CATAPULT Digital

A new perspective on cyber risk, applied to the UK FMI ecosystem

Applying the benefit harm index (BHI) - (A new approach to modelling risk assessment of cyber ecosystems and their socio-economic impacts) to the UK's evolving financial markets infrastructure (FMI) ecosystem.

CHARLES FOX – SECURITY LEAD, DIGITAL CATAPULT
BRIAN MACAULAY – LEAD ECONOMIST, DIGITAL CATAPULT

Executive summary

Within an increasingly complex and interconnected world, the way in which cyber threats are perceived and responded to needs to be reconsidered. Traditional risk models rely heavily on probabilistic approaches, which demand stable distribution and almost complete knowledge of all possible states.

New advances in digital technologies, combining huge data, rapidly evolving automated algorithms and the prospect of a generational shift in network speed and capacity, pose serious challenges to traditional risk modelling. Through the Hermeneut project (part of the European Community's Horizon 2020 programme) Digital Catapult has proposed a new approach to understanding dynamic and emergent threats: the benefit harm index (BHI), which integrates ideas from both economics and complexity science.

This report shows how this exciting new perspective on cyber risk modelling can be applied to the cyber ecosystems that form many of today's critical national infrastructures (CNI) - complex systems of systems that exhibit emergent behaviour and require a new approach to cyber risk assessment.

This study looks at the systemic socio-economic impacts that can result from cyber attacks associated with emergent threats to CNI cyber ecosystems, and uses the UK financial markets infrastructure (FMI) ecosystem as a case study for the new BHI approach.

The UK FMI ecosystem is part of the UK economy and is one of the UK's 13 CNI components. Systemically important FMIs play an essential role in the financial system, and the disorderly failure of such an FMI could lead to severe systemic disruption if it caused markets to cease to operate effectively.

A high-level ecosystem for 2020-30 has been modelled to focus on the critical FMI operational systems domain, and on the associated domains of UK governance, the supply chain and wider non-critical core services.

Finally, there is a description of the approach that can be used to mitigate the growth of harm within these complex systems of systems, and highlights the use of Implication Wheel™[1] methodology to uncover emergent systemic threats to the UK FMI cyber ecosystem.

In this report

Introducing the Benefit Harm Index: A new perspective on cyber risk 4

Using BHI to mitigate to emergent threats 5-6

Using BHI to mitigate to the growth of harm 7

Applying BHI to cyber ecosystems 8

High-level model of the UK FMI ecosystem 9

The UK critical FMI operational system COI 10

The FMI central counterparties 10

The payment systems 10

The securities settlement systems 10

The UK FMI governance COI 11

The wider financial services COI 11

The FMI supply chain COI 11

Complexity and evolution of the UK FMI ecosystem 12-13

Applying the BHI using an illustrative cyber attack scenario 14

The illustrative cyber attack scenario 15-17

Exploring the vulnerability and control aspects of the system 17-18

Exploring the liklihood of attack scenario 18

Exploring the potential impact/harm of the attack scenario 19

The benefit to harm index perspective on the scenario 20-21

An approach to mitigating emergent risk/radical ignorance 22-24

Conclusion: Find out more about BHI and the Hermeneut project 25

Glossary 26

References 27

INTRODUCING THE BHI – A NEW PERSPECTIVE TO CYBER RISK

BHI modelling methodology is designed to provide new insights into the potential risks associated with the cyber ecosystems which underpin complex and dynamic markets that are driven by the exploitation of emerging technologies. These rapidly evolving markets typically contribute significantly to national and international economies, and often form an integral part of CNI.

Unlike a controlled (deterministic) system with a known set of risks and a well-defined future state, a complex system features many unknown risks and will evolve in ways that cannot be fully predetermined. For example, within the biological ecosystem, microscopic changes can propagate rapidly and create a huge-scale impact, such as when a single virus mutates, evolves and spreads to cause a pandemic. This demonstrates Cyber ecosystems are also complex dynamic environments that evolve rapidly and feature high levels of uncertainty. They can generate emergent behaviours which cannot always be predicted by studying the way in which constituent parts interact. Emergent behaviours manifest themselves in many forms (as seen in the murmurations of birds in the biosphere, and new socio-political collective behaviours through social media use online).

Traditional risk assessment methodologies - which assume a complete knowledge of all possible states of the system being assessed and that a mathematical likelihood can be applied to each event - cannot address the complex dynamics, emergency behaviours and associated uncertainties of cyber ecosystems.

The Hermeneut BHI introduces a new approach to risk assessment, by modelling the growth of benefits and risks in the context of complex cyber ecosystems. It also features event driven scenario analysis methods, recognising the evolution of such systems over time.

Modelling dynamic complexity provides a perspective for exploring the rate of growth of socio-economic benefits generated by an evolving cyber ecosystem over time. It also provides a perspective for exploring the rate of growth of threats to that ecosystem, and the associated socio-economic harm that those threats could generate over time. The difference between the level of benefit and the level of harm at any given time period is a key output of the BHI model.

An event-driven scenario approach enables exploration of the implications of cyber chain reactions, helping to identify hidden risks (and benefits) using tools such as the Implication Wheel™[1]. This helps mitigate the fact that the risks for complex dynamic systems cannot fully be predicted as some will be emergent, and could be significant.

The BHI methodology applies many of the principles used in the latest economics research[2], recognising that the economy is a complex system within other systems. When the BHI methodology is applied to a cyber ecosystem, the balance between benefit and harm, and how that balance changes over time, can be explored. BHI is used to identify and mitigate emergent threats, and then to explore ecosystem-level mitigation strategies for those scenarios where the socio-economic harm outweighs the benefits. Any residual risks can then be managed using traditional risk assessment methodologies.

USING BHI TO MITIGATE TO EMERGENT THREATS

Cyber ecosystems are complex, and therefore exhibit emergent behaviour. As the level of complexity increases, different types of emergent behaviour will appear:

- Simple dynamic behaviour (such as a clock measuring time)
- Weak emergent behaviour (such as the flocking of birds or shoaling of fish)
- Strong emergent behaviour (such as bubbles within financial markets)
- Spooky emergent behaviour (such as conscious thought in humans or AI)

The first two emergent behaviour types are associated with deterministic systems, and can be easily reproduced using system simulations. The third and fourth are associated with stochastic (random interactions defined by probability distribution) systems. Stochastic systems can exhibit strong emergent behaviour that cannot be fully reproduced in simulations; spooky emergent behaviour cannot be reproduced by even the most detailed simulation.

The extent to which a cyber ecosystem can be controlled - and defended - is intrinsically linked to its level of complexity. The stability of the system is also related to its level of complexity, and changes at micro level can result in dramatic change at macro level. Therefore, an attack on a cyber ecosystem can trigger a significant chain reaction that will appear as emergent behaviour.

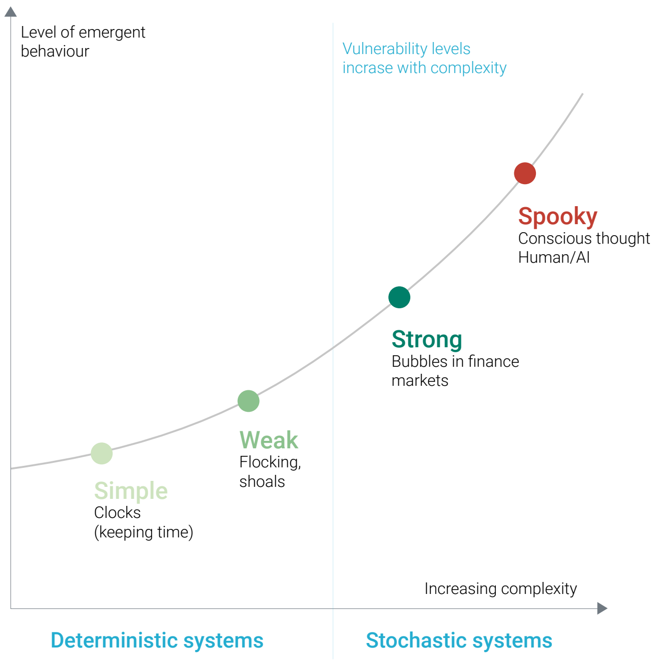

Figure 1 – Complexity and emergent behaviour

In the case of strong and spooky emergence, the stochastic systems) the system is fundamentally uncontrollable.

Table 1 shows how threats and vulnerabilities to components in a system will vary by class. Each vulnerability level (VL) requires a different type of mitigation.

The VL of a component may be changed by reconfiguring other components in the system. Some levels of vulnerability must be mitigated across the ecosystem.

Vulnerability level (VL)	Threat class	Attacker's control
● 5	Emergent system	The system can show emergent behaviour and cannot be controlled, since its phase space changes as emergent behaviours manifest themselves.
● 4	Stochastic system	The system cannot be controlled, but vulnerabilities can be reliably modelled using closed-form probability distributions over a fixed (and finite) set of state variables in the system's phase space.
● 3	Uncontrolled system	The system is not under control, but could be controlled in principle.
● 2	Uncontrolled inputs	An attacker uses a legitimate control input within the system's scope, but outside its expected or normal range.
● 1	Unauthorised activities	An attacker uses legitimate and in-scope control inputs within the control system.

Table 1 – Vulnerability levels and their associated class of threat

One of the key components of the BHI approach to dynamic risk involves mitigating emergent threats in complex ecosystems. Figure 2 illustrates the BHI process for doing this.

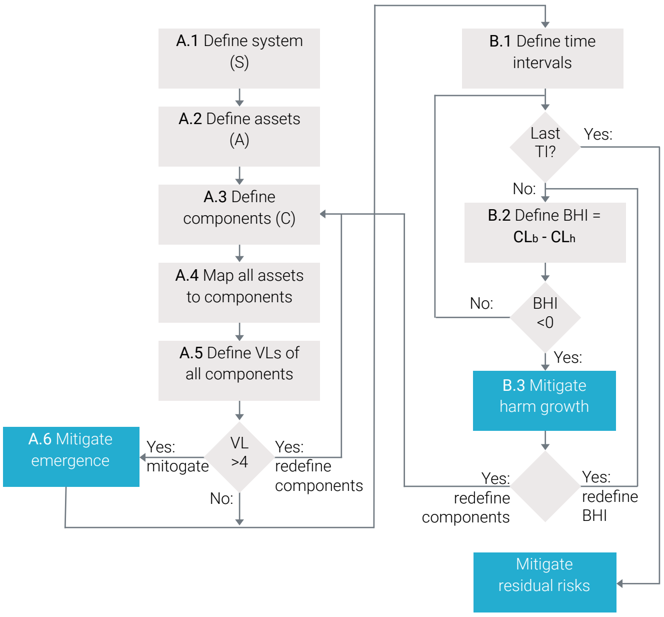

Figure 2 - BHI process for mitigating emergent threats

As shown in Figure 2, the first steps for addressing emergent threats (A.1 to A.5) are to define:

A.1: The ecosystem being considered.

A.2: The set of assets (the sensitivity of which is such that their loss or compromise would cause significant harm, and which - as a whole or in part - may be of interest to a threat agent for malicious, fraudulent and criminal activities).

A.3: The set of components which comprise the system - a component must contain hardware and may contain software and data (it is assumed that components can communicate with each other using sufficiently secure protocols).

A.4: The association between each asset and any component that directly influences its security.

A.5: The VL for each component.

