



NATIONAL  
**PREPAREDNESS**  
COMMISSION



# Building Better Resilience

Paul Martin and Jordan Giddings

December 2020

**The Covid-19 pandemic has exposed major shortfalls in the preparedness of businesses, organisations and nations. Many have turned out to be less well prepared than they thought they were and less well prepared than they should have been. In this paper we highlight three areas that we believe are crucial to doing it better next time:**

**Active resilience**

**Human psychology**

**Complex systems**



---

## CONTENTS

---

ACTIVE RESILIENCE	<b>4</b>
Building passive resilience	5
Building active resilience	6
Personal resilience	8

---

HUMAN PSYCHOLOGY	<b>10</b>
Misreading risk	11
Mishandling crises	13
Aiming off	14

---

COMPLEX SYSTEMS	<b>16</b>
-----------------	-----------

---

CONCLUSION	<b>21</b>
------------	-----------

---

---

## ACTIVE RESILIENCE

Resilient people, businesses and institutions cope well when things go wrong. They roll with the blows, deal effectively with the adverse consequences and return quickly to a stable equilibrium. That, at least, is the conventional view. However, resilience should mean something more substantial than recovering from disruption, desirable though that is. Resilience comes in two distinct forms, which we refer to as passive and active.

**Passive resilience** is the ability to absorb disturbance, recover quickly from a setback and return to normality. This is the colloquial sense of the word, which equates roughly with being robust. Expressed in terms of risk, passive resilience is about reducing the impact of a disruptive incident by reducing the size or duration of its harmful consequences. Passive resilience is clearly a good thing and we should all aspire to have more of it. However, there is – or should be – more to resilience than absorbing blows.

**Active resilience** means growing progressively tougher by learning from adversity and becoming better able to avoid and manage future stresses. Actively resilient people or organisations do more than just return to their prior state after an adverse event: they continually learn from their experience and develop stronger defences, making them better able to resist the next time. They are less likely to experience a crisis and cope better if they do.

The concept of active resilience is similar to what the writer Nassim Nicholas Taleb calls antifragility. Taleb argues that complex systems, including national economies and living organisms, become tougher as a result of coping successfully with moderate stress. Passively resilient people or organisations absorb shocks and return to their previous state; actively resilient ones become tougher. A biological analogy is physical exercise. Strenuous exercise causes mild damage to muscle tissues, which respond by repairing themselves. But the muscles do not merely recover to their prior state: rather, they over-compensate and grow stronger, and hence strenuous exercise makes us physically stronger. Writers and philosophers have been making essentially the same point for millennia. ‘Difficulties strengthen the mind, as labour does the body’, said Seneca two thousand years ago. Nietzsche put it more starkly: ‘That which does not kill us, makes us stronger’.



## Building passive resilience

Passive resilience amounts to reducing the impact element of risk by reducing the severity or duration of harm. There are many ways of achieving this, depending on the nature of the risk.

The impact element of disruptive risks generally involves some form of loss, whether that is lost lives, lost money, lost data, lost business, lost reputation, or some combination of those. The impact will be worse if the victim responds too slowly or makes bad decisions during and after the event. By following the temporal chain of events leading to loss, we can see that impact could be reduced by:

- reducing the amount of time that natural hazards or malicious threat actors have to cause harm (early detection and rapid response);
- building-in alternative ways of working and minimising single points of failure (redundancy)
- keeping copies of assets, so that losing the asset causes less harm (backup);
- sharing the loss with others (insurance);
- improving the ability to cope effectively with the disruption (incident management and crisis management); and
- accelerating the process of returning to normality (business continuity planning and disaster recovery)

The first line of defence in reducing impact is early detection and rapid response. If an unfolding risk is spotted and mitigated quickly, less harm is likely to ensue. Early detection is especially relevant in the cyber domain, where speed is vital. The same principle applies to unfolding natural disasters such as disease pandemics.

A general design principle for reducing impact is to build-in redundancy in order to reduce the severity of failure. The presence of multiple alternative ways of working, or simply a stock of spare parts, can reduce the risk that a system will fail badly following a disruptive event. Redundancy improves reliability under pressure, whereas the presence of multiple single points of failure does the opposite.

Another basic strategy for reducing impact is insurance, whereby financial loss is reduced by compensation. Insurance is the commonest means of reducing the impact of ordinary crime. It works well when the impact of a damaging incident is mainly financial. But other dimensions of impact, such as reputational damage, loss of confidence, interruption of business and psychological harm, are less easily washed away with cash.

Finally, we have the familiar capabilities of incident management, crisis management, business continuity planning and disaster recovery. These should be tested regularly to

ensure that they do in fact work as intended. As the boxer Mike Tyson put it: ‘Everyone has a plan, until they get punched in the mouth.’ The best way of finding and fixing the flaws is through regular live exercising under realistic conditions. Another, less intrusive, way of honing crisis management is through computer simulation. Practising responses in an immersive virtual environment allows a larger range of scenarios to be explored than would be practicable with live exercises.

