified exponentially.

Figure 3.5
Sustainability as a function of efficiency and resilience



Source: Reproduced from *Quantifying economic sustainability*, Goerner et al..

Box 1 WannaCry Ransomware (May 2017)

In early 2017, the US National Security Agency became aware of a security breach which had resulted in the exfiltration of hacking tools that had remained secret for more than five years⁴⁷. Two extracts from this data theft were deemed to be particularly potent:

⁴⁷ Nakashima, E. and Timberg, C (2017). "NSA officials worried about the day its potent hacking tool would get loose. Then it did.", *Washington Post*, 16 May [Online].



- **EternalBlue** – an exploit which takes advantage of a vulnerability in the Microsoft Windows Server Message Block (SMB), allowing malicious code to spread within compromised networks without user interaction, seeking out adjacent vulnerable hosts;
- **DoublePulsar** – a sophisticated memory-based kernel exploit that hooks onto Windows systems, granting a high level of control to execute malicious code.

Upon discovering the vulnerabilities, Microsoft issued a security bulletin (MS17-010) and critical patches on 14 March 2017 for all supported versions of Windows. On 14 April 2017, a hacking group known as **The Shadow Brokers** released details of these exploits and other tools into the public domain. Less than one month later, a threat actor had repurposed these tools and, on 12 May 2017, launched a ransomware campaign now commonly known as **WannaCry**. Fortunately, WannaCry's spread was cut short through the accidental discovery of a kill-switch by an independent researcher later that same day.

Even though it was quickly halted, the malicious code is reported to have affected more than 300,000 computers, across 30,000 firms in 150 countries, including the United Kingdom's National Health Service, FedEx, Telefonica, Nissan and Renault⁴⁸. Microsoft also took the unusual step of issuing software patches for unsupported versions of Windows (e.g. XP), although 98% of infected computers were running Windows 7⁴⁹. Estimates of the economic losses caused by WannaCry range from the hundreds of millions to USD 8 billion.

This cyber incident highlights the homogeneous nature of the technological environment upon which many sectors, including finance, rely. The extensive and sometimes critical dependency on Windows was compounded in this instance by a widespread lack of basic cyber hygiene, demonstrated by a failure to perform patch management on a timely basis.

Complexity

Given its varied application across different scientific fields, a unified definition of **complexity** is yet to be achieved, although most scholars would support the view that it describes *an emergent property from a collection of interacting objects*. Section 3.1 introduced the concept of complexity through the lens of complex systems and their properties. When considered as a system amplifier, complexity has an inverse relationship to transparency. As interconnectedness increases across and between complex systems, the observed behaviours become less understood or predictable.

It is important to distinguish between *complex* and *complicated*. As previously stated, complex systems have feedback loops and behave in an unpredictable non-linear fashion. On the other hand, complicated systems typically consist of many parts, some of which have a rationale which is difficult to understand, and could be defined or described more simply.

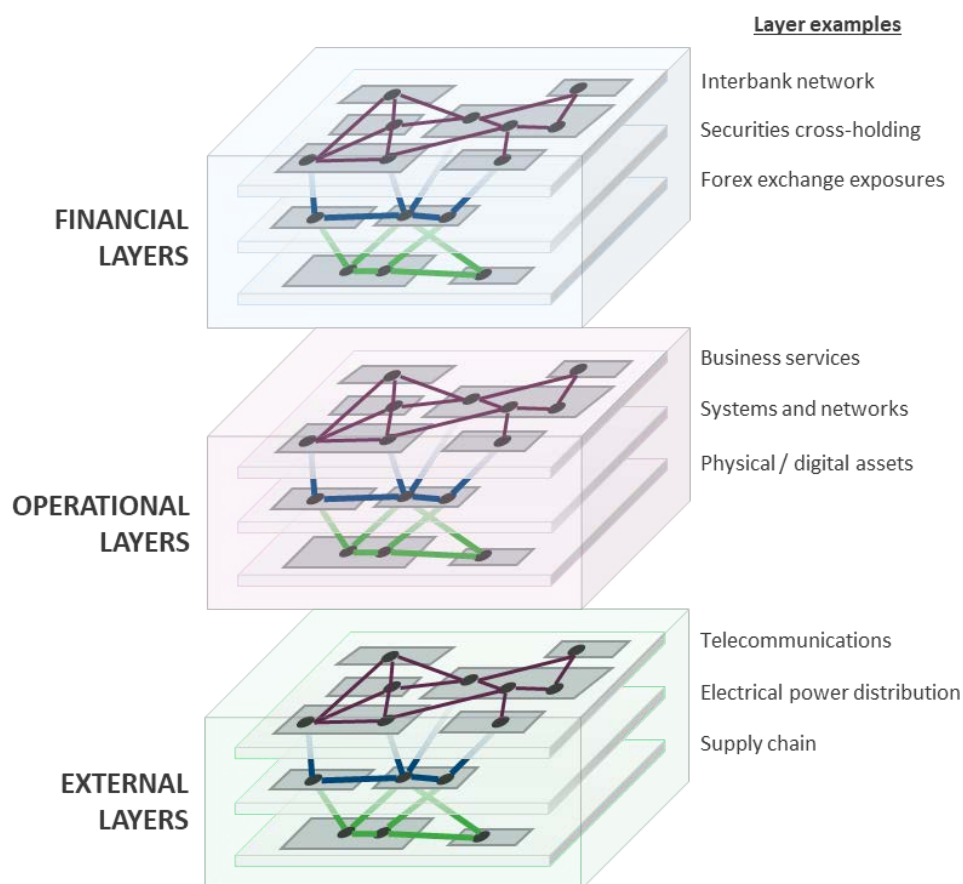
⁴⁸ Benfield, Aon (2017), *Cyber Event Briefing: Wannacry Ransomware Attack*.

⁴⁹ Goodin, D. (2017), "Windows 7, not XP, was the reason last week's WCry worm spread so widely", *Ars Technica*, 20 May [Online].



Figure 3.6

The financial sector as a multi-layered network of complex systems



Source: Inspired by the Office of Financial Research briefing on "A Multilayer Map of the Financial System"⁵⁰.

Notes: This figure is a stylised image which aims to convey the complexity of interconnections which exist between and across a non-exhaustive list of disparate complex systems.

Figure 3.6 depicts the financial sector as system of systems, made up of individual layers that have evolved collectively over time, and that have become more complex and complicated in equal measure. The ability of any individual or entity to understand the ramifications of disruption when a cyber incident occurs becomes exponentially challenging as these layers are traversed.

Equally, this inherent complexity makes cyber threats asymmetric. Cyber incidents will inevitably occur at various levels within the stack, making them impossible to prevent. This is especially true from a software perspective, given that all code will contain vulnerabilities that are both *unprovable* and *unknowable*, as demonstrated by Alan Turing in his Halting Problem theorem⁵¹.

⁵⁰ Bookstaber, R. and Kennet, D.Y. (2016), "Looking Deeper, Seeing More: A Multilayer Map of the Financial System", Office of Financial Research.

⁵¹ Armstrong, R.C. and Mayo, J.R. (2009), *Leveraging Complexity in Software for Cybersecurity*.

Securing arbitrary code is therefore not just hard – it is impossible. When scaled up to consider not just isolated code fragments, but the vast interplay within a technological ecosystem, traditional vulnerability management techniques may no longer be sufficient to understand all of the possible failure states in a way that adequate protection can be achieved.

Time sensitivity

The timing of a disruption has implications in terms of both its starting point (affecting the size of the initial shock) and its propagation through a system. In complex systems, structural transitions from one state to another are referred to as **tipping points**⁵². If a shock were to occur in the lead up to or during a state change in a system, the repercussions could be amplified during this heightened state of susceptibility. For the financial system, and for financial market infrastructure in particular, the timing of a shock can lead to a significant amplification effect, e.g. during critical time periods or outside core operating hours.

Box 2

Critical time points for the United Kingdom's RTGS and CHAPS systems

As shown in Figure 3.7, the daily timetable for systems operated by the Bank of England (i.e. the Real Time Gross Settlement system and CHAPS) indicates the end of settlement day times for different transaction types. These times represent significant tipping points that, unless contingency measures are taken, could amplify a shock if any operational issues were to occur towards the end of the day. To mitigate this issue, the Bank of England has instituted (at least) two operating procedures: (i) CHAPS Direct Participants are expected to settle 50% of payments, by value, by 12 p.m., 75% by 3 p.m. and 90% by 5 p.m. (as indicated in **orange**); and (ii) a two-hour extension window is available as a contingency for unexpected disruptions.

⁵² Jurczyk, J., Rehberg, T., Eckrot, A. and Morgenstern, I. (2017), "Measuring critical transitions in financial markets", Nature, Vol. Scientific Reports, No 7.



Figure 3.7

Summary of the Bank of England's RTGS daily timetable

RTGS, CHAPS and Net Settlement Events	Time	CREST Event	Time
Transfers between own accounts and Enquiry Link access enabled; and Notes Circulation Scheme settlement	05:15		
Start of CHAPS settlement	06:00	Start of Delivery vs. Payment (DvP) / Free of Payment (FOP) settlement	06:00
Hourly CLS pay-in and/or pay-out deadlines	07:00 to 11:00		
Faster Payments settlement	07:05		
Bacs settlement	09:30		
Cheque & Credit settlements (up to six settlements between 10:40 and 11:10)	10:40		
LINK settlement	11:00		
CHAPS 50% Throughput Target	12:00		
Faster Payments settlement	13:05		
Visa settlement	14:00		
		End of equity and gilt DvP settlement	14:55
CHAPS 75% Throughput Target	15:00	Start of Delivery By Value (DBV) settlement	15:00
Image Clearing System settlement	16:30		
CHAPS 90% Throughput Target	17:00		
Faster Payments settlement	17:05		
		End of DBV settlement	17:30
End of CHAPS settlement for customer payments (MT103)	17:40		
End of CHAPS settlement for interbank payments (MT202)	18:00	End of FOP settlement	18:00
Notes Circulation Scheme settlement	18:30		
Latest end of contingency extension	20:00	Latest end of contingency extension	20:00

Source: Bank of England⁵³.

Uncertainty

As a general concept, uncertainty relates to situations involving imperfect or unknown information. Famously described by Donald Rumsfeld, uncertainty can be separated into “*known unknowns*” and “*unknown unknowns*”⁵⁴. The latter, known as **Knightian uncertainty**, is considered to have no delimiting parameters, and is based on a fundamental degree of ignorance, a limit to knowledge, or the unpredictability of future events⁵⁵. When applied to complex systems, uncertainty can be classified as follows⁵⁶:

- **ambiguity** – uncertainty over probability, created by missing information that is relevant and could be known;

⁵³ Bank of England (2018), *Bank of England's RTGS and CHAPS services: Service Description*.

⁵⁴ Rumsfeld, D. (2002), Department of Defense News Briefing [Interview], 12 February.

⁵⁵ Knight, F.H. (1921), *Risk, uncertainty and profit*, Houghton Mifflin, Boston.

⁵⁶ Thunnissen, D.P. (2003), *Uncertainty Classification for the Design and Development of Complex Systems*.



