

RESILIENCE STRATEGIES AND APPROACHES TO CONTAIN SYSTEMIC THREATS

17-18 September 2019, OECD Conference Centre

This paper defines concepts related to systemic threats and reviews the analytical and governance approaches and strategies to manage these threats and build resilience to contain them. This aims to help policymakers build safeguards, buffers and ultimately resilience to physical, economic, social and environmental shocks. Recovery and adaptation in the aftermath of disruptions is a requirement for interconnected 21st Century economic, industrial, social, and health-based systems and resilience is an increasingly important theme and a crucial part of strategies to avoid systemic collapse.

It builds on the report *Resilience Activities at the OECD: Current Practice and Future Directions* and was produced with the support of the Slovak Republic. Comments and suggestions from the NAEC Group are welcome.

William Hynes – william.hynes@oecd.org

JT03450724

Resilience-based Strategies and Policies to Address Systemic Risks

Igor Linkov, Carnegie Mellon University and US Army Engineer Research and Development Center;

Benjamin D. Trump, University of Michigan and US Army Engineer Research and Development Center

William Hynes, Organisation for Economic Co-operation and Development (OECD) and the Johns Hopkins University School of Advanced International Studies

Executive Summary

Modern society proceeds on the assumption that a number of complex systems will work reliably, both individually and in their interactions with other human and natural systems. This bold assumption generally holds true. Yet this assumption is being challenged by the scientific community and, increasingly by policymakers too, on the basis and improved understanding of how systems work, and by drawing on the lessons of disruptive historical events. Examples include how expanding trade routes and changing demography provided the conditions for the Black Death to wipe out up to 60 percent of Europe's population in the 14th century. The way agricultural, social and political factors combined to turn a crop disease into the potato famine that devastated 19th century Ireland. Or more recently, the fact that the world financial system almost collapsed in 2008 because of the negative impacts of a number of innovations that were supposed to improve efficiency.

The immense variety of complex systems gives rise to an equally vast array of systemic risks, identifying which is made more difficult by the characteristics of the systems in question. It is possible though to characterise the nature of these threats as involving a process of contagion that spreads individual failures to the system as a whole, and where disruption in one area can cascade through the system as a whole. The key insight here is that interconnectedness brings many benefits, but it also poses new problems and can intensify the danger from existing threats. Moreover, although the consequences of a systemic threat being realised may be dramatic, the probability of this event occurring may be small. This combination of high impact, low probability, and propagation that is hard to predict makes the task of the policymaker trying to address systemic risk arduous. It requires a different approach to dealing with risks and potential system failures.

Resilience can provide a philosophical and methodological basis to address systemic risk in a more useful way than traditional approaches based on risk management. Risk assessment and management are used to harden components of the systems affected by specific threats, yet such approaches are often prohibitively expensive to implement, and do not address cascading effects of system failure. Resilience approaches emphasise the characteristics and capabilities that allow a system to recover from and adapt to disruption, and a commonly-accepted definition describes resilience as the ability of a system to perform four functions with respect to adverse events: planning and preparation; absorption; recovery; and adaptation.

While risk management and resilience approaches share some of the same concerns, there is a fundamental difference in the timeframes in which they operate. Risk is concerned with what happens before an event – preparing the system for a given threat; resilience is concerned with how the system behaves after the event occurs and how this is affected by and network effects other systems. Resilience

can also be said to be “threat agnostic”. It does not try to identify particular threats, but rather assumes that at some stage, some threat or combination of threats will materialise and disrupt the system. The goal therefore is to be ready for whatever happens, even if it cannot be anticipated and has never happened before. Furthermore, since the system is complex, a resilience approach accepts that transitions to new phases are part of its nature and the system will not return to some previous equilibrium. New normals are normal. This dynamic nature of systems means that they are constantly evolving, and strategies that improve resilience to certain stresses at one point in time may render the system more brittle in the face of shocks at another time. Seen like this, resilience of a system is less a singular moment when a disruption incurs losses, but is instead a process of how a system operates before, during, and after the threat arriving.

For a resilience approach to be useful to policymaking, the domains of resilience have to be identified, along with the potential sources of system collapse. Four main domains can be identified: physical (sensors, facilities, equipment, system states and capabilities); information (creation, manipulation, and storage of data); cognitive (understanding, mental models, preconceptions, biases, and values); and social (interaction, collaboration and self-synchronisation between individuals and entities). Sources of collapse can be as varied as the domains, ranging from failure of political authority to economic or social crises.

Such a wide spectrum of parameters complicates the identification and management of systemic threats in practical terms. A certain number of criteria can guide the process, two of which are particularly important. First, regardless of the type of data collected (qualitative, quantitative, or a mix of the two), the dataset must have a clear and indisputable connection with the resilience project in mind. Second, a crucial step in conducting resilience classification includes the imposition of predetermined notions of system success or failure. The resilience matrix is a powerful tool in this process. One axis consists of the four domains (physical, information, cognitive, social); the other has the four phases (prepare, absorb, recover, adapt). Resilience is assessed by providing a value in each cell that summarises the capacity of the system to perform within that domain and period of time. This helps to identify those parts of the system that need to be made more resilient. It also helps to identify where a solution to one problem may create other problems elsewhere. For example, in dealing with floods, mobile pumps that discharge untreated water may be a source of health problems later.

It is worth noting, however, that resilience matrices do not intrinsically capture the various temporal characteristics within decision making that could cause shifts in preferences or needs over time, requiring users to continually update their matrix outlook at regular intervals. Under such circumstances, resilience matrices may currently be considered an optimal tool to gain a general view of system resilience that could be further refined with a more thorough decision analytic tool or modelling effort. Decision-makers therefore need a suite of tools that could include methods inspired by network science in addition to matrix-based approaches, and combining qualitative and quantitative approaches. Institutional issues have to be considered too, notably the lack of an accepted formal definition or centralised governing body, and considerable efforts are needed to integrate quantitative, qualitative and semi-quantitative approaches.

This is becoming urgent as the risks associated with the trends most likely to influence our economies and societies in the future, notably AI and digitalisation, are not amenable to traditional approaches and will present challenges that we cannot anticipate or prepare to meet using traditional tools.

Outline

- (1) Systemic Threats: Definitions and a Brief Review of the Literature
 - a. Defining Systemic Threats
 - b. The diverse nature of systemic threats – the need for recovery and adaptation
- (2) Resilience as a Philosophy and Tool to Understand and Address Systemic Threats
 - a. Operating philosophy
 - b. How resilience addresses systemic threats

- (3) Using History as a Lens for Civilisational Collapse or Survival Amidst Systemic Threats
 - a. Defining and Understanding Collapse in History
 - b. Systemic Threats in History
- (4) Methods for Identification and Management of Systemic Threats
 - a. Methodological Input Requirements for Risk and Resilience of Systemic Threats
 - b. Semi-Quantitative Approaches: The Resilience Matrix
 - c. Quantitative Approaches: Network Science
- (5) Ideas for Methodological Practices for Resilience: Making Resilience Useful for Decision Makers
 - a. Resilience and Efficiency
 - b. Future Systemic Challenges Facing Society

1. Systemic Threats: Defining and Understanding a Growing Global Concern

Modern society relies upon complex systems for virtually every major activity. From water and plumbing, to energy and electricity, to even the provision of public health and international finance. Such systems allow society to grow and diversify by more efficiently delivering goods and services than was historically possible. However, such interconnection comes at a cost.

Threats to such complex systems may yield disruptions to how such life is able to function. This section details how a phenomenon of growing concern, known as *systemic threats* or *systemic risks*, can contribute to dramatic shifts in the stability of modern complex systems.

1a. Defining Systemic Threats

Systemic threats represent growing challenges within the international landscape. Functionally, systemic threats have been defined in scholarly literature in numerous ways such as Centeno et al. (2015) (“the threat that individual failures, accidents, or disruptions present to a system through the process of contagion”). The International Risk Governance Center contends that systemic threats arise when “systems [...]are highly interconnected and intertwined with one another”, where a disruption to one area triggers cascading damages to other nested or dependent nodes (IRGC 2018). Further, IRGC (2018) states that “external shocks to interconnected systems, or unsustainable stresses, may cause uncontrolled feedback and cascading effects, extreme events, and unwanted side effects”, implying that the potential for cascading disruption is a growing and critical concern for many facets of daily life.

Systems susceptible to systemic risks are intertwined with one another in a series of nested relationships. Renn (2016) describes how such interconnectivity facilitates stochastic, non-linear, and spatially interspersed causal structures that, if triggered, contribute to a ‘domino effect’ that can permanently alter a broader infrastructural, environmental, or social system. Due to this nature, “systemic risks overwhelmingly do not follow normal risk distributions but tend to be fat-tailed, i.e. there is a high likelihood of catastrophic events once the risk arrives” (IRGC 2018).

Upon arrival, systemic threats have the potential to completely disrupt the original configuration or order of a given system – replacing it with something altogether new and sustainable in its own way. Normatively, such outcomes in many fields are generally constructed to be negative, such as the potential collapse of the global financial system, the depopulation of biodiversity on land or at sea, or a major international public health crisis. Overall, however, is that systemic threats push systems from one stable configuration onto another.

Though the nested infrastructural, social, environmental political, and economic systems of modern society made the consequences of systemic risk more pronounced and globally distributed, such threats have contributed to significant disruptions in historical societies as well. One famous example includes the bubonic plague, or The Black Death, which in the mid-14th Century was transferred from Asia towards Europe through a growing trade network that had connected states and peoples by land and sea. Within a matter of decades (and in some countries, years), huge proportions of the European and Asian populations were killed by the Black Death, and triggered massive changes in the economic and social lives of contemporary peoples.

Certain triggers of systemic threats can be violent and forceful, jarring a relatively stable and sustainable system into an altogether different configuration. Generally rapid and external shocks, examples include the disruptive potential of solar storms or space weather upon digital systems, or the sudden release of toxic chemicals into the natural environment. Generally ‘low-probability yet high-consequence’, such events are difficult to predict via conventional modeling techniques. Yet, once the acute disruption occurs, a chain reaction of systemic shift occurs until a new stasis is achieved.

Other triggers are more chronic in nature. Such events, as with gradual climactic shift or slight overfishing within a given ocean or sea, are initially limited in impact, yet can eventually be overwhelming and unstoppable in their effect. IRGC's *Guidelines for the Governance of Systemic Risk* (2018) states that such threats are best addressed early through the detection and interpretation of weak signals, although further notes that such slow-moving chronic systemic threats can be nearly imperceptible in their earliest stages.

Key to such slow-moving chronic threats is the notion of *transitions*. Lucas et al (2018) and Pelling (2012) frame such transitions as 'tipping points' where a system edges towards a critical inflection point that may foster a transformation from one system permutation in favor of another. If breached, such tipping points can foster feedback loops and nonlinear effects that cause a system to shift and change in an increasingly dramatic form. For the case of collapse of economic or social systems, identifying tipping points as well as the drivers or influencers that drive a system towards such points is a difficult yet critical requirement in order to preserve core system functions (i.e., financial markets, labor forces, international trade, etc.).

Systemic threats are characterised by their capacity to percolate across complex interconnected systems – either through an abrupt shock, or gradual stress (IRGC, 2018). Systemic threats are particularly difficult to model and calculate via a risk-based approach due to a mixture of the weak signals, which herald a potential upcoming systemic risk event, as well as the nested interaction effects by which a systemic threat incurs disruption to a system in an indirect manner. For example, the Financial Crisis of the previous decade began as a collection of relatively contained failures of financial firms, which ended in a substantial financial collapse across much of the world.

Helbing (2012) notes that the consequences of failing to appreciate and manage the characteristics of complex global systems and problems can be immense.

1b. The diverse nature of systemic threats – the need for Resilience

Resilience serves as a more helpful alternative philosophy and methodology to analyse complex adaptive systems and systemic risks, which are difficult to analyse via conventional risk assessment methodologies. Many modern systems would benefit from a resilience-based approach, particularly for systems with inherent nested interdependencies with others, or those which are prone to low-probability, high-consequence events that are difficult to accurately predict or model. Resilience helps these systems prepare for disruptions, cope with and recover from them if they occur, and adapt to new context conditions (National Academy of Sciences (NAS), 2012^[1]; Linkov and Trump, 2019^[2]).

In a general sense, resilience has been used as a metaphor that seeks to describe how systems absorb threats and maintain their inherent structure and behaviour. More specifically, resilience is used as a global state of preparedness, where targeted systems can absorb unexpected and potentially high consequence shocks and stresses (Larkin et al. 2015). Such definitions are certainly helpful for policymakers to make sense of their systems, and identify opportunities to improve system function and capacity to counter disruptions. However, to emphasise the role of system function, recovery, and adaptation to disruption, it is necessary to adopt a *system's view*.