## Building active resilience

An organisation that implements the measures outlined above should be more resilient than one that relies on luck. If an incident does occur, the damage should be less severe and normal functioning should be restored more rapidly. But there would still be some damage and there would be more incidents to come in the future. So, what more can be done to reduce the risk? The answer lies in building active resilience.

An actively resilient person, organisation or nation is one that continually learns from experience and applies the lessons to make itself progressively tougher. Each time it deals with an incident it adapts by improving its ability to prevent or manage the next one. Like a muscle growing stronger through exercise, it grows stronger through adversity.

Building active resilience is a cyclical process in which the organisation detects a disruptive incident, responds to it, recovers its critical functionality, learns from the experience and applies the lessons. The virtuous cycle incrementally strengthens its ability to resist the next disruptive incident and reduces the impact if the disruption occurs. Many organisations do the detection, response, and recovery parts but neglect the learning part, leaving them no stronger than they were before.

Fortunately, there is no need to suffer an actual crisis in order to learn from experience. Active resilience can be developed through simulated experience acquired through testing, exercising and red teaming. Another relatively painless way of building active resilience is by learning from others through mechanisms such as information-sharing forums. As a wise person once observed, smart people learn from their own mistakes, but wise people learn from the mistakes of others.

The value of learning from others means that building active resilience works better as a cooperative process. Resilience is not a zero-sum game in which my gain is your loss and vice versa. Everyone stands to benefit from cooperating. Each organisation benefits directly from becoming more resilient and gains even more if others do the same. In a highly interdependent world, organisations are safer if their suppliers, partners, customers and even their competitors are resilient. Furthermore, collective resilience can create a sort of herd immunity that makes it harder for malicious threat actors, such as criminals or hackers, to shop around for the most vulnerable.



Resilience depends as much on people and relationships as it does on plans or infrastructure. An organisation is better able to weather a storm if it has trusted leaders, a motivated workforce, loyal partners and crisis managers who understand how people actually behave in stressful situations. A well-run organisation will take care to nurture its people during a crisis, knowing that their ingenuity and persistence will be crucial. The effects of a major crisis and its aftermath may extend over weeks, months or even years. The wise thing to do at the outset may be to resist the urge for heroic endeavour and send some people home to rest. The critical phase of a crisis may come after days or weeks of intense activity, by which time the sleepless heroes are too exhausted to think straight.

Personal and professional relationships are another crucial component of resilience. We all need help from others, especially when the going gets tough. During a crisis we are critically dependent on our existing relationships and networks. It is too late by then to start developing new ones. Therefore, one good way of building active resilience is to invest in relationships. Regrettably, the short-term imperatives of cost-saving and efficiency often get in the way. Frequent reorganisations, rapid turnover of personnel and so-called transformation programmes can have the unintended effect of breaking up informal networks and reducing organisational resilience. Economic pressures can result in whole cities or nations becoming less resilient. For example, the high cost of living in some large cities makes it hard for public service workers on low wages to afford to live there. A 2016 study found that more than half of police officers, firefighters and ambulance paramedics serving London were living outside the capital, where rents and property prices are lower. In a crisis, particularly one affecting public transport, it is harder and takes longer for them to get to work helping others.

A further benefit of resilience – the attribute that just keeps on giving – is that it can act as a deterrent and hence reduce the threat. It does this by increasing the cost to malicious threat actors, such as criminals and terrorists, and reducing the benefit to them. Actively resilient targets are less vulnerable to attack and a successful attack has less impact. The discerning threat actor will prefer to go after a less resilient victim.

When assessing resilience, it is crucial to take account of the resilience of suppliers and other third parties. It is all very well having your own sophisticated fall-back facilities, but if your business depends on a fragile small supplier for critical services then their lack of resilience will become your lack of resilience. Supply chains should be resilient too.

Like buses, crises do not arrive one at a time at neatly spaced intervals. Even if they were entirely independent of one another, they would still exhibit some clustering for purely statistical reasons. In reality, crises tend to be causally inter-related to some extent, making clustering even more likely. A basic precautionary principle is that you should be prepared to cope with two or more crises at the same time (otherwise known as a cluster\*\*\*\*).

## Personal resilience

The concept of resilience is relevant to individuals as well as organisations or nations. Moreover, personal resilience is an ingredient of organisational resilience because the resilience of key individuals will have a critical bearing on the successful handling of any crisis. An organisation that aspires to be more resilient should think about the resilience of its people, not just its infrastructure.

What does psychological research tell us about personal resilience? Among the most common psychological attributes of resilient individuals are self-confidence, realistic optimism (as distinct from delusional optimism), a sense of humour, the ability to stay focused under pressure, persistence and finding meaning even in negative experiences. Other factors that contribute to personal resilience are expertise, supportive relationships and adequate sleep. You are more likely to cope with stressful situations if you know what you are doing, have friends and colleagues to support you and are not debilitated by sleep deprivation. There are multiple pathways to resilience and individuals can be resilient in different ways. The same is true for organisations.