These definitions should take into account the nature of each component and its vulnerabilities, as well as the threats from the environment and other components. If any component has VL greater than 4 (corresponding with emergent threat), the process takes one of two paths:

- Redefinition of the components, for example, to localise an associated asset in a component that has a lower VL value - this requires in reiteration over steps A.3 to A.5
- Mitigate emergence (A.6) by designing a set of security controls that seek to mitigate associated risks - these controls need to detect, and potentially isolate and neutralise the impact of an attack

Using BHI, characteristics that can be localised need to be distinguished from those which cannot. Organisations cannot be expected to mitigate non local characteristics, so other classes of intervention must be applied to safeguard the ecosystem. For the latter class, mitigations must be a set of governance, standard, and other interventions across the ecosystems, and key criteria for adoption must seek to minimise impact on the individual organisations adopting such recommendations.

Once this iterative process is complete, the process of considering emergent threat is also complete, and analysis passes to using BHI to mitigate threats from growth.

USING BHI TO MITIGATE THE GROWTH OF HARM

Modelling the dynamic complexity provides a perspective for exploring the rate of growth of the socio-economic benefits generated by an evolving cyber ecosystem over time. It also provides a perspective for exploring the rate of growth of threats to that ecosystem and the associated socio-economic harm they could generate over time. The difference between the level of benefit and the level of harm at any given time period is a key output of the BHI model.

Benefit and harm can grow at different rates within a cyber ecosystem. There are two key features of complex ecosystems that help to refine understanding of these growth rates.

1. Each ecosystem will evolve through a number of distinct phase transitions as it evolves.

For example, the introduction of a new product or class of products that penetrates a market. Initially there is near exponential growth, often modelled as compound growth in business plans, with a constant or slowly varying compound annual growth (CAGR) parameter. As penetration of the market occurs and saturation approaches, the Bass diffusion distribution eventually manifests its asymptotic growth complexity at constant of 0.

It is therefore appropriate to consider the BHI in three distinct time intervals:

TI0: From product introduction to when the complexity level is four (exponential)
TI1: From when the complexity level transitions from four to zero
TI3: From market saturation onwards, when the complexity level is zero (constant)

2. Each ecosystem will typically have multiple domains, each of which can feature different levels of complexity and associated growth rates.

The right-hand side of Figure 2 shows the process for using BHI to mitigate threats from growth.

The first step (B.1) defines the set of time intervals relevant to the various developments of both benefit and harm.

In particular, these time intervals will consider:

- Times of events marking the start and end of relevant changes, such as investment rounds or the introduction of new products
- Times at which the distribution of growth is likely to be discontinuous, for example as a result of some material event such as a change in product or the channel it uses to access the market

The second step (B.2) iterates over the intervals to compute the benefit to harm index (BHI) for each sub-interval, by determining the complexity index (CI) for each growth distribution. If the BHI is negative, indicating that the CI for growth of harm exceeds that of benefit, the process proceeds to mitigate harm growth (B.3), which specifies security controls. If a plausible mitigation is found, the process re-computes the BHI value and iterates to the next time interval.

In some cases, for example where an effective mitigation cannot be found, it may be appropriate to redefine the components. In this case, the process returns to the right-hand side of the diagram at step (A.3).

Level of benefit	Level of harm 0	1	2	3	4
0	0	-1	-2	-3	-4
1	1	0	-1	-2	-3
2	2	1	0	-1	-2
3	3	2	1	0	-1
4	4	3	2	1	0

Figure 3 – The BHI for distinct time intervals (TI)

For any BHI greater than 0 systemic (ecosystem) level mitigations are required.

Once all members of CI have been processed, mitigation of risks from growth is complete and the process can continue by using traditional risk management techniques to address any residual risks.

APPLYING BHI TO CYBER ECOSYSTEMS

To apply the BHI methodology to a target cyber ecosystem, the following high level ecosystem domain model is used.

A cyber ecosystem is a complex system of systems, where each system can be modelled in terms of a set of interacting components. Each ecosystem will have a scope/ system boundary and will typically be embedded in a wider environment. Political, economic, social, technological, environmental and legal (PESTL) influences from this wider environment affect the ecosystem's operation and growth.

Each cyber ecosystem is structured into a number of domains that support different dynamic communities of interest (COI). As shown in Figure 4, these domains reflect the distinction between operational systems within the ecosystem and the supply chain systems that support the manufacture and production of the components that will eventually populate that operational system's domain.

The other domains shown include the command and control systems domain, and the underlying system components, processes and interactions that comprise them. The governance and regulatory processes domain contains the governance systems and regulatory frameworks used to set and police the policies, rules and standards associated with governing the cyber ecosystem. The final domain is the value-added services domain, which includes the systems and processes associated with services that add value to the operational services, for example, insurance services.

All cyber system domains will have vulnerabilities. Threats to the ecosystem will exploit these vulnerabilities through attack vectors originating from threat sources (for example, hostile states), and attacking via threat actors (external and internal), as shown schematically in Figure 4. Through multiple iterations, the BHI approach exploits methodologies such as the Implementation Wheel™1 to investigate the vulnerability levels of components and cyber chain reactions being generated in complex systems. Targeted scenario analysis is used to help identify such events by systematically exploring the implications of interaction/contagion through multiple first, second, and nth order interaction flows.

The BHI dynamic approach to risks also enables the construction of multiple phase states of each cyber ecosystem model to reflect its different evolutionary states. This is then used to help create the BHI growth model across those different time intervals, resulting in an output of the form shown earlier in Figure 3.

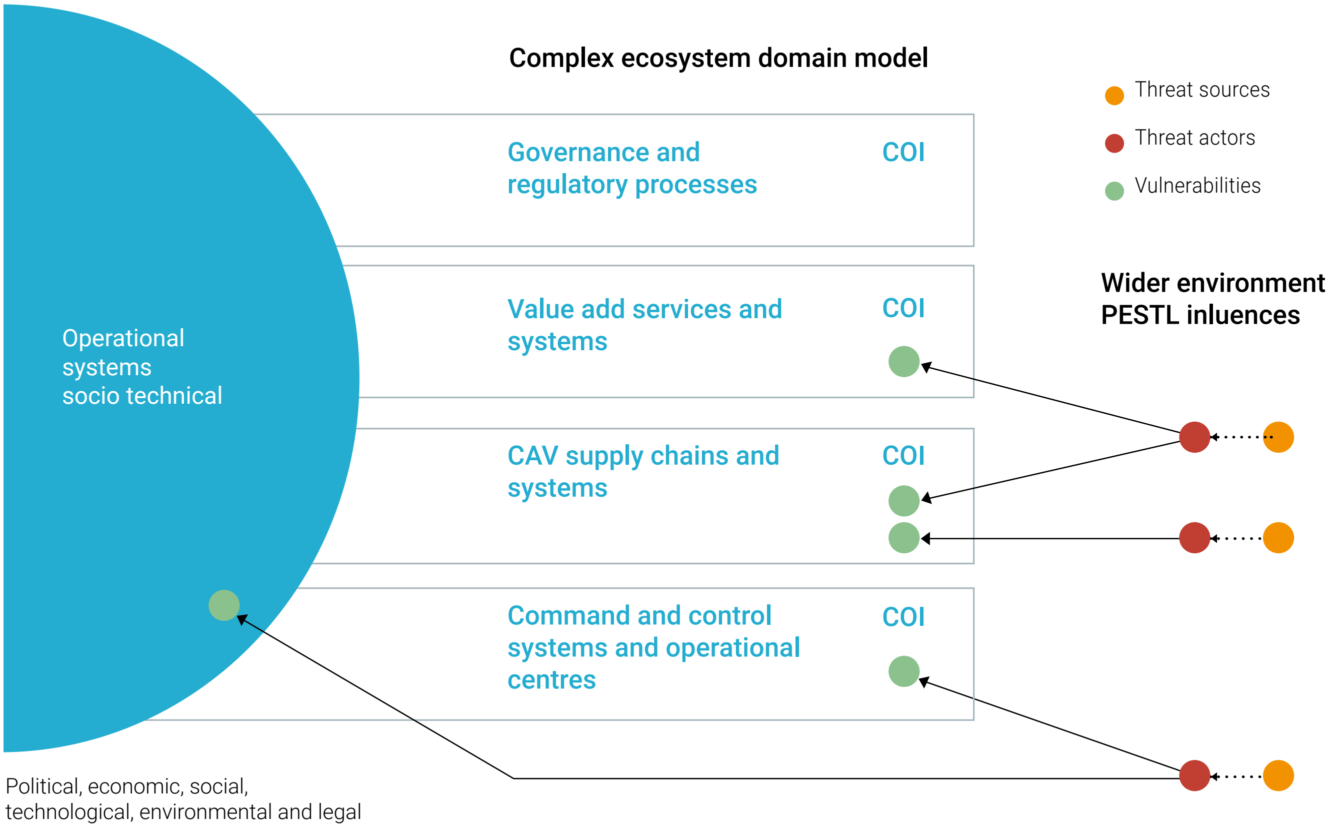
Complex ecosystem domain model
Threat sources
Threat actors
Vulnerabilities
Governance and regulatory processes
COI
Value add services and systems
COI
CAV supply chains and systems
COI
Command and control systems and operational centres
COI
Operational systems socio technical
Wider environment PESTL inluences
Political, economic, social, technological, environmental and legal

Figure 4 - Cyber ecosystem high level domain model

HIGH-LEVEL MODEL OF THE UK FMI ECOSYSTEM

The UK's financial markets infrastructure (FMI) is one of the UK's thirteen CNI components. This paper focuses on the critical FMIs that underpin the resilience of the UK financial sector.

Using our ecosystem domain model, the UK FMI ecosystem can be represented at a conceptual level, as shown in Figure 5.

Each of the domains shown in Figure 5 represents a distinct dynamic socio-technical community of interest (COI) within the UK FMI ecosystem. The central core in Figure 5 contains common infrastructure services, such as SwiftNet, which is managed by the Society for Worldwide Interbank Financial Telecommunication (SWIFT).

The UK FMI ecosystem is embedded in the wider global finance ecosystem and is subject to global political, economic, social, technical and legal (PESTL) influences.

