The Ratio Club Revisited

Professor Daniel Gideon Dresner, PhD, FInstISP and… see *V Acknowledgements*

Abstract — **Inspired by the multidisciplinary collective of the original and its allegiance to Wiener's cybernetics, the University of Manchester's Digital Trust and Security team has re-established the Ratio Club as an instrument to wade through complex matters and attempt some clear thinking on selected topics on that contemporary cyber challenge of cyber security.**

This paper sets out some the thoughts and discussions of the second outing of the reconvened 'Club' and its proving as a format-for-hire for thought leadership and problem solving with an air of academic rigour amidst the nudges of contemporary relevance. A kind of 'Delphi-technique on tap' one might say.

The Ratio Club is working on a diverse membership profile to retain a topic-agnostic freedom and promote a degree of verisimilitude to whatever the challenge it is set for scrutiny.

The winter air of Manchester was warmed by the passionate commitment of the participants in the discussions that followed the keynote view of artificial intelligence which settled into the evening's theme of 'Resilience'[1].

Index Terms — **Artificial Intelligence. Cyber Security. Cybernetics. Interdisciplinary. Machine learning. Multidisciplinary. Resilience. Systems thinking.**

I. Introduction

The University of Manchester reconvened The Ratio Club in December 2018, its second outing since 1958. Following in the tradition of the original, guest speaker Dr Robert Hercock, Chief Scientist at BT, set the scene with a keynote about artificial intelligence and concepts of mind and thinking. Vibrant discussions ensued amongst the evening's ratiocinators and we have attempted to capture as many of these herein. The challenge of this exercise is to give structure to a debate that one wishes to deliver conclusions and outcomes but one does not want to constrain.

So this paper includes both thoughts from the evening, and interpreted thoughts after the matter. Although we had envisioned the evening taking the challenge of answering six questions (*see below*), these were largely lost in the mix of discussion and the enthusiasm of participation. Part of the catalyst for participation and engagement – apart from a good dinner – was to ensure that the variety of disciplines had the opportunity to trade their world views within the constraints of time and space in the venue. To this end, some 'ratiocinators' changed places between courses of the evening's dinner to keep the conversations moving.

'Questions are a burden to others…'[2]

We had set out six questions for discussion – in keeping with the theme of 'resilience'. However, by the end of the evening it became apparent that few – if any – of our guests had paid them much attention! This left us with a body of notes taken by the facilitators and a challenge to map these to the questions which formed the direction we – as organisers – had hoped the evening would take. Three participants heeded the call to address them in written responses with only the lightest of touches to spruce them up for publication (as appendices).

The questions were always in the back of the organisers' minds, at least, and comprised:

(1) What characterises resilience?

(2) How do you measure resilience before it's tested by real world events?

(3) Is there a balance to be struck – and if so how – between the resilience of people in, or affected by, the system and the estate(s) or technology(ies) that comprises the rest of the system?

(4) Should we treat an attack on our digital lives like the harm of a physical attack? When we lose data or access to 'digital' are we bereaved? Do we suffer grief? Will we react to the next bleep from our devices with symptoms of PTSD?

(5) Does 'cyber resilience' exist as a thing or – taking a systemic approach – does it become the way of articulating resilience?

(6) Can you adjust a state of resilience to encourage more favourable outcomes? How might we compensate for expectations of recovery to the old state when we'll be dealing with a system that has come through a change in itself?

In this paper, we have retrofitted the records of the discussions not as answers to the questions but rather as data for the future consideration of how the questions may be

[1] A topic embraced by the complimentary IAAC Cyber Leadership Forum.

[2] '…Answers, a prison for oneself.' The curiosity-stifling mantra from the village in Patrick McGoohan's The Prisoner, ITC Entertainment, (1967 – 1968)

answered. This makes the paper a sort of epistemological pot-pourri. From hereon in – to the Conclusions – we note the discussion with an air of the Chatham House rule. This is not an attempt to withdraw attribution to often original and thought provoking remarks, but rather an attempt to capture as much wisdom as we can without resorting to a format of, 'He said…/she said…'.

WHAT CHARACTERISES RESILIENCE?

Without examining the whole lyric, it would seem that the anthem of resilience might be, 'I'm still standing' (Elton John, 1983). Socially, resilience is 'learning to live' and from a cyber security aspect, learning to operate with bad people on the system. This suggests that – as with cyber security – context is important to understand what is be done to at least feel resilient, and what can we measure to give us that feeling. What threatens our resilience and is your resilience different to mine? If we are to protect the key assets – what are they?