Overcoming adversity often has positive consequences. Individuals who repeatedly cope with moderately stressful events tend to become psychologically more robust and better able to cope with future stress. To put it another way, the experience of coping with challenging situations helps to build active resilience. Highly resilient people typically have a track record of dealing successfully with stressful situations.

The ‘good’ forms of stress that help to strengthen personal resilience are generally acute (i.e., relatively short in duration), controllable (i.e., you can do something to alleviate the problem) and moderate in intensity. In contrast, stress that is chronic (enduring), uncontrollable and severe is more likely to cause psychological and physical harm and weaken the individual’s capacity to cope with further stress. ‘Good stress’ makes us more resilient whereas ‘bad stress’ does the reverse.

When selecting and developing its people, an organisation should value personal resilience. Being clever, ambitious and hardworking is not enough. Personal resilience has more bearing on ultimate success or failure than most other characteristics. An organisation that wants to become more resilient should try to create an environment in which its people can become more resilient. Among other things, this could involve regular testing and exercising, to give people the experience of coping successfully with demanding situations, and training to equip them with relevant expertise. A resilient organisation will also respect people’s need for adequate sleep, foster supportive interpersonal relationships and value humour.

Is it possible for individuals or organisations to be too resilient? Well, yes. Excessive resilience of the wrong kind can be an obstacle to necessary change and consequently





destructive. An unhealthy passive resilience built on delusional optimism, over-confidence and unbending persistence can keep individuals plugging away in hopeless situations where a better course of action would be to stop and change track. Seemingly resilient superheroes sometimes self-destruct. The findings of psychology are in pleasing harmony with complex systems science, which tells us that systems that are too strongly anchored in a stable equilibrium can be resistant to evolutionary change and consequently more vulnerable to big shocks. Unhealthy passive resilience can prevent individuals and organisations from recognising that they need to change. Those that resist change for too long can end up breaking.

---

## HUMAN PSYCHOLOGY

We humans are not rational calculating machines in possession of all the facts – a reality that economists eventually came to recognise. Rather, we are emotional and social animals. Our feelings, relationships, personalities and experiences have pervasive influences on our judgment and decision-making, even more so under conditions of stress and uncertainty. Judgments about risk and crises are prone to systematic distortion by a range of psychological predispositions and cognitive biases. It is better to understand these features of human psychology and aim off for them.

Each of us is equipped with a suite of psychological predispositions that systematically influence how we think and behave in different situations. These predispositions may be regarded as unconscious heuristics or rules of thumb. Evolution has equipped us with them because they help in making potentially life-saving decisions when there is insufficient information and not enough time to think. They work remarkably well in a wide range of situations. In some circumstances, however, they can lead us to make systematic errors, both in how we perceive risks and how we respond to those risks.

Responding intuitively is obviously beneficial when there is no time to review all the options and devise an optimal plan. It is doubly beneficial because acute stress impairs our ability to remember procedures. During the long history of our species, our ability to react unthinkingly has made the difference between surviving and dying for countless of our ancestors. However, intuitive responses can be problematic in novel situations where a different and counter-intuitive response may be required.

The best way to avoid doing the wrong thing in a dangerous situation is to keep practising the correct response until it becomes semi-automatic. Military operators refer to this as developing muscle memory. Another way to stop people doing the wrong thing in an emergency is to design the technology so that their untrained intuitive response is also the correct response – in other words, fitting technology around people rather than the reverse. This principle lies behind the design of the panic bars (or crash bars) that are fitted as standard to fire exit doors. People do not have to be trained to use them because their intuitive behaviour in an emergency is also the right way to open the door.



## Misreading risk

Psychological research has uncovered an assortment of predispositions that can lead us astray when confronted with the sorts of novel, complex and protracted challenges presented by contemporary risks. These predispositions, or cognitive biases, systematically affect the way we think.

For a start, we are inclined to overestimate the likelihood of exotic, attention-grabbing risks and underestimate the likelihood of mundane risks, even when the mundane risks are objectively more likely to materialise. A well-known example is the belief that sharks are more dangerous than stairs, even though the opposite is true by a wide margin. This distortion in our perception of risk is fuelled by a phenomenon called **availability bias**, whereby we find it easier to believe that something might happen if it comes easily to mind. The more easily we can picture an event, the higher our intuitive estimate of its likelihood. Events come more easily to mind if we have seen memorable accounts of them in social media, for example. Thus, we are inclined to over-estimate the likelihood of bad events that we can easily imagine and under-estimate risks that we find hard to imagine.

Availability bias lies behind the recurring cycle of crisis, anxious reaction and complacent inaction that typifies responses to the biggest risks. Following a major incident, we act to strengthen our defences because the event is still fresh in our minds and we are anxious to avoid a repeat. But as time passes the anxiety fades, and so too does our sense of how likely another crisis will be. This false sense of security persists until the next crisis dispels it. And so on.

