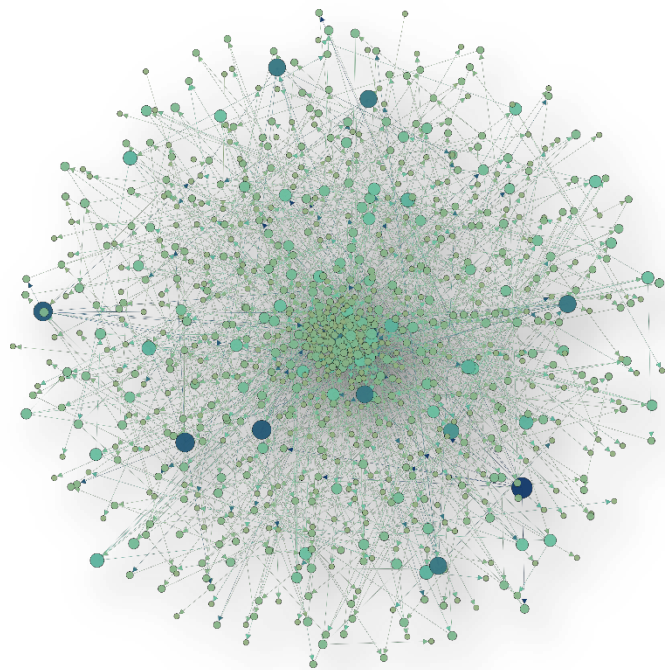

Systemic Risk: How connectivity impacts risk management practice

An Awareness Paper



EXECUTIVE SUMMARY

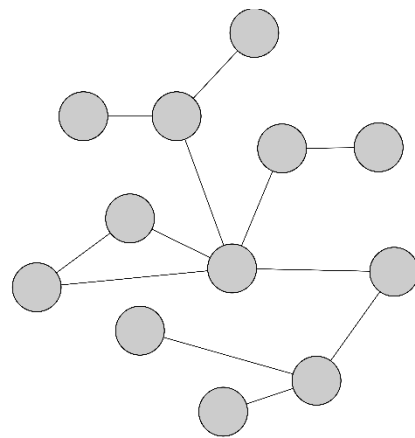
Low-frequency, high-impact events are occurring at a much faster rate than anticipated by traditional risk management assessment techniques, disrupting businesses, their performance and long-term viability. Recent examples of such systemic events include 2003 US Power Grid failure (est. cost= 10 billion USD; BP's 2006 Deepwater Horizon oil spill (est. cost = 40 billion USD); 2008 Lehmann Brothers collapse (est. cost = 2.2 billion USD); 2010 Iceland volcano impacting air flights (est. cost= 2.13 billion USD) and 2013 Tesco horsemeat scandal.

Despite contextual differences, complex network science provides a framework for understanding such large-scale, systemic events. In doing so, the *interdependent* nature of risk is highlighted, compared to the traditional view of risk independence. By shifting the focus from assessment of individual risks towards understanding the interdependency of the underlying network, exposure to systemic risk can be assessed and subsequently minimised. Network analysis provides the key in:

- Assessing the robustness and resilience of an organisation to systemic risk, based on the extent, and nature of its underlying network.
- Measuring the susceptibility of an organisation to systemic risk in an auditable, objective and quantitative way.
- Identification of key drivers of systemic risk and mitigations measures that can be applied at a local and/or global level.
- Monitoring organisational changes that may alter the network architecture in such a way as to raise the exposure of the organisation's to systemic risk.

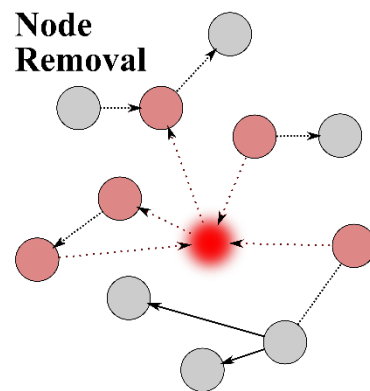
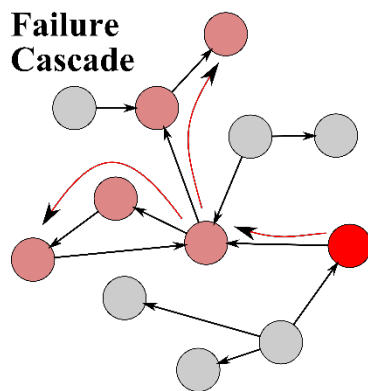
The following document is complementary and is intended for circulation within your organisation. It expands on the aforementioned points by providing a non-technical overview, along with examples on how these techniques have been applied in practice. It further strives to ignite a debate around the role of individual organisations in stabilising the environment in which they operate, an appeal to both regulators and individual business.

Part I: Systemic Risk Drivers



Extended Enterprise:

Each node represents entities of interest; links capture their interdependencies



Exposure to systemic risk is dependent on the capacity of a system to sustain a number of dynamical processes. These dynamical processes can be considered as the drivers of systemic risk and come in two distinct flavours: (a) individual node removal and (b) failure cascades – see upper right and lower right figures above respectively.

1. INTRODUCTION

Systemic risk is the potential of having interdependent failures, which emerge through the interconnectivity within a networked system, whether it is a collection of people, processes or technological artefacts. As such, their potential damage on an organisations performance is immense, potentially threatening its survival. Interestingly, singular systemic events are capable of initiating spectacular cascading failures, often with unpredictable and catastrophic impacts, such as electricity grid black outs or collapse of financial markets. Importantly however, interconnectivity is a necessary but insufficient condition for the emergence and propagation of systemic risks – as with all, the devil is in the detail. To further stress this point, it is not whether a system exhibits a degree of interconnectivity, but rather the specific form that this interconnectivity takes.