UK FMI ecosystem domain model

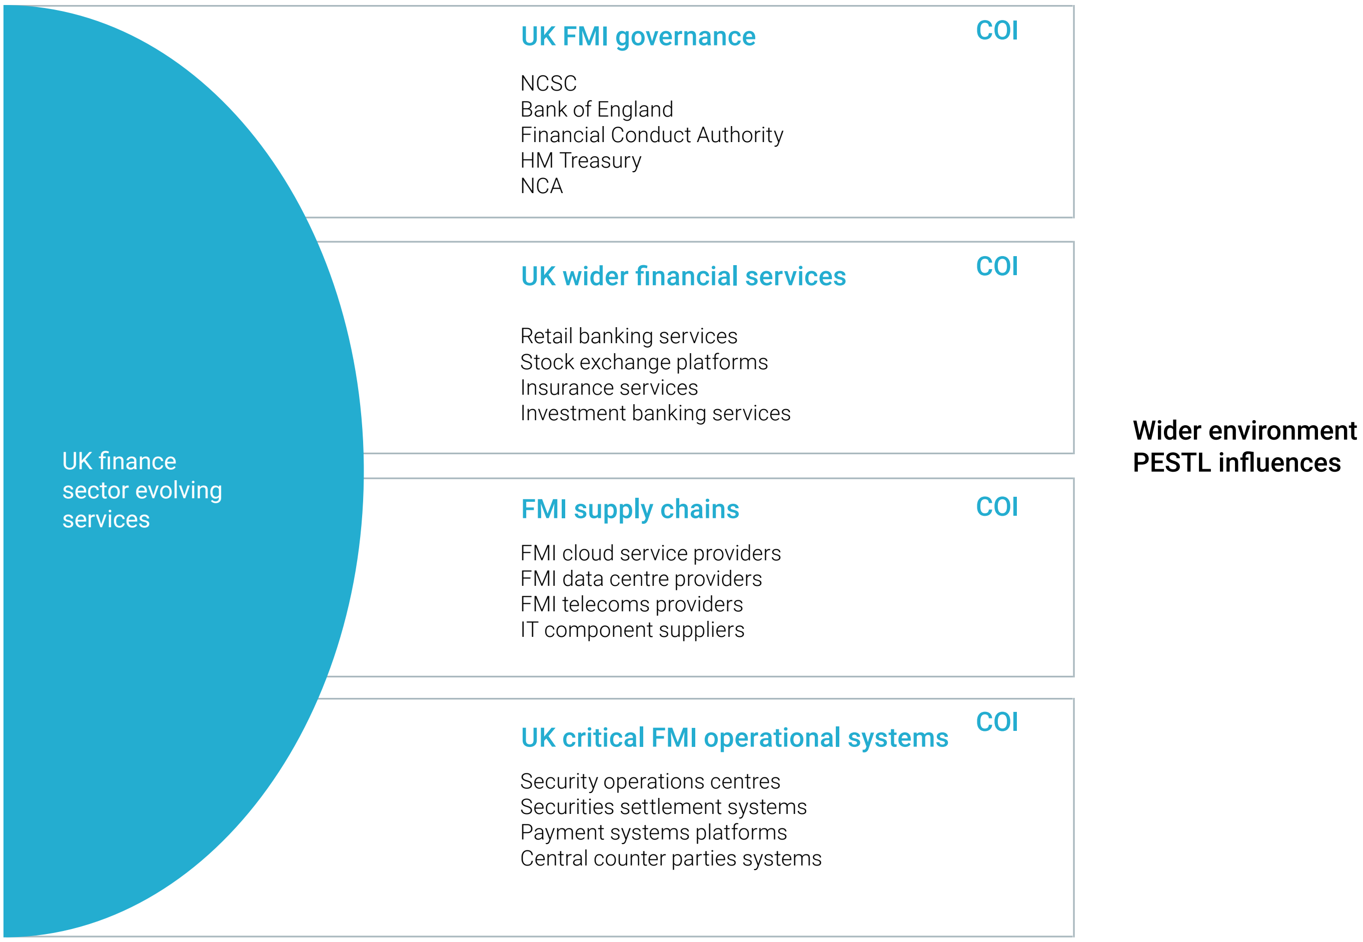

UK finance sector evolving services

UK FMI governance — COI

NCSC
Bank of England
Financial Conduct Authority
HM Treasury
NCA

UK wider financial services — COI

Retail banking services
Stock exchange platforms
Insurance services
Investment banking services

FMI supply chains — COI

FMI cloud service providers
FMI data centre providers
FMI telecoms providers
IT component suppliers

UK critical FMI operational systems — COI

Security operations centres
Securities settlement systems
Payment systems platforms
Central counter parties systems

Wider environment PESTL influences

Figure 5 – The UK FMI ecosystem domain model

THE UK CRITICAL FMI OPERATIONAL SYSTEM COI

The COI depicted by the left hand domain of Figure 5 comprises the operational systems of each of the critical members of the UK FMI. In this high-level model, these critical FMI members fall into three types:

- Central counterparties (CCPs)
- Payment systems
- Securities settlement systems

These critical infrastructures represent a single point of failure, therefore any cyber attack able to successfully disrupt them for a significant period of time would potentially have a systemic impact on the UK economy.

THE FMI CENTRAL COUNTERPARTIES

The FMI central counterparties (CCPs) shown in Figure 5 are recognised clearing houses supervised by the Bank of England. There are three critical UK CCPs:

- **ICE Clear Europe Limited**, which clears a range of exchange-traded derivatives and OTC credit default swaps. It handles half of the world's oil futures contracts, and is typically the only place to go for those linked to the Brent benchmark

- **LCH Limited**, which clears a range of repos, exchange-traded and OTC securities, and derivatives. As London's largest clearing house, it clears interest-rate swaps with a notional value of over $340 trillion, up to 95% global total

- **LME Clear Limited**, which clears a range of metal derivatives traded on the London Metal Exchange and OTC metal contracts. As part of the London Metal Exchange, it is virtually irreplaceable for precious metals traders

Given the vast amounts of derivatives that these clearing houses handle, and that much of the trading cannot be done elsewhere, they are critical components of the UK FMI ecosystem.

THE PAYMENT SYSTEMS

Payment systems that meet defined criteria may be recognised by HM Treasury, and supervised according to the Banking Act 2009. The payments system being:

- BACS, operated by BACS Payment Schemes Limited (BPSL), processes payments of varying values, and is often used for PAYE, direct credit and direct debit payments

- CHAPS, operated by the Bank of England, is the United Kingdom's high-value payment system, providing real-time gross settlement of sterling transfers between participants

- CLS operates the world's largest multi-currency cash settlement system for foreign exchange transactions in 18 currencies, including sterling

- Faster Payments Service (FPS), operated by Faster Payments Scheme Limited (FPSL), processes standing orders and electronic retail transactions, including transactions generated through internet, mobile and telephone banking

- LINK is a network of card issuers and cash machine (ATM) providers that enables cardholders to withdraw cash at any LINK-connected ATM

- Visa Europe is a four-party cards payments processor operating in the EEA, Israel, Turkey and Switzerland, and offers debit, credit, deferred debit and prepaid card products

THE SECURITIES SETTLEMENT SYSTEMS

Securities settlement systems may be regulated under the Financial Services and Markets Act (FSMA) as recognised clearing Houses (RCHs) and are subject to the Uncertificated Securities Regulations 2001 in the UK.

Euroclear UK and Ireland Limited (EUI) operates the Certificateless Registry for Electronic Share Transfer (CREST) system, which is also a recognised payment system under the Banking Act 2009. CREST is the securities settlement system for UK gilts and money market instruments, as well as UK equities, which settles on a gross delivery versus payment basis. EUI also operates CREST for the purposes of settling Irish equities.

THE UK FMI GOVERNANCE COI

The UK FMI governance COI depicted in Figure 5 comprises the key UK government agencies responsible for, or involved in, the regulation and governance of UK financial services. The focus here is on the governance of the critical FMI entities.

The Bank of England (the bank) is committed to ensuring UK FMIs are operating safely, and to working co-operatively with global regulators. Supervision of FMIs is central to the bank's objective of maintaining financial stability, and the bank seeks to ensure that FMIs are operating in a safe way and finding ways to reduce systemic risks.

The bank has developed a supervisory approach based on the Principles for Financial Market Infrastructures (PFMI). These principles set out the international standards that FMIs should follow for governance arrangements, financial resources. and the management of certain types of risk.

The Bank of England co-operates closely with both the Financial Conduct Authority (FCA) and the Payment Systems Regulator (PSR), in relation to supervising market infrastructure and payment systems respectively. The Bank of Engalnd considers its CBEST framework (controlled, bespoke intelligence-led cyber security testing) to be an important addition to FMI in-house testing and vulnerability assessment.

CBEST testing replicates the behaviours of threat actors identified by government and commercial intelligence providers as posing a genuine threat to systemically important financial institutions.

The UK financial authorities (HM Treasury, the Bank of England and FCA) have a single mechanism to coordinate a response to incidents that have affected, or have the potential to affect, the financial sector. This also includes the National Cyber Security Centre (NCSC) and, when appropriate, the National Crime Agency (NCA) for cyber incidents.

THE WIDER FINANCIAL SERVICES COI

The wider financial services COI comprises the many diverse entities that provide additional financial services outside the core critical FMI services identified. Members of this COI include large investment banks, retail banks, insurance providers and a rapidly growing number of fintech companies.

The entities in this COI will typically use critical FMI operational services (for example, retail banks use CCPs), although many fintech companies also provide peer-to-peer (P2P) services, such as loans that do not require established financial institutions to act as an intermediary.

THE FMI SUPPLY CHAIN COI

The FMI supply chain COI shown in Figure 5 comprises the global supply chains associated with the operational systems of members of both the critical FMI and wider financial services COIs. Members of the supply chain COI provide the hardware and software components of financial service providers' operational platforms, as well as FMI telecommunications services, cloud services and data centres.

The Bank of England has set standards for its suppliers that effectively apply to the FMI supply chain COI, requiring compliance with data management, physical, operational and cyber security requirements.

COMPLEXITY AND EVOLUTION OF THE UK FMI ECOSYSTEM

The FMI ecosystem is a complex system of systems already exhibiting emergent behaviours, such as bubbles in the financial markets and flash crashes caused by the interaction of multiple high-speed trading algorithms. The stability of the FMI relies on a simple concept: trust. If trust in financial services is lost, then those services rapidly lose their viability. The increasing complexity of the UK finance cyber ecosystem makes cyber resilience a fundamental requirement for maintaining that trust. This fundamental relationship is core to managing cyber risk to the UK FMI ecosystem.

The services provided by the FMI ecosystem are undergoing a transformation that is being driven by a number of emerging technologies at various stages of maturity, including:

- Artificial intelligence and machine learning
- Big data analytics
- Blockchain-based distributed ledgers
- Open APIs
- Cloud services

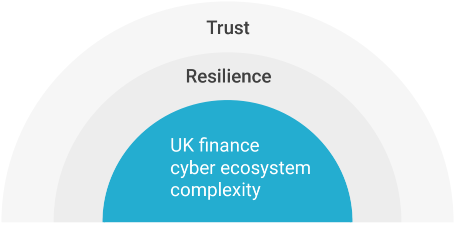

Figure 6 - Trust in relation to cyber resilience and cyber complexity

These technologies are driving new disruptive fintech services, often driven by SME startups. These financial innovators are using new technologies to provide more collaborative, agile and customer-centric financial services.

This paper explores how emerging technologies can bring with them not only significant potential benefits, but also a whole new threat surface, with associated cyber risks. This is especially valid in the case of for established financial institutions migrating from major legacy IT systems.

Figure 7 provides a high-level view of the conceptual architecture of the UK FMI ecosystem.