Another source of distortion is **optimism bias**, which is our tendency to regard the world as a kindlier place than it really is and ourselves as more capable than we really are. Optimism bias helps to explain why big projects over-run and why some people are dangerously over-confident about taking big risks.

Optimism bias has pervasive effects on our judgments about risk. It predisposes us to be over-optimistic about the risk of something bad happening and over-confident about our ability to cope if the risk does materialise. At worst, this can result in outright denial. A recurring feature of protective security and resilience planning is the need to experience a major attack or disaster before converting thought into action. Moreover, the movers and shakers in life tend to be the most confident and optimistic individuals, because these are attributes that society rewards.

Optimism is not all bad, of course. Apart from making us feel good, one of its benefits is helping us to be more persistent in the face of difficulties. If we blithely assume that everything will work out fine, we are less likely to give up. In that sense, a healthy degree of optimism can strengthen personal resilience.

A further source of distortion is our predisposition to worry less about risks that are likely to materialise in the more distant future, compared to more immediate risks. This form of cognitive bias, known as **present bias or future discounting**, makes it harder to form sound judgments about how much effort to invest in preventing low likelihood/very high impact risks such as catastrophic terrorist attacks or natural disasters. A related issue is our inability to understand intuitively how risk accumulates over time as we repeatedly expose ourselves to the same relatively small risk. The risk of dying in a single car journey or from smoking a single cigarette is tiny. But the cumulative risk from a lifetime of car journeys or smoking is surprisingly large.

A pervasive influence on our perception of risk is a bias known as **loss aversion**. We humans are inherently more sensitive to potential losses than we are to potential gains of equivalent size. When making decisions that might lead to gains or losses, our judgments are consistently biased towards the avoidance of losses. We have evolved to be more sensitive to situations that threaten our wellbeing (potential losses) than we are to situations that might bring potential gains. Threats are more cogent than opportunities and therefore bad news takes priority over good news, because it is better to live to fight another day.

One consequence of loss aversion is that we need bigger perceived gains to offset small perceived losses in time or effort. In cyber security, for example, the users of computer systems immediately perceive the losses associated with having to perform irksome security tasks such as updating software, whereas the gains are intangible. Loss aversion means the perceived gain must be considerably bigger in order to win out over the loss.

Our asymmetric response to loss and gain is another reason why resilience and preparedness receive less attention during periods of peaceful stability, when everyone feels safe, but immediately become the overriding priority when something bad happens. When we feel threatened, our normal desire to improve our situation is overridden by a stronger desire to preserve what we already have. To put it crudely, fear trumps greed.

Another consequence of loss aversion is a tendency to stick with the status quo rather than making changes, even when those changes are known to be beneficial. A significant change, such as improving resilience, has costs and almost invariably creates some losers, even when most people are winners and the net effect is overwhelmingly beneficial. Thanks to loss aversion, those who stand to lose are more highly motivated to avoid their loss than the winners are to acquire their benefit. The net result is resistance to change.

In addition to these and other cognitive biases, some individuals have personality traits that make them more than averagely inclined to take excessive risks. The personality traits that have the biggest bearing on risk-taking include sensation seeking and impulsivity. Individuals who score highly on these traits are statistically more likely to smoke, abuse drugs, drink too much, gamble, drive too fast, not wear a seat belt and engage in risky



sex. Psychological experiments have shown that even something as ephemeral as the presence of an attractive member of the opposite sex can have a measurable short-term effect on people's propensity to take risks. The influence of personality traits and situational factors should be borne in mind when choosing people for leadership roles and designing the environments in which they operate. A reckless individual who enjoys behaving dangerously may not be ideally suited to taking responsibility for the safety of others.

## Mishandling crises

The psychological predispositions described above affect our perception of risk and our reaction to immediate threats. Our subsequent ability to manage those risks and deal effectively with crises is influenced by a range of other predispositions and biases.

Four cognitive biases are particularly likely to affect behaviour in a crisis. **Confirmation bias** is the universal tendency to pay attention to information that supports our existing beliefs while ignoring information that contradicts them, thereby entrenching preconceptions. Obviously, this can be hazardous in major incidents, where the situation is rarely clear-cut and existing beliefs sometimes turn out to be wrong. The ability to heed contradictory information and modify our beliefs in the light of new information is vital.

Then there is **groupthink** – the inclination to follow the pack and conform to the majority view, even when we suspect the majority view is wrong. During crises, the people trying to manage the situation rarely have a full understanding of what is going on and how best to deal with it. Groupthink can increase the risk of uniting around the wrong course of action and failing to consider better options.

Next is **sunk-cost bias**, which inclines us to persist with a course of action because we have already invested heavily in it, even if objective evidence suggests it would be better to cut our losses and desist. The investment need not be financial, of course. Sunk costs include time and emotional investment, and hence the bias keeps us plugging away at failing projects or persisting with the wrong policies.