Interestingly, networks described by highly ordered, or conversely, random architectures are less likely to be affected by systemic events materialising. Nonetheless, such resilience comes at the cost of reduced efficiency due to the increased levels of redundancy that defines ordered or random networks (Wang and Chen, 2002). To counter this loss in efficiency, the majority of real world networks are (either by design or evolution (Barabási, 2012)) finely tuned between the two extremes, leading to the enhanced sensitivity of real-world networks to systemic threats. These are described by complex networks, where normal distributions go out the window and heavy tails becoming the norm (see Figure 1). Under these conditions, systemic risk is a real threat, where a single tree falling can induce extensive black outs; a lightning strike can change the status quo of an entire market; a factory fire can bring down an automotive manufacturer and the failure of a single financial institution can trigger a financial meltdown.

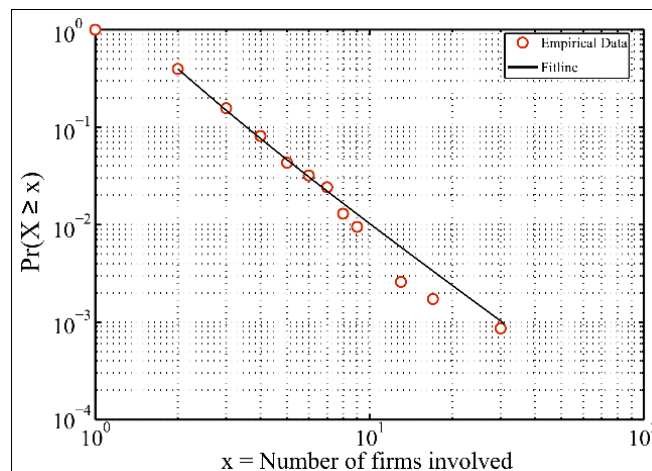


Figure 1: Number of involved firms on single loss events over a period of 6 years. A straight line on a log-log scale highlights the heavy-tail nature of the distribution. This is contrast to what one would normally expect, where on average, a firm’s exposure would be limited to very few loss events (Ellinas et al., 2015a).

Since the 2007-2008 financial crisis, the Bank of England has been actively involved in understanding the drivers of systemic risks, along with impact to both national and global economic networks in the hope of introducing effective governance and built-in robustness (Haldane and May, 2011). The real challenge to such work comes from understanding the system-level contribution of individual businesses in terms of

the probability of triggering systemic risk. The issue for regulators therefore is not whether an individual company is too big to fail, but whether it is too systemic to fail, where its failure can set-off a domino of subsequent failures.

Similarly, large organisations have the same governance dilemmas, where individual, yet tightly-coupled, business units can trigger systemic risk. Modern businesses are increasingly interconnected in sophisticated ways through supply networks, customers, and technologies, to name just a few. Whilst Enterprise Risk Management (ERM) initiatives have been successful in identifying where and when connections arise, an evaluation of the consequent systemic risk is a continuous challenge. A novel way of tackling this challenge requires an appreciation of the emergence of tipping points in the operation of organisations, along with evaluating the effectiveness/efficiency of a wide range of possible mitigation measures (e.g. appropriate culture, containment mechanisms, supplier management etc.).

Network science has emerged as the main tool of scientific enquiry around the effect of interconnectivity and interdependence within business-related networks (Barabási, 2009). As such, pragmatic guidance can now be provided on systemic risk management, adding new approaches for risk practitioners and consequently results in added value by:

- Highlighting key aspects of interconnectivity and how they dictate the exposure of an organisation to systemic risk.
- Providing background information on the construction of network models.
- Helping to understand the nature of systemic risk.
- Recommending ways for managing sources of systemic risk.

2. SYSTEMIC RISK

Through a network lens, systemic risk is a macroscopic property that emerges due to the non-linear interactions of components (or agents) at a microscopic level. These interactions are captured by the networks structure, and can be the result of empirical observations (e.g. insurance claims, contractual links etc.) or artificially generated the use of standard network models (see Part II).

To provide a quantitative measure, one can shock each individual node and evaluate its impact on the overall network, based on how many subsequent nodes are affected. The nature of the perturbation is important, with two distinct mechanisms dominating:

1. 'Removal' of individual nodes, manifesting in a number consecutive instances e.g. consider a malicious hacker targeting and shutting down parts of an IT system; a factory being taken out of a supply chain due to a fire etc.
2. 'Contamination' of individual nodes e.g. consider the impact of a single task failing during a project delivery; a contaminant entering a supply chain; a disease spreading through a population; a piece of information triggering a reputation threat etc.

In the first case ('removal'), an individual node is selected (either by a random or in an informed way) and is removed from the network. As a result of its removal, all links that were attached to it are also removed,

creating a 'structural hole' into the network. Its impact can be measured by using various indicators, including the number of remaining components and/or nodes left within the network. By applying this method, Albert et al. (2000) uncovered the Achilles' heel of the Internet and WWW – a property commonly referred to as 'Robust-Yet-Fragile' (RYF); a direct consequence of its highly-heterogeneous network structure (what is referred to as scale-free, see Part II). Several recent systemic failures can be attributed to this 'RYF' property – examples include the loss of 400 million USD in sales by Ericsson following a lightning strike to a single factory. Under this view, the effect of such 'Acts of God' can be minimised by focusing on the capacity of a network to sustain this node removal process rather than trying to predict the occurrence of this improbable cause.