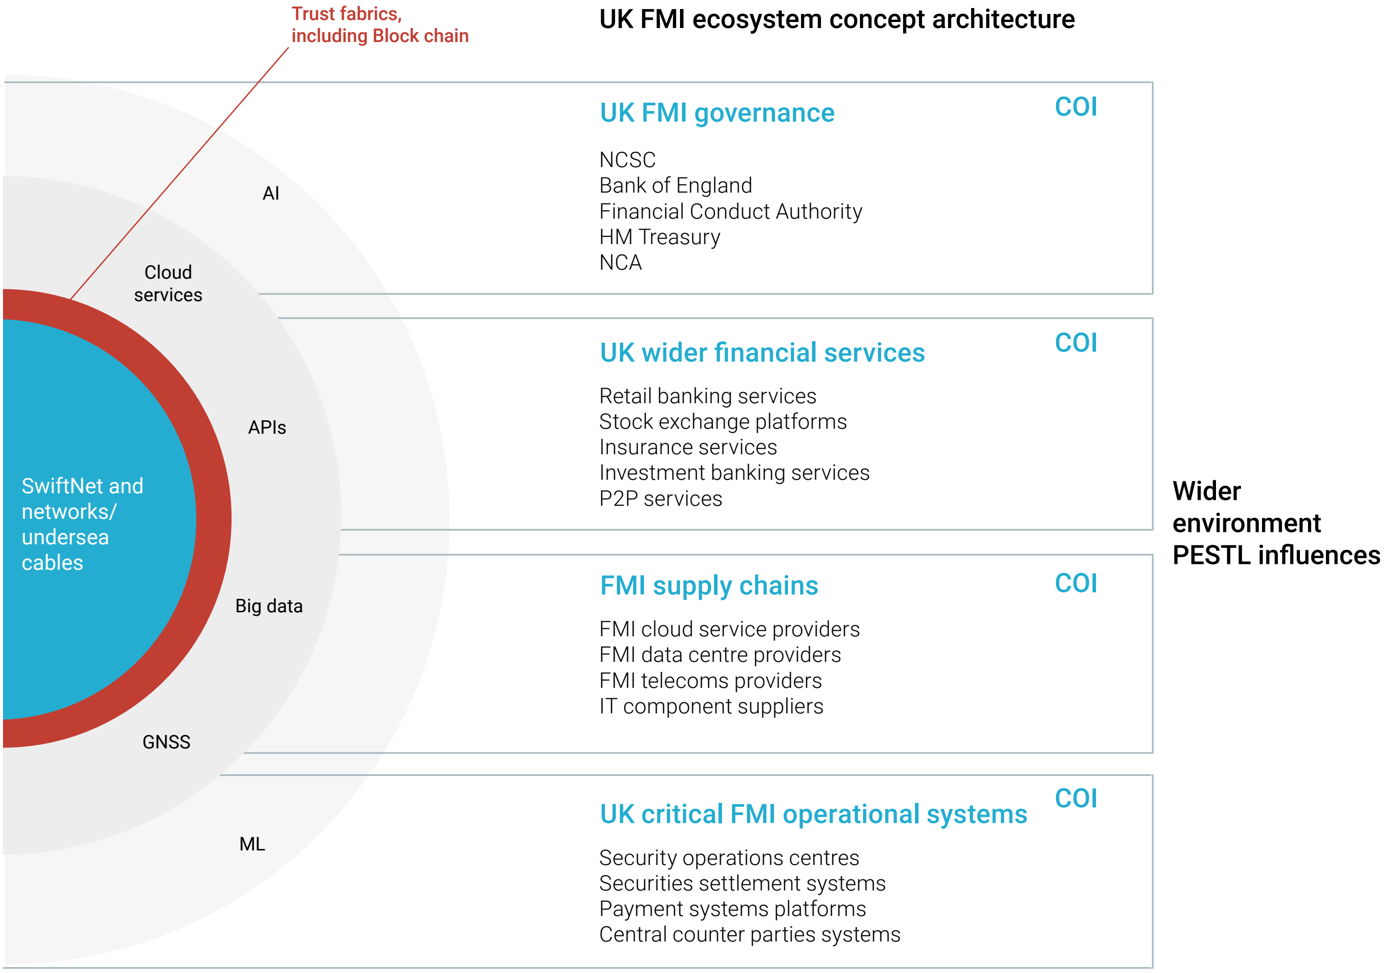
Trust fabrics, including Block chain
AI
Cloud services
APIs
SwiftNet and networks/ undersea cables
Big data
GNSS
ML
COI
COI
COI
COI
Wider environment PESTL influences

UK FMI ecosystem concept architecture

UK FMI governance

NCSC
Bank of England
Financial Conduct Authority
HM Treasury
NCA

UK wider financial services

Retail banking services
Stock exchange platforms
Insurance services
Investment banking services
P2P services

FMI supply chains

FMI cloud service providers
FMI data centre providers
FMI telecoms providers
IT component suppliers

UK critical FMI operational systems

Security operations centres
Securities settlement systems
Payment systems platforms
Central counter parties systems

Figure 7 - UK FMI ecosystem conceptual architecture view

The core communication infrastructure

As shown in Figure 7, the physical network lies at the heart of the FMI ecosystem, as the infrastructure through which all financial transactions between its members are transmitted.

The fibre optic cable links connecting FMI sites within the UK form one component of this physical network infrastructure. However, the most important component is the undersea fibre optic cable network that provides high data volume low-latency connectivity to the rest of the world, and the global financial market infrastructure.

Above this low-level physical network infrastructure sits the international interbank messaging and routing system (MSR), SwiftNet. SWIFT is the Belgium-based co-operative society that links over 11,000 financial institutions, including 193 central banks, in more than 200 countries. There are also low latency IP networks, for example, Secure Financial Transaction Infrastructure (SFTI).

The trust fabric

Another part of the communications core is the trust fabric, which is represented conceptually by the red semi-circle in Figure 7. There are a number of components to the trust fabric, including cryptographic capabilities for aspects of authentication, confidentiality and data integrity trust.

However, there is also implicit trust: for example, in the way SwiftNet trusts participating financial institutions and provides no further authorisation control on payment messages entering or exiting the SWIFT network. Block chain is also introducing disruptive new trust models that remove the need for intermediaries, such as cryptocurrencies like Bitcoin.

Disruptive technology services

The conceptual architecture of the communication core of the UK FMI ecosystem also features a layer of technology capabilities highlighted by the darker grey semi-cirle in Figure 7. These FMI supporting technology capabilities include:

- Cloud services
- Open APIs for collaborative financial services
- Big data analytics services
- GNSS (timing services)

Cloud services

Financial services firms are rapidly adopting cloud services, which provide a powerful set of tools to manage data needs. This shift introduces new opportunities for combining public and proprietary data into big data, which can be used to generate innovative new analytical insights. Over time, increasing numbers of core services are likely to migrate to cloud hosting.

Open APIs

The use of open APIs has introduced a new era in financial services. However, the security of API-based collaboration between established, evolving and new financial entities cannot be maintained and managed by financial institutions alone - it requires a collaborative effort across the entire value chain.

GNSS

Accurate timing (down to the millisecond) is a fundamental requirement underpinning most FMI ecosystem transactions, including high-speed trading and settlement integrity. Financial institutions are legally committed to record their operations against a consistent and accurate timescale.

Banks:
Global navigation satellite system (GNSS) is used for time stamping functions, to log events in a chronological manner and therefore be able to recreate causal links.

Stock exchanges:
Sock GNSS plays a key role in the way that stock exchange servers apply time stamps to the trades they execute and the quotes they issue.

Use of big data analytics in the FMI ecosystem is enabling transformation in a number of areas, including fraud detection and investigation, and trading and investment decisions. However, the emerging technologies truly delivering such transformation are the AI and ML algorithms that feed on big data.

Artificial intelligence (AI) & machine learning (ML)

The use cases for AI and ML are constantly changing, but banks today are focusing on three main applications:

- Building a better customer experience
- Reducing costs, not headcount
- Streamlining risk operations

AI and ML can be developed to exploit big data to make informed real-time investment/trading decisions based on not only buying and selling prices data but on diverse and related data, such as socio-political trends.

However, although AI and ML can provide financial entities with many advantages, it also presents new challenges. The governance of and reporting on AI and ML driven transactions that can out-perform human understanding in ways that are opaque as is the case with deep neural networks is one such challenge.

While banks are using AI and ML to improve their identity authentication processes and to better detect suspicious activity, malicious actors are also using AI to create new cyber threats, for example, by injecting biased data into the training sets of ML algorithms that can then be exploited by an attacker.

APPLYING THE BHI USING AN ILLUSTRATIVE CYBER ATTACK SCENARIO

The BHI approach models the growth of benefits and harm in the context of complex cyber ecosystems. The high-level model of the UK FMI ecosystem can be used as a context for showing how this approach can be applied.

The first step is to model the growth of benefits over a time period. The time period selected here is 2019 to 2030, which corresponds to the period during which the adoption emerging technologies will complete the digital transformation of financial services.

A simple model of the overall benefits growth associated with the digital transformation of the FMI is shown in Figure 8, and assumes that the benefits grow as a Bass diffusion in line with the projected rate at which emergent technologies are adopted into the UK FMI ecosystem.

The second step is to model the growth of risks (as a combination of likelihood and adverse impact) over that same period. To do this we use an illustrative multi-vector cyber attack scenario on the UK FMI ecosystem and explore its associated risks as a function of time during the evolution of the ecosystem.

Firstly, the evolution of the risk likelihood (threat level) is explored by assuming the threat source is a nation state (with associated capabilities); by factoring in the growth of the threat surface (vulnerabilities such as opportunities) over time; and by modelling variations in PESTL influences, such as motivation. Figure 8 illustrates this risk likelihood in red for example, the threat level assuming constant motivation but an exponential growth in vulnerabilities during the transformation period.

The potential impact is then modelled by exploring the potential of the cyber attack to generate a cyber chain reaction that poses a systemic risk to the UK. A systemic risk is generally seen as the potential for a major financial crisis adversely affecting the real economy. The vulnerability level and stochastic nature of the UK FMI ecosystem during the transition period exposes it to so-called 'black swan' events that can result in systemic impacts.

In 2017, the financial services sector contributed £119 billion to the UK economy, 6.5% of total economic output. The sector was largest in London, where 50% of the sector's output was generated. Any successful cyber attack generating a systemic impact to the UK FMI ecosystem would result in a downturn in its contribution to the UK economy, which could easily result in a £ multi-billion loss to the UK. This does not take into account impacts on intangible assets, such as brand equity, reputation and trust.

The third step in the BHI process, therefore, evaluates the difference between the growth in benefits and growth in harm during the transformation period.

Bass diffusion distributions of UK FMI distruptive technologies benefit growth and the associated growth of liklihood of systemic cyber attacks

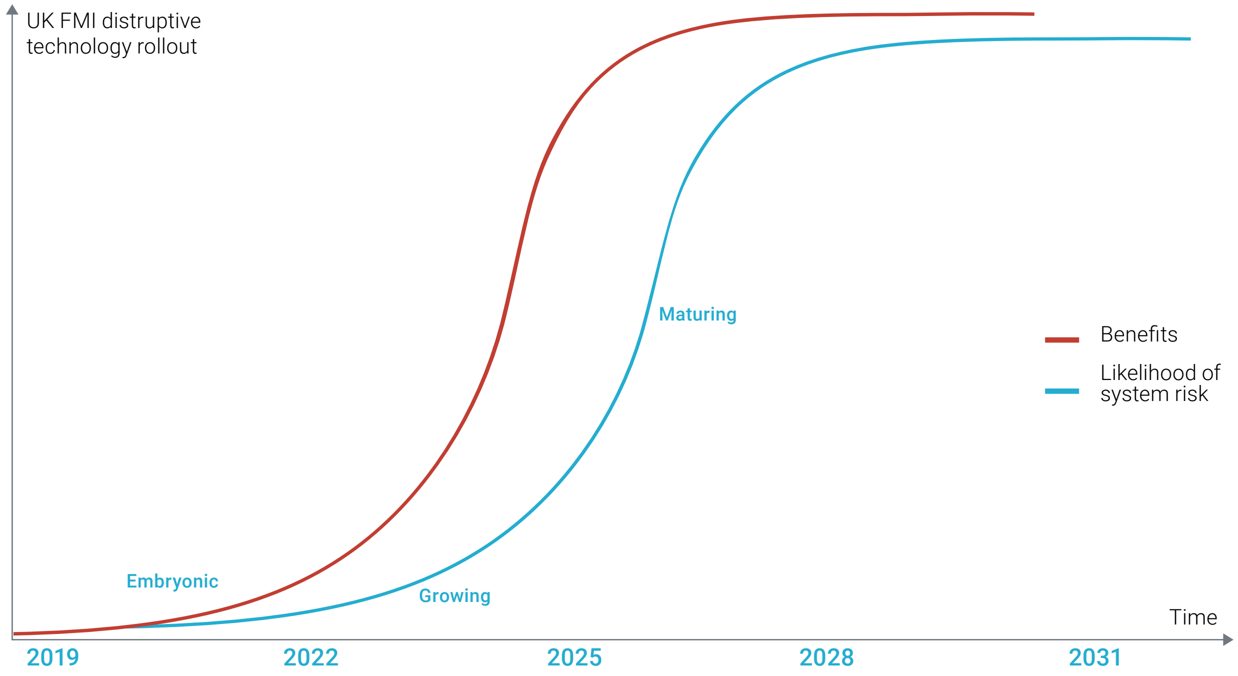

Figure 8 - Model of the growth in FMI benefits generated by disruptive technologies, versus the associated growth of likelihood of systemic risk

THE ILLUSTRATIVE CYBER ATTACK SCENARIO

The illustrative and hypothetical cyber attack scenario assumes that the threat source is a nation state entity, and for this sake of this exercise that fictional role has been given to Russia and its Federal Security Service (FSB) acting through one or more proxy APT groups, including APT 39 in Iran.