Perhaps there is a lack of objective thought about the diversity of expectations to be found for resilience. Technology provision is as much now in the hands of the general populace as the colour of the Model-T Ford was for early motor car customers. The democratisation of connectivity – in certain nations at least – leads us to assume that connectivity is ubiquitous, leaving very few unconnected. We race to design for Mr Spock rather than Homer Simpson in a way that every driver is expected to be a car mechanic too.

We swing between the paradigms of withstanding nation state on nation state cyber attacks and the 'Clapham Omnibus' empathy of the citizen facing theft and fraud. Resilience seems to come in layers. Despite the moribund information security standard of government which failed to take us into a secure cyber age, the concept of the business impact table has almost appeared again in the Cyber *Attack categorisation system to improve UK response to incidents*[3] (NCSC, 11 April 2018). This appears to question the resilience model of the citizen and the SME who are – to all intents and purposes – reliant on their own capabilities and resources in the event of attack. Herd immunity is there, providing that either you are in the right herd or the immunity is cast for the greater good.

So what does good resilience look like? Can it be modeled? Do we have to make choices between (analogously) hard-shelled coconuts or tough-centred avocados? The latter is becoming prevalent in telecommunications – light protection outside, a soft middle, and a highly protected core.

Can we rely on the quirks of our adversaries? In warzones (as observed in Afghanistan), people don't attack communications towers because they need them too. Criminals don't want to destroy banking, they want to live off it. This is parasitic rather than symbiotic, although observing their actions and intent might hint at how to at least achieve a state of antifragility against a perceived desire for robustness. Can resilience be seen at how expectations are managed? For example, power resilience in India is a function of the availability of generators – the power fails so often so generators are commonplace.

The dependence on interconnectivity is past the point of reversal. The proverbial genie won't fit back into the bottle; we must adapt and join with the systems. But human values aren't universal; they differ from person to person and cannot be programmed.

HOW DO YOU MEASURE RESILIENCE?

It was perhaps not a surprise that there seems to be little in the discussions that could be mapped to a question of metrics for cyber resilience – or indeed even those aspects which fit the clearly cyber *security* label. What would be the indicators of an oncoming storm that may rally activity to prevent the 'big attack' happening or cope with it whilst it did?

Part of the problem still lies in the much wanted datasets that reveal what happened and how. Here we are looking for leading metrics and the proverbial rear-view mirror of lagging metrics is covered up. Despite the beefing up of mandatory reporting under the General Data Protection Regulation (GDPR – 2016/679) and the perhaps more prescient EU Network and Information Security Directive (NIS – EU 2016/1148), it would seem that we must still find ways to avoid companies actively suppressing information about security breaches. Some of this may be down to cognitive dissonance; change is resisted just because it's a lot of work to unlearn old patterns. The secrecy may be a habit that's hard to break.

Perhaps we can look forward – metaphorically – to suggest where on the road to resilience we are. We might do this by looking at certain areas that surround the invisibility of cyber security. The presence of these activities, working with diversity in the face of adversity, may help us to detect the undetectable. On the table for consideration are:

- How symbiotic is the system design? Have we stopped swinging on the pendulum at the weakest link yet?
- How active is our monitoring and detection? Cyber security has commoditized intelligence. There is no excuse for not using it.
- Training and education is important but as part of wider programme (*see above and below*). The ability to recognise a phishing e-mail is good, but is less likely to happen months after training or when under pressure. Combining it with better network segmentation and tools to reduce the phishing messages arriving is more sensible. Finding ways to measure the impact of training is more interesting.
- How well the segmentation of devices, networks, and people key to threat-resistance is done.

[3] https://www.ncsc.gov.uk/news/new-cyber-attack-categorisation-system-improve-uk-response-incidents

Cyber security advice is skewed towards identify, detect and protect (prevent) – as in the NIST framework. There is very little about response and recovery.

This approach to measurement proposes a move beyond the culture of 'waiting for failure' and so prompts adjustment in response changes in the leading measurements that are collected. It might support a view of cyber antifragility.

WHAT'S THE BALANCE BETWEEN THE RESILIENCE OF PEOPLE AND THE TECHNOLOGY?