Common usage of resilience causes scholars to infer several principles of what resilience actually means. The first such principle includes the positivity of resilience, or the notion that resilience is an inherently beneficial goal to achieve. The second includes the measurement of resilience by characteristics believed to apply to a given system – effectively driving an inductive approach to resilience thinking (Bené et al. 2012). Lastly, resilience thinking is often viewed in a context-agnostic framework, where principles of resilience can be applied to various situations and cases interchangeably.

We define resilience as the capability of a system to recover in the midst of shocks or stresses *over time*. Recovery implicates multiple interactions between factors, and across scales and sub-systems, that are usually unexpected and complex in nature. Given such concerns, resilience differs from traditional methodological approaches of protecting against risk, where these uncertain and complex shocks and

stresses that affect targeted systems are inherently outside of the design of the system's intended purpose. As such, preparation for such events contains only limited available guidance, and promoting traditional risk approaches such as bolstering system hardness is prohibitively difficult and excessively expensive. Resilience allows us to address these concerns within a framework of resource constraints and the need to protect against low probability, high consequence events more recently described as 'black swans.' In other words, resilience is preferred to traditional risk management strategies where a systems-theory of protecting against risk is required, and where the potential risks in question are highly unlikely yet potentially catastrophic in nature.

Resilience affords greater clarity over such threats (particularly systemic threats) by focusing upon the inherent structure of the system, its core characteristics, and the relationship that various sub-systems have with one another to generate an ecosystem's baseline state of health (IRGC, 2018). Walker et al. (2004) define ecosystem equilibria as a characteristic of "basins of attraction", where the components and characteristics of a system drive it towards a baseline state of health and performance. For example, the Pacific Ocean is a huge and complex ecosystem with a tremendous diversity of flora and fauna whose roles in complex food webs have been reinforced by millions of years of evolution and adaptivity; a localised oil spill may damage small points of ecosystem health but is unlikely to dramatically and permanently shift the species dynamics and food webs which currently prevail across most of the Ocean. However, through constant exposure to trillions of microplastics (i.e., the Pacific Trash Vortex) or continuous chemical and radiological contaminants (i.e. bleaching of the Great Barrier Reef, radiological runoff from the Fukushima Daiichi Nuclear Power Plant), system equilibria can be jolted in a manner that favors a differing basin of attraction. Unfortunately, we are moving in that direction already, where huge regions of oxygen-depletion in the Pacific Ocean are contributing to 'dead zones' where virtually no marine life can survive.




Figure 1: Illustration of a complex interconnected environmental system (Linkov & Trump, 2019)

Both the initial Pacific Ocean baseline state of health (a larger, global system), as well as these dead zones (sub-system within the Pacific Ocean global system), are fundamentally resilient systems that are defined by basins of attraction which possess characteristics that reinforce the status quo in the absence of a shock or disruption. On one hand, a biologically rich and diverse Pacific biological system has recovered from a tremendous array of disruptions over the past century, or adapted system processes in a way that address incoming challenges such as with large-scale commercial fishing. However, continuous overfishing along with chemical and radiological runoff are disrupting the Pacific ecosystem enough to potentially transition

it towards a new and less biologically complex basin of attraction. This will be further discussed in the last section.

More than just a metaphor, ecosystem resilience describes the intensity of a given ecological basin of attraction to preserve the baseline state of health and activity within a given area, whether that state is optimum, desirable or not. It could be normatively positive (i.e., a complex and biodiverse Pacific Ocean) or negative (i.e., a Pacific Ocean comprised of huge dead zones and limited biodiversity). Methodologically, such basins of attraction are comprised of complex interconnected and adaptive systems that are constantly under stress, yet only shift to a new equilibrium if a tipping point has been breached and the system is trending towards a new basin. More than just system recovery and adaptation, ecosystem resilience is a property of natural selection and organism interaction within their broader environment in a manner that produces some sustainable end-state. Such resilience-based approaches can help us understand when and how certain ecosystems might shift from one steady-state to another (Linkov et al., 2018), as well as define the biological and ecological drivers which cause an ecosystem to arrive at a steady equilibrium altogether.

In the modern age, there is no socioecological system that is not influenced by human behaviour or activity. Increasingly, many research organisations find that human activity is directly or indirectly pushing environmental and ecological systems from an initial condition of high biodiversity and systemic complexity, towards simpler, less diverse, and less hospitable climates and food webs. Some human-derived disruptions are relatively abrupt (e.g., industrial logging in tropical rainforests) or more gradual in effect (e.g., ocean pollution), yet both tend to drive such environmental and ecological systems towards a tipping point that limits the potential for diverse environmental life. Such a system is resilient yet normatively unfavorable, where significant energy and resources would have to be dedicated towards returning an at-risk environment to its original basin of attraction.

2. Resilience as a Philosophy and Tool to Understand and Address Systemic Threats

2a. Resilience as a Philosophy and an Approach for Complex Systems

As a term, resilience has centuries of use as a descriptor in fields as diverse as military operations, to psychology, to civil and environmental engineering. Its synonyms are vast and varied, ranging from insinuations of toughness to elasticity. While it finds its roots in these early ideas, the modern application of resilience has centered upon analysing how systems bounce back from disruption. This seems simple enough at first glance, yet the methodological application and analysis of how systems do, in fact, bounce back post-disruption can be quite challenging.

A 2012 National Academy of Sciences (NAS) report on “disaster resilience” defines resilience as the ability of a system to perform four functions with respect to adverse events: (i) planning and preparation, (ii) absorption, (iii) recovery, and (iv) adaptation. Nevertheless, quantitative approaches to resilience in the context of system processes have neglected to combine those aspects of the NAS understanding that focus on management processes (i.e., planning/preparation and adaptation) with those that focus on performance under extreme loadings or shocks (i.e., absorption and recovery). Advancing the fundamental understanding and practical application of resilience requires greater attention to the development of resilience process metrics, as well as comparison of resilience approaches in multiple engineering contexts for the purpose of extracting generalisable principles.

A core problem here is that risk and resilience are two fundamentally different concepts, yet are being conflated as one and the same. The Oxford Dictionary defines risk as “a situation involving exposure to danger [threat]”, while resilience is defined as “the capacity to recover quickly from difficulties.” The risk framework considers all efforts to prevent or absorb threats *before* they occur, while resilience is focuses

on recovery from losses *after* a shock has occurred. However, the National Academy (2012) and many others define resilience as “the ability to anticipate, prepare for, and adapt to changing conditions and withstand, respond to, and recover rapidly from disruptions.” In this definition, adapt and recover are resilience concepts, while withstand and respond to are risk concepts, thus the risk component is clearly added to the definition of resilience. Further, approaches to risk and resilience quantification differ. Risk assessment quantifies the likelihood and consequences of an event to identify critical components of a system vulnerable to specific threat, and to harden them to avoid losses. In contrast, resilience-based methods adopt a ‘threat agnostic’ viewpoint.

We understand resilience as the property of a system *and* a network, where it is imperative for systems planners to understand the complex and interconnected nature within which most individuals, organisations, and activities operate. Risk-based approaches can be helpful to understand how specific threats have an impact upon a system, yet often lack the necessary characteristic of reviewing how linkages and nested relationships with other systems leave one vulnerable to cascading failure and systemic threat. Resilience-based approaches, which inherently review how the structure and activities of systems influence one another, serves as an avenue to understand and even quantify a web of complex interconnected networks and their potential for disruption via cascading systemic threat. Such an approach is one of increasing prominence and focus at the international level, where the need to better protect complex systems from systemic threat becomes a matter not only of whether a system can survive disruption, but importantly in what state would it find itself in the aftermath of such a disruption.

Resilience is both a philosophy and a methodological practice that emphasises the role of *recovery* post-disruption as much as *absorption* of a threat and its consequences. Philosophically, this mindset is one that is grounded upon ensuring system survival, as well as a general acceptance that it is virtually impossible to prevent or mitigate all categories of risk simultaneously, and before they occur. Methodologically, resilience practitioners seek to make use of limited financial and labor resources to prepare their system for a wide variety of threats – all the while acknowledging that, at some point in the future and regardless of how well the system plans for such threats, disruption will happen. While the more conventional practice of risk assessment and management is very concerned with accounting for systemic threats, this exercise is typically undertaken on a threat-by-threat basis in order to derive a precise quantitative understanding of how a given threat exploits a system’s vulnerabilities and generates harmful consequences. As will be discussed later in this chapter, such an exercise works well when the universe of relevant threats is thoroughly categorised and understood, yet develops limitations when reviewing systemic risk to complex interconnected systems. Building from this limitation, resilience complements traditional risk-based approaches by reviewing how systems perform and function in a variety of scenarios, agnostic of any specific threat.

There are some theoretical and empirical implications of the above definition of resilience that have to be taken in consideration. They seldom are, or are not explicitly included in assessment. We unpack these implications below.

i) The dimension of time and experiential learning

The dimension of “time” is not only important to shorten the recovery phase (Linkov et al. 2014), as an indicator of resilience, but also implies the understanding how the system coped with previous stress and what were the dynamics of those changes. Linkov et al. (2014) outlined resilience as a function of system performance over time, which we extend to argue that such system resilience includes the past experiences that a given system has encountered that have stressed its capacities for service delivery or normal function. In other words, exposure to previous shocks in stresses in various capacities can have a direct effect upon the system’s ability to recover from future shocks and stresses. Coupled with the ability of a system to absorb shocks and stresses while still maintaining important functions, recovery serves as

an essential component to judge whether a system is resilient in the face of challenges (see Table 1 for a typology of system absorption and recovery capabilities under stress).

	High Absorption	Low Absorption
High Recovery	Ideal State; high adaptive capacity	Resilient but challenged system; moderate adaptive capacity
Low Recovery	Hard but brittle system; Low adaptive capacity	Significant threat to long-term survival; Low adaptive capacity

Table 1. Typology of System Resilience by Adaptive Capacities for Absorption and Recovery

This phenomenon is driven by the *adaptive capacity* of a system. Whether biological, cognitive, or infrastructural, systems that have been previously exposed to shocks and stresses are more likely to have the experience and memory to adapt in the face of new and emerging challenges. Likewise, those systems with limited exposure to historical challenges may have less capacity to recover to future threats due to their limited experience of having to adapt system capabilities in order to overcome such shocks. For example, the human body is better able to more quickly absorb and/or recover from certain illnesses if it was exposed to similar diseases in the past through the creation of antibodies in the blood together with improved knowledge regarding how to best treat such illnesses with medication and best practices.

ii) The shifting capacity of a system

Since a system is dynamic (it changes over time), system stresses can occur throughout the system’s development. As such, individual strategies can *both* augment an individual system’s resilience to certain stresses *while also* increasing the system’s brittleness in the face of certain shocks. In other words, it is possible for a system to become increasingly adaptive, yet also become increasingly brittle and susceptible to disruptions from shocks and stresses.

Given this idea, it is essential to understand that strategies to promote resilience may also make the system brittle or susceptible to collapse. A recurring example of this includes economic markets, which continually adapt to emerging market conditions through new strategies of investment, capitalisation, and debt transfer. Specifically, while investment markets continually adapt and develop resilience to external shocks, they become increasingly brittle through growing system complexity and appetite for risky investments. These actions are individually rational (i.e. investors seek to grow profits by approving riskier trades that are generally sound but have a higher chance of failure), yet can increase the potential for the stock market to enter recession as a large enough aggregate of investments fail and companies enter default. In this way, the stock market slowly trends towards brittleness in a rational manner over time.

2b. How resilience addresses systemic threats

The key question that resilience practitioners seek to answer is “how can I make sure my system performs as well as possible during disruption, and recovers quickly when disruption does occur” (Figure 2)? More specifically, how can I make sure my system is not vulnerable to cascading disruption posed by systemic threats?

These questions are particularly salient for the study of complex systems, where large organisations like hospitals rely upon the smooth operation of various connected systems and sub-systems to function properly (i.e. the energy grid, secure and efficient information systems, simplified patient intake, medical supply chains, and various others). Resilience is also an important question to tackle threats of very low

probability yet disastrous consequences, where no clear strategy exists to mitigate or prevent such threats from happening in the first place. Regardless of the situation to which it is applied, resilience requires one to think in terms of how to manage systemic, cascading threats, where a disruption to one sub-system can trigger dramatic changes to other connected systems. This is a complex task with few formalised answers, yet a helpful beginning is to operationalise resilience in a meaningful and methodological focus.