The potential capability of this nation threat source is high, although the actual motivation to carry out a major cyber attack on the UK FMI ecosystem is assumed by default to be low. This fictional scenario is of a hypothetical escalation of any Anglo-Russian geo-political tension. This hypothetical escalation is imagined to be a result of Russia and Iran being targeted by America and the UK with a major new round of sanctions in the year 2020, as a result of an escalation in the weaponisation of gas and oil supplies.

The cyber attack scenario models here are based on a sustained multivector cyber attack targeting the UK FMI ecosystem with the objective of causing a systemic impact on the UK. In NCSC terms, this would equate to a category 1 or 2 cyber attack, as defined in Table 2.

The threat source intends to cause a systemic impact on the UK FMI ecosystem by causing a breakdown of trust in the FMI in the area of energy trading, and in particular oil futures contracts. The UK FMI entity being targeted to achieve this is ICE Clear Europe.

From an external perspective, ICE appears to operate strong cyber-resilient critical FMI services, making it a good example for this hypothetical scenario, as its defences will be typical of all critical FMI services providers.

ICE Futures Europe, formerly the International Petroleum Exchange (IPE), was formed in 1980 and is the home of the benchmark Brent and Gasoil futures and options contracts. ICE Clear Europe provides clearing services for futures and options contracts traded on ICE Futures Europe, ICE Endex, and ICE Futures US Energy Division; it handles half of the world's oil-futures contracts; and is typically the only place to go for those linked to the Brent benchmark.

As part of the threat intelligence approach, it is noted that Russia views cyber attacks as a sub-component of information warfare, which covers a broad range (including computer network operations, electronic warfare, psychological operations and disinformation operations). The fictional cyber attack scenario has been designed around this broader multi-vector approach.

The first cyber attack vector will target service disruption and the timing integrity of ICE trading transactions. It exploits the physical vulnerability of the fibre optic cables serving the ICE Clear Basildon data centre located in Langdon Hills. Once these ducts leave the site boundary, they merge with the main multi-tenant cable ducts running along nearby main roads. By cutting all the cables in these ducts at the nearest road-based maintenance point (for example, by using fake telecoms provider vans or roadworks as a cover for the operation) the ICE Clear Basildon data centre can be effectively put offline for a period of hours. Microwave channels are not attacked in this scenario, since they have limited data capacity.

Category level	Category definition	Who responds?	What do they do?
Category 1 National cyber emergency	A cyber attack which causes sustained disruption of UK essential services or affects UK national security, leading to severe economic or social consequences or to loss of life.	Immediate, rapid and co-ordinated cross-government response. Strategic leadership from ministers/Cabinet Office (COBR), tactical cross-government co-ordination by NCSC, working closely with law enforcement.	Co-ordinated on-site presence for evidence gathering, forensic acquisition and support. Co-location of NCSC, law enforcement, lead government departments and others where possible for enhanced response.
Category 2 Highly significant incident	A cyber attack which has a serious impact on central government, UK essential services, a large proportion of the UK population, or the UK economy.	Response typically led by NCSC (escalated to COBR if necessary), working closely with law enforcement (typically NCA) as required. Cross-government response coordinated by NCSC.	NCSC will often provide on-site response, investigation and analysis, aligned with law enforcement and criminal investigation.

ICE Clear Europe's timing integrity is provided at source from the National Physical Laboratory (NPL) in Teddington, UK, which provides a certified precise time signal by fibre, directly traceable to co-ordinated universal time (UTC) and independent of GPS, which is susceptible to vulnerabilities. The NPL timing service underpins time stamping, latency monitoring and synchronisation, in compliance with the Markets in Financial Instruments Directive (MiFID II). This level of precision also aids forensics and audit, thereby improving the functioning of financial markets and strengthening investor protection.

During the downtime created by the initial hypothetical cyber attack, GPS spoofing attacks would also be used across London to corrupt timestamps, and to corrupt any residual processing going on inside the Basildon data centre (assuming there is a fallback to GPS, as the NPL time source will also be/ go down). These secondary attacks would also spoof at least two of the three legacy GPS time references used by the SFTI network.

The Basildon data centre is part of a global network with built-in resilience, including another primary ICE Clear data centre in Chicago and a secondary data centre in Atlanta. In a disaster recovery situation, the Basildon centre will attempt to failover to the secondary data centre in Atlanta.

In order to disrupt failover activities associated with the Basildon data centre being taken offline for at least two hours, the hypothetical threat actor also takes down some of the key transatlantic cables during this period. By targeting key cables, the threat actor could cause re-routing of massive volumes of internet traffic, which would severely impact latency and slow down the Basildon data centre's failover to the Atlanta data centre.

In this hypothetical scenario, the target undersea cable system is GTT Express (formerly Hibernia Express). Other targets could include America Europe Connect (AEC-1 and AEC-2) and MAREA.

GTT Express is a 4,600km and 6-pair transatlantic submarine cable system linking North America and the United Kingdom. It was built with the state-of-the-art submarine network technology, specifically designed for the financial community, and offers the lowest latency route from New York to London (with a 58.55 millisecond round trip delay). Russia operates a number of very small, nuclear powered submarines that are capable of diving in excess of several thousand metres; this includes the AS-12 Losharik deep-sea submarine. In our fictionalised attack scenario, Russia uses this capability to cut the target undersea cables at a significant depth, making repairs more time-consuming. For the sake of our example, it can safely be assumed that Russia had already located and mapped all key transatlantic undersea cables using a combination of easily-obtained maps and high resolution sonar.

The resulting disruption to physical networks would keep the various FMI security operations teams busy trying to restore and maintain ICE Clear's services, and would be likely to involve government agencies, such as the NCSC. These initial attacks are designed to disrupt FMI services and lower confidence, but while they are significant, they are not true systemic attacks.

However, the second hypothetical attack vector has the potential to escalate FMI disruption to a systemic level. It would be launched during the disruption resulting from the network attacks. This vector exploits vulnerabilities in the end-to-end SwiftNet interbank payment system, and would transfer large amounts of money to a number of innocent trading parties and targeted individuals.

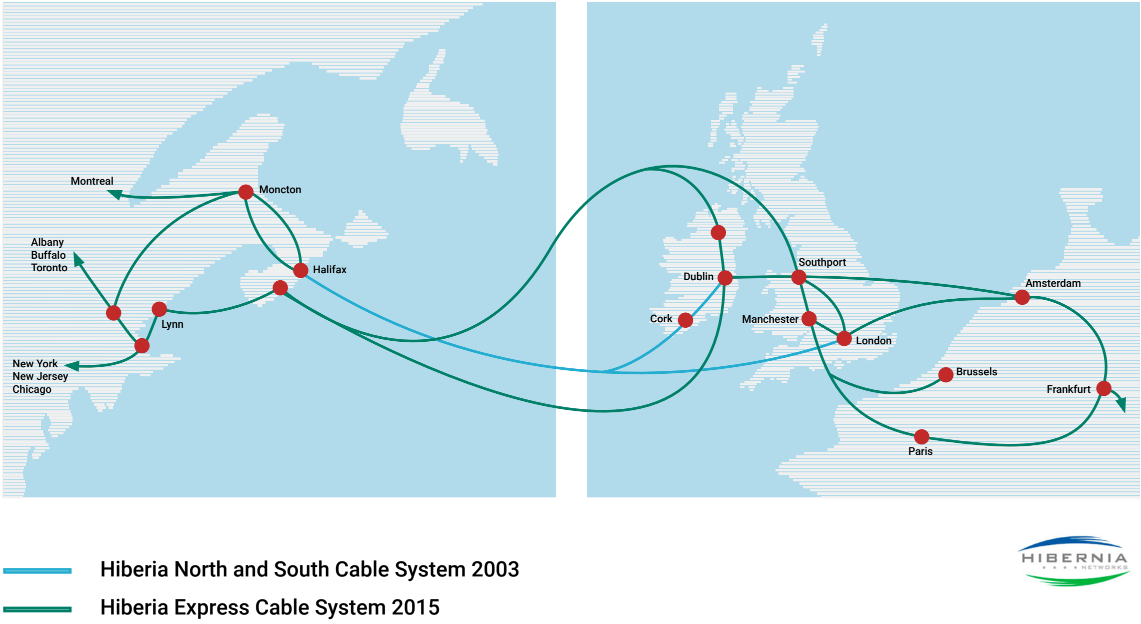

The purpose of such an attack would not be financial gain but to support a negative news narrative designed to destroy trust in the energy/oil futures markets. The narrative could be one of many variations, based on false accusations being made about key individuals and institutions. In this fictionalised scenario, the narrative is that certain governments are manipulating the price of oil through trading parties such as hedge funds and large banks.

The hypothetical disinformation campaign starts by spreading rumours that oil prices are being manipulated by named parties, and associating this fictional activity with the service disruptions. The attackers then build on this narrative by revealing evidence based on the funds transferred through rogue SwiftNet transactions to these innocent parties.

Eventual rejection of this false narrative would come at a cost, since it would involve a public admission that all the rogue transactions had been created through a successful cyber attack on SwiftNet. This admission would, in turn, lower the level of trust in the UK FMI.

EXPLORING THE VULNERABILITY AND CONTROL ASPECTS OF THE SYSTEM

The complexity of the FMI ecosystem is growing as it undergoes digital transformation driven by the emerging technologies described in Section 2 (such as AI and ML, big data analytics, block chain, open APIs and cloud services).

The vulnerability of the FMI ecosystem is increasing in line with its complexity. This is driven in part by the rapidly evolving interconnectivity and collaboration enabled by these technologies. The illustrative cyber attack scenario shows how these evolving vulnerabilities can be exploited in different ways, via multi-vector attack paths. The vulnerabilities associated with the physical network and timing attacks, for example, are relatively well known and constant.