It is often observed that both the lessons of history and the concepts of science fiction have much to teach us about cyber security[4]. But however many times we seem to go around the security block with our imaginations leading us, our adversaries seem to be there waiting to exploit something new or something we might have expected to be long-since sorted out. Despite the speed of our technology, we still seem to be challenged to navigate innovation through the politics of inertia without a high degree of damage and harm that is jauntily marketed as 'disruption.'

Companies die; cities are immortal. Companies become monocultures. They lose diversity. Companies fail variety – the protective measure that Ashby observed systemically – not at the heart of – cybernetics. Monocultured societies or systems seem to embrace problems like a clenched hand around a hot pan handle. The more order you put in, the more vulnerable it becomes. Why are cities relatively immortal? How do they differ from the adaptive systems we hope to create for our digital communities? Perhaps it is because of scaling. Computer systems morph and become mono-cultured; they try to diversify but they can't adapt...yet. What is the cross over point for cyber resilience? Where are the transitions made from individuals to communities, from regional to national scales? Looking backward to the individual, that has never been a single unit – brain and other organs are presented in diverse and changing forms through aging, or disease, and through changes to health and wellbeing. Which factors that contribute to human immunity could be adapted to cyber immunity? How do we immunizse against behaviour – and behaviour changes – which bring harm to the individual, the communities, supply chains, and undermine the positive risks of achieving a positive state of antifragility?

WHEN WE LOSE DATA OR ACCESS TO 'DIGITAL' ARE WE BEREAVED?

Cyber impact tables (*see above*) do not give forewarning of the risk landscape. They may be symptomatic of risk being in the eye of the beholder[5] making it difficult to create a levelled, universal assessment. However, if there is one dock leaf amongst the nettles of cyber threats, it is that the Internet has led to the commoditisation of intelligence and an understanding of the localised exposure to risk is accessible. For example, IASME's risk profiling[6] enables a small to medium-sized enterprise (SME) to gauge the level of protective measures that a business should take to orchestrate an acceptable level of information assurance. This includes the fundamental recovery methods that the business should have rehearsed to assure that a cyber attack will not have lasting, detrimental effects.

There is a growing data set that may show what is apparent from anecdotes from the initial security review of these small businesses. The idea that risk can be transferred to anything other than underwriting costs to insurance is debatable. It manifests as businesses transfer the risk of their resilience – described by preparations for disaster recovery and business continuity – to third-party service providers. Decisions are based on the relative size of the provider (a significant cloud computing brand) and the SME. SMEs consider the cloud architecture providing access to commensurately secure data or applications with integrity to be sufficient to dispel with long-held ideas of back-up and data restoration. It will be interesting to see cloud providers – as the Digital Service Providers of the NIS Directive – being held to account for a business' consequential loss in the face of a cyber attack.

Like NCSC's cyber attack categorisation, the individual or SME is unlikely to get the attention they deserve as they firefight and move from the state of 'who you gonna call?' to 'who you gonna sue?' The provider will doubtless be protected by the long-since clicked-through terms and conditions.

We are in a state of uncertainty where it is not the likelihood of cyber attack that is unknown. We live in an environment of persistent attack (even from automated malware released into the wild long ago[7]). It is the victims' resilience to the effects of an attack that is most likely to hold the challenge. Until we have trustworthy systems in the widest sense[8], we require a coping strategy whilst in this period of inevitable risk: an ability to recover despite a lack of preparation. This should extend not only to the technical cleansing of devices and their resident software and data but also to the post-traumatic stress of the people who will have suffered.

Threat perception will be found along a variable scale with viewpoints including international, national, local, community, and personal (including family). A proliferating quagmire of prevention advice is often difficult to navigate, conflicting, and ironically assumes that the person needing it will have Internet access, when in practice they may have lost all safe access or may be too nervous to log back online.