For the second case ('contamination'), consider a domino effect metaphor where, each domino (a node) is affected in some way and, under a given set of conditions, is capable of affecting its neighbouring dominos, resulting in a possible failure cascade. The larger the failure cascade is, the greater the systemic role of the node. By applying this methodology, recent work highlight that

- a) Large failures may follow the exact same dynamics as small failures, making them unrecognisable before their full impact has been unravelled (Bak and Paczuski, 1995). This does emphasise the importance of path dependency and a shift of focus from trying to predict the impact towards understanding the exposure of a system.
- b) A single node failure is capable of inducing a significant amount of damage to the entire network, regardless of size. Theoretically, the magnitude of this damage can be infinite; practically, this indicates that the possible size of the largest failure is limited by the size of the network rather than the probability of occurrence.
- c) The probability of an increased amount of systemic risk being materialise is exceedingly high.

These elements become particularly important where traditional risk mitigation strategies provide limited support for mitigating these particular forms of risk. For example, consider the case of project risk management, where risk is typically mitigated via the deployment and analysis of project schedules (Zwikael and Sadeh, 2007) e.g. using variants of Critical Path Method (CPM) and Program Evaluation and Review Technique (PERT). These techniques assume linearity (Williams, 1999), yet recent work has shown that projects are capable of sustaining non-linear failures (Ellinas et al., 2016, Ellinas et al., 2015b), and in effect, suffering from the aforementioned points (a) – (c).

The following section will use a number of real world examples to elaborate on the distinct nature of these two mechanisms.

2.1. SYSTEMIC RISK: THE NODE REMOVAL CASE

Numerous networks achieve a great deal of their functionality merely by their capacity to connect e.g. supply chains, IT systems, communication systems etc. For these systems, nodes (components) of the network can become non-operational in a number of ways – factories can become non-functioning; servers can be taken offline for maintenance (or taken out maliciously) etc. As a consequence, the impact of such a local event may have a disproportional effect and accordingly, lead to systemic risk materialising. Such events may arise under a variety of scenarios: a node may be removed randomly (e.g. corresponding

to regular maintenance in which any given component is routinely taken out of the system) or through careful attack (e.g. a hacker targets the most connected part of an IT system hoping to induce maximum damage). To model the random removal scenario, a node is randomly chosen and subsequently removed from the network – this process is repeated until all nodes are removed. In the targeted removal, the node to be removed is based on the number of connections that it possesses (i.e. most connected node is removed first, followed by the second most connected etc.) – again this process is repeated until all nodes are removed. For both cases, a robust system would experience a drop in its performance in a linear way, suggesting a linear relationship between cause and effect. If this was not the case, then one would observe the disproportional deterioration of the system as the node removal process was taking place – in other words systemic risk would be unravelling. Let us explore this concept further through the means of a network that captures the tasks responsible for delivering a large project.

Figure 3 captures the impact of the node removal process by monitoring the connectivity of the network, in terms of connected components (Figure 3, left) and connected nodes (Figure 3, right). In this case, the random removal case is represented by the red plot; the targeted removal is captured by the black plot.

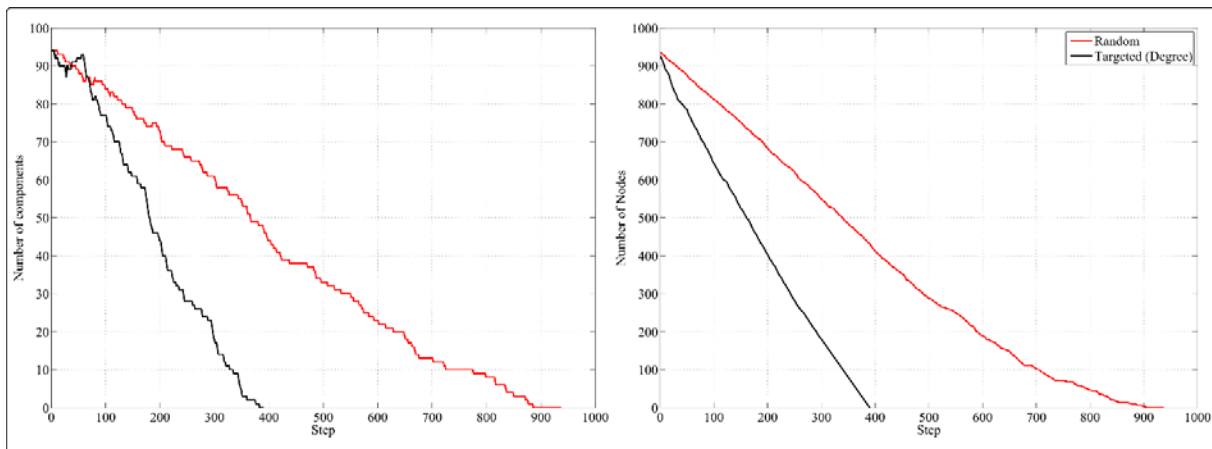


Figure 3: The impact of node removal is captured by monitoring the number of connected components (left subplot) and connected nodes (right) that remain in the network. Note how the system is increasingly robust under the random removal process (black plot), showing a slow, linear-like reduction in its performance. Conversely, the network is increasingly fragile under a targeted node removal, leading to the complete disintegration of the network less than half time steps.

Two key features are worth highlighting:

1. Under the random case, the response of the network is approximately linear, indicating a cause and effect (e.g. by taking out 100 nodes, the overall number of connected nodes and forms the network is approximately 100 less, Figure 3, right). In other words, the network is quite robust under random nodal removals, and no systemic risk materialises. In the case of a targeted node removal, the converse is true indicating that the system can also be rather fragile (and thus the term “Robust-Yet-Fragile”). Specifically, by removing 100 nodes, the overall number of nodes decreases to approximately 600 – that is roughly 300% increase in damage. Taken to the extreme case, the network is completely dismantled with less than half effort required by the random case.