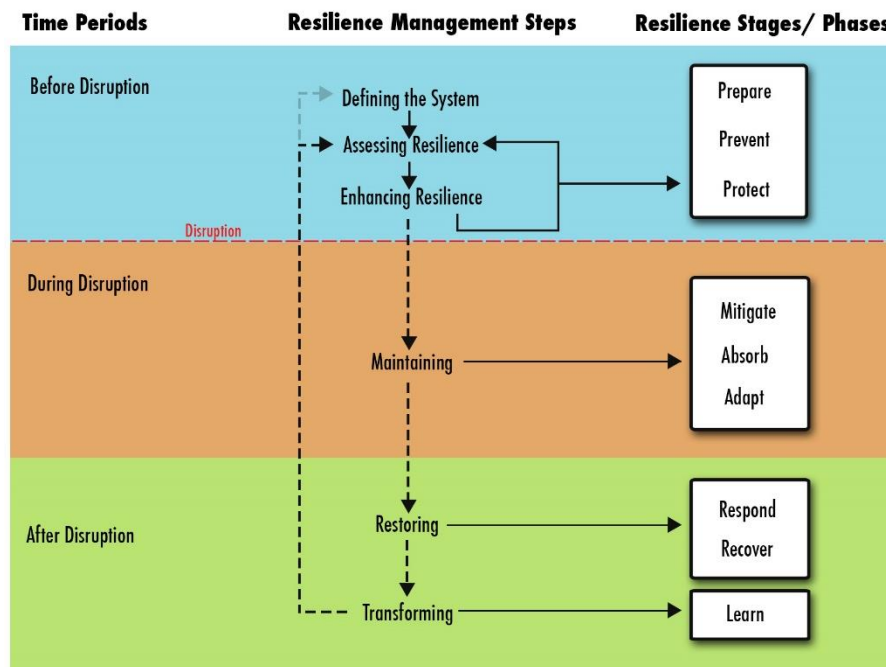


Figure 2. Role of Resilience in Systems, Emphasising Importance of Combating Disruptions

A central requirement for analysts is to frame resilience as a function of both *time* and *space*. We emphasise these considerations due to the multi-temporal and cross-disciplinary view by which one must review systemic threats.

Stages of Resilience

With respect to time, resilience of a system is less of a singular moment when a disruption incurs losses, but is instead a process of how a system operates before, during, and after the threat arrives. No single definition has been formalised in this area, yet the National Academy of Sciences’ 2012 report on Disaster Resilience notes resilience as involving how a system plan and prepares for, withstand and absorbs, recovers from, and adapts to various disruptions and threats (Figure 3) (NAS 2012). In this approach, system resilience is an ever-changing characteristic whereby a system’s core functions are constantly shifting to deal with threats.




Figure 3. Stages of Resilience as Proposed by US NAS.

Most conventional, risk-based approaches emphasise the plan/prepare and withstand/absorb phases the identify, assess, and prevent/mitigate threat (Linkov et al. 2018). Regardless of whether a specific threat is considered, these stages focus upon (a) identifying and interpreting signals associated with threats to a system, (b) exploring the structure and connections that a system has with others, and (c) identifying strategies that preserve a system's core capacity to function regardless of the disruption that occurs (Patriarca et al. 2018; Park et al. 2013). Signals include statistics and other information that might indicate a pending systemic threat, i.e. early reports of new and virulent disease as an indicator of a pending epidemic and public health crisis (Scheffer et al. 2012). Signal detection is a difficult and recurring task, but can be the only avenue to better understand the variety of systemic threats that may arise at different points in the future. Likewise, mapping of the various connections and dependencies within one's system can help identify critical functions that, if taken offline, could generate cascading systemic failure.

If possible, system preparation and absorption of threat is accomplished via a prevention-based approach where a threat is avoided altogether. However, when this is not possible, emphasis is placed upon the capacity of a normatively beneficial system to avoid total collapse. This can be accomplished by a more gradual approach known as 'graceful degradation', where the core operations of a system are prioritised over non-essential services for as long as possible. By limiting the extent and scope of disruption to a system, it becomes easier to keep system functions online and avoid the sudden shocks or losses in critical system capabilities. Often, this is accomplished by 'hardening' different functions of a system so that they will not break under pressure.

While the plan/prepare and absorb/withstand stages are important to help a system address systemic threats before they occur and as they arise, resilience approaches also must place importance upon how a system performs after the threat has arrived. This includes (i) recovery, and (ii) adaptation. Recovery includes all efforts to regain lost system function as quickly, cheaply, and efficiently as possible, while adaptation centers upon the capacity of a system to change and better deal with future threats of a similar nature. Dealing with recovery and adaptation constitute the particularly novel additions by resilience to the broader fields of risk analysis, assessment, and management, and force stakeholders to take account of percolation effects due to disruptions. The role of adaptation and recovery is the primary point of focus for any resilience analyst, where a system with a robust capacity for recovery can efficiently weather serious disruptions that would otherwise break even the most hardened of system components.

Domains of Resilience

Going beyond the US National Academy of Sciences' (NAS) stages of resilience, the spatial component of resilience requires one to consider how a disruption to one system can trigger consequences in others – including those that have indirect or in apparent linkages to the disrupted system.

Alberts & Hayes (2003) identify four different Network-Centric Operation (NCO) domains important to a system's agility, or what they later define as “the ability to successfully effect, cope with, and/or exploit changes in circumstances” (Alberts & Hayes 2006). While a relatively early contribution, this effort at resilience thinking is intended to force its users to consider the wide breadth of characteristics and decision inputs that may factor into system performance. Each domain is impacted in a different yet equally important manner when a critical event or disruptive arises and success in one domain may not guarantee the same outcome in the other areas. Additionally, it is important to note that the greatest resilience and ability to recover from adverse events is achievable only when all domains are considered and resolved in a resilience analysis policy problem. These domains include: (Hayes 2004; Alberts 2007).

1. Physical: sensors, facilities, equipment, system states and capabilities.
2. Information: creation, manipulation, and storage of data.
3. Cognitive: understanding, mental models, preconceptions, biases, and values.
4. Social: interaction, collaboration and self-synchronisation between individuals and entities.

These domains are important to decision making for complex systems in general and resilience in particular (Roeger et al. 2014; Collier & Linkov 2014). The physical domain represents where the event and responses occur across the environment, and is typically the most obviously compromised system in the midst and aftermath of an external shock or critical risk event. Elements here can include infrastructural characteristics ranging from transportation (roads, highways, railways, airports, etc.) to energy or cyber networks that deliver services to public and private entities alike (DiMase et al 2015). As such, the physical domain of resilience thinking generally includes those infrastructural factors that are most directly impacted by a hazardous event, where the other domains include outcomes and actions that are a response to damage to physical capabilities and assets. Threats to such infrastructure can range from environmental (i.e. a catastrophic storm) to human (i.e. terrorist violence or military attack). In this domain, the objective of resilience analysis is to bring the infrastructural or systems asset back to full efficiency and functionality for use by its original owner or user.

The information domain is where knowledge and data exists, changes, and is shared. Such elements here can include public or private databases, which are increasingly under potential attack from private hackers and other aggressive opponents (Osawa 2011; Zhao & Zhao 2010). Another growing target for information domain-type risks includes stored online communications and e-mails, which if acquired by a nefarious third party could generate individual embarrassment or even national security risks (Murray et al. 2014; Berghel et al. 2015; Petrie & Roth 2015). Where such aggressions are a growing reality in the Information Age (Kaur et al. 2015), adequately protecting against these risks and bolstering information systems to be resilient and robust under attack is of paramount importance to government agencies and private companies alike (Lino 2014). For this domain, the objectives of resilience management are to both prepare information assets for a variety of potential attacks while also assuring that these systems will react quickly and securely to such threats in the immediate aftermath. In this way, risk preparedness, risk absorption, and risk adaptation make information and cybersecurity resilience a growing priority for a variety of governmental and business stakeholders (Linkov et al. 2013; Collier et al. 2014; Björck et al. 2015).

The cognitive domain includes perceptions, beliefs, values, and levels of awareness, which inform decision-making (Linkov et al. 2013; Eisenberg et al. 2014). Along with the social domain, the cognitive domain is the “locus of meaning, where people make sense of the data accessed from the information domain.” (Linkov et al. 2016). Such factors are easy to overlook or dismiss due to a reliance upon physical infrastructure and communication systems to organise the public in response to a disaster, yet such

perceptions, values, and level of awareness of publics to strategies to overcome shocks and stresses are essential to the successful implementation of resilience operations (Wood et al 2012). In other words, without clear, transparent, and sensible policy recommendations that acknowledge established beliefs, values, and perceptions, even the best-laid plans of resilience will fall into disrepair. A robust accounting for the cognitive domain is particularly important for instances where policymakers and risk managers may have a disconnect with the local population, such as with international infrastructure development projects of health-based interventions. For such cases, policy solutions which may seem sensible and common-sense to the policymaker or risk manager might be assumed to be robust, yet rejected by locals as contrary to established custom or practice.

The social domain is characterised by interactions between and within entities involved. The social domain also provides an area to which careful attention should be paid in overall community resilience. Social aspects of society have impacts on physical health (Ebi & Semenza 2008). For example, individuals or communities can have better recovery in the face of epidemic when they also have strong social support and social cohesion. The social domain also ties into the information domain in regards to trust in information. When the community does not trust the source of information, they often do not trust the information itself or have to take the time to verify it, leading to a need for community engagement by the authority or organisation to increase their social relations and therefore trust within the community (Longstaff 2005).

While the physical and cognitive domains attract a lot of attention in both overall resilience and hazard-specific resilience, the information domain is of great importance for overall functioning. In more than just health events, information has a huge impact on citizen response. Not all individuals understand and interpret information the same way. This leads to a need for attention to be paid on how to get information out effectively and in a timely fashion during a crisis. Information is important to more than just the citizens, however. Adequate information is crucial in real time for authorities to make informed and appropriate decisions. As important as information is, however, it is equally important to account for the role of human decision making. Specifically, human interpretation of data is important as raw numbers can be misleading if not considered in context of a given environmental setting or policy application. This ties back into needing to disperse tailored information for understanding that presents data pertinent to a threat in a manner that is not obscure for its recipient. The way in which authorities and citizens handle information should be evaluated with careful consideration for the communities being discussed.

Social resilience within this context may apply to societies and communities of various size, ranging from local neighborhoods and towns to more regional or national governments. For smaller communities, organisations, and businesses, discussions of resilience may center on the ability of local governments and set communities to address long-term concerns such as the impact of climate change (Berkes & Jolly 2002; Karvetski et al 2011), ecological disasters (Adger et al. 2005; Cross 2001), earthquakes (Bruneau et al. 2003), and cybersecurity concerns (Williams & Manheke 2010), as well as other manmade hazards such as transnational wars, civil wars, terrorism, migration, and industrial hazards. For larger communities and governments, such concerns are similar yet often more complex and varied in nature, where they involve hundreds to potentially thousands of stakeholders and include the interaction of various infrastructural systems.

These domains often overlap and exist in all systems, such as for messages from the information domain to be shared, infrastructure in the physical domain or interactions in the social domain must support dissemination. At its core, a focus upon domains ensures that a policymaker or risk manager acquires a holistic understanding of their policy realm, and are able to understand how a shock or stress could trigger cascading consequences that were previously difficult to comprehend. For example, the collapse of Lehman Brothers triggered a worldwide economic recession in 2008 due to the inherent interconnectivity of various economic and financial systems at that time.

3. Using History as a Lens for Civilisational Collapse or Survival Amidst Systemic Threats

Though systemic threats and civilisational disruption are framed as 21st Century concerns, they also help frame reasons why certain civilisations collapsed or survived in the midst of extreme disruption in certain eras of history. Learning from past events can allow us to identify (a) reasons why societies, economies, or ecologies persisted despite disruption, and (b) why others collapsed in the face of an adverse event. Lessons gained from such an exercise may help prepare modern society for similar disruptions, or at least position policies in a manner that limits collapse on similar lines from previous societies.

Such notions of history as an avenue to understand and explain civilisational disruption and collapse was a focal point of Princeton University's PIIRS (Princeton Institute for International and Regional Studies) Global Systemic Risk (GSR) research community, which explores the underlying fragility of the contemporary world. By reviewing the economic, political, social, environmental, and military systems of past societies, it is possible to trace how interconnections and interdependencies from past societies triggered cascading failures, albeit on a slower timeframe than would generally be experienced in a digitally-interconnected 21st Century world. PIIRS' GSR applies qualitative and quantitative indicators to assess civilisational outcomes through a variety of disruptive triggers, with hopes of developing models or frameworks to trace systemic collapse over time.