- **epistemic** – incomplete information or incomplete knowledge of some characteristic of a system or the environment;
- **aleatory** – inherent variation associated with a physical system or environment under consideration;
- **interaction** – unanticipated interaction of many events and/or disciplines, each of which might, in principle, be (or should have been) foreseeable.
- Uncertainty as a systemic amplifier has a strong link to the confidence contagion channel. For example, agents may feel less confident in their actions when operating with incomplete information, or where there is ambiguity regarding the shock or its consequences. They have to rely on their subjective sentiment and perception of the future. Agents are therefore more likely to be overly risk-averse because they lack the ability to assess risks, which makes them prone to erratic choices and herding behaviour. As a result, a system may experience increasing volatility in times of panic.
- The financial system's dependence on information itself creates a new target for malicious threat actors. As part of a malicious cyber incident, a threat actor may attempt to control a target's information, misdirecting their situational awareness or decision-making processes⁵⁷. Thus, the threat actor may lead its target to make the decision they want them to make and/or to seed confusion. The breadth, complexity and pace of improvements in technology indicates that the challenges posed by such situations may overwhelm decision-makers. It is expected that future malicious cyber incidents will increasingly incorporate the introduction of false information, target individuals with regard to information degradation, or seek to specifically corrupt the information that reaches decision-makers. These deliberate acts will be coordinated in campaigns in order to maximise uncertainty and maximally exploit cyber successes.

Box 3

Large-scale watering hole campaign (February 2017)

In early February 2017, financial institutions (primarily based in Poland) began to report unusual activity on their user workstations, including network traffic to foreign locations and malware installation. Subsequent investigations led to suspicions that websites had been compromised so that visitors using the IP addresses of targeted financial organisations were redirected and infected (a technique known as a **watering hole**).⁵⁸

In a campaign that ultimately spread to victim organisations in at least 31 countries, uncertainty exacerbated the issue in two distinct ways:

- uncertainty over which third parties could be considered “safe” and the extent to which the issue had spread. The compromising of legitimate websites (including those belonging to

⁵⁷ Styzt, M.R. and Banks S.B. (2014), “Cyber Warfare Simulation to Prepare to Control Cyber Space,” National Cybersecurity Institute Journal, Vol. 1, No 2, pp. 9-25.

⁵⁸ Trend Micro (2017), **RATANKBA: Delving into Large-scale Watering Holes against Enterprises**, 27 February.



national authorities) known to be popular with the targeted visitors undermined confidence in pre-existing relationships of trust.

- uncertainty over attribution. The subsequent forensic analysis concluded that those responsible had deliberately inserted Russian words into their malware in order to confuse investigators and obfuscate their origin.

Erosion

In Section 3.1, the financial system is referred to as a CAS. The system characteristics described have so far related to concepts of complexity. However, the adaptive qualities of the financial system also have susceptibilities that could lead to a system becoming more vulnerable.

CASs are dynamic – they change over time. They actively anticipate and react to the consequences of certain responses based on history and feedback, even if they do not materialise in reality. These systems show emergent aggregate behaviour that is not simply derived from the sum of its parts, and cannot be predicted from their original state, i.e. it is non-deterministic.

A system “remembers” through the persistence of internal structure. The memory of a financial system exists in various locations, such as in business plans, the experience of individuals, or even regulation. If the future state of a system is in part based on its past states, adverse feedback loops could feed a progressive build-up of fragility and aggregate risks. These risks are not directly attributable to the activities of a single institution, but derive instead from collective behaviour, which leads to an amplification of volatility in the financial sector and in the real economy, thus leading to **procyclicality**⁵⁹. The build-up of risks over time that are hidden and under-priced weakens a system, making it increasingly susceptible to failure.

Adverse conditions

Complex systems exist in varying states of order, from equilibrium to the “edge of chaos”⁶⁰. Systems that reach this latter threshold are particularly prone to sudden, nonlinear transitions from one state to another. Such critical transitions can be the result of either external perturbations (such as a cyber incident) or the endogenous functioning of a system itself. They are both difficult to forecast and potentially irreversible.

The presence of adverse conditions places a system under greater stress and a degree of instability, which could amplify an external shock. Examples of such conditions for the financial system include, but are not restricted to:

- extreme market volatility;

⁵⁹ Caruana, J. (2010), *Systemic risk: how to deal with it?*, BIS.

⁶⁰ Waldrop, M. (1992), *Complexity: The Emerging Science at the Edge of Order and Chaos*, Simon and Schuster, New York.



- asset bubbles;
- high interest rates;
- excessive off-balance sheet debt;
- political instability;
- low profitability / recession;
- trade tariff wars.


Detection time

The relationship between incident detection rates and the impact experienced is not a new concept. In any operational setting, the longer an incident goes undetected, the greater the consequences. The same is true for cyber incidents, whether malicious or not.

In their 2018 global study of data breaches, IBM and the Ponemon Institute's conclusions supported this well-founded link between incident detection and financial implications⁶¹. The study involved almost 500 diverse contributors, and illustrated the link between an organisation's mean time to identify (MTTI) an incident, and the average total cost of the data breach (see Figure 3.8). For the finance sector, the average MTTI figure was reported as **163 days to detect that a data breach had occurred**. This average increased significantly for deliberate acts in comparison with accidental circumstances.

Figure 3.8
Relationship between mean time to identify (MTTI) and average total cost

(USD millions)



Source: Reproduced from IBM/Ponemon's 2018 Cost of Data Breach Survey.

⁶¹ IBM and Ponemon Institute (2018), *2018 Cost of a Data Breach Study: Global Overview*.



Unlike human error, malicious cyber incidents can be purposefully orchestrated so that they stay under the detection threshold across all security layers (from the network to the human layer). Instead of being “loud and proud”, an increasing proportion of threat actors are adopting a more sophisticated **low-and-slow** approach, as shown in the Capital One example overleaf. This tactic involves remaining invisible for as long as possible, while stealthily moving between compromised hosts without generating regular or predictable network traffic patterns. This approach aims to bypass traditional security tools because each of the individual actions that make up the larger threat is too small to detect. These campaigns are designed to operate over a longer period of time, and by minimising disruption to any data transfer or connectivity levels, they blend into legitimate data traffic. Equally, one of the threat actor's goals may be to maintain long-term (persistent) access, in contrast to threats which only need access to execute a specific task.

The ramifications of this type of stealthy activity could be significant both in cases of extensive data exfiltration or in subtle integrity compromises that remain undetected for months or even years. The scale of the shock to a system at the point of discovery, followed by uncertainty over the extent of assets affected or knowledge of the last known uncompromised position, contributes to the overall amplification effect.

Box 4 Capital One data breach (July 2019)

In March and April of 2019, Capital One experienced a multi-stage data breach, exposing the personal information of nearly 106 million individuals in the United States and Canada. The data accessed related to credit card applications from 2005 to early 2019 for consumers, individual applicants and small businesses, and included names, addresses, dates of birth, credit scores, transaction data, social security numbers and linked bank account numbers.⁶²

Capital One became aware of the problem on 17 July (four months after the initial intrusion), when the organisation received an anonymous responsible disclosure email alerting them to the presence of sensitive Capital One data on a public forum. Although the incident was quickly addressed once the company had been notified, it highlighted multiple lapses that had led to the successful breach. These would have gone unnoticed if not for the tip-off.

Early detection of the unauthorised access or exfiltration might have limited the scale of, or altogether prevented, the intrusion. Instead, Capital One faced the amplifying effect of a prolonged and undetected breach, momentarily undermining confidence in the institution. In addition, although the impact on customers appears to be contained, the full extent of the consequences for Capital One from a legal or regulatory perspective is yet to be determined.

⁶² Capital One (2019), **2019 Capital One Cyber Incident**, 19 July.



Ineffective response

The financial sector is increasingly confronted by a number of cyber incidents which have the potential to spread beyond national borders and create significant economic knock-on effects. It is therefore incumbent upon the sector to manage these disruptive events effectively, as trust in a system is directly affected by how swiftly and how efficiently affected stakeholders react in a crisis situation.

The complexities of modern crises often require the involvement of many agents, and this demands effective coordination if a successful outcome is to be achieved. The need for coordination also raises significant multi-agent governance challenges, especially when spanning an institutional-level response with that of the sector as a whole, national authorities, other jurisdictions and non-financial sectors. The capacity to coordinate crisis management is a fundamental element of good governance, as it tests a system's capacity to provide the appropriate responses at the right time, in order to protect the real economy and mitigate the impact of disasters.

A 2006 study carried out by Donahue and Tuohy explored after-action findings from recent disasters to identify where deficiencies in collective response had worsened the outcome⁶³. These included:

- **poor communications** – unwillingness to agree specifications or commit resources to shared systems, or to share information in a comprehensive manner, due to a lack of trust between responders.
- **uncoordinated leadership** – unclear, multiple, conflicting, uncooperative and isolated command structures were cited as major problems.
- **weak planning** – multi-party commitment to plans that might be watered down to permit compromise.
- **resource constraints** – in a cost-pressured environment, competition for key resources is tough. In a major incident, this becomes much tougher and resources become even more scarce. As a result they exert extra pressure at the time they are needed most.
- **a lack of training and practice to embed lessons** – there was insufficient training or practice to implement plans. As a result, responders did not work well together or understand each other and their needs in an emergency.

For example, in the wake of the 2008 financial crisis, European authorities issued a paper highlighting *“the EU’s lack of effective crisis management for cross-border financial institutions.”*⁶⁴. Although in the context of resolution, the paper called on all competent authorities to *“effectively coordinate their actions and have the appropriate tools for intervening quickly to manage failure”* and for existing reforms to be *“complemented by a clear framework that will, in future, enable*

⁶³ Donahue, A.K. and Tuohy, R.V. (2006), *Lessons We Don’t Learn: A Study of the Lessons of Disasters, Why We Repeat Them, and How We Can Learn Them*.

⁶⁴ Commission of the European Communities (2009), *An EU Framework for Cross-Border Crisis Management in the Banking Sector*.



authorities to stabilise and control the systemic impact of failing cross-border financial institutions.” When transposed to a cyber incident that could potentially have systemic consequences, the lack of effective coordination and/or tools among respondents could become a factor contributing to shock amplification, as seen during the Equifax data breach example below.

Box 5

Equifax data breach (2017)

The official congressional report into the 2017 Equifax data breach provides a detailed end-to-end account of this incident, including aspects of the firm's incident response which may have aggravated the situation⁶⁵. In particular, the following actions have received criticism.

- After discovering the breach on 29 July, Equifax took a further six weeks to notify the public on 7 September. The firm needed this time to develop a dedicated site and call-handling facility to deal with the high volumes of inquiries, with a total of 148 million consumers affected. It is not evident that the firm had prepared for or rehearsed for an incident of this magnitude, and internal resources came under significant pressure to create new incident-handling channels in very challenging conditions.
- As part of its response, Equifax made the decision to direct its customer base to a dedicated portal (www.equifaxsecurity2017.com) which was separate from its main website. Not only did this link appear suspicious, or at least confusing, but the firm's social media team redirected customers to an incorrect address (www.securityequifax2017.com) over a period of two weeks. The incorrect address was subsequently registered by a security researcher to highlight the fact that over 200,000 individuals had navigated to what might otherwise have been a phishing website, thereby underlining the dangers of the company diverging from its trusted equifax.com domain.
- The firm's leadership during the incident was called into question over the decisions they had taken. These included the CIO's two-week vacation in August in the midst of the crisis response; withholding the true identity of the affected party (Equifax) from the response team developing the consumer-facing portal; the planned but accelerated “retirement” of its CIO, its CSO and its CEO; and the termination of a senior IT executive for failing to forward a patching alert that had been received by 430 other employees.

This incident is now viewed not only as one of the largest corporate breaches in history, but also as a textbook example of how a lack of preparedness and the failure to handle an incident effectively can make matters worse.

⁶⁵ US House of Representatives – Committee on Oversight and Government Reform, [The Equifax Data Breach – Majority Staff Report](#).