In other words, if nodes are removed in a given order (as identified using various levels of information such as its node degree) an entire system can experience rapid deterioration – a clear indication of systemic risk at play.

2. The ability to construct early-warning signs of systemic risk is highly dependent on the measure used to capture the state of the network. For example, consider the case where a decision maker can only monitor the number of components of a system (Figure 3, left). Interestingly, red and black plot lines only show significant divergence after about 8% of the system has already been removed. Hence, if simulations steps were translated to time, monitoring the number of components will provide little information on whether a random (non-systemic) or targeted (systemic) case of node removal was taking place. As a result, one would be left to the dark with respect to the extent upon which mitigation action should be deployed. In contrast, by better understanding the architecture of the network, one would be able to appreciate the exposure of the network to such systemic events, and consequently deduct whether the event is minor or major.

In summary, node removal can provide a possible mechanism in which systemic risk can emerge. Importantly, the exact same system can exhibit significantly different causal relationships. As a result, knowing whether failure can lead to systemic risk becomes exceedingly hard when the focus is on observe-and-react approach. In response, the concept of limiting the exposure to systemic failures (e.g. by focusing on the uniformity of a network's features) can provide a more useful perspective in shielding a network from such systemic failures. Importantly, it is not a question of when the next systemic failure may take place, but whether the structure of a network allows for a systemic failure to arise – a shift in focus from risk prediction to risk exposure.

2.2. SYSTEMIC RISK: THE FAILURE CASCADE CASE

Consider a set of dominos, arranged in an intrigued way. Clearly, the number of dominos that can be affected by the fall of any one domino is a function of its location, with the first domino having the capacity to affect all dominos and the last being able to effect none. This “domino effect” underpins a number of phenomena in various complex socio-technical systems, ranging from the power-grid failures to financial system meltdowns and the emergence of social norms (Vespignani, 2012). This cascading behaviour has been attributed to a threshold mechanism (Granovetter, 1978) and can provide a modelling framework for capturing this complex behaviour in a tractable way (Ellinas et al., 2015a).

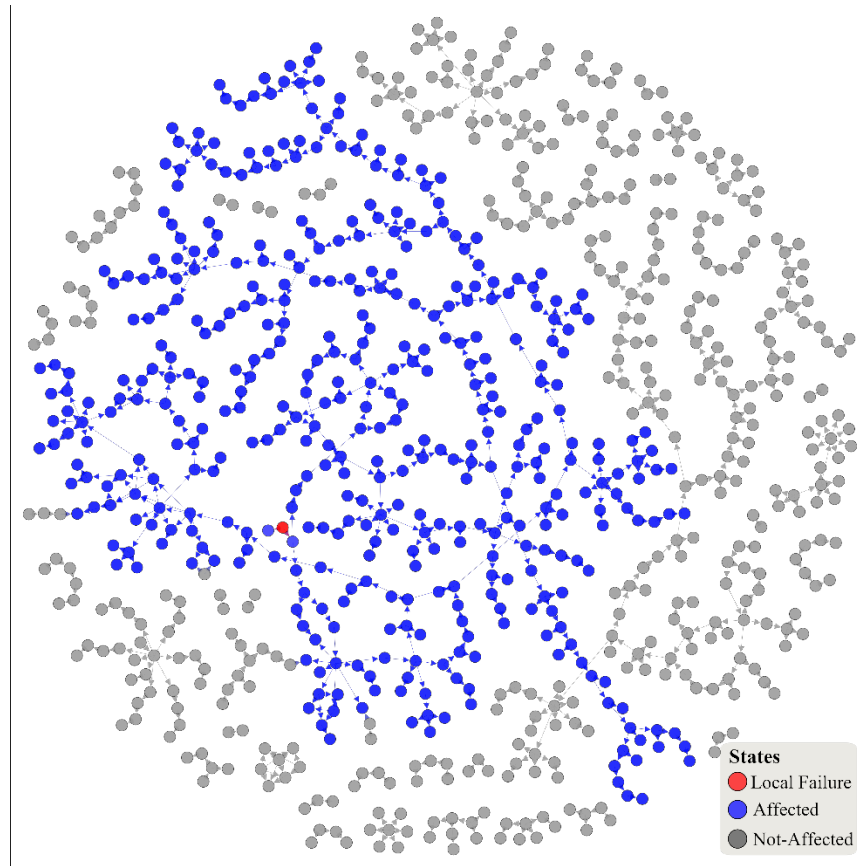


Figure 4: Example on how the failure of a single node (red) can propagate throughout the network, subsequently impacting other nodes (blue) which on first sight, are not directly linked. Figure corresponds to the largest cascade possible in a real-world engineering project (Ellinas et al., 2015b).

Under this view, systemic risk is captured by the number of nodes that can be affected by the failure of a single node, as a ratio of the total number of nodes. By doing so, single nodes responsible for triggering large cascades, or nodes that are more likely to be affected by one, can be identified. Figure 4 precisely captures this effect across a real-world engineering project, abstracted in the form of an activity network. In this case, every node corresponds to an activity, with every link indicating a dependency (e.g. node i being connected to node j means that the start of activity i depends on finishing activity j). By shifting the focus of analysis from the local to the global scale, the capacity of a network to sustain systemic risk, along with its probability of occurring, can also be obtained – see Figure 5. As such, the effect of local/global mitigation (e.g. increase in resource allocation/change in network structure) can be quantitatively assessed, providing insight on how such failure cascades can be contained in an effective and efficient way.