Finally, when things have gone badly wrong and the official enquiries are underway, we are inclined to exhibit **hindsight bias** – the unattractive propensity to be wise after the event. When we retrospectively blame individuals for making the wrong decision, we assume that the error lay in their poor judgment or folly. A more likely explanation is that they did not know then what we know now. Many well-made decisions turn out to be wrong for reasons that were unforeseeable at the time, and many of those who cast blame from the lofty position of hindsight would probably have made the same decision had they been in the same situation. We apply hindsight bias to our own judgments as well, with the result that we react to a surprising event as though we had been expecting it all along. We

unconsciously revise our beliefs in the light of what we have just experienced and forget what we previously believed.

Hindsight bias is an obstacle to learning from adverse experience and therefore a barrier to developing active resilience. If we knew all along that something bad was going to happen then what more is there to learn? Hindsight bias also fosters an unhealthy culture of risk aversion. Officials soon come to realise that if something bad happens they will be judged by the outcome, not the quality of their prior decisions. This encourages them to adopt standard procedures that would be harder to criticise after the event, regardless of whether they are the best procedures.

## Aiming off

This baleful catalogue of psychological snares and delusions should not be read as a counsel of despair. There is much we can do to counteract their less helpful effects. Simply being aware of their existence is a start and consciously bringing them to mind is a good practice. We can further enhance our ability to make better judgments by listening attentively to people with relevant experience and running a lessons-learned exercise after every significant incident. Beyond that lies a menu of methods for countering our cognitive biases.

A simple technique for confronting confirmation bias and groupthink is Devil's advocacy. This entails making someone explicitly responsible for challenging the consensus and advocating an alternative view. Devil's advocacy serves as a check against a dominant opinion being accepted too readily. It also helps to build confidence that an eventual conclusion will withstand robust scrutiny.

A more elaborate technique that counteracts groupthink is the **Delphi method** – a structured process for distilling the views of experts. One problem with putting experts around a table is the intrusion of group dynamics, which can result in the conclusion being less than the sum of its parts. The Delphi method provides an antidote. It works by eliciting the experts' views individually, anonymously and iteratively. Each expert gives their opinion, usually by means of an online questionnaire. The various opinions are collated and the participants receive moderated feedback, which summarises the emerging views without revealing whose views they are. The process is repeated, usually for three rounds. The aim is normally to reach a consensus, although the exposure of dissent can be equally valuable. A practical benefit of the Delphi method is that the experts do not have to meet physically. Its main benefit, however, is diluting the powerful interpersonal dynamics that often play out when opinionated experts convene.

The distorting effects of optimism bias and groupthink can be resisted with a technique known as the premortem. The premortem is held just before an organisation commits itself to an important decision to proceed with a plan. The individuals most closely



associated with the decision are gathered together and told to imagine a future in which they went ahead with their plan, which turned out to be a disaster. Each of them is then asked to write their own brief history of the imagined disaster, describing what went wrong. The premortem gives the decision-makers licence to express doubts and forces them to contemplate, if only briefly, the possibility that they might be wrong.

Other techniques for countering biases and improving the quality of analysis include checking key assumptions, reviewing the quality of the evidence on which the conclusions are based, analysing alternative hypotheses and conducting a “What If?” analysis. When trying to assess complex situations, it is a good idea to list all the key assumptions on which the final judgment will rest and check that they are valid. Another good discipline is scrutinising the quality of evidence from which the conclusions are derived. Some strands of evidence will be stronger than others and it is all too easy to forget caveats about their reliability. We can improve our confidence in a judgment by systematically identifying and evaluating alternative hypotheses (a more elaborate version of Devil’s advocacy). This helps to avoid becoming too firmly wedded to the first solution we thought of – a tendency reinforced by confirmation bias.

Probably the single most effective defence against making bad decisions is expertise – the combination of relevant knowledge, skills and wisdom needed to cope with demanding situations. Experts pay attention to the right things, whereas novices focus on irrelevant details and become overwhelmed with information. Experts also cope better with ambiguity. They ask the right questions and know what to do.

Expertise forms the bedrock of sound intuitive judgment. In a stressful and uncertain situation, the ability to make quick, intuitive judgments is crucial. Intuition is often portrayed as a semi-mystical property born of age and wisdom. However, psychological research has shown it to be something more mundane. In simple terms, intuition is a form of pattern recognition – the ability to spot patterns or signals in uncertain situations, based on past experience. That is why individuals who are highly experienced in a particular field, whether it be mountaineering or protective security, tend to have more accurate intuitions. Their experience has equipped them with a richer set of reference points against which they can recognise patterns and quickly judge how a situation is likely to develop. They spot a cue in their environment and can link this to relevant information in their memory. One practical implication is that we should pay more heed to people’s intuition (including our own) if they have more relevant experience. The surest way of acquiring relevant experience is through regular practice, training, testing and exercising.