3.5 Cyber-specific amplifiers

Persistence

Section 1.2 describes intent as one of the factors to be considered when examining cyber threats through the lens of deliberate acts of human origin. There are varying degrees of intent, e.g. acting negligently, recklessly, knowingly or purposefully. At the upper end of the intent spectrum, **persistence** encapsulates a commitment to higher aspirations or objectives, as well as acting in a determined manner, even in the face of difficulty or opposition. Cyber incidents which demonstrate this characteristic are clearly differentiated from acts which are opportunistic, indiscriminate, speculative or isolated in nature.

The presence of persistence infers an amplifying effect, where a threat actor seeks to achieve their aims using any means available. This characteristic embodies the intelligent and adaptive nature of cyber threat presented by a human adversary, through specific and prolonged targeting, customisation, and by seeking to maximise the opportunities of the intrusion. Threat actors with this type of intent will relentlessly pursue their motives, be these disruptive or for financial gain, while minimising the level of interaction required from them to execute their objectives and evade detection (see APT1 example overleaf).

Prevention is no match for persistent cyber threats. Countermeasures based on previously identified tactics, techniques or procedures may help to block – but will not necessarily stop – persistent adversaries as they evolve the approaches they use to seek out exploitable weaknesses. No organisation is immune to this cyber threat – an idea reinforced by McAfee researcher Dmitri Alperovitch in 2011, who claimed that there are only two kinds of organisations: “*those that know they’ve been compromised, and those that don’t know yet*”.⁶⁶

Box 6 Mandiant APT1 report (Feb 2013)

In February 2013, the US cybersecurity firm Mandiant (since acquired by FireEye) released an extensive report [60] into the activities of a threat actor known as “APT1”. With 141 known victim organisations spanning 20 sectors (see Figure 3.9), APT1 is believed to have conducted one of the most far-reaching cyber espionage campaigns in terms of the sheer quantity of information stolen.

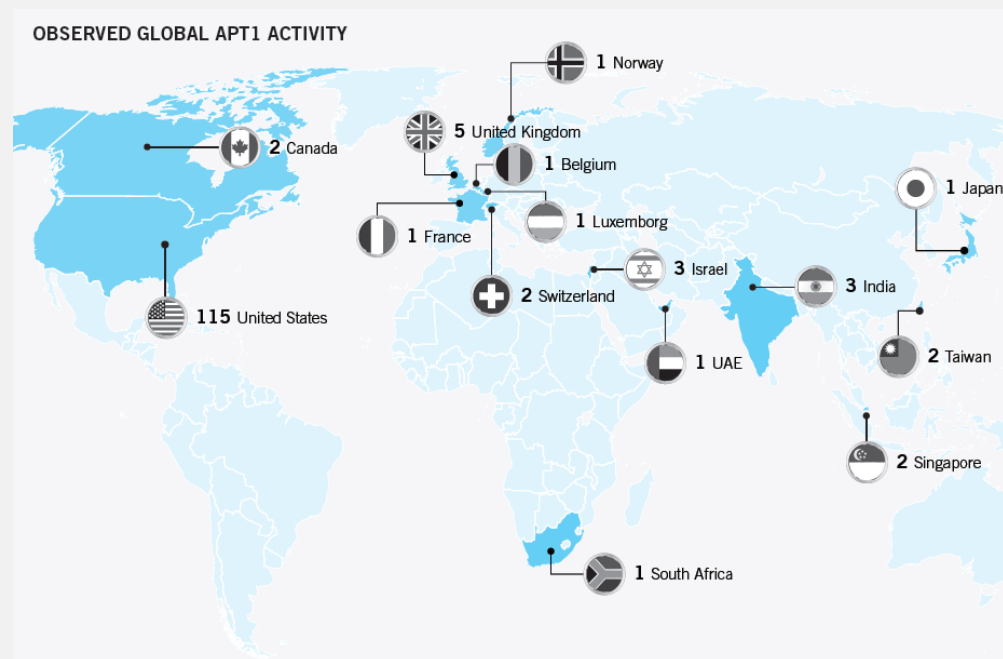
Over a seven year period from 2006, Mandiant’s analysis concluded that APT1 had systematically stolen hundreds of terabytes of data, although this may only have reflected a small fraction of the group’s activities. With well-defined and honed approaches, APT1 would periodically revisit victims’ networks to steal large volumes of intellectual property, including technology blueprints, proprietary manufacturing processes, test results, business plans, pricing documents, partnership agreements, and emails and contact lists from victim organisations’ leadership.

⁶⁶ Alperovitch, D. (2011), **Revealed: Operation Shady RAT**, August.



APT1 demonstrated persistence not only in the way they pursued their objectives, but also in their ability to dwell within victims' infrastructure to extract information over time. In one case, APT1 successfully maintained access to a victim's network for 1,764 days (four years and ten months), with an average dwell time of 356 days. In another case, the group was observed stealing 6.5 terabytes of compressed data from a single organisation over a period of ten months.

Figure 3.9
Geographical location of APT1's victims



Source: Reproduced from Mandiant APT1 Report⁶⁷.

Notes: In the case of victims with a multinational presence, the location shown reflects either the branch of the organisation that APT1 compromised (when known), or else the location of the organisation's headquarters.

Concurrency

Until this point, the model has only explored the shock arising from a single cyber incident. However, one of the unique characteristics of cyber risk is its potential for **concurrency** in forming a compound shock, typically only as a result of deliberate acts. This characteristic can be expressed in two distinct forms: (i) using multiple vectors or techniques as part of a coordinated attack on a single entity; or (ii) where (nearly) simultaneous, but apparently unrelated, disruption occurs in different parts of a system.

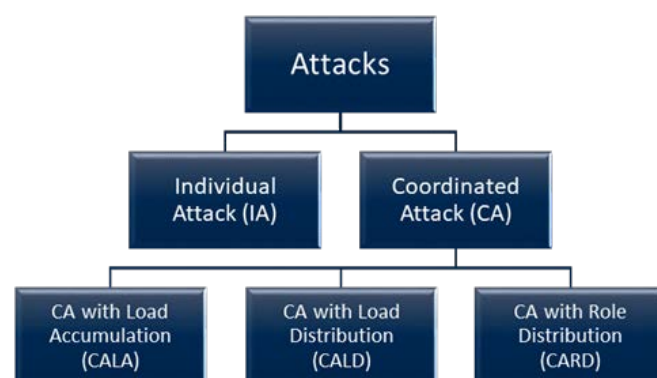
⁶⁷ Mandiant (acquired by FireEye) (2013), **APT1 – Exposing One of China's Cyber Espionage Units**, 19 February.

In the first instance, a coordinated attack is a multi-step scenario where the collaboration of several attacking sources is needed to achieve a common goal. This technique may be used to distract or overwhelm a target, so that the overall objective can be reached. The great danger of coordinated attacks is that they exceed the power of single attacking sources. Hence, coordinated actions can amplify shocks in a system in a way that would not be possible if they were performed individually.

Samarji [61] proposes a taxonomy which differentiates between different forms of concurrent malicious cyber activity (see Figure 3.10):

- **Individual Attack (IA)** – an elementary action executed by a single attacking source;
- **Coordinated Attack (CA)** – an action consisting of joint individual actions executed by several collaborating attackers;
- **CA with Load Accumulation (CALA)** – attackers accumulate their capacities, executing an action in a distributed and simultaneous manner (e.g. Distributed Denial of Service);
- **CA with Load Distribution (CALD)** – a shareable attack accomplished by a group of attackers (e.g. a coordinated password cracking attack);
- **CA with Role Distribution (CARD)** – a multi-task action in which attacking sources distribute roles or tasks to accomplish their goal.

Figure 3.10
Coordinated cyber attack taxonomy



Source: Reproduced from Samarji, Léa El⁶⁸.

Box 7 APT10 Campaign (2016-17)

One of the most prolific examples of a sustained CARD-based campaign is the multi-year espionage operation conducted by the threat actor known as “APT10”. Joint research by PwC and

⁶⁸ Samarji, L.E. (2015), “Risk-aware Decision Support System to Counter Coordinated and Simultaneous Attacks”, *Télécom Bretagne*, Université de Rennes.




BAE Systems has highlighted the ambitious scale of APT10's activities seeking to compromise multiple managed IT service providers (MSPs) in 2016-17⁶⁹.

APT10's objectives were believed to centre on intellectual property theft and information collection on a massive scale that is closely aligned to strategic national interests. MSPs represent ideal high payoff targets for espionage-focused threat actors, given their direct and unfettered access to client networks.

The scale of APT10's activities is unprecedented, affecting nine different sectors in 15 jurisdictions across five continents. The number of individual operations number in the thousands, carried out by multiple teams, each responsible for different aspects of day-to-day operations, during 08:00 to 19:00 business hours (see Figure 3.11). Threat actors such as APT10 demonstrate an ability to trigger concurrent and targeted cyber incidents with global and cross-sectoral reach.

Figure 3.11
Operational times for APT10 in UTC+8



Source: Reproduced from PwC/BAE Systems – Operation Cloud Hopper.

Notes: The time zone for this data set is UTC+8, i.e. 8 hours ahead of Coordinated Universal Time (UTC).

⁶⁹ PwC / BAE Systems (2017), *Operation Cloud Hopper – Exposing a systematic hacking operation with an unprecedented web of global victims*.

Pervasiveness

The concept of **pervasiveness** as an amplifier, i.e. existing in or spreading through every part, is intrinsically linked to the property of “scaling”, which can differentiate cyber incidents from other types of operational events. As the pervasiveness of cyberspace grows, so does the potential for asymmetric disturbances to occur within that medium.

In 1997, the economist and journalist Frances Cairncross coined the term **death of distance** to describe how telecoms, the internet and wireless technology were overcoming geography as a barrier to communication, creating an environment in which every actor in cyberspace theoretically has global reach⁷⁰. In parallel, the fourth industrial revolution has driven a surge in ubiquitous computing, fuelling demand for always-on, always-connected everyday devices with some form of processing capability. Coupled with an internet which was designed based on principles of openness and accessibility, this digital explosion favours those of malicious intent as it increases the volume of endpoints that could be exposed to threats and the likelihood of their exploitation.

In this cyber ecosystem, it is possible for relatively small actors to have a disproportionate effect, whether they intend to or not – mass is no longer important. In an adversarial context, the decreasing barrier to entry around the globe can be used as an effective force multiplier, where waging cyber campaigns is significantly more simplistic than defending them. Pervasive cyber incidents are typified by their rapid spread, affecting numerous victims in a very short period of time, as the NotPetya example below illustrates. Their seemingly random dispersion makes these incidents difficult to control, often overwhelming response teams whose bandwidths are strained as they try to put out many fires at once. Consequently, the pervasiveness of a cyber incident can have an amplifying effect which transcends geographical or sectoral boundaries, thereby increasing the complexity of response coordination and its ability to collectively maintain confidence.

Box 8 NotPetya (June 2017)

The hyper-connected nature of cyberspace provides a pervasive medium within which the financial system operates. However, when this medium is compromised through a combination of speed and destructive intent the consequences can be devastating, as seen during the NotPetya incident⁷¹.

Within hours of its first appearance, a destructive worm raced beyond Ukraine to cripple multinational companies across the globe, causing **damage in excess of USD 10 billion**. The maritime giant, Maersk, which accounts for one-fifth of the world’s shipping capacity, had to rebuild its entire network of 4,000 servers and 45,000 PCs from the ground up. Merck, whose ability to manufacture a number of drugs was temporarily halted, told shareholders it had lost USD 870 million due to the malware. French construction giant Saint-Gobain lost around USD 400 million, while Mondelēz, the owner of chocolate-maker Cadbury, took a USD 188 million hit. FedEx’s European subsidiary TNT Express was also crippled, requiring months to recover some data. The

⁷⁰ Cairncross, F. (1997), *The Death of Distance: How the Communications Revolution Will Change Our Lives*, Harvard Business School Press.