3a. Defining and Understanding Collapse in History

Collapse is a particularly difficult term to succinctly define in a manner that various disciplines and scholars would appreciate. For example, collapse could refer to the complete and total elimination of a nation-state from historical record (i.e., the Austro-Hungarian Empire in the aftermath of World War I), or it could also refer to civilisations that were fundamentally changed by disruption, despite their persistence and ability to survive in a differing form (i.e., the Eastern Roman Empire surviving for nearly a thousand years beyond its more famous Western counterpart).

For our purposes in this document, collapse refers to the permanent breakdown of complexity in the socioeconomic network of a given state. Societies are made up of networks and interconnected systems that (a) receive resources, energy, or participation from a broader network, and (b) yield some benefit, function, or purpose back towards that same society. Significant disruptions such as economic collapse, invasion, epidemic disease, or environmental catastrophe represent some of the possible situations where activities requiring significant resources or energy to operate (i.e., economic markets, extensions of geopolitical influence, growth-based demographic and economic policy) are not able to continue. In such situations, societal complexity decreases in the face of disruption until it reaches a sustainable inflection point – thereby reforming and rebuilding into an entirely new configuration.

In the face of severe disruption, the order and stability provided by a complex society is lost and replaced by increased levels of disorder and anarchy. In a workshop sponsored by Princeton University's PIIRS GSR, such loss of civilisational complexity and order was not necessarily due to the extreme degradation of the nodes of civilisational systems (i.e., centers of commerce, religious houses, or centers of legal and judicial authority), but may also be triggered by a disruption of the *linkages* between those nodes (i.e., transportation networks, trade routes, communications systems, environmental conditions that prevent collaboration of civilisations due to changing climactic conditions). As such, historical civilisational collapse is viewed as a systemic exercise, where such disruptions understood as a disruption of a basic nested system requirement (i.e., a link between nodes) that the society cannot survive without in its current form. Disruption to that nested system requirement cascades into other systemic losses and even collapse, and contributes to a reformation of a society or civilisation in an entirely new manner.

To address questions related to historical systemic threats and their potential to trigger societal or civilisational collapses, core questions discussed within Princeton's PIIRS GSR Workshop included:

- a) How fragile are the connections between all the system nodes?
- b) How could these be disrupted?
- c) How could such disruptions lead to a catastrophic breakdown?
- d) What would the costs associated be and for whom?

These questions require one to acknowledge that societal collapse is almost always due to a multivariate causal explanation – no single event or cause is responsible for the society's disruption and fall. Instead, multiple interlocking disruptions trigger feedback loops that amplify the effect of the disruption upon societal systems (environmental, sociopolitical, economic, etc).

Sources of collapse

Princeton's PIIRS GSR notes that one of the most common is the failure of political authority. This can take two stages: In its most obvious stage, is the breakdown of the monopolisation of control over the means of violence. In other words, authority and the administration of justice is increasingly localised to a regional or even household level, increasing the uncertainty that an encounter would lead to a legitimised violent outcome with few options for adjudication by a higher power.

In its most Hobbesian stage, it is literally "all against all" where "mere anarchy is loosed upon the world." In most fictional accounts of a global collapse the predominant political feature is rule by the strongest and best armed. In the most optimistic utopias, the only danger can come from outside the aggregated system. We can imagine levels of aggregation of the control of the means of violence from extended family through villages, small regions, nation states, and empires. Renfrew associates collapse with the decline of hierarchical authority structures, Marcus writes about the cyclical states of Mesoamerica.

A much more important breakdown might precede or succeed this: the diminution of disunion of the idea of communal legitimacy. All systems rely on some understanding of "the rules that need not be spoken." This can range from the epistemological authority, to ethical authority (what is good), to most mundanely bureaucratic authority.

Another source of collapse might be crises in economic production and consumption. Imagine a spectrum beginning with complete and total individual/household autarky and self-reliance (Robinson Crusoe). On the other, we are currently living through the most extensive economic aggregation of world history where virtually all societies depend on some sort of merchandise or capital flow. This can be measured by distance traveled by basic goods. The transition in the area of the Western Empire between the 4th and 13th C is a classic example of a cycle (Pirenne). You might also have battles about the distribution of social product. Here it is no longer the supply that is the cause of the crisis but disputes about who gets what and how much.

Another form of collapse might involve the cultural and physical segregation of individual. How much does any person have any contact with others? Or hear about others? How far away? We can imagine a society with limits of the sound of a human voice, a horn, or a church bell. Today, we possess the technology to make it almost impossible for people not be aware of what is going on thousands of miles away and to have at least intermittent contact with people around the world. While multiple connections are usually a sign of order, it can also lead to contagion.

An obvious source of collapse can come from the simple failure of the infrastructure of command and control of the military, the maintenance of roads (quality, distance, quantity), and of communications. Water systems are always critical (Ouellette; Orloff). We can also estimate the degree of a society's sophistication by examining its dependence on a certain technology or infrastructure. The most primitive one depends in the end of human motion through space, while today's depends on the maintenance of the internet.

The most common form of collapse might be summarised as the Biblical four horsemen: war, conquest, famine, and plague. Here we have to distinguish between the appearance of the horsemen that is due to cosmic bad luck and HOW a society responds and reacts.

What collapses when?

One of the most common critiques of the “collapse” literature is that it has a binary bias—civilisation and barbarie, that it occurs very quickly, and that once it does, it leaves little comprehensible residue. (Thus, the common Armageddon trope of post-apocalyptic savages not understanding what a car is or what buildings were for.) We can think of these as critical variables of the “timing” of collapse:

What survives? Just because a large system with a clear hierarchy collapses does not mean that daily life at the household level is disrupted. Those at the social bottom may not even recognise that anything has happened for some time. The key task here is to measure the territorial/demographic scope of fragmentation and aggregation. This is particularly useful for comparing parts of a previous system that undergoes dramatic change, but with very different consequences (e.g., The West/East Roman Empire(s)).

How long does it survive? With the new interest in resilience, we can also imagine a series of cycles of fragmentation and aggregation as the system reorganises itself. Thus, a “civilisation” or a cultural system may persist in its individual branches following a systemic “collapse.”

When we think of a collapse, we often imagine it in a short time span—a day as in the case of nuclear aftermath, a generation in terms of an eroding empire. But the velocity of a collapse may vary across time and space and we need to recognise “stages” or “tipping points.” We can also think of cycles. Yoffee discusses intervals of urbanism, ruralism, and local autonomy. Schwartz et al. focus on what they call “second generation” states following some initial disintegration. Holling provides a theoretical model of the cycles of ecosystem organisation that attest to “adaptive capacity.”

Why Does it Collapse?

We can divide these causes into two rough groups—exogenous and endogenous—but the line between these can become quite fuzzy. Cumming et al. suggest that we should pay attention to the interaction between pre-existing structures and the challenges faced. (See Avin et al. for a contemporary classification.)

The most obvious forms of these are exogenous to the system in question: invasions or natural catastrophes. In these instances, nothing about the society would lead one to expect a collapse, but the entry of some other factor (conqueror) or event (volcanic eruption) destroys the basis of the system. Weiss and Bradley and Diamond suggest that many historical collapses can be traced to significant climactic changes.

On the other side of causality, the system collapses because of its own endogenous qualities. Either it depends too much on a tightly knit and complex base, which cannot endure, or it creates an atmosphere where an elite becomes corrupt and no longer does the system maintenance required by its function in a society. The classic case here is what Tainter calls the decreasing marginal utility of complexity. Helbing focuses on the degree of interdependence and the possible limits of organisation. Downey et al. have identified a “boom-bust” cycle in the European Neolithic that seems to indicate a loss of resilience due to the introduction of agriculture. Diamond speaks of a hubris not allowing the necessary adaptation. Turchin et al., identify “principal components” of societal organisation that may explain crises and outcomes.

The most interesting cases result from the interaction of some exogenous factor (a new disease vector) AND the existing social structures. The most obvious example here are the many cases where social structures could not meet the challenge of an environmental change. Butzer focuses on how identified options, understanding and other aspects of resilience may explain different outcomes to similar crises.

Homer-Dixon, et al., focus on “synchronous” failure. Scheffer writes that societies may tend to resist change in environment until too late. Downey et al. do the same for Neolithic societies.

3b. Systemic Threats in History

While contemporary technology and the level of global integration may be new, many of the systems, mechanisms, dynamics, and fundamental foundations of civilisation (food, water, health/epidemiology, trade/transportation, political peace/security, and dependence on technologies) are the same. Seemingly different historical failures may have systemic commonalities that have not yet been studied from an interdisciplinary point of view.

We may begin with Joseph Tainter’s definition of social complexity: “the size of a society, the number and distinctiveness of its parts, the variety of specialised roles that it incorporates, the number of distinct social personalities present, and the variety of mechanisms for organising these into a coherent, functioning whole” (1988). The maintenance of this complexity requires ever more amounts of physical and social energy to maintain, and this in and of itself becomes an increasing strain on society. Peter Turchin has a similar fascination with what he might call “organisation” as described in the *Ultrasociety*. This may be best expressed by the exponential increase in both population and per capita energy use that has endangered our survival as a species. In short, we take for granted an unprecedented level of social organisation in the modern world, the fragility of which is a critical topic of study.

The second motivation follows the work by Kai Erikson and his belief that social life may sometimes be best understood through the prism of catastrophe. The argument is simple: if we wish to understand the most important social structures, we might best analyse what happens when these and their supporting institutions disappear. How much crime without police, how much illness without medicine, how much exchange without markets? When significant aspects of society come apart, we can better appreciate what they contributed to the status quo ante and how societies evolve to deal with their development. We have significant amounts of historical analysis of catastrophes, but have mined relatively little of this for sociological insights.

4. Methods for Identification and Management of Systemic Threats

4a. Methodological Input Requirements for Risk and Resilience of Systemic Threats

Risk quantification is an essential element of any risk or resilience management tool. Classification can be based on a variety of categorical or ordinal metrics, all based upon the type of risk quantification method utilised. Along with a consideration of the scope and severity of the hazards that may accrue from a given activity, such classification efforts are all largely dependent on the type and abundance of information available. With regards to consequences and severity, it is unlikely that a decision maker would place significant time and resources to the task of promoting a costly and time-intensive classification effort for a project with a small and inconsequential universe of potential negative outcomes. Likewise, decision makers would be less hasty in their efforts to push early-stage risk classification forward without thorough analysis (although there are tragic and famous cases to the contrary). While early stage classification efforts are imperfect in focus due to their unavoidably subjective nature, they generally serve as a reflection of the realities facing decision makers and stakeholders within their given field.

For cases where more objective information is available, risk quantification allows for greater precision with the risk classification effort (assuming, of course, that the data and model used are both relevant and rigorous). This precision is derived from robust sources of lab or field data that, if produced in a transparent and scientifically defensible manner, indicate statistically significant trends or indications of risk and hazard. Over time, multiple trials and datasets with similar indications ultimately contribute to risk profiles that establish best practices in the research, production, commercialisation, and disposal of various

substances, materials, and processes. As new studies and information become available, these best practices may be improved or updated to sharpen existing perceptions of risk. Because of this, the process of improving risk classification for a given project or material is continuously evolving. Just because quantitative information provides for more objective judgment, however, does not indicate a total absence of subjectivity in the risk classification process.

Conversely, there are many applications where objective data is not available. There are various reasons for this: acquiring such information may be legally or morally irresponsible, the application in question may be too novel for rigorous experimentation to have taken place, or the available data may be outdated or irrelevant for the particular risk application at hand. Such concerns are common with respect to new or emerging technologies or futuristic risks – greatly complicating the risk perception and management process in the midst of extreme uncertainty. Under such limitations, risk and resilience managers are required to turn to qualitative information such as through expert elicitation – a process which can set risk priorities in order if done correctly.

Regardless of the approach chosen to classify risk, resilience analysis quantification requires several component parts for it to be conducted in a transparent and defensible manner. While what follows is not an exhaustive list (particularly due to the fact that requirements of data reporting and project guidelines will differ strongly across infrastructural characteristics and security needs), it does offer some idea of the challenges faced by aspiring resilience analysts. With this in mind, these characteristics include (Merad & Trump 2020): (i) the availability of an outlined and transparent dataset, derived qualitatively or quantitatively, (ii) a framework or approach to process such data in a scientifically defensible and easily replicable fashion, (iii) pre-established notions of resilience success and failure, or various gradients of both, and (iv) considerations of temporal shifts that may strengthen or weaken system resilience in the midst of a variety of factors, including those which are highly unlikely yet particularly consequential.