[4] See CyberTalk, 6, The SciFi Issue, Autumn 2014

[5] Bernstein, P. L., (1998) Against the Gods: The Remarkable Story of Risk, John Wiley and Sons

[6] The IASME Consortium, *The IASME Governance Standard for Information and Cyber Security*, Issue 5.0, January 2018

[7] Shanmughapriya, M., Sumathi G., Aarthi, K.C. (2018) Bot Net of Things – A Survey, International Journal Of Engineering And Computer Science, Vol 7 No 05

[8] BS 10754-1:2018, *Information technology. Systems trustworthiness. Governance and management specification*

What actions might help deal with the technology involved? This will involve the disciplines of computer science and engineering but immediately connects with the forensic integrity demanded by law enforcement and the intelligence service. There will be information assets that may need to be recovered that are common, if not to all then to many: bank accounts; social media accounts; e-mail accounts; computers (whatever their guise: laptops, tablets, smartphones), as well as the specialised items from the so-called 'Internet of Things' ranging from cameras and optional health monitoring devices through to critical items such as pacemakers and insulin pumps. The sheer volume of recovery needed means that we shall have to call on the very systems that are giving us heartache. Getting autonomous mitigation and self-defending systems right (Jones and Dresner, 2014) is our next big 'moonshot' (according to Nicola Whiting, CSO, Titania Ltd).

IS 'CYBER RESILIENCE' A THING OR JUST A WAY OF ARTICULATING RESILIENCE?

The discussions pertaining to the challenge of defining what resilience – or cyber resilience – might be characterised by pervaded the evening. They were particularly bound by the attempts to drawn on the metrics of resilience (*see above*) which may not define it, but may at least help us define what 'good' looks like.

Much is described in terms of the experience of what has not worked. Programmes that have nothing to do with cyber security may bring useful analogies. Challenges in German forestry were referenced[9]. But it was out of the positive that the dark twins of duality emerge to evolve from the monocultures. Examples vary from the shanty town in the thrall of the affluent (or is it vice versa?) and the jailbreaking of technology to create freedoms at high risk.

Not learning almost seems to be an active rather passive activity (oil spills were mentioned here) in that what could be learnt is suppressed. There are opportunities for learning from giant mistakes such as Y2K and CO_2 suppression. Despite the foreseeability of catastrophe, the attempts to forestall it tend to favour waiting to seeing success or failure. The interim action tends to doing something to feel safe and happy rather than solving the problem. We are just starting to retract the blame and shame culture because of the damage it does to suppressing knowledge. We might not have the requisite variety to control the network effects. But we might be able to see the tipping points where we can't program change but can influence the directions taken,rather like the use of phenomenology in architecture. Cyber – or information – security was locked into the targeting of creating a known state, described in the *Risk Management and Accreditation Document Set* (RMADS) with periodic inspections back to that space that struggled with the evolving context of the environment in which the system operated.

CAN YOU ADJUST A STATE OF RESILIENCE TO ENCOURAGE MORE FAVOURABLE OUTCOMES?

Conventional wisdom has it that when people's values change, they will reliably change their behaviour. We discussed ways of starting with behaviour change where this in turn might drive a change in values. Successful examples of this include smoking bans, drink-driving, and seatbelt wearing (*see above*). These met resistance when they were introduced but have now become strongly embedded in the culture and values of society. In the case of cyber security, is the government strong enough or clear enough about what is needed to impose this? As drivers will hurt others and smokers will take medical resources from those unable to counter the risk of disease threatening them, so too do poorly-protected Internet resources threaten others on the network or the network itself.

Government initiatives take credit for what are really cultural successes. Smoking bans work because of culture, not fines and cancer. They create the 'policeman in the head' overseeing the cultural incentives.

Is resilience a state of feeling safe and happy? When government comes up with advice for resilience – think about the 'Protect and Survive' mantra of the 1980s – we see conflicting messages of wanting government guidance and feeling that government ought to be doing more for us. What are the values that we want to see preserved? If we put values in, do we get resilience out? When is harm more important than share price? Has a generation of cyber talent been suppressed by the parents who saw War Games (1983, United Artists) in the cinema and have steered their children away from being experimental with computers?

By of way of an analogy in evolution, this could be like developing a frontal cortex to control behaviour - pervasive behaviour-change strategies. A good example is the COM-B framework for understanding behaviour (Michie, van Stralen and West, 2011). This might form a toolbox of cognition to be programmed into systems to encourage human-machine symbiosis. Change the environment to make people behave better; use dark design to produce a user interface that nudges users into doing safe things. This would include designing interventions so that it is harder work to do something unsafe, stopping things getting through to people in the first place, or mounting a campaign such as People-Like-You-Act-Like-This. Such schemes require careful management of their effects to avoid nurturing digital inequality amongst those quick to adapt and those not.