Technology can help by providing automated decision-support tools that guide the thinking processes of decision-makers when they are under pressure. However, there is a world of difference between automated decision-support tools and automated decisions. Automated systems cannot be expected to make optimal, unbiased decisions in complex and uncertain situations, not least because the algorithms that underpin them are the products of imperfect human judgments.

---

## COMPLEX SYSTEMS

As noted above, we humans are often guilty of oversimplifying situations and not identifying the long-term effects or unintended consequences of decision making. History is littered with examples of where humans have not considered the broader picture when making decisions. For example, the Beeching cuts to the UK's railways in the 1960s failed to identify many broader, longer term consequences. When the severe snows of December 2010 hit London, much of the rail network ground to a halt. One small but important contributory factor was that rail personnel were unable to travel to the starting stations to crew or dispatch trains because they lived some considerable distances from stations and depots and were unable to travel due to the snow. In the pre-Beeching world, many of them would have lived in railway-owned accommodation situated near stations and depots. The sale of the real estate following implementation of the Beeching recommendations had led to an unintended consequence nearly 50 years later. Knowing this today can enable train operating companies to build resilience by prepositioning their drivers when forecasting indicates a high probability of inclement weather that may stop railway employees getting to their destinations for work.

Nature is littered with examples of where the complexity of interlinked systems themselves builds inherent resilience. Symbiotic systems, where individual parts of the system work with each other for mutual benefit, abound at every scale on our planet. From the relationship between the clownfish and the anemone, through the complex dynamics between trees, fungi and insects in a forest, to the global interactions of fauna and flora within oceans and its interactions with the atmosphere, we see the ability of the system to veer and haul during times of feast, famine and external challenge. And in each case, if the symbiotic relationship is pushed too far from its natural equilibrium by external factors, such as diseases, pollutants or natural disasters, then we may see a system failure – a coral reef dies, a forest is overrun with an invading non-native species, or climate change develops into a runaway crisis. In every case the system will initially counteract the change from equilibrium by adapting to emerging behaviours. But in many instances, that initial system failure can cause a cascade of impacts – a coral reef dying can start to precipitate local, wider ecosystem failures which eventually cascade to the planetary level. For the want of a nail... Knowing enough about how the complex systems are working, their interactions and, most importantly, the conditions in which they can remain stable, are core starting points for understanding their broader resilience.





There is often confusion between systems that are complicated, as distinct from systems that are complex or exhibit complexity. Something that is complicated can be designed and ultimately has predictable, controllable behaviour. Not so with complex systems. The point is well illustrated by Professor Chris Rapley, a leading climate change scientist. During his time as Director of the British Antarctic Survey, he had a discussion with a field ecologist. The ecologist was studying a small patch of ground that was an oasis of activity in the otherwise barren landscape. Within this patch were various plants, fungi and bacteria, all interacting in a complex system. Rapley suggested that the ecologist could perhaps identify the key species in the system and separate them to simplify it, making it easier to understand. To this the ecologist bluntly replied that physicists are always trying to isolate things and make them simpler – to him the objective was to study the complexity.

When multiple systems are connected, the way in which they collectively behave can become unpredictable and is said to exhibit emergent properties. An example of emergent behaviours relevant to the science of resilience concerns human culture. Professor Eve Mitleton-Kelly of the LSE notes that you cannot design a culture, despite the best efforts of organisations to do so; you can only provide a framework in which a culture can develop. Culture is something that emerges from complex human behaviours.

Emergence is a term used in systems science, biology, philosophy and art to describe properties of an entity that arise from the interactions of smaller entities. The emergent property is not found in the parts from which it arises and, as such, emergence represents more than the sum of these parts. Consequently, making assumptions based on a view of only a subset of the overall system can lead to poor decision making.

***One of the first principles to establish when designing a resilient system is to recognise the limits on understanding that complexity brings and ensure that the solutions are engineered accordingly.*** For example, there have been instances of run-away stock markets driven by High Frequency Trading – that is, computer-to-computer trading driven by algorithms. The algorithms are developed by humans, driven by data and make very quick trading decisions without human intervention. And if they are driven outside the bounds of their assumptions then a cascade of trading events can take place, potentially leading to major shifts in the stock markets driven by algorithms and not reflecting economic factors.

Trying to understand systems, their relationships with other systems, and their collective behaviour remains at the cutting edge of research. But we can gain some understanding of core features of interdependent systems by looking at simpler systems with which we are more familiar. For example, physical systems can undergo phase changes. Water in a glass (the system) freezes when cooled below zero degrees Celsius, for example. It changes from one phase (water) to another (ice). But it also has another, more catastrophic mode of changing phase, exemplified by supercooled water. In certain circumstances water can be cooled well below its normal freezing point of zero degrees Celsius and remain liquid. This requires the water to be very pure and undisturbed.

However, as soon as an imperfection is introduced – say, a particle of dust lands on the surface or the side of the glass is tapped – an instability is introduced and the water immediately and violently changes state from water to ice. Moreover, it does so in a way that, unlike ordinary freezing, cannot be stopped once the process has initiated. It is an example of a catastrophic phase change.