With regards to dataset needs, such requirements do not strongly differ from more traditional risk assessment or other decision analytical methods. In this vein, risk and resilience analysis require a dataset with a clear connection to the host infrastructure or system and the various adverse events that could arise to threaten the said system, along with a consideration of how recent and relevant such data may be to decision making in the immediate or near future timeframe. In other words, regardless of the type of data collected (qualitative, quantitative, or a mix of the two), the dataset must have a clear and indisputable connection with the resilience project in mind. Unclear or fuzzy relationships with targeted projects, goods, or infrastructure may yield inaccurate or inefficient recommendations to bolster system resilience – effectively leaving stakeholders worse off than they were before they sought help from a resilience analyst. This should be a relatively simple exercise, however, as the dataset is either collected directly by the resilience analysts for their given project or is acquired from a similar project's data collection activities. Where no quantitative data is available, qualitative information may be classified in such a way as to serve as a temporary placeholder to allow analysis to continue (Vugrin et al. 2011). In this way, the type or quality of data available can directly inform the method chosen to process available information for resilience analysis (Francis & Bekera 2014; Ayyub 2014).

After acquiring a dataset relevant to the resilience issue at hand, the next requirement of resilience quantification is the need for a framework to process that data. Method selection is driven by a variety of factors, which most notably include the quality and robustness of available data as well as statutory requirements for output and transparency (Francis & Bekera 2014; Linkov et al. 2014). Additionally, certain clients may request specific methodologies to process their resilience analysis inquiry, yet this should only be undertaken if the necessary data are available and if the computation of such data with a particular method makes sense.

Throughout the process of conducting resilience thinking for a variety of projects, a crucial step in conducting resilience classification includes the imposition of pre-determined notions of system success or failure. This is consistent with virtually any other branch of scientific inquiry, where users must establish

some notion of ‘goodness’ or ‘badness’ that they seek to identify prior to data manipulation and results classification. Such efforts to establish system resilience success and failure could be generalist in nature (as in, deploying categorical variables or quantitative cutoff points that signify a positive or negative performance under certain stressors) or specific (as with the use of extensive quantitative data to inform precise points of system failure, as with the use of levees in flood management). Opting either way is at the discretion of the resilience analyst and their stakeholders along with a reflection of the degree of precision truly needed to assess system resilience. In other words, some stakeholders may be satisfied with answers of ‘there is a moderate probability that system x could fail under condition y’, yet for other cases, such stakeholders may need to know the exact conditions and points at which degradation and/or failure occurs. Generally, with more information available and the more potential the system has of incurring damage to society if broken, the greater precision is needed to assess systemic resilience success or failure.

Lastly, any resilience quantification effort must be considered under a wide array of temporal shifts that could dramatically alter system operations, efficiency, and challenges. The great challenge of resilience analysis and decision making is to consider a wide breadth of time horizons over which hazards and challenges could arise to shock a system, project, or infrastructural asset, where such events may not be projected to be possible until several years, decades, or even centuries into the future. Given this, an analyst should seek to discuss shifting preferences, threats, and system capabilities over time with stakeholders and managers in order to gain a more accurate view of how such a system may be challenged and behave in the midst of an external shock, with additional considerations with regards to how those systems could evolve and become more resilient over time.

One solution is to use a structured framework for selection metrics and organising the assessment. The individual performance factors are kept separate for more easy interpretation but can be aggregated to a single score, if relevant. The Resilience Matrix, described in the next section, provides a two-dimensional approach to selecting metrics, rather than a one-dimensional list of factors. More specifically, the Resilience Matrix explicitly incorporates the temporal phases of the event cycle, identified in the National Academy of Sciences definition of resilience: prepare, absorb, recover, and adapt.

4b. A Semi-Quantitative Approach: Resilience Matrix

A matrix assessment methodology affords users the capability to construct a framework that compares various decision metrics on a broad, ‘big picture’ level of resilience thinking and decision making. Given this, resilience matrix approaches will assist those local level stakeholders and policymakers focused on resilience performance along with broad and regional emergency response teams who seek to institute resilience thinking to “adopt a more holistic view of resilience necessary to reduce the impact of an adverse event” (Linkov et al. 2013). Collectively, the development and execution of such resilience matrices will provide robust and transparent policy guidance for national policy goals, while also offering improvements to large-scale system resilience for areas ranging from industry to energy to medicine (Kelic et al. 2013; Rosati et al. 2015; Roeger et al. 2014).

	PREPARE	ABSORB	RECOVER	ADAPT
Physical				
Information				
Cognitive				
Social				

Figure 4. Demonstration of a Resilience Matrix. The Y-axis includes domains of resilience, and the x-axis includes stages of resilience as established by NAS

Resilience is a property of any system and corresponds to the system's ability to perform critical functions in the midst of catastrophic and unexpected happenings. Described by Linkov et al. (2013), a resilience matrix collectively provides a unifying framework to assess system resilience which may be applied productively to societies and groups, when seen as systems (Figure 4). Linkov et al. (2013)'s formal Resilience Matrix (RM) classifies four general resilience domains of complex systems that include a mixture of physical infrastructure and more abstract capabilities, and takes into account the performance of these domains throughout the event's occurrence and disruption. The RM does not define specific metrics or attributes to use, but it gives guidelines to select the appropriate measurements to judge functionality from the perspective of a broader system. The RM guidelines diverge from the accomplishments of different community resilience progressions by taking advantage of a stakeholder-driven approach to characterise signs and ranges of system progression that are directly related to the aforementioned community. Through this method, the progression is characterised in relation to the necessities of the local environment rather than against the advancement of some generalised or national goal, which could or could not be acceptable in the local setting.

Cutter et al. (2014) reflect the difficulty to specify values of community resilience that are accepted nationally and no clear formulation for the approval of an external source of values of community resilience is given at this time (Cutter et al. 2014). As a consequence, the acceptability and usability of any resilience judgment can only be assessed by the very community in which it is being utilised. Even better, stakeholders are prompted to incorporate values from those identified by other resilience assessment strategies, where accessible, as important signals which connects the RM with other formulations to balance the strength of both approaches. The RM's simplified guidelines promote other strong attributes as well. Interdependences are ubiquitous in all systems, particularly in fields such as urbanisation, globalisation, and technological advances in which they are considered as particularly valid for communities. This being said, the time and cost it takes usually prohibits the investigation and modelling of all of these dependencies. The basic idea underlying the use of the RM is that in order to create resilience, achievement in all sectors of the system must be identified. This is different from the methodology of solutions in the past, which have maximised singular factors of the system. A consequence of such a narrow focus is that failures in the system can lead to cascading effects; the collapses of communities in light of calamities are frequently an effect of overflowing collapses from critical components in the system that are not identified as such. To be resilient on any scale, singular time steps cannot be relied upon to restore functionality. Even though the real relationship between system factors may not be revealed, by improving the resilience of all aspects of the system, performance can be kept or quickly restored. The Resilience Matrix methodology includes a general set of guidelines for the resilience

judgment for systems that has already been produced and idealised in the use of cyber, energy, engineering, and ecological.

The Resilience Matrix consists of a framework to conduct assessments regarding the performance of complex and of incorporated systems or projects across varying focal points. Generally, risk matrix frameworks consist of a 4x4 matrix, “where one axis contains the major subcomponents of any system and the other axis lists the stages of a disruptive event” (Fox-Lent et al. 2015). Next, matrix rows include the four primary domains to be considered within any systemic evaluation project, including physical, information, cognitive, and social (as noted in the US Army’s Network-Centric Warfare document) (Alberts & Hayes 2003). Additionally, matrix columns illustrate the four steps of disaster management, including the plan/prepare, absorb, recover, and adapt phases of resilience management as outlined by the National Academies of Science (Committee on Increasing National Resilience to Hazards and Disasters et al. 2012). Altogether, these sixteen cells give a basic description of the performance of the system throughout an adverse event.

Resilience is assessed by providing a value in each cell that summarises the capacity of the system to perform within that domain and period of time. As an example, the Information-Recover cell is provided with a score which reflects the functionality of the system to garner (monitor and detect) and share (analyse and disseminate) information that will provide relief in recovery. The Social-Adapt cell is provided with a score according to the ability of the system utilisers to change behaviour and keep changes beyond the reply to the initial incident. The matrix of values provides a visualisation, a snapshot of the overall resilience of the system, which can be tracked through a period of time, used to compare against comparable systems, or which may under closer inspection reveal holes in the performance, preparation, or organisation of the system (Eisenberg et al. 2014). In order to begin a resilience assessment utilising the matrix approach, Fox-Lent et al. (2015) recommend: (1) clearly outlining the system or project’s boundaries along with an array of hazard and threat scenarios that could impact the system, (2) enumerating critical system functions and capabilities that must be maintained throughout a crisis or shock, (3) selecting indicators for each critical function and subsequently compute performance scores in each matrix cell, and (4) aggregate all cells of the matrix—if necessary—to provide an overall system resilience rating, which will provide information about the system’s ability to respond to and overcome the effects of an external shock (Fox-Lent et al. 2015).

The RM method can be scaled to any observable system (from local to national to international). The system can be portrayed as business, a neighborhood community, a city, or even broader as an entire region. The system demarcation should be idealised on a geographic level, and the size of the boundaries will then direct how precise the indicators must be. Also, the scope of the threats being contemplated should be specified and recorded. Among others, this could include climatological disasters, human disasters (cyber-attacks, terrorism, spills, massive electrical grid failures, etc), or societal concerns (infectious disease, economic troubles, etc).

Looking at an extreme event, it is not necessary to always have every activity happening in a region to proceed undisturbed. Necessary aspects are the ones that have to be managed at nearly full scope to endure giving the necessary assistances of the system through the affair and to support the restoration of other activities after the affair. Most activities or note will be placed into divisions in relation to residents, economy, or ecosystem. Conceivable necessary functionalities for communities and areas are: housing/shelter, food and clean water, medical services, transportation, electricity, sewage, industry/commerce, ecosystem services, education, and recreation. In the RM method, each necessary functionality of the system is singularly judged by using the matrix. By practicing a judgment at this stage, the values may show that the system is resilient enough for one function but not as resilient for other functions, which services more useful information to advance improvement than a standardised community resilience value. The magnitude of necessary functions selected by those who instantiate the method should be kept down to 3-5 to ensure that the inquiry is of an acceptable size. Necessary functions will be different and will be dependent on the location, scale, history, and values of the community.

The Resilience Matrix utilises a civilian- or local expert-informed approach. The best utilisation is to bring together a group of representatives from the community to instantiate the judgment. This group should consist of experts that are well informed about the community itself, such as with municipal representatives from emergency management, community development, local threat management, and others concerned with the general welfare of a local citizenry. Each part of the matrix serves as a signal of the performance of the system’s given necessary function. Rather than figure a set of universally accepted values, the RM receives data based on local experience to find signals that have to do with the local problem. These indicators should be founded while taking into account some of the necessary characteristics of resilient systems that have been proclaimed by others—modularity, dispersion, redundancy, flexibility, adaptability, resourcefulness, robustness, diversity, anticipation, and feedback response (Park et al. 2013; Frazier et al. 2010) —and taking into account where each attribute is most reasonable with the system that is being observed. To act as a screening function, the RM allows for the utilisation of the most convenient and most significant data, whether it involves a numerical aspect or a qualitative aspect. In consequence, signals and values for each cell can be constructed in a number of ways:

1) An individual distinct number may be acceptable when it is a value that leads or can indicate how well the part of the system is doing. In order to realise how this measurable quantity affects the resilience value, the value itself must be put into terms relating to the problem itself; upper and lower ranges must signify sufficient performance and insufficient performance. Fox-Lent et al. (2015) and Linkov and Moberg (2011) indicate that these “two points define a linear utility function (unless sufficient information is available to suggest a nonlinear function), and the metric score is calculated as $(\text{metric value} - \text{lower bound}) / (\text{upper bound} - \text{lower bound})$, which results in a score between 0 and 1 (Linkov & Moberg 2011; Fox-Lent et al. 2015) (Figure 5).

	PREPARE	ABSORB	RECOVER	ADAPT
Physical	0.20	0.15	0.25	0.40
Information	0.32	0.22	0.34	0.12
Cognitive	0.53	0.05	0.26	0.16
Social	0.22	0.39	0.20	0.19

Figure 5. Strawman demonstration of the Resilience Matrix with single-unit outputs for each cell

2) When numerous values are big donors to the system’s functioning but have different values of achievement, an alternative option is to take a mean or weighted amount of these numerous values. The weighted amount should be completed after the individual values have been brought into context using a linear utility function.

3) If characteristics of the system are not fully accepted and an acceptable signal of the magnitude or stage of achievement cannot be seen, a functional checklist method could be utilised to develop a value. Using a broad list of needed values for functioning, the magnitude of items reviewed is a possible value. This method is easily utilised in cognitive domain where the magnitude of plans or range of planning activities and be conceptualised, but the tolerability of the plans is much more difficult to judge. The matrix can be improved as more specific limiting values are foreseen.