Influencing good behaviours to encourage cyber security is not helped by 'IT' crying foul of existential threat when people's immediate experiences are different. Expectations must be managed. Master-slave is a term embedded in the history of computing and communications. How far do we stretch this analogy? Do we see our expectations of mastery over our systems as lost as the slaves revolt? Perhaps a cyber security breach hasn't happened in a bad enough way yet to enough people to change the societal mindset. We

[9] Scott, J. C. (1999) Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed, Yale University Press

may never know how many lives were shortened by the Wannacry infection suffered by the NHS in 2017.

As we discussed above, we are fixated on our rear view of failures with a layer of fear, uncertainty, and doubt as to what can happen in future. We rarely see cyber security risk positively (security product vendors not withstanding) as an opportunity to consider the risk of success by looking at the good things and asking. 'Why did this work?'. Rather than trying to influence people, try to understand how it's done. The slow changes to legislation for the greater good have no capacity or mechanism to consider the probabilistic paths to emergence – good or bad. Changing our attitude to security and risk rather than rushing to change infrastructure might deliver the calmness for resilience to emerge.

The discussion frequently centred around the difference between robustness and resilience. Robustness was described as coming back from a crash, whereas resilience was described as not getting to the worst crash point in the first place by developing coping strategies so as to be able to continue to operate under stress. The example tabled was a solar event taking out the national grid: national resilience may mean helping some people and not others. This takes us back to the NCSC attack categorisation.

The models of cybernetics and the concepts of Ashby's homeostat cry out for attention. Organisations who survive cyber security breaches rarely point to their business continuity plans with pride. They have not recovered from disaster but rather an organism that has survived. 'No campaign plan survives first contact with the enemy.'[10] but the synergy created by the planning process is priceless.

Dresner and Jones' (2014) paper – The Three Laws of Information and Cyber Security – focused on 'security' as a state of resilience to attack (or failure from other causes) that can be recognised by the successful delivery of the 'system' under scrutiny to be delivering on its objectives. This was a useful approach to understanding resilience in the context of cyber security. The system continues to 'be.' Or does it? And if so, by how much? The suggestion of deterministic control objectives for security are replaced by a systemic view of overall objectives. Control models at least suggest that resilience is a worry. Attention to recovery is just starting to emerge. When the other 3 laws don't work – the Asimov's 0th law is the thing.[11] The greater good model again…

II. The Wayward Child of Cyber Security

The Ratio Club meeting centred around the theme of resilience with intellects being fired up by Dr Hercock's view of evolving artificial intelligence.

You could almost label artificial intelligence as the wayward child of cyber security. But this is not necessarily a bad thing. I'm almost tempted to add 'in the right hands' but whose hands those are is a terribly subjective thing and an honest colleague might grab an opportunity (in the words of Francis Bacon) when personal circumstances create pressures that may be alleviated monetarily (or from other forms of relief which we won't discuss here).

Shakespeare might have observed that there are some disciplines that are pure cyber security (think all that design of malware and anti malware), some disciplines which may have analogous algorithms that may inform cyber security (think approaching computer networks as quasi-biological entities that need an immune system), and Malvolio's third view of the disciplines that will have cyber security thrust upon them or lose control to bad actors in a sea of exploited vulnerabilities.

So, what has this got to do with cyber security and artificial intelligence. Leaving aside - in such a short space – the is it or isn't it discussions of AI, machine learning and complementary technologies – we stand before an opportunity that inaction could cause us to lose, and its new owner may not fly a national flag that we like. It may also be flag of iconography which may be more readily associated with forensic anthropology. There is a dark side and a light side to every technology. The challenge will be making sure that AI doesn't come to make its own decisions as to which side it is on. Some AI will be used for good, some for bad, and some will be used to turn one from the other.

Contemporary AI is an augmenting technology. With the petabytes of data that may conceal just a few bytes and nibbles under control of our adversaries, AI may just give us the support we need to identify what's going on and face up to the challenge with well-decided actions. Dresner and Jones – in our 2014 paper – define cyber security as having three components. Two – which have a traditional security feel – are wrapped around the other. They see the operations of computer systems from the (not so) humble tablet computer all the way up to gargantuan data centres as the core of existence. They are there to book train tickets, watch Netflix, deliver Netflix, run power stations, switch traffic lights and so on *ad* – the complexity of – *infinitum*. They are there to operate. Wrapped around this operation are the expected good things to do for cyber security - the protective measures such as HMG's Cyber Essentials, the SME's IASME Governance System, or international standard of good practice number 27001 – and that part of cyber security that cries resilience: the ability of the system to self-preserve.