⁷¹ Greenberg, A. (2018), “The Untold Story of NotPetya, the Most Devastating Cyberattack in History”, WIRED, 8 August.



TNT disruption also affected its ability to operate the CREST Courier and Sorting Service for the deposit, withdrawal and delivery of physical securities.

On a national scale, NotPetya disrupted many parts of Ukraine's infrastructure, affecting four hospitals in Kiev alone, six power companies, two airports, more than 22 Ukrainian banks, ATMs and card payment systems in retail and transport, and practically every federal agency. The speed at which NotPetya crippled its victims' systems was also unprecedented, bringing down the network of a large bank in 45 seconds, and with a part of one major transit hub being fully infected in just 16 seconds. From a national perspective, it might even be reasonable to describe the damage done to the Ukrainian economy, where 10% of all computers in the country were wiped, as a systemic event. On a global scale, however, the incident demonstrates how a pervasive cyber incident can spread through an operational network to affect entities far beyond Ukrainian ground zero.

Speed of transmission

In a 2009 speech on systemic risk, the former president of the ECB, Jean-Claude Trichet, spoke of speed being a *"key element in phases of turmoil and crisis"*⁷². Trichet noted that the speed at which shocks are transmitted through a system had accelerated tremendously in recent decades. Many factors contribute to this, including:

- the process of global financial integration;
- increasing leverage in institutions;
- the accumulation over a long period of time of unsustainable global imbalances;
- the technological advancements that facilitate the instantaneous transmission of information worldwide.

Since the financial crisis, there has been a proliferation of technological capabilities and processes, with increasing levels of cost effectiveness and speed. Consumer and business demand for increased convenience and speed have driven the digitalisation of financial services. However, the increasing speed of operations can, conversely, work against an entity when disruption occurs. An entity's ability to detect and react to incidents could become compressed to the point at which it is unable to respond quickly enough, regardless of how effective its decision-making processes are. Consequently, its ability to contain or absorb the impact of the disruption may be undermined, potentially resulting in the fast and widespread propagation of a shock.

⁷² Trichet, J.-C. (2009), *Text of the Clare Distinguished Lecture in Economics and Public Policy*, Clare College, University of Cambridge.



Box 9

Knight Capital (August 2012)

The Knight Capital incident highlights how the speed of high frequency trading (HFT) can combine disastrously with human error and poor internal processes. As a result of incorrect code deployment, the financial market-maker executed over four million trades in 154 stocks over a 45-minute period, causing the **firm to lose over USD 460 million** from these unwanted positions. The market handled the failure through a relatively organised rescue, although there were wider systemic consequences. A necessary condition for HFT to increase systemic risk is that it creates channels for adverse feedback loops, whereby relatively small events like a micro crisis or the sudden disappearance of liquidity can be amplified into a systemic event. The IOSCO Technical Committee highlighted the issue of algorithms operating across markets, with the potential to transmit shocks rapidly from one market to another, thus amplifying systemic risk⁷³.

3.6 Alignment of amplifiers

In 1986, the ecologist Buzz Holling introduced the **adaptive cycle model**, a useful metaphor for the long-term dynamics of change within CASs. The model is represented by four stages in a Möbius strip diagram: growth (r), conservation (K), creative destruction (Ω) and reorganisation (α). In 2001, this concept was augmented to account for the interconnections between systems which exist at different scales of space, time and social organisation⁷⁴. This interacting and nested set of hierarchical structures, known as a **panarchy**, makes it possible to understand the relative contributions of individual cycles and identify points where they may be vulnerable. In applying panarchy thinking to the financial system, the nested hierarchies range from its component parts at an organisational asset level, right through to it being a subsystem of an overall economic system and human society.

In Figure 3.12, three selected levels of a panarchy are illustrated to show the two connections that are critical to creating and sustaining adaptive capability. One is the *revolt* connection, which can cause a critical change in one cycle to cascade up to a vulnerable stage in a larger and slower cycle. The other is the *remember* connection, which facilitates renewal by drawing on the potential that has been accumulated and stored in a larger, slower cycle.

When a shock occurs that initiates *creative destruction* and the other levels have accumulated vulnerabilities and rigidities within their adaptive cycle, the crisis can spread across a system towards a **panarchical collapse**, in a fashion similar to Lorenz's *butterfly effect*⁷⁵. Extremely large events can overwhelm the sustaining properties of panarchies, destroying levels altogether and triggering destructive cascades within the panarchy. Holling defined these circumstances

⁷³ Technical Committee of the International Organization of Securities Commissions (2011), *Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency – Final Report*.

⁷⁴ Holling, C.S. (2001) "Understanding the Complexity of Economics, Ecological, and Social Systems", *Ecosystems*, Vol. 4, No 5, pp. 390-405.


⁷⁵ Lorenz, E. (1972), "Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?", in 139th meeting of the American Association for the Advancement of Science.



symbolically as an “*alignment of the stars*”, where hazards and vulnerabilities interact at different scales. Unaligned vulnerabilities would minimise the propagation of impacts, whereas aligned vulnerabilities would facilitate it.

Figure 3.12

Panarchy with nested adaptive cycles at multiple scales



Source: Adapted from Davoudi, Simin⁷⁶.


In Reason’s Swiss Cheese model, a system’s defences against failure are modelled as a series of barriers, represented as slices of cheese (see Figure 3.13)⁷⁷. The holes in the slices represent vulnerabilities within a system and are continually varying in size and position across the slices. When disruption occurs from a cyber incident, the alignment of holes in each slice permits “a trajectory of amplification” in such a way that the rate of contagion following an initial shock augments with each hole. The effect of each system vulnerability alters the rate at which impacts accumulate within a system, typically in a non-linear fashion. If sufficient vulnerabilities are present or emerge as a result of further contagion, the combination could lead to a “perfect storm” or systemic event.

⁷⁶ Davoudi, S. (2012), “Resilience: A Bridging Concept or a Dead End?”, Planning Theory & Practice, Vol. 13, No 2, pp. 299-307.

⁷⁷ Reason, J. (1990), “The Contribution of Latent Human Failures to the Breakdown of Complex Systems”, Philosophical Transactions of the Royal Society of London, Vol. 327, No 1241.

Figure 3.13

Contagion amplification (using the Swiss Cheese model)



Source: Adapted from the original James Reason model.

3.7 Contagion channels

In Section 2.3 the localised consequences of a cyber incident are expressed in the form of business impacts (first order impact). As these impacts extend beyond the bounds of the affected institution(s) at the starting point, the transmission and amplification of an impact is described as **contagion** (second order impact). In the context of financial stability, the FSB defines contagion as follows:

“Distress experienced by a single financial institution or sector can be transmitted to other institutions or sectors – owing either to direct exposures between them, or commonalities that lead to a general loss of confidence in those institutions or sectors.”⁷⁸

For the purposes of the model, this concept of contagion is extended to cover amplification and/or the transformation of shocks between or within entities. Contagion types have been broadly classified into three distinct **channels** – operational, confidence and financial – that can interact with each other concurrently and on a many-to-many basis. Aside from the individual characteristics of each contagion channel, which the paper will describe in further detail, the relationship between channels is captured in a **3x3 contagion grid** in Figure 3.14.

⁷⁸ Financial Stability Board (2017), *Financial Stability Implications from FinTech: Supervisory and Regulatory Issues the Merit Authorities' Attention*.

Figure 3.14

Transmission within and between contagion channels (with examples)

	Operational	Confidence	Financial
Operational	O►O : operational disruption in one part of a system results in operational disruption in other parts	O►C : operational disruption results in a loss of confidence or trust in the affected entity, the services it provides (affected or not), or the financial system as a whole	O►F : operational disruption results in failure to perform financial activity in line with stakeholder needs or requirements leading to financial impacts
Confidence	C►O : a loss in confidence results in operations being unintentionally or intentionally degraded or ceased	C►C : a loss of confidence in one part of a system results in loss of confidence in other parts	C►F : a loss of confidence leads to a change in financial activity with negative consequences
Financial	F►O : financial impacts result in operations being unintentionally or intentionally degraded or ceased	F►C : financial impacts result in a loss of confidence or trust in the affected entity, the services it provides (affected or not), or the financial system as a whole	F►F : financial disruption in one part of a system results in financial disruption in other parts

Source: Developed by the ESRB's ESCG.

Operational channel

The operational contagion channel encapsulates the cascade effect of disruption as it arises from, or traverses through, the operational layers of the financial system. Operational contagion can thus be seen as an operational disruption in one part of a system resulting in operational disruption in other parts. To understand operational transmission in a systemic context, it is useful to combine complexity and systems theory with the concepts found in **natural accident theory (NAT)**. NAT serves as a useful lens because it addresses the widespread impacts of single defects by examining two properties:

- **tight coupling** – the speed at which actions propagate in systems that lack slack or buffer between components;
- **interactive complexity** – the extent to which components within a system can interact, especially in unfamiliar, unplanned or unexpected ways.

Perrow argues that systems that are both tightly coupled and highly complex may be more susceptible to system accidents⁷⁹. High complexity in a system means that when disruption occurs, it takes time to work out what has happened and to act appropriately. Tight coupling means that this time is not available. In essence, a system's structure and complexity can impede the detection and correction of failures, while tight coupling may cause a failure to spread faster.

When these two properties are considered in the context of cyber security, information systems are purposefully designed to promote efficiency by being both tightly coupled and interactively complex.

⁷⁹ Perrow, C. (1999), *Normal Accidents: Living with High-Risk Technologies*, Princeton University Press.



As a result, NAT would conclude that cyber incidents are both inevitable and unpredictable. This view is echoed by Zurich, who states that *“the increasingly tight coupling of the internet with the real economy and society means a full-scale cyber shock is far more likely to occur.”*⁸⁰

Confidence channel

Confidence can be described as an emotion of assured expectation⁸¹. In turn, **trust** is a particular form of confidence, forming the emotional basis of cooperation. In financial services, trust can be expressed in two forms⁸²:

- trust between participants in a financial transaction that the parties will honour their side of the agreement, even if it means one or more of the parties incurring unexpected losses;
- trust by the population at large that the financial sector is focusing on its core role of efficiently bringing savers and investors together in ways that optimise the allocation of private savings to financial investment in physical and human capital.

Trust builds up via repeated satisfactory interactions between entities, and morphs over time into a positive reputation, allowing a new party to confidently expect that they can trust the other party without having to perform repeated transactions. Conversely, according to the sage advice of Warren Buffet, *“It takes twenty years to build up a reputation, and five minutes to ruin it.”*

A sufficient level of trust, especially systemic or institutional trust, plays a crucial role in the stability and maintenance of the economic and financial system. For example, market liquidity is critically dependent on confidence in the security and reliability of clearing and settlement arrangements for funds and financial instruments. In previous financial crises, containment policies were implemented with the goal of restoring public trust, in order to minimise the repercussions in the real sector of loss of confidence by depositors and other investors in the financial system.

Uncertainty and confidence are linked by an intrinsically inverse relationship. Therefore, an increase in uncertainty during periods of disruption amplifies impacts via the confidence contagion channel. In a cyber incident, confidence can be quickly eroded or destroyed if there is a lack of certainty over the circumstances which caused the disruption, the actual effect of the disruption, or the possibility of containing or minimising the impact. The destruction of that trust is not necessarily tied to the digital world, can then spill over into a loss of confidence in the delivery of financial services.

⁸⁰ Zurich Insurance Company and Atlantic Council (2014), *Risk Nexus – Beyond data breaches: global interconnections of cyber risk*.