The wider world also exhibits such instabilities whereby a system is pushed a long way from its equilibrium, appears for a while to be normal and then suddenly and unpredictably changes its state. The financial crash of 2007/2008 is arguably an example of a system that had moved a long way from equilibrium, was unstable and suddenly went through a change of phase.

A former Chief Systems Engineer in the Ministry of Defence once noted that ‘Systems are simple until you add people to them’. The trajectory of human civilisation has been one of constantly trying to produce order on the planet by the deployment of energy and natural resources. Left alone, the planet would find its own equilibrium and disorder would gradually increase over time – in the language of physics, entropy always increases. In some senses, humanity has moved the planet away from its natural equilibrium and artificially sustains it in an unstable equilibrium, where a relatively small change to a condition can lead to a major change to the system. Humans are able to sustain a system when times are normal, but when the unexpected happens (Covid, subprime mortgages, etc.), a relatively small change to the system can cause runaway effects that become uncontrollable, leading to unpredictable outcomes with global impacts.

Our interventions as humans on this planet are inherently risky. We never know quite how near to that phase change we are pushing the system, or what levers we might have available to manage system impacts if it does start to go through a phase change. ***So the second principle for designing a resilient system is to ensure that the system has significant latitude to flex in response to a broad range of changes to internal and external conditions without going through a phase change, and that we suitably instrument and measure it to monitor its behaviour.*** While we may not be able to fully understand the complexity of a system, we can put in place diagnostics and instrumentation to measure its behaviour, identify when the system is changing in unexpected ways and respond accordingly.

Part of building resilience is just better measurement and instrumentation of our systems, so that we can understand the direction in which they are heading and have the opportunity to steer them away from more dangerous conditions before it is too late. But in many cases this balance is difficult to achieve, because the more interesting, innovative and often financially beneficial outcomes from human engineered systems occur at, or very near, the boundary of a phase change, and in many instances they actually generate a phase change. The industrial revolution is one such example and many commentators believe that the ubiquitous rise of Artificial Intelligence (AI) is already driving that next major phase change.



AI and its underpinning data science is a good example of how a new capability can be both an enabler for building resilience (such as managing complex supply chains or cyber security defences in real time) and yet also capable of introducing new issues that reduce the resilience of systems (such as the 2010 Flash Crash cited earlier). The vast majority of AI today is classed as ‘black box’: a set of inputs produces a set of outputs without any true understanding of what happens in the middle. AI systems are often trained by copying other systems, notably humans, but the reason why an AI system makes a particular choice can be opaque. It just does so (although the same could be said for humans). Applying such artificial intelligence to systems that need to demonstrate high integrity – for example, financial systems, self-driving cars and autonomous systems more generally – is all well and good until things go wrong. Whose fault was it? What went wrong?

AI may need to become increasingly amenable to audit and reverse engineering to enable us to see what went wrong. An active area of current research is around new forms of artificial intelligence that enable us to understand better the links between cause and effect. Applying AI techniques to complex data sets to help us understand how complex systems evolve and might be managed is increasing in importance all the time – from drug discovery in the world of biopharma to the interactions of global economies. We are only at the beginning of this journey.

Returning to the severe snow of 2010, it emerged that there were starkly different outcomes to the ways in which the two major international airports in the south of the UK responded. Heathrow airport had substantial difficulties during the incident and its runways and operations had to close down. Gatwick, on the other hand, was able to continue running, albeit at a lower capacity. Comparing their different responses and outcomes reveals an important lesson about systems and resilience. The Begg report on the incident noted that Heathrow was ill prepared for snow and recommended that the airport “should adopt an improved resilience target that [it] never closes as a result of circumstances under its control, except for immediate safety or other emergency threats”. Heathrow’s crisis response planning appeared to be focussed on a number of core scenarios whereas Gatwick planning was based around the core functions needed to respond to a broad range of difficult circumstances. While scenarios can be very useful in exercising crisis response plans, they may not be a good mechanism for building resilience into a system and responding to complex, emerging situations. Too much focus on planning for particular scenarios can make it harder to cope when the unexpected happens. What is far more important is to build functions that enable systems to be managed during times of challenge, and to enable those systems to degrade in a managed way under control of the decision makers. Ultimately, resilience is independent of scenarios. ***And thus, the third principle of resilience and complexity is that we should invest in general capabilities (functions) and decision-making abilities, rather than prepare for a small cluster of specific scenarios.***

There is a growing move to build ‘digital twins’ of complex systems in order to explore issues of complexity in a ‘safe space’. Such digital twins are not just simulations of reality, not some sort of a virtual-world representation of reality as popularised in the Matrix film series. Rather, they are an abstract simulation, often driven by real world data, allowing decision makers, designers, operators and others to explore opportunities and choices before interacting with the real world – whether that be cutting steel, changing interest rates, building a new city or intervening in climate change. They often focus on some areas at much higher granularity than others to analyse relevant areas of the decision space before committing. They enable the decision maker to stress-test their area of responsibility in parallel with the real world before formally committing to do so.