And this the idea whose time has come (to play on the words of Stafford Beer), the application of AI as tool to take action at commensurate speed when indicators of compromise rear their ugly binaries. And the irony? Why are we having this discussion now? Why are we treating AI as if it's new and exciting and something to get a career in (Adams, 2002)? Because we bandy the word cyber and pooh-pooh its origins. Cyber – from the Greek 'to steer' – as a way of modeling and then creating the systems of AI we now crave was trampled over years back by a sort of

[10] Field Marshall Helmuth Graf von Moltke

[11] "A robot may not harm humanity, or, by inaction, allow humanity to come to harm." (Asimov, I. (1985) *Robots and Empire*, Doubleday)

academic patricide. The influencers of the day seemed to ride roughshod over the cybernetics of Norbert Wiener. A language which naturally describes the communication in humans and machines we crave to benefit from in a state of cyber security. Some of us are waving this flag. Here at the University of Manchester we've even reawakened one of its organs - the interdisciplinary thought leadership group known as 'The Ratio Club'. But don't worry AI, we love your potential. Our role model – a member of the original club did too. Who was he? You may have heard of Alan Mathison Turing...

III. Conclusions

On structure

It ironic that the roots of The Ratio Club are themselves embedded in a discipline that take its nomenclature from the idea of 'steering' and so suggests that systems can be steered towards the desirable outcomes of at least one 'weltanschauung' of its actual or implied owner (Checkland, 1981) at the time. This irony lies in the challenge of bringing together brilliant minds for 'unfettered' discussion (Dresner and Williams, 2018) and yet giving them the topic that is to receive scrutiny before trying to contain the evening within the bounds of that topic. It is to be observed that few contemporary discussions around cyber security escape attempts to address AI. However, demarcation seems to be rarely made as to whether this is about the use of AI for controlling protective, cyber security measures, the use of AI for overcoming protective cyber security measures, or the need to AI to be duly protected against attacks on its integrity. This second outing of The Ratio Club was tested by a change of keynote speaker – who by contrast to the remit of the resilience theme – could not attend and a lecture on the state of the art of AI replaced the original intentions. This challenged the facilitators to restore the theme to the evening whilst not creating a blinkered and restricted forum for the evening.

Diverse disciplines + Academic freedoms + Target of evaluation = The Ratio Club

Figure 1: The Ratio Club Format

The net effect was a challenge for the notetakers as they deliberately made no effort to stifle the creative discussion and recorded the proceedings in the hope that their recordings could be carried forward to produce this report in the spirit of the original theme.

IV. Further reading

[1] Adams, D. N. (2002) *The Salmon of Doubt,* William Heinemann Ltd.,

[2] Checkland, P. (1981). *Systems Thinking, Systems Practice*. John Wiley and Sons.

[3] Dresner, D. G., and Williams, C. (2018), *Rational thinking (The revival of The Ratio Club)* with Colin Williams, CyberTalk, 10, October 2018

[4] Michie S, van Stralen MM, West R. The behaviour change wheel: a new method for characterising and designing behaviour change interventions. Implement Sci. 2011;6:42. Published 2011 Apr 23. doi:10.1186/1748-5908-6-42

[5] Taleb, N. N. (2013) *Antifragile: Things that Gain from Disorder*, Penguin

[6] Weiner, N. (1948) *Cybernetics: Or the Control and Communication in the Animal and the Machine*, MIT

V. Acknowledgements

Our thanks are extended to the Sociotechnical Security Group (StSG) of the National Cyber Security Centre (NCSC) for its generous support of The Ratio Club.

Figure 2: Some of The Ratio Club 2.0, December 2018

- Professor Colin Williams, University of Warwick
- Anna G., The National Cyber Security Centre
- Professor Chris Hankin, Imperial College London
- Doctor Keith Scott, De Montfort University
- University of Manchester
    - Professor Emma Barrett
    - Carmel Dickinson
    - Professor Chris Taylor
    - Matthew Thomas
    - Gwern Hywel
- Doctor Ben Bowles, University of London
- Doctor Robert Hercock, British Telecom
- Professor Leonard Smith, London School of Economics

And special thanks to

- Doctor Vassilis Galanos, University of Edinburgh
- Nigel Jones, Information Assurance Advisory Council (IAAC)
- Sir Dermot Turing

for their written responsa after the evening – *see the appendices* – as well as their participation on the night, and,

- Evie Hull, University of Manchester

for editing the final manuscript.