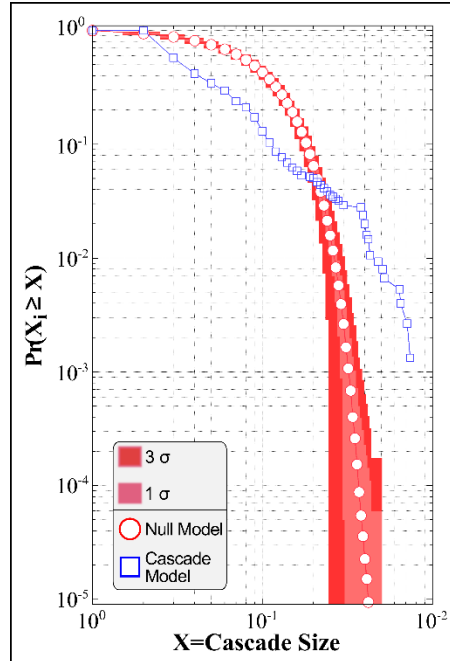


Figure 5: Cumulative probability distribution of failure cascade sizes. The probability (y-axis) of having a failure cascade equal or greater than a given size (x-axis) is shown in blue; the same probability is shown under the case of the cascade sizes being Normally distributed (red). Note that the cascade size is given as a proportion of the overall number of nodes.

Consider Figure 5, which captures the probability of having a cascade of a given size, under the case of a cascade model (blue plot) modelled against of an activity network. A null model which represents failure cascade sizes drawn from a Normal distribution are also shown for comparison (red plot). Through this example, the manifestation of systemic risk can be captured by considering the total number of nodes that can be affected by a single node failing. For example, in Figure 5, left, the failure of single nodes can impact roughly a tenth of the entire system (90 nodes). Clearly, nodes that are capable of triggering such failures are of systemic importance and need to be identified and accounted for. It is worth noting that once such model has been set up, the effectiveness and efficiency of various mitigation actions (applied at either the local or the global level, in a uniform or targeted fashion) can be assessed, supporting a variety of decision making functions.

Three key aspects of the results captured by failure cascades must be highlighted:

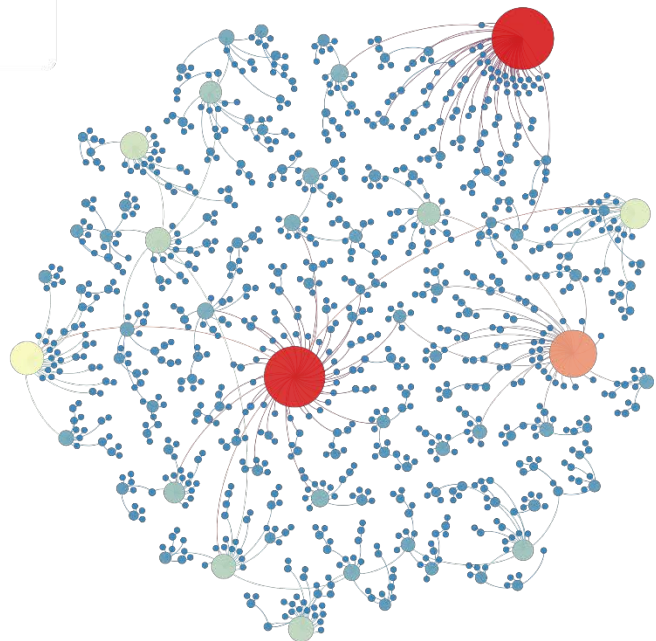
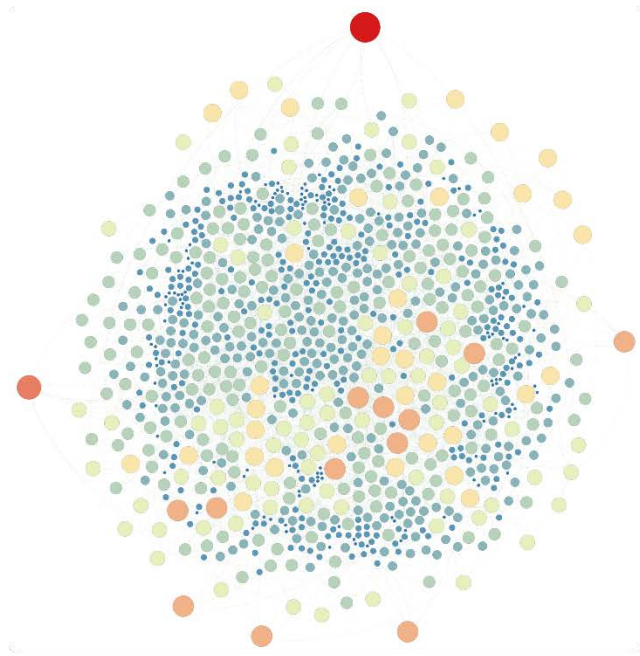
- The process that drives these cascades is essentially self-contained, indicating that no great exogenous force is needed to induce large systemic failure (Bak and Chen, 1991). As a result, both small and large failures can be shown to follow the exact same dynamics, challenging the ability to predict large failures (Ellinas et al., 2015a). Nonetheless, by focusing on the structure of the underlying system, the exposure of the network can be numerically explored and assess whether a network is prone to such failures.
- Regardless of the level of mitigation action, the cumulative probability distribution is heavy-tailed, and resembles a power-law. As a result, the extreme variance that describes this class of distributions, the average size of a failure cascade provides no useful information on the scale of damage

that a network can sustain. In fact, uniform regulatory frameworks and/or risk mitigation strategies should not be applied uniformly to all nodes – some nodes have a limited impact, yet other have an extraordinary capacity to impact the entire network and hence, should be individually treated. Identifying them and their role within the network provides the first key step in shielding against systemic risk.