⁸¹ Barbalet, J.M. (1996), “Social Emotions: Confidence, Trust and Loyalty”, *International Journal of Sociology and Social Policy*, Vol. 16, No 9, pp. 75-96.

⁸² Vanston, N. (2012), *Trust and reputation in financial services*, Government Office for Science.



Financial channel

There is extensive macroeconomic literature on the various forms of financial contagion that a system can become exposed to. This paper will therefore not attempt to provide detailed insight into this topic, but will instead highlight some of the ways that financial contagion may arise within the financial system.

- **Interbank credit shock.** Credit losses are related to counterparty risk, and are incurred by lender banks when their borrower banks default and fail to fulfil their obligations. These losses can then lead to the default of lenders, resulting in another wave of credit shocks. In a seminal paper, Eisenberg and Noe developed an analytical framework to determine the set of payments that clear the network following an initial shock, assessing in this way how losses propagate through a system⁸³. More generally, due to the interconnectedness of the financial system, losses in one financial institution can quickly lead to losses for other financial institutions.
- **Market liquidity shock.** Losses may arise in the value of a financial institution's external assets, caused by a generalised fall in market prices, a rise in expected defaults or the “fire sale” actions of a failing financial institution. Such market liquidity shocks are conventionally and sensibly represented by discount factors which, for a given asset class, are proportional to the number of failing financial institution holding the asset.
- **Funding liquidity shock.** Financial institutions that are subject to regulatory solvency constraints mark their assets to market in order to replace lost funding. If the interbank market is experiencing distress and shrinks, banks short of liquidity may be unable to borrow all the money they need from the market, and may be forced to sell their illiquid assets. In the presence of fire sales, market demand for illiquid assets becomes inelastic, depressing the market prices of these assets and resulting in effective losses for banks. These fire sales spill-overs create an incentive to hoard liquidity, which can in turn induce another wave of sales, activating a liquidity spiral that could cause the interbank market to freeze completely.

⁸³ Eisenberg, L. and Noe, T. (2001), “Systemic Risk in Financial Systems”, *Management Science*, Vol. 47, No 2, pp. 236-249.



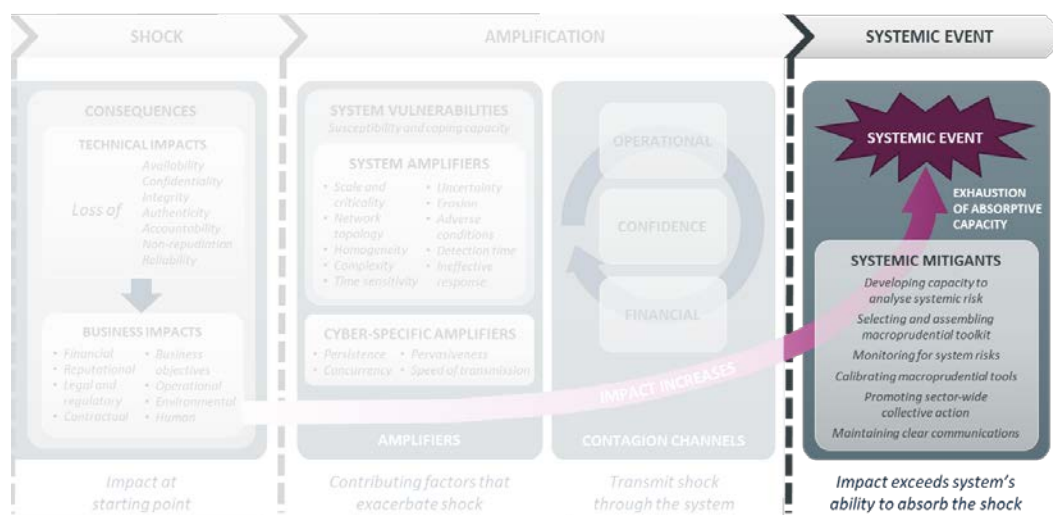
4 Systemic event

The point at which a shock leads to a systemic event

Building on the FSB's definition of systemic risk, the model considers systemic cyber risk as the risk of disruption to financial services that is⁸⁴:

1. caused by an impairment of all or part of the financial system following a cyber incident,
2. has the potential to result in serious negative consequences for the real economy.

While a serious cyber incident may cause system-wide disruption, this need not lead to a systemic event. However, a systemic event could occur when the system no longer has the capacity to absorb the shock and recover so that it can continue to provide key economic functions.



4.1 Impact thresholds

In Section 2 the model describes the degree of localised disruption experienced, based on a set of coexistent and often interdependent impacts. As disruption spreads through a system, other nodes exhibit impact, leading to an overhaul accumulation. At a systemic level, the overall impact experienced can be described as the **aggregate impact**. However, this is not an additive process, as it would be incorrect to express a system view based on a summation of the individual behaviour of constituent elements.

⁸⁴ IMF/BIS/FSB (2009), *Guidance to assess the systemic importance of financial institutions, markets and instruments: Initial considerations*, Report to G20 Finance Ministers and Governors.

Within each system, the point at which the aggregate impact becomes too great to bear, and the impairment of key economic functions has the potential to significantly affect the real economy, can be referred to as the **impact tolerance threshold**.


Figure 4.1A describes the concept of impact tolerance and related principles⁸⁵. A system is portrayed as facing a continuous flow of disruptive events, though virtually all of these have negligible consequences at a systemic level. From a macroprudential perspective, it would not be desirable to experience impacts within a system close to the upper impact tolerance threshold.

Instead, a lower-bound threshold is necessary to define the parameters according to which the individual entities in a system, the services they provide, and the economic functions they support, should operate. This may be referred to as an “impact appetite” or **requirements threshold**.

Upper thresholds can differ between individual systems, whereas lower thresholds may be set in line with the remit of the individual authorities that oversee the safe running of each system. Both thresholds are expressed in terms of impacts, described in either qualitative or quantitative terms.

Figure 4.1

The charts below illustrate the concept of impact tolerance and absorptive capacity (a), a shock being absorbed (b), and disruptions with differing rates of impact amplification (c)



Source: Inspired by the Bank of England Consultation Paper 29/19.

4.2 Absorptive capacity

Figure 4.1a shows the gap or buffer between these two thresholds – this represents the coping capacity within a system to absorb shocks. This **absorptive capacity** provides a degree of mitigation against systemic events. The requirements threshold is therefore a device which conveys **resilience requirements** to the constituent parts of a system (in the form of operating limits to be followed) which do not deplete the systemic buffer. However, there is a trade-off to consider when setting requirements thresholds. An increase in absorptive capacity through a lowering of the threshold drives cost into a system due to more stringent operational requirements.

⁸⁵ Bank of England, Prudential Regulatory Authority; Financial Conduct Authority (2019), *Consultation Paper: Operational Resilience: Impact tolerances for important business services*.



Figure 4.1b depicts a cyber incident where aggregate impact exceeds the requirements threshold but does not breach the impact tolerance threshold. This example highlights a situation where one or more nodes in a system have been unable to operate within their requirements but sufficient absorptive capacity is present to avoid impact tolerance being breached.

Figure 4.1c overlays two independent disruptions, each with different occurrences of system vulnerabilities affecting the outcome. In both cases, the difference in amplifying effect (shown as α) is reflected in the steepness of each impact curve. As CASs are governed by non-linear equations, the rate of impact will tend towards exponential growth.

However, absorptive capacity is not only a function of its constituent parts. Additional measures can be established at a systemic level to improve a system's ability to absorb shocks.

Box 10

Impact tolerance and absorptive capacity: additional considerations

1. Using existing impact analysis methodology in a system context

When considering impact tolerance from a domestic macroprudential perspective, the upper threshold defines the point at which **financial instability** occurs, i.e. failure to provide a consistent supply of critical (economic) functions that the real economy demands from the financial system.

As noted in Section 2.3, BIA has been used in an organisational context for several decades. By transposing these techniques to a system setting, the application of impact tolerance seeks to describe the levels of aggregate impact experienced when the provision of economic functions is disrupted. The upper threshold is similar to concepts such as Maximum Acceptable Outage (MAO) or Maximum Tolerable Period of Disruption (MTPoD) although, importantly, it is impact-based rather than time-based.

Given that it would be undesirable to operate a system close to its impact tolerance threshold, public authorities may seek to set their appetite for disruption, i.e. the requirements threshold, to levels which align with the needs and expectations of the real economy. Deciding on the lower limit (and therefore the degree of absorptive capacity within the system) is likely to involve an iterative process of finding an optimal balance between resilience and efficiency, as described in Figure 3.5.

The process resembles other macroprudential calibration functions performed on a periodic basis, such as setting the countercyclical buffer (CCyB), albeit in a context of building resilience to operational stress. Threshold tuning may be informed by the use of scenarios to estimate the build-up of aggregate impact when different economic functions are disrupted. The process also exposes the relative criticality of economic functions – in other words it puts the “C” into CEF. However, the challenge lies in expressing impact with sufficient accuracy, which may lend itself to qualitative measures at the outset, with the gradual introduction of proven quantitative indicators.

2. Introducing flexibility into resilience requirements

Section 3.1 describes the financial system as a CAS in a constant state of flux. It is therefore natural to conclude that the criticality of any given part of the system is both dynamic and time sensitive. However, prescribing rigid resilience requirements works against this principle. If



individual organisations contribute to different economic functions to differing extents, it is logical to surmise that their individual requirements must also differ.

A notable example which typifies this debate is the two-hour recovery time objective (RTO) designated for financial market infrastructure⁸⁶. However, the requirement's origin dates back to the *Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System (2003)* written following the terrorist attack on 11 September 2001⁸⁷. Even though the threat landscape and the nature of financial services have changed drastically in the intervening years, the resilience requirements for the most critical parts of the system remain unchanged.

Financial authorities using an impact-led approach no longer need to set rigid requirements directly. Instead, impact thresholds relative to the objectives of individual authorities can be established (whether driven by financial stability, or prudential or conduct goals) from which firms can determine the resilience requirements necessary to operate within those thresholds.

3. Linking system and individual organisational views on impact

Currently, financial institutions already perform business impact analyses for their own purposes, with varying degrees of proficiency. If firms are requested to additionally consider impact thresholds based on the role that they play relative to the system, behaviours driven by commercial incentives become **aligned with public interests**. It may not, however, be possible for firms to understand their individual contribution in isolation without recourse to the system-wide perspective the authorities can offer.

This macro-micro relationship and related concepts are explored further by PwC in guidance it has issued to its clients⁸⁸ on impact tolerance, although it should be noted that the central analogy relates only to a single organisation (a bridge) providing a single service (crossing a river).

4. Reinforcing links between risks faced and capability validation

All too often, the testing scenarios used to validate the presence and adequacy of response and recovery capabilities are not reflective of the risk environment within which firms operate. However, the scenarios used to calibrate impact thresholds can also inform **assurance activities** within firms. In a cyber context, a firm can use its situational awareness to drive the customisation of scenarios relative to its circumstances and, subsequently, demonstrate that it can keep within impact thresholds when those stresses are simulated.

There are different philosophies with regard to the severity of scenarios that should be incorporated (e.g. severe but plausible, extreme but existential, or all-hazards) but ultimately the selection is an individual firm's decision. However, the financial authorities may be in a position to review and challenge this selection, or to stretch firms through the use of cyber stress tests informed by the current sectoral threat landscape.

⁸⁶ CPMI-IOSCO (2012), **Principles for financial market infrastructure**, April.

⁸⁷ Board of Governors of the Federal Reserve System; Office of the Comptroller of the Currency; Securities and Exchange Commission (2003), **Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System**, 7 April.

⁸⁸ PricewaterhouseCoopers (2019), **Operational resilience: how to set and test impact tolerances**, October.