The second example, the hypothetical attack on SwiftNet, exploits the SwiftNet end-to-end trust model. SwiftNet is vulnerable because when Bank A sends a transaction through the cross border/international payments system (for example, SWIFT) to Bank B, there are no security controls or authentication checks carried out by the SWIFT Messaging and Routing System. SWIFT trusts Bank A and assumes Bank A has its own solid security controls. This is clearly not a great trust model, given that the vulnerabilities associated with the evolving IT systems of banks and other financial entities that use SwiftNet are growing exponentially during digital transformation.

In 2016, this SWIFT trust model was exploited by cyber attackers, who compromised the in-house security of the Bank of Bangladesh. They then used SWIFT to submit a number of fraudulent payment orders through its accounts in the Federal Reserve Bank of New York, totalling $951 million. Though most of this money was recovered, the attackers successfully laundered $81 million through parties in the Philippines. The SWIFT Customer Security Programme (CSP) and Relationship Management Application have since been put in place to help mitigate this trust model's vulnerability.

The vulnerabilities of the IT system of banks and other financial entities using the SWIFT network is growing as they undergo digital transformation, as shown earlier in Figure 8. Therefore, in our hypothetical example the vulnerability to our SwiftNet attack vector (and in particular to APTs) is growing exponentially during the period 2020 to 2025.

In our fictional cyber attack scenario the ATP attack vector targets the points where the SWIFT system interfaces with other applications within the participant bank. Our ATP's TTPs are designed to detect and compromise defences such as deception traps and tokens.

Classic risk mitigations work on the basis that a system can be controlled in the presence of threat actors, reducing or removing the threat. However, as shown in our example, systems that support business ecosystems, and the information and computer technologies that support a given organisation's activities are increasingly complicated.

Vulnerability Level (VL)	Threat class	Attacker's control	Economic rationale
● 4	Stochastic system	The nature of the system is such that it cannot be controlled, but vulnerabilities can be reliably modelled using closed-form probability distributions over a fixed (and finite) set of state variables in the system's phase state space.	Radical ignorance: black swan events may occur, as preparation for such events is frequently hindered by an assumption of knowledge of all the risks. Scenario modelling using Shackle's potential surprise.

Table 3 - VL4

Control of small systems is a mature discipline: controllability of linear systems is well understood, and understanding for non-linear systems has been developing steadily. In contrast, control of complex systems is poorly understood, and mostly poorly characterised. A threat actor can leverage this lack of knowledge to cause harm to a system in ways that a defender cannot control through prior mitigation.

In the BHI model, the VL of a system – of a given scope and phase space with a given resolution – is a measure of its intrinsic lack of controllability, from the perspective of the defenders who legitimately operate the system.

In the scenario here in 2020 to 2025, the vulnerability level of the UK FMI ecosystem is at level 4, in line with its level of complexity. VL4 is shown from a control perspective in Table 3.

Vulnerability level 4 is representative of the fact that the UK FMI operators, and in particular banks and other entities that use the FMI, have no actual knowledge of, for example, the 'zero day' threats in the armoury of nation state-sponsored APT groups. The expanding threat surface associated with the emerging/collaborative technologies driving digital transformation makes it difficult, if not infeasible, to detect all such latent threats.

This lack of knowledge makes risk decisions far less certain even than gambling, as at least a gambler knows the odds against success. In the context of cyber threats, it is the attacker who holds the knowledge. In other words, the knowledge status supporting risk decisions has moved from rational ignorance to one of radical ignorance.

EXPLORING THE LIKELIHOOD OF AN ATTACK SCENARIO

Classic risk assessments model the likelihood of a cyber attack on a particular target of interest in terms of a threat level assessment at a given point in time. The threat level is typically modelled as a function of the capability of the threat sources/actors and the level of motivation and priority for attacking that target of interest.

Threat level = F(capability of threat source/actor(t), motivation/priority(t))

The capability of the threat source and associated actors in our hypothetical example scenario are those associated with a nation state, in this case the capabilities of the FSB and their APT groups. The capability of such nation state actors for launching sophisticated cyber attacks is high.

The attack vector in the hypothetical scenario exploits significant vulnerabilities (attack opportunities) within the UK FMI ecosystem, and these are growing exponentially in line with digital transformation and the explosive growth in the threat surface. Such vulnerabilities are relatively easy to exploit, so the likely capability of the threat actors relative to the difficulty in exploiting them is high.

Readiness of latent zero day threats to the UK's CNIs would give any hostile nation state the potential to launch a cyber attack with a significant socio-economic impact on the UK.

Therefore, the likelihood of the UK energy grid ecosystem entities potentially being compromised via such latent zero day back doors is very likely.

The likelihood of an actual attack being executed that exploits (and thus exposes) any zero day vulnerability would depend on motivation and priority, which themselves would be driven by the state of the geopolitical relationship between the UK and the hostile nation state in question.

EXPLORING THE POTENTIAL IMPACT/HARM OF THE ATTACK SCENARIO

The illustrative scenario hypothesised in this report assumes that the Russian FSB and their APT groups launch an attack that exploits insider threat and latent zero day vulnerabilities to compromise UK FMI ecosystem components. The objective is to cause economic damage to the UK as part of a geopolitical weaponisation of energy supplies campaign that begins to escalate in the year 2020.

The levels of harm involved are modelled using the examples of the impact on the UK FMI ecosystem, as shown in Table 4.

Level	Impact
🔴 Very high	Overall capability of the UK FMI ecosystem brought to a halt. Significant socio-economic scale disruption. High impact on all intangible assets. Systemic impact (for example, NCSC Category 1 cyber-attack, national cyber emergency).
🟠 High	Total disruption of one or more UK FMI domains (for example, a clearing house, or payment or settlement systems). Significant impact on most intangible assets. Systemic impact (for example, NCSC Category 2 cyber attack)
🟡 Medium	Localised significant intra-FMI domain operational disruption. Minor UK-wide disruption of overall UK FMI ecosystem operations. Minor impact on most intangible assets (for example, NCSC Category 3 or 4 cyber attack)
🟢 Low	Localised intra FMI domain short term operational disruption (for example, NCSC Category 5 cyber attack)

Table 4 - UK FMI ecosystem level impact levels

As shown in Table 4, when assessing the impact of a successful cyber attack on the UK FMI ecosystem, the potential harm to both tangible and intangible assets must be included. For example, brand equity can be lost as a result of the reputational damage caused by succumbing to a successful cyber attack.

If our hypothetical threat actors were to successfully launch the example multivector cyber attack scenario in the year 2020, the impact of such an attack, if successful, could include the following:

- Reputational damage to many of the players in the oil futures market, in particular clearing houses
- Loss of trust in UK FMI and its counterparties systems
- Potential migration of energy/oil futures market exchanges out of the UK, thus impacting the UK economy

A key aspect of this cyber attack scenario is that the impact is on intangible assets - primarily trust. That loss of trust however, is designed to lead to a tangible high level impact on the UK economy. However, the impact level would be different at different points in time, as would the motivation of the attacker. For example the impact could potentially be very high if carried out aggressively in 2025, when there will be significantly more collaborative interconnectivity within the FMI ecosystem. This would enable the attacker to broaden the attack and cause both planned and unplanned cyber chain reactions to propagate across all forms of financial services.

Having explored the illustrative cyber attack scenario in classic risk assessment terms, it can now be explored from the perspective of the BHI.

THE BENEFIT HARM INDEX PERSPECTIVE ON OUR SCENARIO

The overall socio-economic benefits of the UK FMI ecosystem grow over time in line with a Bass diffusion distribution, as shown earlier in Figure 9. As shown in the hypothetical cyber attack scenario, the harm which can be inflicted on the ecosystem by a specific threat can also grow with time, and the associated threat level will vary with time.

Benefits are defined in terms of the positive business and socio-economic impacts multiplied by their likelihood. Harm is defined in terms of the negative business and socio-economic impacts multiplied by their likelihood. A simple discrete formulation of how to calculate the associated growth is shown below:

$$B(t_{n+1}) = B(t_n) + b(t_n) * P_b(t_n) - h(t_n)*P_h(t_n)$$

Where $P_h(t_n)$ is proportional to the Threat Level

Threat level = F(capability of threat source/actor(t), motivation/priority(t))

The benefit harm index relates to differences in the complexity levels of benefit (CLb), and harm (CLh), over a time interval, TIi assuming M distinct threats (j) where j ranges from 1 to M.

BHI = CLb(TIi) - CLh(TIi)

Where:

CLb(TIi) = MAX(Level(Distribution(b(TIi)), Level(Distribution(Pb(TIi))))

CLh(TIi) = MAX(Level(Distribution(h(TIi))), MAX(over all j, {Level(Distribution(capability(j, TIi)), Distribution(priority(j, TIi)))))

In simple terms, for the hypothetical cyber attack scenario on the UK FMI ecosystem there is an overall set of socio-economic benefits that are growing in line with a Bass diffusion distribution curve, as described earlier in Figure 8.

- During the strong growth period 2020 to 2022, the benefit growth rate is embryonic, equating on average to a benefit complexity level of 2 (see Figure 3)
- During the period 2023 to 2028, benefit growth is exponential, equating to complexity level 4
- During the period 2029 to 2031, the benefit growth rate decreases rapidly from exponential to asymptotic, which equates to a benefit complexity level 4 decreasing to 0 during this period

Given the UK government's strategy of ensuring that the UK remains a leading player in financial services, we can assume that the associated probability of following that distribution is high and flat, so for simplicity it is assumed that it is close to 1. The accuracy of the market forecasts is assumed to be high.

If initial assumptions are that the value of benefits and harm over each interval are the same then, in effect, there is a focus on the difference in the growth rates of benefit and risk likelihood, rather than on the actual quantitative benefit and harm multipliers.

For the period 2020 to 2022: **CLb(2020-2022) = 2**

For the period 2023 to 2031:
CLb(2023-2028) = 4 CLb(2029-2030) = 2, CLb(2030-2031) = 0.,

For the hypothetical cyber attack scenario, an attack threat level (likelihood) has been selected which also grows with a Bass diffusion distribution curve. However, the Bass diffusion curve for potential harm grows in advance of the benefits Bass diffusion curve, since nation state threat actors will be targeting the UK's critical national infrastructure, by creating an arsenal of zero day threats for each CNI as part of their cyber warfare readiness capabilities.

The potential for harm associated with the hypothetical threat scenario will grow exponentially (for example, complexity level 4) in line with the growth in complexity of the UK FMI ecosystem during the digital transformation period.

This high level of complexity is associated with the explosive growth in the size of the overall threat surface for the selected attack vector. It can be expected that any nation state threat actors would exploit by further developing their arsenal of associated zero day threats.

As already mentioned, the motivation priority in the hypothetical threat scenario is low until the year 2020 when it becomes high, taking the threat level (likelihood) to a very high value.

So for the period 2020 to 2025: **CLh(2020-2025) = 4**, which reflects the strong growth phase of Bass diffusion distribution of the growth of potential harm. For the period 2026 to 2031 the exponential growth in potential harm asymptotically decreases. The resulting BHI values simply show the difference in growth rate, as outlined in Figure 8: the potential harm is growing faster than the potential benefit in the earlier period 2020 to 2022, before evening out.

These calculations assume that the level of benefit and the level of harm were of equal magnitude for each time interval. However, the level of systemic harm that can be inflicted on the UK FMI is largely relative to the incremental growth in benefit generated by the early stages of the digital transformation of the FMI ecosystem.

Although there is no formal quantitative analysis in this paper, the value the UK financial services sector contributed to the UK economy in 2017 was £119 billion, 6.5% of total economic output.