4) For cases where it is difficult to quickly denote specific decision metrics in order to assess overall systemic risk and performance, qualitative interviews with knowledgeable stakeholders and subject experts may help generate scores based their opinions, beliefs, and perceptions of resilience and risk for the target system (see Figure 6 for an example of a ‘stoplight’ qualitative approach). When adopting this approach to information acquisition and matrix population, these subject experts are required to simultaneously consider various factors in order to rank and score them based upon a pre-defined notion of success, failure, or a gradient somewhere in between. Such expert scoring mechanisms can either be categorical or ordinal based upon our general understanding of the system alongside the expert’s level of comfort giving certain answers in a high stakes environment. If each value is a specified measure but is used as a signal of the performance across the cell itself (system component), announcing the real judgments during the last output may prescribe immoderate precision for the screening purpose of the risk matrix. Instead of this, subsequent findings are contextualised into quintiles and illustrated as a colour-coded heat map of relative system resilience. Through this path, the results granted by the matrix will dial in on topics about the improvement of resilience on what aspects of the system can access scope goals rather than try to only enhance the signal utilised for the screening judgment.

	PREPARE	ABSORB	RECOVER	ADAPT
Physical	Red	Red	Yellow	Green
Information	Green	Yellow	Green	Red
Cognitive	Green	Red	Yellow	Red
Social	Yellow	Green	Yellow	Yellow

Figure 6. Strawman demonstration of the Resilience Matrix via a ‘stoplight’ approach. Red signals a significant lack of resilience with significant threat of system collapse, yellow indicates a minor lack of resilience with potential for system collapse, and green indicates acceptable system resilience and few plausible scenarios where a disruption could trigger system collapse.

For the majority of infrastructural projects and systems in need of resilience thinking, there exists multiple components and characteristics that must be considered throughout the decision making process. This is no different for resilience matrices, which break down the resilience problem for a given system into pre-determined parts. To acquire measures of overall system resilience for a given case, “the scores for each sector can be averaged across the critical functions to create a single matrix reporting general resilience” (Fox-Lent et al. 2015). Quite often, however, the stakeholders or policymakers conducting the resilience matrix shall attribute differing quantities of comparative rank across these various critical functions. Within each function, relative importance should be understood as criteria weights to be included within matrix cell aggregation for our defined critical functions, which ultimately generates a final resilience score that denotes system performance for predetermined management needs, perceptions, and goals.

The immediate use of the matrix method (relative to the youth of the field as it stands in 2015) is to find holes in the scope of the system to aid the most prosperous management of adversities (i.e., resilience). This knowledge can carve a path to prioritise continuous communal functions to provide certainty that the least achieving factors of the system are approached in a timely and just manner. With the breakdown of results, the RM method gives the option to allow for communication and promote connections. The first part of the stakeholder pact engages conversation between local civilians and faculties of the government to bring attention to the beginning development of projects to improve the local area to provide for

accessible alternatives. Coming second, the course of finishing the RM provides attention to the full scope of necessities of the system and provides appreciation to the fact that no government or community faculty has the experience, jurisdiction, or assets to watch over the resilience of the full system. For each necessary responsibility, the exact guideline can be utilised to encourage the development of co-operation and friendships among applicable management faculties by connecting necessary jobs and functions.

As an example, government faculties have usually concentrated on the plan and prepare temporal phase of resilience and usually work alongside each other only towards their own goal (Larkin et al. 2015). Forming these beneficial co-operative groups and partnerships could reduce costs by getting rid of unnecessary attempts and improve resilience by improving co-operation and networks of communication that ensure to be helpful in disaster relief missions. While the USACE can spend more in investing in the strengthening of coastal fortifications on the Atlantic side, without a partner to boost the fortification on the Bay side (which is outside the USACE authority and domain), any project would not be very effective. Many approaches are available for judging the resilience of the community, but not all approaches provide a final value that has been developed specifically to guide future action. The inevitable utilisation of resilience judgment tools such as the one described will be to give a standard achievement value on which the resilience boost future of recommended system adjustments can be analysed. At the very least, the values can be utilised to guide the choice between potential projects that are the simplest to finish or the most visible to the community as a whole. Even though it is not portrayed here, the matrix method can be adequately used to this output through an objective method.

Proposals can be judged by finding which signals of which necessary functions would be changed by instantiating the proposal and going through the calculations of the resilience values over again to compare against the standard. This method gives a procedure to list observed decreases of achievement in some system values with the advancement of achievement in others. As an example, mobile generator-run pumps may help with restoration for citizens in the area by providing another resource. However, unlike defined storm water pumps, mobile pumps discharge water that is untreated, which could reduce the restoration of the overall health of the ecosystem. In listing the effects of planned proposals across many necessary activities, the RM provides for a framework for decision making.

With all of this in mind, the main takeaway here is that the resilience matrix approach described in Linkov et al. (2013) possesses several assets due to the method's relative simplicity, transparency, and ease of use within multi-criteria decision methods for further evaluation and risk assessment needs (Yatsalo et al 2016; Collier & Linkov 2014). Resilience matrices have the capability to help benchmark and assess early stage system resilience through the use of available qualitative information, which could effectively assist stakeholders and policymakers identify gaps within that given system's capabilities that could result in catastrophic failure under certain conditions. It is worth noting, however, that resilience matrices do not intrinsically capture the various temporal characteristics within decision making that could cause shifts in preferences or needs over time, requiring users to continually update their matrix outlook at regular intervals. Under such circumstances, resilience matrices may currently be considered an optimal tool to gain a general view of system resilience that could be further refined with a more thorough decision analytic tool or modelling effort.

Ultimately, the resilience matrix approach offers a potential framework to compare and contrast various decision metrics from multiple disciplines that reside in the same matrix cells. In this way, such an approach will greatly assist those focused upon improving system and infrastructural resilience performance alongside those shareholders and managers required to prepare for and respond to emergencies to the systems in question. In both cases, resilience thinking allows its users to take on a holistic view of the process of bolstering systemic resilience properties by ensuring a) that the given system is adequately prepared for a host of potential challenges and b) that a variety of domain and temporal horizons are considered throughout the resilience evaluation process. Overall, the development and use of such frameworks should offer much needed guidance into policy implementation at the national and local levels for bolstering systemic resilience on a variety of projects and topics.

4c. A Quantitative Approach: Network Science

One example of a more quantitative approach towards modelling complex systems includes *network science*. In an attempt to computationally analyse complex interconnected systems, network science approaches for resilience are rooted in the premise that resilience necessarily has a temporal dimension to it. Indeed, as Holling (2001) points out, there are two conceptual ways to characterise resilience: the first, more traditional paradigm, concentrates on stability near an equilibrium steady state, where resistance to disturbance and speed of return to the equilibrium are used to measure the property; the second paradigm emphasises conditions far from any equilibrium steady state, where instabilities can flip a system into another regime of behaviour. The former group of methods define resilience from the engineering perspective, while the latter are termed ecological resilience. What is common between both paradigms is that they look at the dynamics of a system taking place in time: engineering resilience specifically mentions “speed of return to the equilibrium”, while ecological resilience looks at a “steady state” of a system. Another common property of both paradigms is the assumption that the state of a system needs to be measured at some points in time so that it is possible to determine whether the system has returned to the original equilibrium. Finally, it is important to notice that resilience is defined with respect to a disturbance or instability. With those three prerequisites in mind, quantitative approaches to resilience characterisation aim to investigate the evolution of a system in time both under normal conditions and under stress.

Most complex systems may be decomposed to simpler components with certain relationships between them. For example, transportation infrastructure may be represented as a set of intersections connected by roadways, global population may be mapped to a set of cities connected with airlines, railways, or automobile roads, ecological system can be decomposed to a set of species with ecological food-chain relationship. The decomposition is, of course, can be done in different ways and is carried out with respect to some predefined objective (e.g. traffic improvement for a transportation system, or species diversity preservation in the case of an ecosystem). In cases where such a decomposition is possible, it is often convenient to deploy methods developed in a branch of mathematics called graph theory, or network science. Network science represents a system under study as a set of points, called nodes, connected with relationships referred to as links. Dynamics of a system is then defined as a composition of individual node states, which in turn depend on their neighboring nodes as well as on internal and external factors (an example of this is illustrated in Figure 7).




Figure 7. Example demonstration of linkages between nodes in a directed network graph

When approaching the problem of resilience characterisation from a network science perspective an important question is whether the approach should be threat agnostic or not. Threat agnostic approaches

(Ganin et al. 2016; Linkov et al. 2018) maintain that resilience is defined regardless of a specific threat that hits the system. The rationale here is that it is often impossible to predict what hits the system, how much of a disruption will ensue, and what the likelihood of a threat scenario is. The opposite group⁽⁷⁻¹¹⁾ of methods define resilience by modelling a specific threat. Those frameworks often imply and require that a probability be assigned to each threat as well as that an algorithm be defined to model how a threat affects the network. While in the world of perfect knowledge, the latter approaches may offer a more realistic way to prioritise resilience-enhancing investments, they appear to convolute resilience analysis with risk analysis. Such a convolution is not necessarily beneficial. For example, Ganin et al. (2016), who argue that risk and resilience analyses should be complementary, but separate of each other, claim that resilience analysis is, in part, motivated by the imperfect knowledge about the threat space.

Before we move forward with network resilience approaches, we introduce some terminology and network classifications. Specifically, networks may be directed and undirected. In undirected networks, links do not define a node serving as an origin or a destination, both nodes are equivalent with respect to the relationship defined by the link. An example of such a relationship is friendship in a social network. In a directed network, nodes are not equivalent with respect to the link, e.g. in a transportation network, a road goes from one point to the other and, not necessarily, the other way. Mixed networks may contain both directed and undirected links. Another important class of networks are interconnected networks. In the case of two interconnected (coupled) networks, system nodes may be logically separated into two sets or layers corresponding to each network. Consider for example, a power distribution system and an information network controlling the former. In this case, links among power nodes may represent transmission lines and links among computers may define cyber connections. Notably, computers need power to function while the power distribution is controlled by the cyber system. This interdependency may be defined with links going from nodes in one layer to nodes in the other and vice versa.

One of the most resilience-relevant problems studied in graph theory is the connectivity in graphs. It is argued, that the state of the system is defined by the size of the network's largest connected component. The larger the set of connected nodes is, the better the system is able to function. Indeed, while there may be many better mappings from the graph representation to the state of a system, connectivity is one of the most universal ways of such a mapping. Moreover, many networks are engineered to be connected on purpose, which also justifies this mapping. Well-known classical results here are based on percolation theory, where nodes and/or links are removed at random and the connectivity of the remaining sub-network is analysed. Percolation theory establishes that the distribution of links among network nodes (degree distribution) is a key characteristic in determining network robustness (Kitsak et al. 2010; Linkov et al. 2013). Yet, percolation modelling typically results in a point estimation of connectivity, or robustness, after a random removal of a certain number of nodes or links and does not look at system dynamics. One step towards resilience quantification is to look at the stability of connectivity between nodes in multiple percolations. As nodes or links are removed at random, different disruptions disconnect different sets of nodes. Based on a graph's degree distribution, Kitsak et al. (2010) answer the questions of what nodes will be connected in multiple percolations, what the likelihood is, and how the size of the persistently connected component changes with the number of percolations. The authors look at both single and coupled undirected networks.

As we said earlier, resilience can be either defined as engineering or ecological. The next group of approaches we look at builds on the engineering definition. Specifically, those frameworks aim at modelling both a disruption and a system's recovery process and define a performance function of the system. For example, Ganin et al. (2016) define a system's critical functionality as "a metric of system performance set by the stakeholders, to derive an integrated measure of resilience". Critical functionality serves as a function of time $K(t)$ characterising the state of the system. Resilience is evaluated with respect to a class of adverse events E (or potential attacks on targeted nodes or links) over a certain time interval $[0, T_C]$ where T_C is the *control time* (Kitsak et al. 2010) <https://www.nature.com/articles/srep19540 - ref41> which can be set *a priori*, for instance, by stakeholders or estimated as the mean time between adverse events.

Provided that $K^{nominal}(t)$ represents the system's normal critical functionality, resilience is mapped to a value R between 0 and 1 as follows (Figure 8).

$$R \equiv \frac{\frac{1}{|E|} \sum_E \int_{t=0}^{T_c} K(t)}{\int_{t=0}^{T_c} K^{nominal}(t)}$$




Figure 8: A generalised resilience profile, where a system's resilience is equal to the area below the critical functionality curve (reproduced from Ganin et al. 2016).

Using one potential network science-driven approach, the authors demonstrated the formulation on two classes of models: 1) multi-level directed acyclic graphs, and 2) interdependent coupled networks. For both models, synthetic case studies were used to explore trends. For the first class, the approach was also applied to the Linux operating system. Results indicated that desired resilience and robustness levels were achievable by trading off different design parameters, such as redundancy, node recovery time, and backup supply available. The nonlinear relationship between network parameters and resilience levels confirmed the utility of the proposed approach.