5. Challenging the availability-centric norm

Resilience requirements are the derived product of impact tolerance and, in the context of cyber security, describe the operating parameters through which the properties in Figure 2.1 are preserved. Of all the properties, availability is the simplest to comprehend, given its human connection with time. However, this can lead to an over-emphasis of availability-based objectives such as RTOs, to the detriment of others. In recent times, the **integrity and the veracity of information** have become increasingly prominent concerns, as these are both difficult to define and to maintain.

Meanwhile, RTOs themselves represent a double-edged sword – easy to convey, but overly simplistic. “*How long?*” may depend on a combination of factors for which a static answer lacks realism. A single figure does not allow for variance across time of day/week/month or point-in-time criticality. The use of worst-case scenarios is recommended from a business continuity planning and preparedness perspective, although this rarely reflects the complex dynamics encountered during incident response. There is ultimately an implementation trade-off between accuracy and effort and, to paraphrase Einstein, the concept should be *as simple as possible, but not simpler*.⁸⁹

4.3 Systemic mitigants

Macroprudential policy is defined by the IMF as *the use of primarily prudential tools to limit systemic risk*⁹⁰, by:

- increasing the resilience of the financial system to aggregate shocks;
- containing the build-up of systemic vulnerabilities over time;
- controlling structural vulnerabilities within the financial system.

In support of this objective, authorities need to select and assemble a set of macroprudential instruments that can help address the key potential sources and dimensions of systemic risk. For systemic cyber risk, the macroprudential toolkit is comprised of **systemic mitigants** that can help offset the effect of amplifiers. Under ideal conditions, these mitigants compensate, to the point that the aggregate impact is contained below the requirements threshold for a system. Based on IMF guidance, Figure 4.2 provides an example of how systemic mitigants could be arranged to form a macroprudential toolkit for systemic cyber risk – each of these mitigants is described further below.


⁸⁹ Calaprice, A. (2000), *The Expanded Quotable Einstein*, Princeton University Press, p. 314.

⁹⁰ International Monetary Fund (2013), *Key Aspects of Macroprudential Policy*.



Figure 4.2

Examples of systemic mitigants that could constitute a macroprudential toolkit for systemic cyber risk and the amplifiers they may offset



Source: Developed by the ESRB's ESCG.

Developing the capacity to analyse systemic risk

In order to facilitate the identification of system vulnerabilities, an understanding is required of the interconnectedness of a system, and the operational flows therein. When this is achieved jointly by public authorities and private entities, the sector as a whole benefits from the shared clarity with regard to a system and its component parts.

This common understanding helps to reduce opacity and uncertainty in the networks' operations, and provides an opportunity to reduce complexity through simplification. This outcome can be achieved through several complementary activities:

- **Structural mapping** – a functional and/or operational mapping of the interlinkages between nodes in a system. This can take a top-down approach, which describes a system from an economic function perspective, e.g. following the FSB's critical function analysis, becoming

steadily more granular⁹¹. Alternatively, a bottom-up approach can be adopted, based on the sector's collective knowledge.

- **Agent-based modelling (ABM)** – an analytical or computational abstraction of dynamic flows between system nodes. This technique allows for more realistic modelling of CASSs.
- **Time-series analysis** – a critical path analysis of end-to-end services to identify key times or events that represent periods of increased sensitivity to disruption.

Monitoring for system risks

Macroprudential authorities are typically responsible for **horizon scanning** to detect emerging risks to the financial system as a whole. This concept persists for cyber risk where a view of the threat landscape and key vulnerabilities affecting the financial sector can be established in collaboration with sector participants and relevant national competent authorities. **Collective vulnerability** information can be pooled across diverse sources and peers to better inform microprudential and macroprudential interventions. Where system vulnerabilities arise from key dependencies outside the financial sector, authorities may explore the implications for their **regulatory perimeters**⁹².

Promoting sector-wide collective action

Crisis response arrangements, both between domestic authorities and collectively with the private sector, should be in place to respond to an escalating cyber incident. This **collective incident response** capability can also be practised and validated through programmes of exercises, which can also aid vulnerability discovery. Domestic arrangements for **information sharing** between market participants are also critical both prior to and during cyber incidents. Authorities can also play a key role in shaping solutions led by the private sector to aid market-wide recovery, such as **data vaulting** and other contingent service arrangements⁹³.

Selecting and assembling a macroprudential toolkit

Stress testing emerged in the 1990s as a tool employed by financial institutions to assess their exposure to large risks⁹⁴. Stress testing is intended to test the resilience of an individual institution or an entire financial system to exogenous and endogenous shocks. **Cyber stress testing** adopts a similar approach by requiring firms to demonstrate their ability to operate within impact thresholds in the face of severe but plausible cyber incident scenarios. The use of cyber stress testing is still

⁹¹ Financial Stability Board (2013), *Recovery and Resolution Planning for Systemically Important Financial Institutions: Guidance on Identification of Critical Functions and Critical Shared Services*.

⁹² IMF/BIS/FSB (2016), *Elements of Effective Macroprudential Policies: Lessons from International Experience*.

⁹³ DTCC and Oliver Wyman (2018), *Large-Scale Cyber Attacks on the Financial System: A case for better coordinated response and recovery strategies*.

⁹⁴ Ron, A., Danielsson, J., Baba, C., Das, U., Kang, H. and Segoviano, M. (2018), *Macroprudential Stress Tests and Policies: Searching for Robust and Implementable Frameworks*, International Monetary Fund.



nascent, although jurisdictions such as the United Kingdom are actively developing their approach⁹⁵.

Where network analysis identifies potential fragilities (e.g. points of concentration), a **structural intervention** may be considered to reduce the potential amplifying effect. A recent example of systemic risk reduction as a result of historical network concentration is the multi-year programme to de-tier the United Kingdom's high value sterling payment system (CHAPS)⁹⁶. The operational risk reduction could equally be associated with a disruption arising from a cyber incident at a direct participant.

In a similar vein to network segmentation at an institutional level, the public authorities may wish to assess whether appropriate **circuit breakers** are needed to contain the operational contagion of a cyber incident spreading across a system⁹⁷.

Calibrating macroprudential tools

As explained in Section 4.1, public authorities that adopt an approach based on impact tolerance will have the ability to adjust **impact thresholds** periodically, in line with changing conditions. However, the downstream consequences of making adjustments need to be carefully considered, given the impact on the sector's investment choices. Authorities may also perform a **scenario setting** role, especially in the context of cyber stress testing. For each test, authorities may choose to adjust the parameters of the scenario (informed by the threat landscape and vulnerability information), as well as the firms in scope and the economic functions tested.

Maintaining clear communications

Public authorities are also able to exert influence through their communications, to achieve a desired policy outcome. On the premise of seeking to offset the systemic amplifier, authorities could choose to address market participants on topics such as:

- a broader adoption of the concept of systemic risk manager as found in financial market infrastructure;
- a focus on substitutability where possible;
- initiating a discussion on the viability of zero-trust architectures for the finance sector.

⁹⁵ Bank of England (2018), *Financial Stability Report*, No 43, June.

⁹⁶ Bank of England (2013), "Tiering in Chaps", Quarterly Bulletin 2013 Q4, pp. 371-378.

⁹⁷ Yong, H.K. and Yang, J.J. (2004), "What Makes Circuit Breakers Attractive to Financial Markets? A Survey," *Financial Markets, Institutions & Instruments*, Vol. 13, No 3, pp. 109-146.



5 Using the model

The use of historical and theoretical scenarios through the model

This section of the paper describes how real-life or hypothetical cyber incidents can be expressed using the conceptual model, and provides a structured approach to scenario analysis. By stepping through the phases of the model, the various elements of a cyber incident and its subsequent effects can be articulated and observations drawn. In addition, the scenario walkthroughs provide a mechanism for testing the model for accuracy and completeness.

5.1 Modelling scenarios

Scenarios can be used to unlock insights and serve as a means of exploring potentially extreme events. However, the viability of the analysis process demands a degree of rigour to ensure that results are consistent, comparable and repeatable. The model's phased and parameterised structure provides a logical sequence for capturing the relevant aspects of cyber incidents to facilitate analysis. For the benefit of those who wish to perform scenario analysis using the model, a separate [data capture template](#) has been developed to accompany this paper⁹⁸.

Preparation

The first scenario decision is to choose whether a historical or a theoretical example will form the basis of the analysis. The use of past cyber incidents requires access to sufficient knowledge of the circumstances surrounding the event if the analysis is to be useful. Conversely, hypothetical scenarios offer flexibility in terms of using the full range of options provided by the conceptual model, although this relies heavily on judgement and experience-based predictions. As this is an experiential exercise, the conclusions drawn should clarify the assumptions made at each stage of the walkthrough.

Regardless of the selection, scenario preparation involves identifying the base parameters to facilitate the subsequent walkthrough. Essential "Context" components include the selection of the incident starting point, the affected asset types, the vulnerabilities to be exploited, and an understanding of impact thresholds for the system. Optionally, the analysis can include a description of the threat as background information or to add a degree of realism which may help ground the judgement formed in later phases. The inclusion of countermeasures that would offset the initial risk crystallisation is only required if a worst-case scenario walkthrough is not performed.

⁹⁸ European Systemic Risk Board (2020), *Scenario Analysis Template*.



Walkthrough

Once the preparatory steps have been taken, the scenario contains all the prerequisite elements needed to describe the circumstances through which the incident arose.

The walkthrough begins by describing the technical and business impacts associated with the “Shock” phase of the model. Although impact types are relatively simple to identify, impact measurement is likely to be qualitative in nature without access to historical impact data or a proven predictive capability.

Based on the elements selected in the “Context” phase, the walkthrough process will need to identify the systemic amplifiers and the contagion channels which are likely to be triggered in each successive cycle of contagion. Using the 3x3 contagion grid in Figure 3.14, one or more transmission mechanisms can be reflected as active in any given cycle.

The final step in the scenario walkthrough is to form a judgement as to whether the aggregate impact experienced at the peak of the cyber incident exceeded the pre-set threshold, i.e. whether the scenario resulted in a systemic event.

Walkthrough variants

In addition to using the model under worst-case circumstances, two further variants are proposed:

- **Central estimation.** In addition to gross risk evaluation, in which all countermeasures and systemic mitigants are either missing or ineffective, the walkthrough process can be repeated to perform a net risk evaluation (or best-case scenario) which assumes that all reasonable measures are in place. During the re-run, the effects noted during the “Shock” and “Amplification” phases are adjusted based on the presence of effective countermeasures and systemic mitigants.

With worst and best-case scenarios described, the aggregate impact profiles of each scenario walkthrough provide the outer bounds for a central estimate. Further walkthroughs of the scenario can be performed by altering any of the variables to capture different outcomes, and are subsequently illustrated using a fan chart (see Figure 5.1). Each walkthrough can provide further insight or a confidence factor with regard to the most likely path of aggregate impact.


- **Reverse walkthrough.** Alternatively, a scenario can also be examined by working backwards from a pre-determined endpoint. Instead of identifying the specific assets affected to understand the downstream effects, this approach begins by identifying the economic function(s) affected, and reverse engineering the events that might have led to that outcome.

Scenario designers may wish to experiment with each of these techniques as part of an initial discovery phase, before fixing the parameters for a scenario walkthrough.



Figure 5.1

Fan chart of aggregate impacts from multiple scenario walkthroughs, with a central estimate of the most likely outcome



Source: For illustration purposes only.

5.2 Scenario examples

In order to validate the conceptual model and determine the extent to which a cyber incident could lead to a systemic event, a mix of historical and hypothetical scenarios were used from which observations could be drawn. The six scenarios reflected in this paper are a subset of a broader selection considered, but have been consolidated to demonstrate different paths through the model. The full analysis behind the summary tables in Figures 5.2 and 5.3 can be found in the ESCG's report on *Systemic Cyber Risk*⁹⁹.