Any single successful cyber attack generating a systemic impact on the UK FMI ecosystem would result in a downturn in its contribution to the UK economy, which could easily result in a multi-billion pound loss to the UK. This does not take into account impacts on intangible assets such as brand equity, reputation and trust.

In the 2015 report 'UK Fintech on the Cutting Edge', EY estimates that the UK fintech sector represented around £6.6 billion in revenue in 2015 and generated around £524 million in investment. HM Treasury has reported fintech growth of around 20% per annum since then.

This indicates that a sustained series of systemic cyber attacks on the UK FMI ecosystem resulting in, for example, a 10% (£12 billion) loss to the UK economy would wipe out the incremental benefits derived from fintech during most of the transition period.

The CLh values would need to be multiplied to reflect this impact. The resulting BHI values are reflected in Table 5 where the deeper red indicates increasingly negative BHI values.

	2020 - 2022	2023 - 2035	2026 - 2028	2028 - 2031
BHI value	🔴 < 0	🟠 < 0	🟡 < 0	🟢 < 0

Table 5 - of the hypothetical cyber attack scenario

In the case of BHI <= 0, the growth order (CL) of the harm exceeds the growth order of benefit. In such a case, unless there is mitigation, it is reasonable to expect that however the benefit grows, it will be overtaken by harm.

Even for just one hypothetical cyber threat scenario, the complexity of the ecosystem and the vulnerability levels of the components at these negative BHI time intervals make it hard to predict the full spectrum of associated cyber chain reactions. In Section 4 of this paper there is an illustration of this scenario in more detail, showing how the Implication Wheel™[1] can be used to try and detect emergent systemic threats in this context.

In applying the BHI formally, it can now be looked at systematically across a significant number of risks rather than just the one hypothetical example explored in this report.

AN APPROACH TO MITIGATING EMERGENT RISK/ RADICAL IGNORANCE

The approach highlighted here uses the Implication Wheel™1 methodology to help uncover emergent threats. Figure 9 illustrates the context that will be used to introduce the Implication Wheel™1 methodology. It features a hypothetical illustrative threat scenario as it could unfold in the UK FMI ecosystem.

Cyber ecosystems are complex systems of systems like the UK FMI ecosystem explored in this paper. As described earlier, such ecosystems are constantly changing, often in surprising ways.

Cyber attacks on such systems can cause cascading cyber chain reactions of indirect and unanticipated consequences. The direct first order effects are often relatively easy to predict and mitigate. However, the second and third order effects are much less obvious and may contain surprises, some of which will pose a systemic risk. These are referred to as 'black swan' events.

The Implication Wheel™1 is participatory 'smart group' methodology that uses a structured brainstorming process to uncover multiple levels of consequences, and which can lead to the discovery of black swan events. Each smart group comprises a diverse set of individuals who will bring different perspectives to the task.

Each smart group starts by considering an initial event, such as the hostile state actor (FSB) launching a hypothetical multi-vector cyber attack.

The example threat actor is shown in the red outlined squares on the left of Figure 9. The initial event resulting from the launch of the multi-vector attack is represented by the set of white boxed activities in the grey highlighted column of Figure 9.

Illustrating a hypothetical cyber chain reaction leading to systemic risk in UK FMI ecosystem

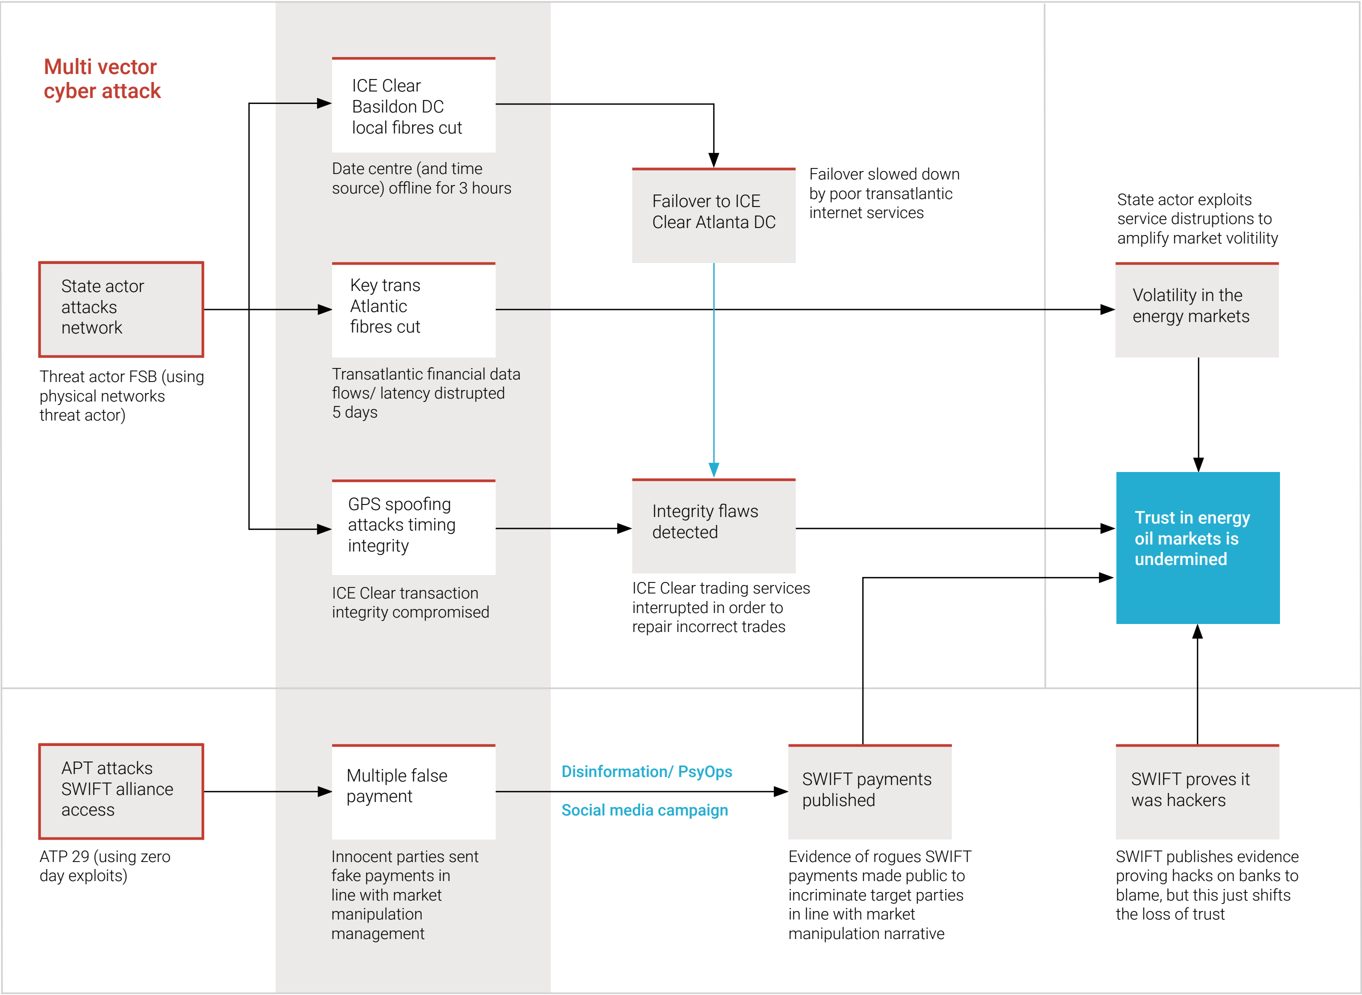
Multi vector cyber attack
State actor attacks network
Threat actor FSB (using physical networks threat actor)
ICE Clear Basildon DC local fibres cut
Date centre (and time source) offline for 3 hours
Key trans Atlantic fibres cut
Transatlantic financial data flows/ latency distrupted 5 days
GPS spoofing attacks timing integrity
ICE Clear transaction integrity compromised
Failover to ICE Clear Atlanta DC
Failover slowed down by poor transatlantic internet services
Integrity flaws detected
ICE Clear trading services interrupted in order to repair incorrect trades
State actor exploits service distruptions to amplify market volitility
Volatility in the energy markets
Trust in energy oil markets is undermined
APT attacks SWIFT alliance access
ATP 29 (using zero day exploits)
Multiple false payment
Innocent parties sent fake payments in line with market manipulation management
Disinformation/ PsyOps
Social media campaign
SWIFT payments published
Evidence of rogues SWIFT payments made public to incriminate target parties in line with market manipulation narrative
SWIFT proves it was hackers
SWIFT publishes evidence proving hacks on banks to blame, but this just shifts the loss of trust

Figure 9 - The Implication Wheel™1 approach applied to the UK FMI ecosystem

Each smart group is then asked "What might happen next?" This generates the direct first order consequences.

In Figure 9, we illustrate some potential first order consequences that propagate from the initial event in the yellow column. These first order consequences, as described earlier, include the loss of the ICE Clear data centre for three hours, five days of disruption to the specialist cable network providing specialist low-latency high-capacity services for transatlantic financial traffic.

This "What might happen next?" process is then repeated by the smart groups for each first order consequence, creating an associated set of second order consequences. This process can be repeated to explore third order consequences, and so on.

For illustration in Figure 9, we have shown a second order black swan event, equating to a significant breakdown in trust in the UK's financial market services as a consequence of this multi-vector cyber-attack.

When the Implication Wheel™1 is used more formally in this context, a layered structure is produced, like the wheel is produced, as shown in Figure 10. This illustrates just one second order effect and its associated third order effects.

The Implication Wheel™1 methodology permits smart group participants to propose levels of impacts and importance, and the likelihood of each consequence. For example, the likelihood of trust in the FMI being damaged by claims of collusion and market manipulation by geo-political players from a nation state threat actor might be low. However, if those claims were supported by evidence of large/suspicious SWIFT financial transactions between those parties, then that likelihood might change.

The smart group should include people with different perspectives and expertise. The chain reaction involves not just technical aspects but socio technical aspects all of which have consequences.