- The probability of encountering small failure cascades is significantly overestimated by traditional models (what corresponds to the red plot in Figure 5). As a result, decisions made on models that do not explicitly capture the intricacies of the underlying networks (e.g. CPM, PERT) can lead to overspending. Perhaps more importantly, the probability of having large scale failures (i.e. systemic risk) is significantly underestimated by traditional models – examples of divergence of over two orders of magnitude are typical (i.e. a systemic failure being underestimated by a factor of 100). Furthermore, the size of the largest possible failure cascade is significantly underestimated by models that fail to capture the underlying interconnectivity

In summary, failure cascades can be used as a second mechanism by which systemic risk can be modelled. Evidently, the probability and impact of these types of systemic failures is surprisingly high; models that ignore the effect of interconnectivity substantially underestimate both. In search for mitigation against such events, the efficiency of local mitigation (applied uniformly or in a targeted manner) can be assessed; increasingly involved measures for mitigation can be deducted by considering specific features of the network and subsequently engineering them in order to efficiently reduce exposure to systemic risk.

Part II: Networks



Everything is connected with everything else – this is an often cited explanation on how large-scale, systemic events manifest themselves. Yet, this is not entirely true, as the architecture of the resulting networks is at the very core of these events. As such, a brief overview of the three main network models ('random', 'small-world' and 'scale-free') is presented.

1. A BRIEF OVERVIEW OF NETWORKS

In 1963, Stanley Milgram set out an experiment to assess an increasingly popular conjecture – that the social world was a rapidly shrinking medium due to the increased social connectedness, as enabled by technological advances. To his great surprise, Milgram reached a stunning conclusion – everyone was connected to everyone else on an average of 6 intermediate steps (Milgram, 1967). This result is now widely known as the “Six Degrees of separation”, after the name of a theatrical play that popularised the idea. This surprisingly low number has captivated the imagination of mathematicians, sociologists, physicists and computer scientist up to this date. In fact, a similar experiment that took place in Facebook (a widely used social network platform), involving 5.8 million users concluded that the average degree of separation was 5.73 – a value surprisingly close to the one estimated by Milgram over 40 years ago.

Widely popularised by social platforms, networks play a crucial role in our lives – examples range from the human brain and global markets to supply chains and projects. The study of these networks (or *graphs*) is historically attributed to the mathematician Leonhard Euler and his attempt in devising a method for crossing all seven bridges of Königsberg, published in 1735. The adopted approach in tackling this problem was what is now commonly referred to as **graph theory**. Despite its numerous branches, random graph theory as introduced in the seminal work of Erdos and Rényi (1960) is closer to its modern successor. In their work, the concept of the **random network** (or *ER networks*, after the initials of the two authors) was introduced.

A random network is composed of a set of nodes (or vertices), connected by a set of links (or edges). Each node receives a link with a probability p , independent from any other node. As a result, the degree distribution of the network follows a Gaussian distribution, with the average degree serving as an adequate description of the state of the network. An important aspect of a random network is its relatively low clustering, resulting in large average path length (or what was previously referred to as degrees of separation). Clearly, this is in contrast to what has been noted by Milgram, so why do we bother with random graphs to begin with?

Some Formal Definitions

The system of interest (e.g. supply chain, infrastructure, project, firms etc.) can be mapped as a network (or graph). It is defined as $G = \{N\{E\}$ where every component i can be abstracted as node i , where $i \in N$. Similarly, a dependency between component i and j are captured by an (un)directed link $e_{i,j}$, $e_{i,j} \in E$

The structure of the entire network can be captured by the so-called adjacency matrix A , where $A(i, j) = 1$ if there is a link between node i and j ; and 0 otherwise. Finally, note that the entry $A^k(i, j)$ will provide the total number of paths between i and j of length k .

For expansive technical reviews, see (Albert and Barabási, 2002, Boccaletti et al., 2006, Dorogovtsev and Mendes, 2002, Newman, 2009, Newman et al., 2006, Newman, 2003)

The value of random networks lies at this very distance from real-world, purposeful networks – an explicit illustration of pure randomness. If one is able to show that an aspect of a network significantly deviates from its random equivalent, then it is worth further exploration. In fact, the findings of Milgram are one such example - the noted value for the degrees of separation (or average path length) is surprisingly low, compared to what is predicted by its random counterpart. Importantly, similar behaviour has been noted in several real-world systems (Watts and Strogatz, 1998), highlighting the practical relevance of the concept – these networks were subsequently called ‘**small-world**’ networks (in tribute to Milgram’s original observation).

This ‘small-world’ effect emerges as the structure of the network transitions from complete order (what is referred to as a lattice structure) to disorder (represented by a random graph). Found within the middle, ‘small-world’ networks have high clustering, yet low average path lengths – a result that leads to the apparent connectedness of a social network. Importantly, the degree distribution of these networks is equivalent to that of random networks – a Normal distribution.

This assumption of Normal degree distributions was subsequently challenged by the seminal work of Barabási and Albert (1999), who used empirical observations to illustrate that the Internet significantly deviates from this assumption, by being heavily heterogeneous. Specifically, nodes with a degree several orders of magnitude greater than the average degree were observed – clearly contradicting its random (and ‘small-world’) counterpart networks. In fact, these observations were noted throughout numerous systems, with heavy tail¹ distributions being the rule rather than the exception.

Network Measures

Network measures can be used to meaningfully capture, and subsequently quantify, important network measures.