Historical scenarios

Past cyber incidents affecting the financial sector provide a useful source of historical data from which lessons can be drawn, and for which "What if?" permutations may be considered. Indeed, the financial sector has experienced, and continues to experience, cyber incidents in varying guises, as captured by the **incident timeline** maintained by the Carnegie Endowment for International Peace and BAE Systems¹⁰⁰.

When studying a real-life incident, a selection of credible sources should be used to underpin the analysis. Public reporting and second-hand media accounts of events tend to contain less insight

⁹⁹ European Systemic Risk Board (2020), *Systemic Cyber Risk*.

¹⁰⁰ Carnegie Endowment for International Peace, **Timeline of Cyber Incidents Involving Financial Institutions**.

into the build-up of impacts or transmission mechanisms, focusing instead on the threat or eventual outcome. Equally, supply and access to relevant incident information may be limited by the level of disclosure from market participants, or fragmented across public authorities. It may therefore be necessary to augment sourced material with reasonable assumptions, to complete all aspects of the scenario analysis.

Figure 5.2

Summary of historical cyber incidents used in scenario analysis

	WannaCry (2017)	NonPetya (2017)	Cosmos Bank (2018)
Background	<ul style="list-style-type: none"> Large-scale crypto ransomware incident See Homogeneity box in Section 3.4 for details 	<ul style="list-style-type: none"> Destructive wiper malware See Pervasiveness box in Section 3.5 for details 	<ul style="list-style-type: none"> Coordinated and sophisticated cyber-enabled fraud Perpetrated across 28 countries in two hours
Context	<ul style="list-style-type: none"> Multiple firms affected Affecting unpatched Windows-based IT assets Indiscriminate malicious intent for financial gain 	<ul style="list-style-type: none"> Multiple firms and supply chain affected Infected accounting software (M.E.Doc) Targeted malicious intent on national disruption 	<ul style="list-style-type: none"> Single firm affected Altered core banking functionality to enable fraudulent transactions Targeted malicious intent for financial gain
Shock	<ul style="list-style-type: none"> Loss of availability and integrity resulting from data encryption Simultaneous service incapacitation across sectors and jurisdictions 	<ul style="list-style-type: none"> Loss of availability and integrity from permanent data destruction Primarily affected local and international firms operating in Ukraine 	<ul style="list-style-type: none"> Loss of authenticity, accountability and non-repudiation Fraudulent ATM withdrawals, resulting in a USD 13.5 million loss
Amplification	<ul style="list-style-type: none"> Homogeneity Pervasiveness Speed of transmission Concurrency 	<ul style="list-style-type: none"> Pervasiveness Speed of transmission Concurrency Uncertainty 	<ul style="list-style-type: none"> Concurrency Persistence
Systemic Event	<ul style="list-style-type: none"> Averted through quick discovery of 'kill switch' Loss estimates up to USD 8 billion Financial sector unaffected Overall confidence not sufficiently affected to put financial stability at risk 	<ul style="list-style-type: none"> Global impact limited to specific geography or affected institutions, as a result of specific targeting If directed at a global financial centre, effects could have been far larger 	<ul style="list-style-type: none"> Insufficiently large financial or confidence loss to trigger a systemic event Level of penetration achieved suggests that significantly more damage could have been inflicted had the threat actors intended to do so

Source: Summarised from ESCG report on Systemic Cyber Risk.

Hypothetical scenarios

Theoretical scenarios provide an opportunity to create a chain of events which has the potential to cross the systemic event threshold in a way that real-life incidents have yet to fully demonstrate. However, the challenge here lies in maintaining a degree of plausibility. It is very easy to create a doomsday scenario which unquestionably leads to a systemic crisis, yet misses an opportunity to use scenario analysis as a tool for vulnerability identification and learning. Designers may look to incorporate elements of past incidents into their scenario, but then augment this to amplify the impacts to a desired level.



Impact estimation for hypothetical scenarios requires many assumptions to be made at the outset, e.g. the nature of the financial institution(s) affected, the business services it provides, its role within the financial system (domestically / globally). The walkthrough process also requires a deep understanding of the interplay between different parts of the financial system, so as to capture the different forms of contagion.

Figure 5.3

Summary of hypothetical cyber incidents used in scenario analysis

	D-SIB Payment Disruption	Account Data Destruction	Price Feed Manipulation
Background	<ul style="list-style-type: none"> Accidental disruption of payment functions at large domestic bank (D-SIB) Significant contributor to retail payments systems Crisis compounded by fake news on social media 	<ul style="list-style-type: none"> Permanent destruction of account balance information and other data related to value at a large multi-national bank, with severe impacts on both wholesale and retail clients 	<ul style="list-style-type: none"> Manipulation of price feeds from commodities and futures markets, as well as the trade and position information that market participants receive from the market's CCP
Context	<ul style="list-style-type: none"> Single firm affected Planned IT change accidentally repurposes redundant code in batch scheduler software, corrupting all payments data in the batch jobs Attempts to manually re-load batch fail to meet key cut-off times, causing further backlog and cascading effect 	<ul style="list-style-type: none"> Supply chain affected Compromise of outsourced IT contractors Pre-emptive corruption of data backup and restore processes, and prolonged extraction of sensitive data Large-scale fraudulent payments and destruction of payment software and account balance data at critical time 	<ul style="list-style-type: none"> Multiple firms affected Market data providers and a CCP are simultaneously targeted Malicious code inserted into the infrastructure used for processing and outputting of price, trade and position data Enables selective modification of information being received or sent
Shock	<ul style="list-style-type: none"> Loss of availability and integrity related to payment services and transaction data D-SIB temporarily forced to shut retail operations, with inability to reconcile account balances, or provide access to online banking and cash points 	<ul style="list-style-type: none"> Loss of availability and integrity of account balance data Operations suspended, with manual workarounds offering short-term relief Bank subsidiaries in other countries also affected due to centralised nature of the group's IT systems 	<ul style="list-style-type: none"> Loss of integrity and accountability of price feed information Results in random errors for entered trades, trade rejections, errors in the reporting of current positions and conflicting market prices observed by market participants
Amplification	<ul style="list-style-type: none"> Scale and criticality Complexity Uncertainty Ineffective response 	<ul style="list-style-type: none"> Persistence Detection time Time sensitivity 	<ul style="list-style-type: none"> Concurrency Network topology Uncertainty
Systemic Event	<ul style="list-style-type: none"> Prolonged disruption of a significant fraction of a country's payment system combined with uncertainty could trigger large-scale financial instability Loss of confidence in one financial institution may quickly spread to a general loss in confidence and liquidity problems across the sector 	<ul style="list-style-type: none"> Realisation that data recovery is not possible, requiring systems to be completely rebuilt Affected bank and several counterparties report liquidity problems, whilst national payment, clearing and settlement systems also disrupted Key concern is whether account data is (perceived to be) permanently lost 	<ul style="list-style-type: none"> Market makers seek to exit positions, further depressing prices, leading to distressed liquidations Ensuing market panic takes on a self-reinforcing and self-sustaining dynamic with multiple 'liquidity spirals' Conceivable that the CCP will incur losses exceeding its default fund, triggering the default of the CCP

Source: Summarised from the ESCG's report on Systemic Cyber Risk.



5.3 Observations from scenario analysis

The following observations, grouped by model phase, were recorded either: (i) while performing scenario analysis; or (ii) in retrospect, when seeking to identifying commonalities across the walkthroughs.

Context

The nature of the threat has a significant bearing on systemic outcome. Scenario analysis suggests that the greatest financial or confidence loss occurs when there is deliberate intent to disrupt or destroy parts of the financial system. Malicious acts motivated by financial gain (e.g. cyber-enabled fraud) are not perceived as generating sufficient losses to trigger a systemic crisis. Furthermore, cyber incidents which are accidental in nature are much less likely to provoke the uncertainty amplifier, as the circumstances behind the incident are more quickly understood.

Shock

Aside from availability and integrity, it is unlikely that the isolated loss of other cybersecurity properties will result in a systemic event. Loss of integrity from deliberate acts is also likely only to result from the initial loss of another property, e.g. loss of authenticity through a successful spear-phishing attempt, allowing a foothold to be gained in the environment. However, perceived (or actual) permanent loss of integrity is most likely to lead to systemic consequences, especially if the information affected is a store of value.

Amplification

The scenarios with most systemic amplifiers present have the potential to escalate most rapidly and generate the greatest levels of aggregate impact. Cyber incidents which possess all the three characteristics of speed, scale and malicious intent tend to demonstrate the most significant repercussions.

In terms of contagion channels, scenarios where **loss of confidence coincides with the expectation of significant financial loss** dramatically increase the prospects of the cyber incident becoming systemic. Significant operational impacts on their own may have system-wide effects, but are unlikely to endanger financial stability as long as trust endures among counterparties. The loss of this trust marks a key transition point in a worsening situation, e.g. when institutions transition from “*not being not able to lend to each other*” to “*not willing to lend to each other*”.

Systemic event

Adequate preparedness, rapid coordination and proactive communication on the part of affected parties and relevant authorities can be key systemic mitigants. The velocity of cyber



incidents dictates that there is unlikely to be sufficient time or capacity to draw on an ad hoc and untested incident response capability at the point of need.

Cyber risk can crystallise in many forms, with the vast majority of cyber incidents handled according to stakeholder expectations. **However, an alignment of systemic amplifiers, further exacerbated by deliberate intent, could conceivably result in a systemic event, endangering financial stability and materially impacting on the real economy.**



References

- Alperovitch, D. (2011), **Revealed: Operation Shady RAT**, August.
- Armstrong, R.C. and Mayo, J.R. (2009), *Leveraging Complexity in Software for Cybersecurity*.
- Baker, M. and Wurgler, J. (2007), "Investor Sentiment in the Stock Market", *Journal of Economic Perspectives*, Vol. 21, No 2, pp. 129-151.
- Bank of England (2013), "Tiering in Chaps", *Quarterly Bulletin* 2013 Q4, pp. 371-378.
- Bank of England (2018), *Bank of England's RTGS and CHAPS services: Service Description*.
- Bank of England (2018), *Financial Stability Report*, No 43, June.
- Bank of England, Prudential Regulatory Authority; Financial Conduct Authority (2019), *Consultation Paper: Operational Resilience: Impact tolerances for important business services*.
- Barabási, A.-L. (2002), *Linked: The New Science of Networks*.
- Barbalet, J.M. (1996), "Social Emotions: Confidence, Trust and Loyalty", *International Journal of Sociology and Social Policy*, Vol. 16, No 9, pp. 75-96.
- Basel Committee on Banking Supervision (2011), *Global systemically important banks: Assessment methodology and additional loss absorbency requirement*.
- Basel Committee on Banking Supervision (2011), *Principles for the Sound Management of Operational Risk*.
- Benfield, Aon (2017), *Cyber Event Briefing: Wannacry Ransomware Attack*.
- Birkmann, J. and Wisner, B. (2006), "Measuring the un-measurable: The challenge of vulnerability", *United Nations University Institute for Environment and Human Security*.
- Board of Governors of the Federal Reserve System; Office of the Comptroller of the Currency; Securities and Exchange Commission (2003), **Interagency Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System**, 7 April.
- Bookstaber, R. and Kennet, D.Y. (2016), "Looking Deeper, Seeing More: A Multilayer Map of the Financial System", *Office of Financial Research*.
- Borgatti, S.P. (2005), "Centrality and network flow", *Social Networks*, No 27, pp. 55-71.
- Cairncross, F. (1997), *The Death of Distance: How the Communications Revolution Will Change Our Lives*, Harvard Business School Press.
- Calaprice, A. (2000), *The Expanded Quotable Einstein*, Princeton University Press, p. 314.
- Capital One (2019), **2019 Capital One Cyber Incident**, 19 July.