Digital twins are not a panacea, however. Care must be taken when building them to ensure that the models are not extended beyond the assumptions embedded within them, and that when multiple digital twins are federated the assumptions are compatible. And that itself becomes an area of difficulty. How do you know that your model is valid?

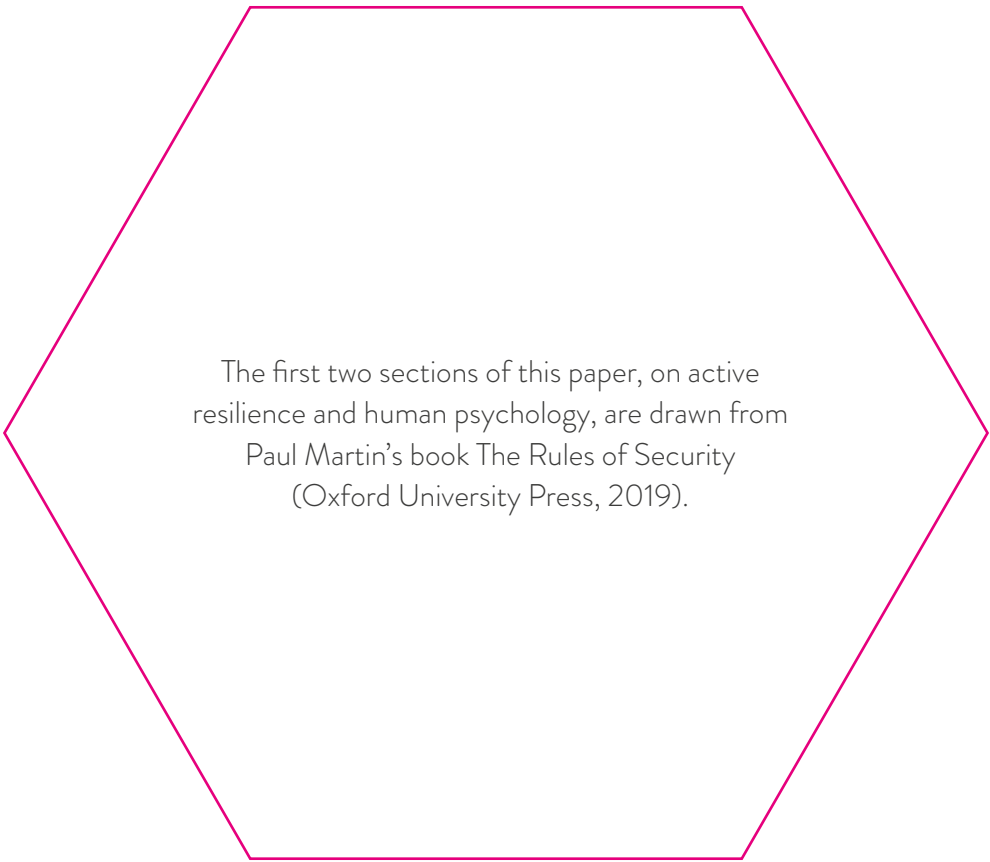
There is increasing interest in driving digital twins or simulations in real time and measuring their capability against real world behaviours. When something happens in the real world that knocks it off course (for example, a closed motorway, a meteorological event, a pandemic disease) a digital twin can be used to experiment with possible real-world interventions to manage the situation. By understanding diagnostics in the real world, feeding them back into the model and making decisions accordingly, in real time or near to real time, it is possible to both validate and verify. These models can be used to help decision makers look for the best ways of managing an evolving situation – and identifying the gaps in knowledge.

Much of the important work around digital twins is about developing common frameworks, or a lingua franca, to enable communities of interest to share, integrate and collaborate and advance by learning in a far more effective way. And that is the core of the fourth key principle of complexity and resilience – which is the need to provide a common framework for all stakeholders to share, integrate, collaborate and advance by learning. While fully understanding ‘the system’ at every level of complexity and at every level of abstraction may be a pipe dream, the ability to model it, measure it and in some senses to experiment with it enables those involved at all levels of policy, industry government and beyond to explore new ways of operating, new ways of collaborating and ultimately to build more resilient futures.

---

## CONCLUSION

**If you want to do better at building resilience, avoiding crises and managing the crises you cannot avoid, then you need to (1) understand and apply the concept of active resilience, (2) understand and adapt to human psychological factors, and (3) apply the lessons from the science of complex systems.**



The first two sections of this paper, on active resilience and human psychology, are drawn from Paul Martin's book *The Rules of Security* (Oxford University Press, 2019).





[www.nationalpreparednesscommission.uk](http://www.nationalpreparednesscommission.uk)

c/o Resilience First | 85 Tottenham Court Road, London, W1T 4TQ