At this section we will define three widely-used measures: the degree distribution, clustering and average path length.

Degree distribution

A degree distribution is essentially a histogram capturing the number of nodes that have a given number of connections. In the case of directed networks, a degree distribution for both incoming and outgoing connections is needed to capture the underlying connectivity of the network.

In the case of random and ‘small world’ networks the degree distribution decays exponentially while ‘scale-free’ networks are defined by slow-decaying tails (often, power-law).

Note that in the case of scale-free networks, degree distributions are better illustrated through cumulative probability plots on a double log scale, in effect reducing the scattering effect at the tail.

¹ Despite the wide reference to power-laws, the reader is warned that the majority of these claims have been shown to be poor estimates - see CLAUSET, A., SHALIZI, C. R. & NEWMAN, M. E. 2009. Power-law distributions in empirical data. *SIAM review*, 51, 661-703.

These networks have been coined as ‘**scale-free**’ **networks**² due to the absence of a meaningful average³ that can describe the scale of the system.

In summary,

- **Random networks** are the simplest form of networks, described by high clustering, high average path length and Normal degree distributions. They are mainly used as benchmarks to identify aspects of a network worth further exploration, simply because they cannot have emerged by pure randomness.
- **‘Small-world’ networks** were introduced to explain the noted high clustering yet low average path length showcased by real-world networks; their degree distribution is also Normal.
- **‘Scale-free’ networks** were introduced to explain the significant deviations in the degree distribution of real-world networks; they are used to model the heavy-tail degree distributions found in numerous natural and man-made systems.

Table 1 summarises the core difference between the three network models, in terms of the degree distribution, clustering and average path length (see side panel for definitions). Figure 2 further illustrates the core differences graphically.

Table 1: Differences between the three network models

	Random	‘Small-World’	‘Scale-Free’
<i>Degree Distribution</i>	Normal	Normal	Heavy-Tail (often, power law)
<i>Clustering</i>	High	High	Low – High ⁴
<i>Average Path Length</i>	High	Low	Lowest

² It is worth noting that despite their pervasive nature, the mechanism of their occurrence is still debatable, whether being a result of pure luck (e.g. the “rich get richer” effect) or reason (e.g. topology reflects an optimisation attempt) BARABÁSI, A.-L. 2012. Network science: Luck or reason. *Nature*, 489, 507-508..

³ This is due to the extremely broad (in principle, infinite) variance that such distributions exhibit.

⁴ Sensitive to other topological features as well.

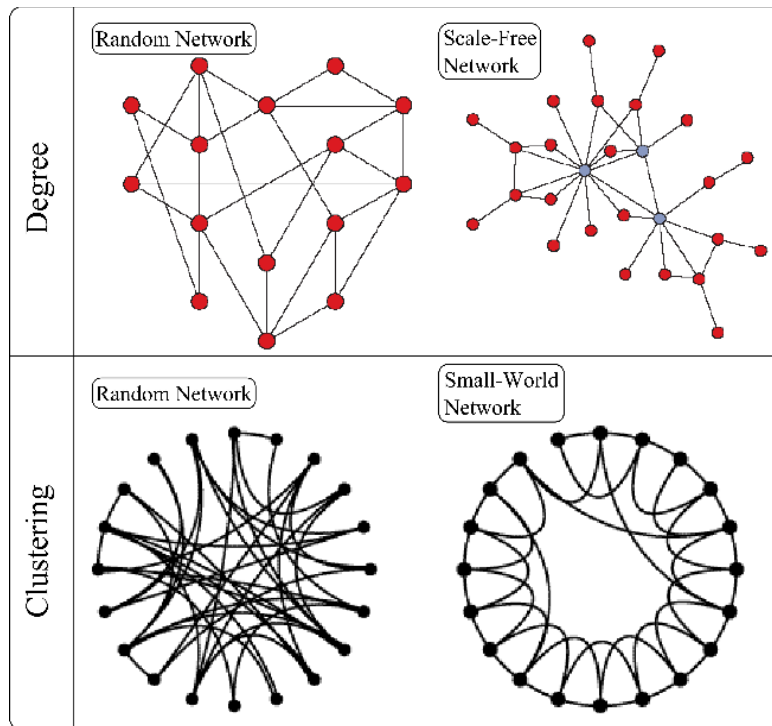


Figure 2: Upper row illustrates the core difference between a random network and a scale-free network, in terms of the degree distribution. In the random case, each node has more or less the same number of connections, while in the 'scale-free' case, some nodes (coloured blue) have a surprisingly high number of connections. Lower row illustrates the core difference between the random network and the 'small-world' network. In the random case, clustering and average path length is small. On the other hand, clustering in the small-world network is high, yet average path length is low due to the emergence of shortcut connections, capable of connecting distant parts of the network. Figures adapted from (Barabasi and Oltvai, 2004, Watts and Strogatz, 1998)

Network Measures

Clustering

Clustering refers to how many closed triangles exist within a network. In the context of social networks, you can think of it as the likelihood of two of your friends to know each other. As such, high clustering suggests a rather dense network, where navigation is rather inefficient due to the large number of links that need to be traversed to range any node – this is typical for a random network.

Average path length

Average path length measures a global property of the network. Specifically, it refers to the number of links that need to be traversed across every pair of connected nodes, averaged across all paths.

Normally, one would expect that increased clustering would also result to high average path length (as more steps are required to traverse through the dense network). However, Watts and Strogatz (1998) have illustrated that this need not be the case, as shortcuts can be introduced to dramatically reduce the average path length whilst maintaining relatively high clustering – this is the celebrated 'small-world' phenomenon.