Carnegie Endowment for International Peace, **Timeline of Cyber Incidents Involving Financial Institutions**.

Caruana, J. (2010), *Systemic risk: how to deal with it?*, BIS.

Cilliers, P. (1998), *Complexity and postmodernism – Understanding complex systems*.

Commission of the European Communities (2009), *An EU Framework for Cross-Border Crisis Management in the Banking Sector*.

CPMI-IOSCO (2012), **Principles for financial market infrastructure**, April.

Dalziel, E.P. and McManus, S. (2004), *Resilience, Vulnerability, and Adaptive Capacity: Implications for System Performance*.

Davoudi, S. (2012), “Resilience: A Bridging Concept or a Dead End?”, *Planning Theory & Practice*, Vol. 13, No 2, pp. 299-307.

Deloitte (2016), *Beneath the surface of a cyberattack – A deeper look at business impacts*.

Donahue, A.K. and Tuohy, R.V. (2006), *Lessons We Don't Learn: A Study of the Lessons of Disasters, Why We Repeat Them, and How We Can Learn Them*.

DTCC and Oliver Wyman (2018), *Large-Scale Cyber Attacks on the Financial System: A case for better coordinated response and recovery strategies*.

Eisenberg, L. and Noe, T. (2001), “Systemic Risk in Financial Systems”, *Management Science*, Vol. 47, No 2, pp. 236-249.

Eisenegger, M. and Imhof, K. (2007), *The True, the Good and the Beautiful: Reputation Management in the Media Society*.

European Union Agency For Network and Information Security (2018), *Threat Landscape Report*.

European Systemic Risk Board (2020), *Scenario Analysis Template*.

European Systemic Risk Board (2020), *Systemic Cyber Risk*.

Financial Stability Board (2013), *Recovery and Resolution Planning for Systemically Important Financial Institutions: Guidance on Identification of Critical Functions and Critical Shared Services*.

Financial Stability Board (2017), *Financial Stability Implications from FinTech: Supervisory and Regulatory Issues the Merit Authorities' Attention*.

Financial Stability Board (2018), *FSB Cyber Lexicon*.

Gai, P. (2013), *Systemic Risk: The Dynamics of Modern Financial Systems*.

Gelbstein, E. (2013), “Quantifying Information Risk and Security”, *ISACA Journal*, Vol. 4.



- Goerner, S.J., Lietaer, B. and Ulanowicz, R. (2009), "Quantifying economic sustainability: Implications for free-enterprise theory, policy and practice", *Ecological Economics*, Vol. 69, No 1, pp 76-81.
- Goodin, D. (2017), "Windows 7, not XP, was the reason last week's WCry worm spread so widely", *Ars Technica*, 20 May [Online].
- Greenberg, A. (2018), "The Untold Story of NotPetya, the Most Devastating Cyberattack in History", *WIRED*, 8 August.
- Haldane, A.G. (2009), *Rethinking the financial network*.
- Haldane, A.G. (2015), *On microscopes and telescopes*.
- Hassel, H. (2007), "Risk and Vulnerability Analysis of Complex Systems: a basis for proactive emergency management", *LUCRAM (Lund University Centre for Risk Analysis and Management)*.
- Holling, C.S. (2001) "Understanding the Complexity of Economics, Ecological, and Social Systems", *Ecosystems*, Vol. 4, No 5, pp. 390-405.
- IBM and Ponemon Institute (2018), *2018 Cost of a Data Breach Study: Global Overview*.
- IMF/BIS/FSB (2009), *Guidance to assess the systemic importance of financial institutions, markets and instruments: Initial considerations, Report to G20 Finance Ministers and Governors*.
- IMF/BIS/FSB (2016), *Elements of Effective Macroprudential Policies: Lessons from International Experience*.
- International Financial Reporting Standards (2018), *Conceptual Framework for Financial Reporting*.
- International Integrated Reporting Council (2013), *The International <IR> Framework*.
- International Monetary Fund (2013), *Key Aspects of Macroprudential Policy*.
- International Monetary Fund (2018), *Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment*.
- International Organization for Standardization (2015), *Societal security – Business continuity management systems – Guidelines for business impact analysis (BIA)*.
- International Organization for Standardization (2018), *Information technology – Security techniques – Information security risk management*.
- International Organization for Standardization (2018), *Security and resilience – Vocabulary*.
- ISACA (2009), *The Risk IT Framework*.
- Jurczyk, J., Rehberg, T., Eckrot, A. and Morgenstern, I. (2017), "Measuring critical transitions in financial markets", *Nature*, Vol. Scientific Reports, No 7.



Kahneman, D. and Tversky, A (1979), "Prospect Theory: An Analysis of Decision under Risk", *Econometrica*, Vol. 47, No 2, pp. 263-291.

Knight, F.H. (1921), *Risk, uncertainty and profit*, Houghton Mifflin, Boston.

Lorenz, E. (1972), "Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?", in 139th meeting of the American Association for the Advancement of Science.

Mandiant (acquired by FireEye) (2013), **APT1 – Exposing One of China's Cyber Espionage Units**, 19 February.

Nakashima, E. and Timberg, C (2017). "NSA officials worried about the day its potent hacking tool would get loose. Then it did.", *Washington Post*, 16 May [Online].

National Institute of Standards and Technology (2012), *Guide for Conducting Risk Assessments*.

National Institute of Standards and Technology (2018), *Risk Management Framework for Information Systems and Organizations*.

Open Web Application Security Project (OWASP), **OWASP Risk Rating Methodology**.

Perrow, C. (1999), *Normal Accidents: Living with High-Risk Technologies*, Princeton University Press.

PricewaterhouseCoopers (2019), **Operational resilience: how to set and test impact tolerances**, October.

PwC / BAE Systems (2017), *Operation Cloud Hopper – Exposing a systematic hacking operation with an unprecedented web of global victims*.

Reason, J. (1990), "The Contribution of Latent Human Failures to the Breakdown of Complex Systems", *Philosophical Transactions of the Royal Society of London*, Vol. 327, No 1241.

Rodriguez, J. and Gasper, P.D. (2015), *Understanding the Value of a Computer Emergency Response Capability for Nuclear Security*, Idaho National Laboratory.

Ron, A., Danielsson, J., Baba, C., Das, U., Kang, H. and Segoviano, M. (2018), *Macroprudential Stress Tests and Policies: Searching for Robust and Implementable Frameworks*, International Monetary Fund.

Ruf, L., Thorn, A., Christen, T., Gruber, B., Portmann, R. and Luzern, H. (2008), "Threat Modeling in Security Architecture – The Nature of Threats", *Information Security Society Switzerland (ISSS) – Working Group on Security Architecture*.

Rumsfeld, D. (2002), Department of Defense News Briefing [Interview], 12 February.

Samarji, L.E. (2015), "Risk-aware Decision Support System to Counter Coordinated and Simultaneous Attacks", *Télécom Bretagne*, Université de Rennes.



Stytz, M.R. and Banks S.B. (2014), "Cyber Warfare Simulation to Prepare to Control Cyber Space," National Cybersecurity Institute Journal, Vol. 1, No 2, pp. 9-25.

Technical Committee of the International Organization of Securities Commissions (2011), *Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency – Final Report*.

The MITRE Corporation (2018), *MITRE ATT&CK: Design and Philosophy*.

The MITRE Corporation (2019), **Common Vulnerabilities and Exposures (CVE)**.

The Open Group (2013), *Technical Standard - Risk Taxonomy (O-RT), Version 2.0*.

Thunnissen, D.P. (2003), *Uncertainty Classification for the Design and Development of Complex Systems*.

Trend Micro (2017), **RATANKBA: Delving into Large-scale Watering Holes against Enterprises**, 27 February.

Trichet, J.-C. (2009), *Text of the Clare Distinguished Lecture in Economics and Public Policy*, Clare College, University of Cambridge.

US House of Representatives – Committee on Oversight and Government Reform, **The Equifax Data Breach – Majority Staff Report**.

Vanston, N. (2012), *Trust and reputation in financial services*, Government Office for Science.

von Bertalanffy, L. (1968), *General Systems Theory: Foundations, Development, Applications*.

Waldrop, M. (1992), *Complexity: The Emerging Science at the Edge of Order and Chaos*, Simon and Schuster, New York.

World Economic Forum (2015), *Partnering for Cyber Resilience - Towards the Quantification of Cyber Threats*.

World Economic Forum (2020), *The Global Risks Report 2020, 15th Edition*.

Yayla, A.A. and Hu, Q. (2010), *The impact of information security events on the stock value of firms: the effect of contingency factors*.

Yong, H.K. and Yang, J.J. (2004), "What Makes Circuit Breakers Attractive to Financial Markets? A Survey," *Financial Markets, Institutions & Instruments*, Vol. 13, No 3, pp. 109-146.

Zurich Insurance Company and Atlantic Council (2014), *Risk Nexus – Beyond data breaches: global interconnections of cyber risk*.



Abbreviations

APT	Advanced Persistent Threat	IIRC	International Integrated Reporting Council
BC	Business Continuity	IMF	International Monetary Fund
BCBS	Basel Committee on Banking Supervision	IOSCO	International Organization of Securities Commissions
BIA	Business Impact Analysis	ISACA	Information System Audit and Control Association
CAS	Complex Adaptive System	ISO	International Organization for Standardization
CHAPS	Clearing House Automated Payment System	MSP	Managed Service Provider
CIA	Confidentiality, Integrity, Availability	MTTI	Mean Time To Identify
CIV	Common Individual Vulnerability	NAT	Natural Accident Theory
CVE	Common Vulnerabilities and Exposures	OT	Operational Technology
DR	Disaster Recovery	OWASP	Open Web Application Security Project
ECB	European Central Bank	PT	Prospect Theory
ENISA	European Union Agency for Network and Information Security	PwC	PricewaterhouseCoopers
ESCG	European Systemic Cyber Group	RDBMS	Relational Database Management System
ESRB	European Systemic Risk Board	RTGS	Real Time Gross Settlement
EUT	Expected Utility Theory	RYF	Robust-Yet-Fragile
FAIR	Factor Analysis for Information Risk	SCADA	Supervisory Control and Data Acquisition
FSB	Financial Stability Board	SCAV	Standing Committee for Assessment of Vulnerabilities
GNSS	Global Navigation Satellite System	UTC	Coordinated Universal Time
HFT	High Frequency Trading	VaR	Value at Risk
ICS	Industrial Control System	WEF	World Economic Forum
ICT	Information and Communications Technology		
IFRS	International Financial Reporting Standards		



Imprint and acknowledgements

This paper was completed under the auspices of the ESRB's European Systemic Cyber Group (ESCG), chaired by Paul Williams, from the Bank of England. The author would like to thank Eric Schaanning for Secretariat support, and other members of the ESCG for their contributions and discussions which influenced this paper.

Greg Ros

Bank of England, London, United Kingdom; email: greg.ros@bankofengland.co.uk

Eric Schaanning

European Systemic Risk Board, Frankfurt am Main, Germany; email: eric.schaanning@esrb.europa.eu

© European Systemic Risk Board, 2020

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.esrb.europa.eu

All rights reserved. Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

The cut-off date for the data included in this report was 19 December 2019.

ISSN 2467-0669 (pdf)
ISBN 978-92-9472-132-7 (pdf)
DOI 10.2849/915512 (pdf)
EU catalogue No DT-AC-20-001-EN-N (pdf)