Appendix A. Sir Dermot Turing on the theme of resilience

I. What is resilience? and some other questions

Perhaps we should start, in the spirit of the Ratio Club, with Ross Ashby's *homeostat*, the purpose of which might be somewhat obscure. It's a self-stabilising device, somewhat akin to the Cerebellum in the human brain. It's self-normalising, or a guarantee of going back to the steady state : it is resilience in machine form.

Some of the characteristics of 'resilience' are thus evident: resilience is a response, not a status. It's a dynamic characteristic: you can't measure or create 'resilience' by looking at a snapshot (try painting go-faster stripes on a stationary car; it may look good, but it isn't speed). The snapshot is probably showing 'preparedness', which is not the same thing as resilience. To test resilience, you need to expose the system to a shock, or a change of some sort, and see how the system responds. It is a resilient system if it springs back to a steady state (not necessarily the starting state) without using excessive resources (time, electricity, people, money, whatever). Resource usage relates to severity of the shock. To attempt an equation:

$$\text{Resilience} = \frac{\Delta \text{ resource usage}}{\Delta \text{ system effectiveness}}$$

(though it's likely the equation is non-linear for certain types of system and certain types of shock).

As to whether a resilient system should return to the old state, it's not a necessary or even desirable outcome of a dynamic process. First of all, the shock may make it unsuitable to return to the old state: other systems may not function properly anymore (they're less resilient) and 'our' system needs to be open to change as well as resilient. Secondly, the data being used by 'our' system should be upgraded to include the fat-tail shock (cyber-attack, financial price move, earthquake, assassination, discovery of huge mineral deposit) as that will affect parameters the system uses for effectively delivering its output. (We should upgrade the bestiary to include the black swan.)

How do we try to assure resilience? Factors at work include distribution and redundancy; use of multiple micro components rather than one big centre; dispersal and dampening. These things all sound like network characteristics (see *Sidetrack into networks* below). See also the centralise, decentralised, and distributed network diagrams of Baran (1962)[12].

II. Sidetrack into networks

There is an immense and rather un-joined-up collection of literature on network design in various fields, for example, competition economics, military logistics, telecommunications, social network studies, insect biology, financial markets etc. This sizeable corpus shows us that topology is vital. Perhaps less suspected is that the topologies observed are dependent on the purpose for which the network exists; real-life topologies tend towards the scale-free ones where the impact of shocks – especially targeted shocks – can be devastating.

III. What can we learn from all this stuff?

(1) Resilience is typically about efficient transmission of data across a network when nodes fail.

(2) Although we may be just as concerned about static data integrity within a node – the theoretical analytics may be similar, if you read 'transmission' as 'retrieval', and the structural solutions may be similar too.

(3) Factors at work are topology, integrity, redundancy, and bandwidth.

Appendix B. Vassilis Galanos on the theme of resilience

What characterises resilience?

Resilience is currently characterised by the amazing hype surrounding it. I often joke by saying that if any academic paper's title contains the terms 'resilience/nt,' 'sustainability/ble,' 'inter/trans/crossdisciplinary,' and 'robust' in any order, about any subject (say, technologies, climate change), the more likely it is that such a paper will be accepted for publication. Resilience is a tricky (cryptic) term because of its double, relatively contradictory meaning; toughness through elasticity, or elasticity through toughness, or something in-between or somewhere around. Its Latin etymology, which one may mostly relate to words such as 'salient' does not really help. Given that the present survey is occasioned by the meeting of the novel instalment of the Ratio Club, and for the above reasons, I will reflect on 'resilience' by returning to the basics of Cybernetics. While there are some excellent papers aiming to relate current literature on resilience with cybernetic thinking (for example, Walker and Cooper 2011: Genealogies of resilience: from systems ecology to the political economy of crisis adaptation; Dijkistra 2007: Cybernetics and Resilience Engineering: Can Cybernetics and the Viable System Model Advance Resilience Engineering?) and to my knowledge, Prof Charles Raab (who has a long background in Cybernetics) from Edinburgh works extensively on the concept of resilience in governance and political science, to keep things simple, I will refer to a single source, Ross W. Ashby's *Design for a Brain* (the second, corrected 1954 edition) which, in my view, contains most principles, definitions, and concepts that help current researchers compare those with 'real-world' problems; all in-parenthesis