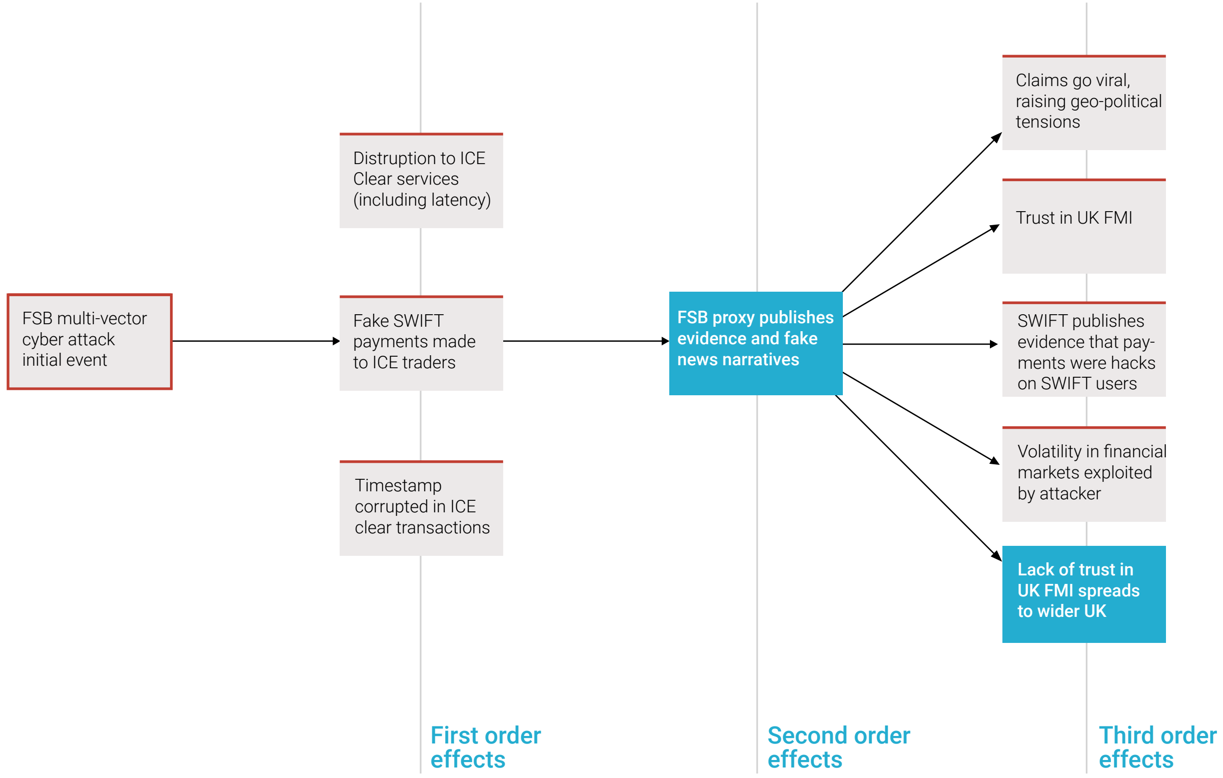
FSB multi-vector cyber attack initial event
Distruption to ICE Clear services (including latency)
Fake SWIFT payments made to ICE traders
Timestamp corrupted in ICE clear transactions
FSB proxy publishes evidence and fake news narratives
Claims go viral, raising geo-political tensions
Trust in UK FMI
SWIFT publishes evidence that pay-ments were hacks on SWIFT users
Volatility in financial markets exploited by attacker
Lack of trust in UK FMI spreads to wider UK
First order effects
Second order effects
Third order effects

Figure 10 - Formal Implications Wheel™1 example showing layered structure

Impacts can range from macro level (the entire UK FMI ecosystem and beyond) down to small localised consequences for a specific member entity.

As mentioned previously, when exploring the impacts of attacks on cyber ecosystems, the impact on both tangible and intangible assets needs to be included, as illustrated in Figure 11.

Cyber attacks can impact on intangible assets but are less likely to impact on physical assets, such as plant and machinery.

The intangible assets associated with the FMI ecosystem include the brand equity of each of the participants, and in particular the critical FMI service operators. Crucially, intangibles include trust in the UK FMI ecosystem overall, which is fundamental to the viability of the financial sector and a significant contributor to the UK economy.

The black swan event hypothesised in the example is therefore significant, since it demonstrates how such trust could be damaged, and how the illustrative multi-vector cyber attack could pose a systemic risk.

A whole spectrum of cyber attacks would need to be modelled in this way to help discover some of the many emergent systemic risks associated with the complex system of systems that forms the UK FMI ecosystem.

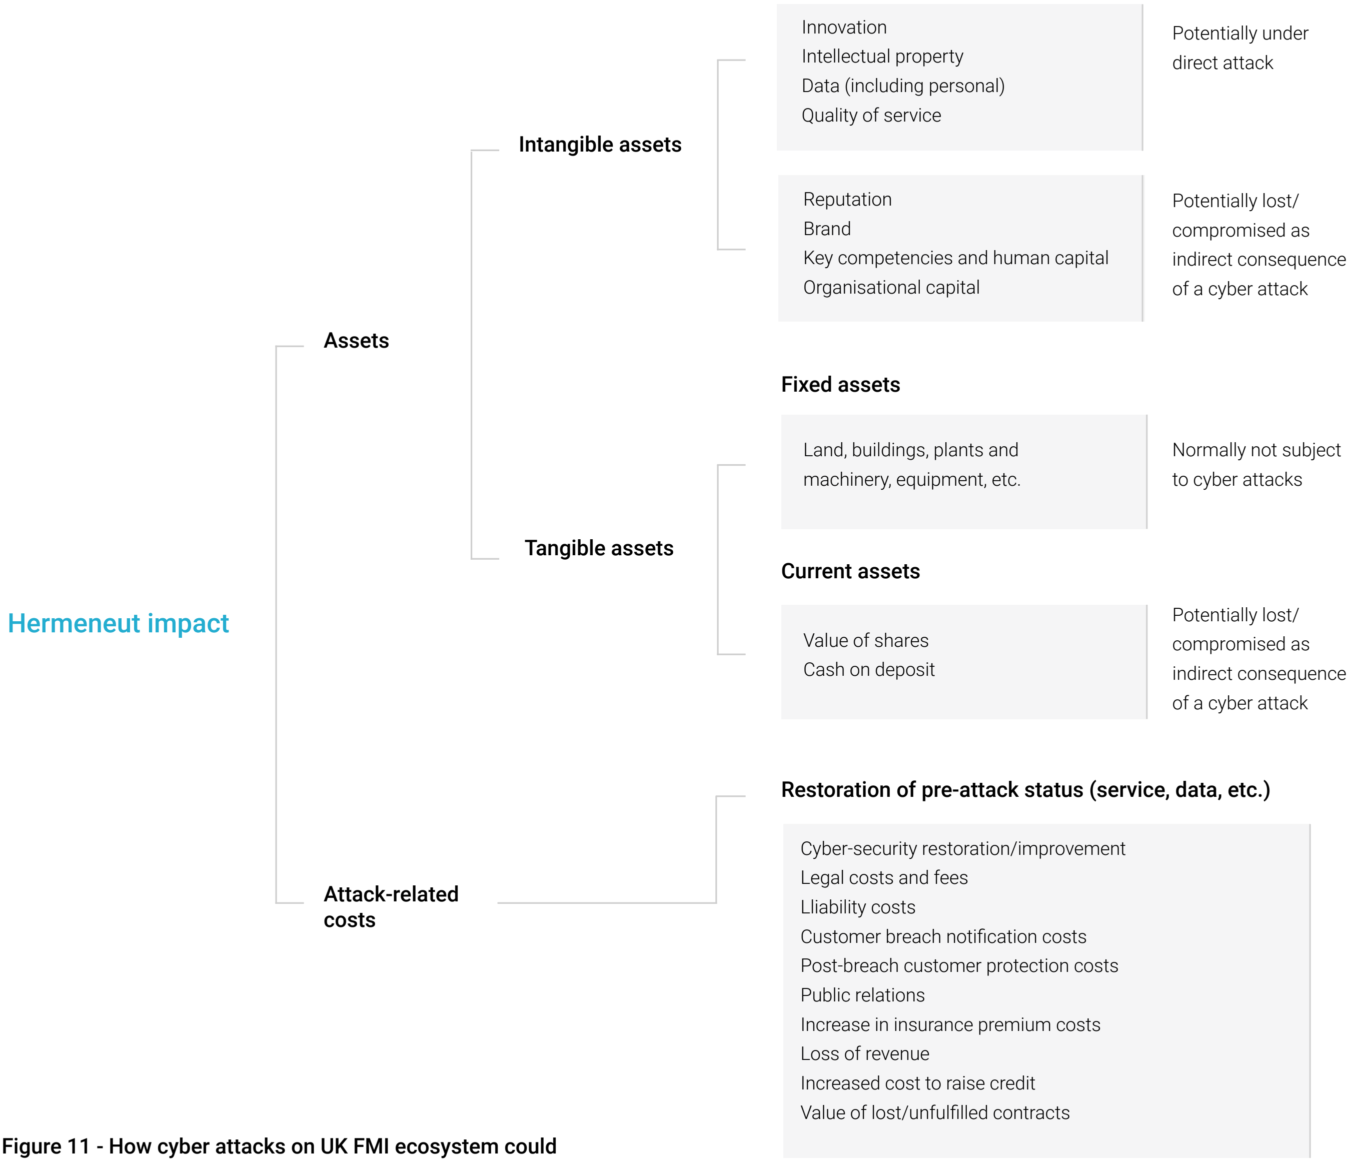
Hermeneut impact
Assets
Intangible assets
Innovation
Intellectual property
Data (including personal)
Quality of service
Potentially under direct attack
Reputation
Brand
Key competencies and human capital
Organisational capital
Potentially lost/ compromised as indirect consequence of a cyber attack
Tangible assets
Fixed assets
Land, buildings, plants and machinery, equipment, etc.
Normally not subject to cyber attacks
Current assets
Value of shares
Cash on deposit
Potentially lost/ compromised as indirect consequence of a cyber attack
Attack-related costs
Restoration of pre-attack status (service, data, etc.)
Cyber-security restoration/improvement
Legal costs and fees
Lliability costs
Customer breach notification costs
Post-breach customer protection costs
Public relations
Increase in insurance premium costs
Loss of revenue
Increased cost to raise credit
Value of lost/unfulfilled contracts

Figure 11 - How cyber attacks on UK FMI ecosystem could impact on tangible and intangible assets

Conclusion

Find out more about the BHI and the Hermeneut project

This report shows how to apply the BHI to CNI cyber ecosystems, and uses a hypothetical cyber attack scenario to illustrate the process. The formal application of the BHI to CNI cyber ecosystems would uncover potentially significant emergent threats in advance of exploitation by hostile nation state actors and their proxies, as well as threat actors such as terrorists.

Digital Catapult welcomes further discussion with CNI stakeholders on the potential benefits of such projects.

The BHI approach is described in full technical detail in EU Hermeneut project document: D4.2 BHI (Benefit Harm Index) Report. This is available on the Hermeneut site at the following link: www.hermeneut.eu/resources/

Hermeneut's cybersecurity cost-benefit approach to risk assessment combines integrated assessment of vulnerabilities and their likelihoods with an innovative macro and micro economic model for intangible costs, delivering a quantitative estimation of the risks for individual organisations or a business sector, and investment guidelines for mitigation measures.

Learn more about the wider Hermeneut project here: www.hermeneut.eu/about/

CATAPULT Digital

Hermeneut

HORIZON 2020

Glossary

ATM	Automated teller machine
BHI	Business harm index
FCA	Financial conduct authority
FMI	Financial markets infastructure
FSMA	Financial services and markets act
CNI	Critical national infrastructure
CREST	Certificateless registry for electronic share
GNSS	Global navigation satellite system
NCA	National crime agency
NGSC	National cyber security centre
PESTL	Political, economic, social, technical, legal

References

The following research papers and other resources are referenced by this white paper:

Endnotes

1. Countryeconomy.com https://countryeconomy.com/gdp/uk

2. Rogers, E.M. 1962, 'Diffusion of Innovations', New York: The Free Press

CATAPULT
Digital

Digital Catapult is the UK's leading advanced digital technology innovation centre, driving early adoption of technologies to make UK businesses more competitive and productive to grow the country's economy.

We connect large established companies, startup and scaleup businesses and researchers to discover new ways to solve big challenges in the manufacturing and creative industries. Through this collaboration businesses are supported to develop the right technologies to solve problems, increase productivity and open up new markets faster.

Digital Catapult provides physical and digital facilities for experimentation and testing that would otherwise not be accessible for smaller companies.

As well as breaking down barriers to technology adoption for startups and scaleup, our work de-risks innovation for large enterprises and uncovers new commercial applications in immersive, future networks, and artificial intelligence technologies.

For more info please visit **digicatapult.org.uk**

Photography Alina Grubnyak