Performance functions capture both absorption and recovery resilience phases defined by the National Academy of Sciences, but stop short of a giving a straightforward way to address planning and adaptation. Moreover, these approaches often need to be tailored to a specific system so that its disruption response is captured meaningfully. Finally, it is not always possible to enumerate, let alone, model, all possible disruptions in the class of adverse events considered. As a way to address the last issue, it is proposed to perform a series of Monte-Carlo simulations.

Applications of performance function methods to realistic systems include studies of malware spreading in a computer network (Linkov et al. 2019), where the authors studied the tradeoff between overregulation and underregulation of computer users arguing that too many rules may result in some rules being neglected and, in fact, result in a lower resilience. Another example of the above methods is epidemic modelling in a metapopulation network (Massaro et al. 2018), where it was found that travel restrictions may be harmful to a system's resilience. Specifically, critical functionality was defined based on the number of people infected and the number of people restricted from travel. Insufficient travel restrictions were shown to only slow down the epidemics without significant changes to the final number of infections. As people movement was diminished for a longer period of time, resilience was lower than that without any restrictions.

An example of an ecological resilience inspired approach is the analytical framework proposed by Gao et al. (2012). The authors focused on the stable points of the equation characterising the system state. They start by looking at a one-dimensional (1D) nonlinear dynamic equation $\frac{dx}{dt} = f(\beta, x)$ where x is a system's

state, t is time, parameter β captures the external conditions and function f defines the dynamics of the system. Moving from the 1D system to a network, the authors change the equation as follows:

$$\frac{dx_i}{dt} = F(x_i) + \sum_{j=1}^N A_{ij}G(x_i, x_j)$$

Above, x_i define node states, N is the number of nodes, function $F(x_i)$ shows the internal dynamics of nodes, function $G(x_i, x_j)$ and matrix $A_{ij} > 0$ represent interactions between node pairs. The authors proposed a way to reduce the N -dimensional state of the system to a 1D model:

$$\frac{dx_{eff}}{dt} = F(x_{eff}) + \beta_{eff}G(x_{eff}, x_{eff})$$

The one-dimensional state x_{eff} is defined based on the nodal degree distribution as well as individual node states x_i , while β_{eff} depends only on the degree distribution and is mathematically easy to evaluate. Thus, thresholds corresponding to β_{eff} values leading to the system moving to a new equilibrium state may be found from the above 1D equation instead of a system of N equations. The approach was illustrated on gene regulatory networks, ecological species, and multiple other networks (Ganin et al. 2016). The analytical results unveiled the network characteristics that can enhance or diminish resilience, offering ways to prevent the collapse of ecological, biological or economic systems, and guiding the design of technological systems resilient to both internal failures and environmental changes.

An example of a mixed approach to resilience evaluation is given by Ganin et al. (2017). The authors have looked at the resilience of an urban transportation network. Specifically, the authors quantified a transportation system's efficiency through delays experienced by auto commuters under normal conditions and resilience as additional delays ensuing from a roadways disruption. The approach borrowed from engineering resilience by allowing traffic redistribution, which may be viewed as recovery, and from ecological resilience by studying the resulting steady state equilibrium achieved by the system. The authors evaluated resilience and efficiency in 40 real urban areas in the United States. Networks were built by mapping intersections to nodes and roadways to links. The authors proposed graph theory inspired metrics to quantify traffic loads on links. Based on the loads they evaluated delays. The results demonstrated that many urban road systems that operate inefficiently under normal conditions are nevertheless resilient to disruption, whereas some more efficient cities are more fragile. The implication was that resilience, not just efficiency, should be considered explicitly in roadway project selection and justify investment opportunities related to disaster and other disruptions (Ganin et al. 2017).

5. Ideas for Methodological Practices for Resilience: Making Resilience Useful for Decision Makers

Due to its relative infancy as a modern method of determining risk and system robustness, no single method has been solidified as the 'go-to' approach for conducting resilience analysis. Generally speaking, this would be a significant limitation for the quantitatively and methodologically driven, which view standardisation as the ability of resilience, like other ideologies, to be generalised to a variety of fields and cases seamlessly. In this, proponents of standardisation are not entirely wrong, as the sheer diversity of cases in which resilience thinking is proposed requires some movement towards consistency and method objectivity. However, rather than developing and implementing a single standard for resiliency across disciplines and countries, a more effective approach would be to further a suite of methods and tools that can be utilised and modified based upon a country, company, or discipline's institutional, political, economic, and cultural incentives and needs.

The main barrier to furthering this suite of methods and analytical tools currently facing the resilience field is its lack of a formal definition or centralised governing body. Resilience thinking and resilience analysis possesses different meanings for differing disciplines, which will only become more entrenched and divided as time passes. Should no consensus definition be reached, there still remains the possibility that some shared meaning may be held to different types of resilience methodologies – specifically in qualitative, semi-quantitative, and fully quantitative work. This is driven by the mathematical and logical backgrounds of these methods, which require something of a shared language across users (albeit with differing flairs and twists based upon stakeholder need) in order to convince their audience that their method’s findings are legitimate and acceptable.

For qualitative methods, tool development is a bit easier due to a reduced reliance upon strict mathematical tools and more upon the need to acquire information for an emerging topic of high uncertainty and risk. Despite the different and specified needs of various disciplines when utilising qualitative approaches to resilience thinking, the general approach expressed by most users is one of user-defined categorical metrics which are filled out by a pre-determined list of subject experts or lay stakeholders. In such an exercise, the opinions of such experts serve as indicators of risk and system resilience, and offer a context-rich view of resilience decision making for a particular case. In this way, qualitative methods share a common root function of eliciting feedback from the world at large and processing results in a transparent and meaningful way, making qualitative methodologies inherently generalisable despite intellectual differences across disciplines.

We would argue that the main hurdle towards resilience-based method and tool development for qualitative methods centers on their overall acceptance by the quantitative community. Across various disciplines, arguments have been levied at qualitative methodologies’ lack of objectivity in pursuit of scientific understanding, with quantitative methods and mathematical approaches being easier to accept and verify (Mahoney & Goertz 2006; King et al. 1994). However, we contend that as resilience is used to tackle cutting edge emerging systems and systemic threats, quantitative information may not always be available or useful to resolve a context poor situation (Ritchie et al. 2013; Trump et al., 2018). In this way, qualitative research in resilience thinking and analysis will help bridge initial gaps in risk understanding by offering an expert-driven view of a system’s resilience for a given array of external shocks and challenges.

Semi-quantitative methods may help assuage the concerns of the quantitatively driven due to the use of mixed qualitative and quantitative data in evaluating resilience decisions. Specifically, risk matrices categorise available objective data into a small set of classification factors that inform overall resilience decision making – allowing for a transparent and scientifically defensible method of conducting resilience analysis. While some information is lost in the transformation of quantitative data to qualitative categorical metrics, this method can simplify resilience-based decision making by breaking down systemic factors into a small number of easily understood subsets. Additionally, this method allows its users to integrate qualitative information and elicited expert opinion alongside available data, bringing some additional context to the available dataset. This approach has a slightly steeper learning curve than traditional qualitative research methods; however, resilience matrices require some fundamental understanding of the math behind matrices as well as understanding of proper use (to avoid the garbage-in, garbage-out problem that haunts any decision analytical tool), which may prevent some from placing such matrices in their resilience tool kit.

Quantitative methods such as network science enjoy perhaps the greatest level of trust amongst lay stakeholders due to the perception of objectivity and raw scientific explanatory power in a variety of applications. Under the assumption of correct math, lay stakeholders can physically witness the transformation of data into rigorous findings of risk and benefit, and may ultimately help pave the way to notions of causality for a particular application of resilience management. Where valid data is plentiful, quantitative methods can go a long ways towards advancing most fields in science, let alone resilience thinking. However, incompleteness or lack of clarity in existing data can put a damper on such research, and even the mathematical method used to generate objective outcomes may in itself be inherently

subjective. An additional concern includes the even steeper learning curve than qualitative or semi-quantitative methods; users are frequently tasked with mastering advanced formulas or computer programs even prior to looking at a dataset. Often, implementation of such methods will require a model to be custom-built by an external consultant or academic. This is not to discourage the use of such methods – their contribution to science is extensive and frequently proven across virtually all fields – yet we cannot leave this section without noting the drawbacks of quantitative-only approaches to resilience along with the complications that their users will face in the midst of high uncertainty and context-poor information limitations.

Perhaps the most important difficulty which is not treated here is in the definition of a graph in a particular context. For example, in Haldane and May's well-known papers on systemic risks in networks, a node is a bank or a country and the links represent net indebtedness with the other node. Yet, this masks the fact that the bank itself is a network of interacting agents and desks as Bookstaber has pointed out in his "End of Theory" and that the lack of resilience may well come through the fragility of the internal network as the failure of Lehman Brothers shows (Zigrand 2015).

By its nature, resilience classification is difficult. If it were not so, there would be little need to discuss the pros and cons of differing ideologies of resilience practice, let alone write a book on how each method functions in this growing field. However, when used properly, these methods can do much good through their ability to inform complex and uncertain resilience analysis and decision making by providing some structure for any methodological venture (Jackson 2018). More methods have and will undoubtedly continue to creep into the field as more disciplines come to embrace resilience thinking, making full method standardisation unlikely. Yet this may open the door for shared fundamental concepts of resilience analysis across ideological and theoretical divides. In other words, each individual discipline will transform resilience to fit its own needs, yet these methods will serve as the cornerstones of a structure that will allow resilience thinkers to have a shared philosophical discussion.

5a. Resilience and Efficiency

Most systems are under considerable stress to improve their *efficiency*. Colloquially, efficiency may be understood as achieving maximum productivity with minimum wasted effort or expense. Where waste is treated as 'lost opportunity' or 'unrealised potential', organisations ranging from private businesses, to industrial production, to transportation and urban planners, to economic policy analysts all seek to identify ways of increasing systemic efficiency by either reducing waste or increasing output per unit of energy invested within a given activity.

Generally, efficiency measures are predicated upon eliminating unneeded redundant systems or resources that have little to no discernable value in the short to intermediate term. For example, for a postal service, why should I purchase, operate, and maintain one hundred delivery trucks for a city when I could accomplish the same goal in essentially the same time with fifty? The resource costs of doing more than is necessary can be considerable, and individuals, companies, and governments alike have rational incentives to operate in a lean manner that minimises waste for maximum productive return.

This may be a fine strategy under normal operating conditions, when externalities are low and disruptions are minimal and predictable. Yet, when situational and environmental conditions shift, or a sudden disruption occurs, a lack of redundant capacity or alternative systemic configurations may leave my business, government, household, or other organisational unit unable to cope with losses in core system functions. Returning to our postal example, say a blizzard or other storm arrives in a manner that slows down or renders inoperable a number of my postal trucks. With one hundred such trucks, we have enough redundant capacity to absorb the disruption and ensure that mail service will continue with minimal delay. However, at fully efficient operations with fifty trucks, any additional disruption will lead to system losses (i.e., the inability to deliver mail to certain parts of my network in a required timeframe). This example is one of relatively low consequence, yet other examples may illustrate the dangers of hyper-efficiency and

a lack of preparedness for system absorption of a disruptive catalyst. For example, hospital configuration and design allows them to be able to meet the increased needs of a certain number of patients in the midst of an epidemic crisis or other disaster, and thereby intake and treat far more patients than their normal operating procedures account for.

All systems require a certain level of resilience to function, with that level changing based upon the needs and degree of importance that such systems play into modern society. Modern aircraft are designed for considerable resiliency, as a lack of quick recovery from an adverse event could contribute to the deaths of dozens or hundreds of people. Similarly, economic supply chains require some element of resiliency to ensure the macroeconomy is sustainable in the event that one node (i.e., producer of materials, goods, or services) or one link (i.e., transportation or delivery mechanisms of materials, goods, or services) are degraded or disrupted from one of an infinite number of causes. For systemic threats, balancing efficiency and resilience is a matter of survival, where critical systems (i.e., water, energy, communications, security, food, etc) must be maintained above a basic needs level in order to ensure that a disruption in one area does not cascade to many others.

5b. Future Systemic Challenges Facing Society – the Promise and Peril of Digitalisation and AI

Though the nature of systemic threats are difficult to accurately predict, the incoming technological revolution due to artificial intelligence (AI) will likely have radical effects upon the global society. Economically, socially, environmentally, and politically, AI will reshape human activity and behaviour consistent with the technological benefits that AI-based technologies will offer, such as the more efficient provision of services to improved analytics and decision making capabilities. Yet, these improvements come at a considerable cost.