Part III: Bringing it all together

Network models are becoming increasingly prominent in understanding the nature of systemic risk. While conventional tools focus on identifying correlations (e.g. regression models), complex networks provide a framework for understanding the casual mechanism that drives systemic risk. Specifically, two distinct processes – node removal (Section 2.1) and cascading failures (Section 2.2) – have been proposed that can lead to systemic failures from modest, local failures. In practice this means that medium-low impact risks should be re-evaluated for their ability to cause very large impacts and that these cascading type failures have a dramatically reduced likelihood.

Both processes emphasize the inability to confidently assess the probability of such failures from occurring, at least in an *a priori* basis. In other words, global failure of a system or organisation (i.e. systemic risk) follows the exact same dynamics as local ones; making them effectively indistinguishable. As a result, it challenges two key steps in any risk management process – that of identification and prediction. So what do we do when what matters can neither be identified nor predicted?

By focusing on the exposure of the whole system, one can begin to tackle this challenge. Specifically, highly heterogeneous networks (e.g. ‘scale-free’) have been shown to be susceptible to systemic risk to a much greater extent than homogeneous networks (e.g. random; ‘small-world’). Such distinct topological trademarks can provide proxies in which the exposure of a system to systemic risk can be examined. More explicitly, formal models that acknowledge these complex network topologies can be used to numerically explore the exposure to systemic risk, along with examining the efficiency and effectiveness of a large variety of mitigation techniques that can contain the impact of such events. Extensive databases and powerful computational methods to utilise these models are now becoming widely available, enhancing decision making processes across a number of levels within organisations.

REFERENCES

- ALBERT, R. & BARABÁSI, A.-L. 2002. Statistical mechanics of complex networks. *Reviews of modern physics*, 74, 47.
- ALBERT, R., JEONG, H. & BARABÁSI, A.-L. 2000. Error and attack tolerance of complex networks. *Nature*, 406, 378-382.
- BAK, P. & CHEN, K. 1991. Self-organized criticality. *Scientific American*, 264.
- BAK, P. & PACZUSKI, M. 1995. Complexity, contingency, and criticality. *Proceedings of the National Academy of Sciences*, 92, 6689-6696.
- BARABÁSI, A.-L. 2009. Scale-free networks: a decade and beyond. *Science*, 325, 412-413.
- BARABÁSI, A.-L. 2012. Network science: Luck or reason. *Nature*, 489, 507-508.
- BARABÁSI, A.-L. & ALBERT, R. 1999. Emergence of scaling in random networks. *Science*, 286, 509-512.
- BARABASI, A.-L. & OLTVAI, Z. N. 2004. Network biology: understanding the cell's functional organization. *Nature reviews genetics*, 5, 101-113.
- BOCCALETTI, S., LATORA, V., MORENO, Y., CHAVEZ, M. & HWANG, D.-U. 2006. Complex networks: Structure and dynamics. *Physics reports*, 424, 175-308.
- CLAUSET, A., SHALIZI, C. R. & NEWMAN, M. E. 2009. Power-law distributions in empirical data. *SIAM review*, 51, 661-703.
- DOROGOVTSSEV, S. N. & MENDES, J. F. 2002. Evolution of networks. *Advances in physics*, 51, 1079-1187.
- ELLINAS, C., ALLAN, N. & CANTLE, N. How resilient is your organisation? From local failures to systemic risk. ERM Symposium 2015, 2015a.
- ELLINAS, C., ALLAN, N., DURUGBO, C. & JOHANSSON, A. 2015b. How Robust Is Your Project? From Local Failures to Global Catastrophes: A Complex Networks Approach to Project Systemic Risk. *PLoS ONE*, 10, e0142469.
- ELLINAS, C., ALLAN, N. & JOHANSSON, A. 2016. Project systemic risk: Application examples of a network model. *International Journal of Production Economics*, 182, 50-62.
- ERDOS, P. & RÉNYI, A. 1960. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci*, 5, 17-61.
- GRANOVETTER, M. 1978. Threshold models of collective behavior. *American journal of sociology*, 1420-1443.
- HALDANE, A. G. & MAY, R. M. 2011. Systemic risk in banking ecosystems. *Nature*, 469, 351-355.
- MILGRAM, S. 1967. The small world problem. *Psychology today*, 2, 60-67.
- NEWMAN, M. E. 2003. The structure and function of complex networks. *SIAM review*, 45, 167-256.
- NEWMAN, M. E. 2009. *Networks: an introduction*, Oxford University Press.
- NEWMAN, M. E. 2011. Complex systems: A survey. *arXiv preprint arXiv:1112.1440*.
- NEWMAN, M. E., BARABÁSI, A.-L. & WATTS, D. J. 2006. *The structure and dynamics of networks*, Princeton University Press.
- VESPIGNANI, A. 2012. Modelling dynamical processes in complex socio-technical systems. *Nature Physics*, 8, 32-39.

- WANG, X. F. & CHEN, G. 2002. Synchronization in small-world dynamical networks. *International Journal of Bifurcation and Chaos*, 12, 187-192.
- WATTS, D. J. & STROGATZ, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature*, 393, 440-442.
- WILLIAMS, T. M. 1999. The need for new paradigms for complex projects. *International Journal of Project Management*, 17, 269-273.
- ZWIKAEL, O. & SADEH, A. 2007. Planning effort as an effective risk management tool. *Journal of Operations Management*, 25, 755-767.



CONTRIBUTING AUTHORS

Christos Ellinas, Systemic Consult
Neil Allan, RiskIQ
Neil Cante, Milliman

POINT OF CONTACT

Neil Allan (neil@systemicconsult.com)