One major concern includes ‘The Future of Work.’ Where past technological revolutions (Agricultural, Industrial, etc) yielded new jobs for every position that was eliminated through scientific and technological progress, there is less certainty that an AI revolution will do the same. The world is already experiencing significant obsolescence of human labor in various fields (older taxi systems, clerical and records support, other entry-level positions regarding information storage and analysis), and is poised to reduce the need for extensive human labor in others in the near future (i.e., retail, food preparation, transportation/freight, education, and others). The mass exodus of humans from the workplace will trigger substantial societal uncertainty related to socioeconomic sustainability and social identity on one hand, to taxation, policymaking and implementation, and economic growth on the other.

Disruptions posed by AI will arrive at a time where the globalised world is undergoing ‘digitalisation’. Digitalisation, defined as the increased connectivity and networking of digital technologies to enhance communication, services, and trade between people, organisations, and things, has been posited as both an emerging opportunity and as a challenge to the United Nations (UN) Global Sustainable Development Goals (SDGs), comprising 17 goals and 169 targets or objectives. The growth and maturation of the digital world, where an increasing scale of individual and communal activities are being recorded, digitised, and analysed for future technological improvement, is creating unique opportunities to enhance social and environmental well-being, and further improve global standards of living while preserving and improving environmental health for future generations. Nevertheless, digitalisation is increasingly shown to also enhance the likelihood of social and environmental sustainability challenges and threats, including the carbon footprint associated with increased electricity generation demand, cybersecurity vulnerabilities, and social discrepancies posed by the widening gap in access to information and communication technologies, commonly referred to as the “digital divide” between those who benefit from a digital economy, and those who may lose jobs, economic resources, or other social benefits.

Between the increasing digital interconnectedness of peoples and societies on one hand and the emerging AI revolution on the other, social and economic activity within the next 50 years will be largely

unrecognisable using the metrics and frameworks of daily life in the current age. How will people/families provide for themselves? What will the educational and societal priorities be? How will economies grow, function, and provide for the necessities and luxuries demanded by the masses? How will governance be executed, and governmental systems function? What is the future of war, peace, and global prosperity?

These are all questions that have no easy answer. Yet, they are all driven by the common question: How will Digitalisation and AI reshape/disrupt society, and how will society's core systems transform and adapt to this inevitable scientific technological revolution? A cynical answer is one driven by a Hobbesian approach, indicating that older Westphalian and neoliberal economic systems will not be able to provide for the comfort and wellbeing of most individuals, let alone their larger states. Still others view AI and Digitalisation as an opportunity for societies to adapt to technological disruption in a planned and normatively positive manner, and position their societies for improved medical, economic, and social lives and livelihoods that are raised by the benefits afforded by AI. The future is highly uncertain here, and more research is needed to understand how (a) AI/Digitalisation disrupts societal systems and networks, (b) how such systems will transform and reform into a sustainable state given the shift in such technologies on a national and global scale, and (c) what types of measures might shift the trajectory of individual and global societies towards a more normatively positive or negative outcome?

As this is fundamentally a question of systemic threat, societal resilience to the socioeconomic consequences wrought by AI/Digitalisation is a subject worthy of further inquiry. How does an increase in economic efficiency due to AI/Digitalisation affect economic and social resilience of one or more countries? Who are the 'winners and losers' of such transformations? What types of policies and activities increase sociopolitical and socioeconomic resilience over time against such disruptions, and which decrease it? Answers to these questions will differ based upon national and cultural contexts, yet are worthy of exploration due to the considerable technological consequences that are not too distant in the future.

References:

- Adger, W. N., Hughes, T. P., Folke, C., Carpenter, S. R., & Rockström, J. (2005). "Social-ecological resilience to coastal disasters." *Science*, 309(5737), 1036-1039.
- Ayyub, B. M. (2014). "Systems resilience for multihazard environments: definition, metrics, and valuation for decision making." *Risk Analysis*, 34(2), 340-355.
- Berghel, H. (2015). "Cyber Chutzpah: The Sony Hack and the Celebration of Hyperbole." *Computer* 2, 77-80.
- Berkes, F., & Jolly, D. (2002). Adapting to climate change: social-ecological resilience in a Canadian western Arctic community. *Conservation ecology*, 5(2), 18.
- Björck, F., Henkel, M., Stirna, J., & Zdravkovic, J. (2015). "Cyber Resilience—Fundamentals for a Definition." In *New Contributions in Information Systems and Technologies*, Springer International Publishing, 311-316.
- Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., ... & von Winterfeldt, D. (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. *Earthquake spectra*, 19(4), 733-752.
- Centeno, M. A., Nag, M., Patterson, T. S., Shaver, A., & Windawi, A. J. (2015). The emergence of global systemic risk. *Annual Review of Sociology*, 41, 65-85.
- Collier, Z., DiMase, D., Walters, S., Tehranipoor, M. M., Lambert, J. H., & Linkov, I. (2014). "Cybersecurity standards: managing risk and creating resilience." *Computer*, 47(9), 70-76.
- Cross, J. A. (2001). Megacities and small towns: different perspectives on hazard vulnerability. *Global Environmental Change Part B: Environmental Hazards*, 3(2), 63-80.
- Cutter, S. L., Ash, K. D., Emrich, C. T. (2014). The geographies of community disaster resilience. *Glob Environ Change*, 29, 65–77
- Ebi, K. L., & Semenza, J. C. (2008). "Community-based adaptation to the health impacts of climate change." *American journal of preventive medicine*, 35(5), 501-507.
- Eisenberg, D. A., Park, J., Bates, M. E., Fox-Lent, C., Seager, T. P., Linkov, I. (2014). Resilience metrics: lessons from military doctrines. *Solutions* 5(5):76–87
- Fox-Lent, C., Bates, M. E., & Linkov, I. (2015). A matrix approach to community resilience assessment: an illustrative case at Rockaway Peninsula. *Environment Systems and Decisions*, 35(2), 209-218.
- Francis, R., & Bekera, B. (2014). "A metric and frameworks for resilience analysis of engineered and infrastructure systems." *Reliability Engineering & System Safety*, 121, 90-103.
- Ganin A. A., Kitsak, M., Marchese, D., Keisler, J. M., Seager, T., & Linkov, I. (Dec. 2017). "Resilience and efficiency in transportation networks," *Science Advances*, 3(12), e1701079.
- Ganin, A. A., et al. (Jan. 2016). "Operational resilience: concepts, design and analysis," *Scientific Reports*, 6, 19540.
- Gao, J., Buldyrev, S. V., Stanley, H. E., & Havlin, S. (2012). Networks formed from interdependent networks. *Nature physics*, 8(1), 40.
- Helbing, D. (2012). Systemic risks in society and economics. In *Social Self-Organization* (pp. 261-284). Springer, Berlin, Heidelberg.
- Holling, C. S. (2001). Understanding the complexity of economic, ecological, and social systems. *Ecosystems*, 4(5), 390–405.
- Jackson, P. (2018). *The New Systemic Risk: Adoption of New Technology and Reliance on Third-Party Service Providers*. Thompson Reuters.
- Karvetski, C. W., Lambert, J. H., Keisler, J. M., Sexauer, B., & Linkov, I. (2011). Climate change scenarios: Risk and impact analysis for Alaska coastal infrastructure. *International Journal of Risk Assessment and Management*, 15(2-3), 258-274.
- Kaur, S., Sharma, S., & Singh, A. (2015). "Cyber Security: Attacks, Implications and Legitimations across the Globe." *International Journal of Computer Applications*, 114(6).
- King, G., Keohane, R., & Verba, S. (1994). "Scientific inference in qualitative research." Princeton, NJ: Princeton University Press 1, 994.

- Kitsak, M., Gallos, L. K., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. E., & Makse, H. A. (2010). Identification of influential spreaders in complex networks. *Nature physics*, 6(11), 888.
- Larkin, S., Fox-Lent, C., Eisenberg, D. A., Trump, B. D., Wallace, S., Chadderton, C., & Linkov, I. (2015). Benchmarking agency and organisational practices in resilience decision making. *Environment Systems and Decisions*, 35(2), 185-195.
- Linkov, I., & Trump, B. D. (2019). *The science and practice of resilience*. Springer.
- Linkov, I., Moberg, E. (2011). *Multi-criteria decision analysis: environmental applications and case studies*. CRC Press, Boca Raton
- Linkov, I., Poinssatte-Jones, K., Trump, B. D., Ganin, A. A., & Kepner, J. (2019). Rulemaking for Insider Threat Mitigation. In *Cyber Resilience of Systems and Networks* (pp. 265-286). Springer, Cham.
- Linkov, I., Trump, B. D., & Fox-Lent, K. (2016). Resilience: Approaches to risk analysis and governance. IRGC, Resource Guide on Resilience, EPFL International Risk Governance Center, Lausanne, v29-07-2016, available at <https://www.irgc.org/riskgovernance/resilience/> (accessed at 10.10.2016).
- Linkov, I., Trump, B. D., & Keisler, J. (2018). Risk and resilience must be independently managed. *Nature*, 555(7694), 30-30.
- Linkov, I., Trump, B. D., Poinssatte-Jones, K., Love, P., Hynes, W., & Ramos, G. (2018). Resilience at OECD: Current State and Future Directions. *IEEE Engineering Management Review*, 46(4), 128-135.
- Lino, C. (2014). "Cybersecurity in the federal government: Failing to maintain a secure cyber infrastructure." *Bulletin of the American Society for Information Science and Technology*, 41(1), 24-28.
- Longstaff, P. H. (2005). "Security, resilience, and communication in unpredictable environments such as terrorism, natural disasters, and complex technology." Center for Information Policy Research, Harvard University.
- Lucas, K., Renn, O., & Jaeger, C. (2018). Systemic Risks: Theory and Mathematical Modeling. *Advanced Theory and Simulations*, 1(11), 1800051.
- Mahoney, J., & Goertz, G. (2006). "A tale of two cultures: Contrasting quantitative and qualitative research." *Political Analysis*, 14(3), 227-249.
- Massaro, E., Ganin, A., Perra, N., Linkov, I., & Vespignani, A. (2018). Resilience management during large-scale epidemic outbreaks. *Scientific reports*, 8(1), 1859.
- Merad, M., & Trump, B. D. *Expertise Under Scrutiny*. Springer International Publications
- Murray, P., & Michael, K. (2014). "What are the downsides of the government storing metadata for up to 2 years?" 5-30
- National Academy of Sciences (NAS) Committee on Science, Engineering, and Public Policy. (2012). "Disaster Resilience: A National Imperative," The National Academies Press.
- Osawa, J. (2011). "As Sony counts hacking costs, analysts see billion-dollar repair bill." *The Wall Street Journal*.
- Pelling, M. (2012). Resilience and transformation. *Climate Change and the Crisis of Capitalism: A Chance to Reclaim Self, Society and Nature*, 51-65.
- Petrie, C., & Roth, V. (2015). "How Badly Do You Want Privacy?." *IEEE Internet Computing* 2, 92-94.
- Renn, O. (2016). Inclusive resilience: A new approach to risk governance. An edited collection of authored pieces comparing, contrasting, and integrating risk and resilience with an emphasis on ways to measure resilience, 184.
- Ritchie, J., Lewis, J, Nicholls, C. M., & Ormston, R., eds. (2013). *Qualitative research practice: A guide for social science students and researchers*. Sage.
- Trump, B. D., Florin, M. V., & Linkov, I. (2018). IRGC Resource Guide on Resilience (Volume 2) (No. BOOK). International Risk Governance Center (IRGC).
- Vugrin, E. D., Warren, D. E., & Ehlen, M. A. (2011). "A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane." *Process Safety Progress*, 30(3), 280-290.
- Williams, P. A., & Manheke, R. J. (2010). *Small Business-A Cyber Resilience Vulnerability*. Chicago
- Wood, M., Kovacs, D., Bostrom, A., Bridges, T., & Linkov, I. (2012). Flood risk management: US Army Corps of Engineers and layperson perceptions. *Risk Analysis: An International Journal*, 32(8), 1349-1368.

Yatsalo, B., Gritsyuk, S., Sullivan, T., Trump, B., & Linkov, I. (2016). Multi-criteria risk management with the use of DecernsMCDA: methods and case studies. *Environment Systems and Decisions*, 36(3), 266-276.

Zhao, J. J., & Zhao, S. Y. (2010). "Opportunities and threats: A security assessment of state e-government websites." *Government Information Quarterly*, 27(1), 49-56.

Zigrand, J-P. (2015). How Can We Control Systemic Risk? World Economic Forum. <https://www.weforum.org/agenda/2015/08/how-can-we-control-systemic-risk/>