[12] https://www.rand.org/content/dam/rand/pubs/papers/2005/P2626.pdf.

numbers refer to pages from this edition[13]. While the questions posed here have nothing to do with the understanding of a brain's (or other 'thinking' machine's) functions, there lies exactly the charm of cybernetic thinking: the question 'how does the brain produce adaptive behaviour?' (1) can easily be adapted to 'how does a system produce resilient [sic] behaviour?'.

To get rid of the mysticism regarding resilience, it is useful now to provide the following working definition: resilience is characterised by the systematic study of homeostatic processes and adaptive behaviour. Push homeostasis to the extreme and the system will become static, entropic, probable, and an easy target as it does not change. Then one needs to apply adaptive behaviour for the survival of the system. Push adaptive behaviour to the extreme and the system will become fragile and susceptible to extraneous parameters (intentional or unintentional). Resilience appears to denote a golden ratio between homeostasis and adaptability; however, the response to the following question will help clarify the problem.

How do you measure resilience before it's tested by real world events?

It is quite likely that there can be no such 'in-advance' test, at least one that can be reliable against unprecedented events. Machine learning is good at tried experiences, not novel ones. Even for tried experiences, no machine learning system can protect from, say, extreme dangers such as airplane crashes, since no vast database exists with all contained variables leading to a crash, or at least, as vast as the databases containing variables of face recognition or 'did you mean' recommendation systems. First definition to keep in mind: 'A variable is a measurable quantity which at every instant has a definite numerical value' (14) and a 'system is any arbitrarily selected set of variables' (15). Later on, I will reflect on how selection of variables will play a crucial role for the measurement of resilience, but let us get a clearer view first on how such variables relate to the system's survival:

'Every species has a number of variables which are closely related to survival and which are closely linked dynamically so that marked changes in any one leads sooner or later to marked changes in others. […] These […] will be referred to as the essential variables of the animal' (41). It already becomes clear that animal behaviour, yet another systems behaviour, and more specifically animal survival can offer specific insights on the concept of resilience. As Ashby continues: 'We can now define 'survival' objectively and in terms of a field: it occurs when a line of behavior takes no essential variable outside given limits' (42). Therefore, adaptability (or homeostasis) becomes the self-organising principle of system bearing its own end in order to describe resilience: 'I propose the definition that a form of behavior is adaptive if it maintains the essential variables within physiological limits' and this becomes the system's goal. A question arises with regard to the selection of the essential variables.

Although measurability is proven to be arbitrary itself, it goes without saying that when it comes to data about variables, they have to be somehow measurable. A chicken-egg type of question arises: do we first need to find measurable data about the variables and see whether they are essential to our needs of describing the system's resilience? Or, do we need to first define the essential variables of the system that interests us and then see whether we can measure them somehow? Both approaches bear strengths and weaknesses. Several measurable quantities may not be that essential to the system's resilience (and we would be biased if we take them into account just because they are measurable; it would offer a distorted view, guided by rhetoric). On the other hand, we cannot define essential variables if we have not much experience about the system. It seems that a combination of intuition (expert committees, brainstorming, gamification for collecting and selecting variables) and trial-and-error processes (recent and non-so-recent history of science and technology, experimental tests) can be a good start point; although, followed by the humble acceptance of potential failure. But the selection of essential variables (and their separation from other variables) is not enough, as the system exists within a given environment.

'Given an organism, its environment is defined as those variables whose changes affect the organism, and those variables which are changed by the organism's behavior' (35); this mutual shaping of organism and environment is known as 'feedback' (36). While considering these definitions, one is asking: in our case, are the humans the environments of digital systems? Is it vice versa? Do the two form larger system with other environmental factors surrounding it? Ashby acknowledges these difficulties between environment and system: 'As the organism and its environment are to be treated as a single system, the dividing line between 'organism' and 'environment' becomes partly conceptual, and to that extent arbitrary. […] Once this flexibility of division is admitted, almost no bounds can be put to its application. […] Variables within the body may justifiably be regarded as the 'environment' of some other part' (39) – we can think of any problem occurring within a digital system and the way it affects 'external' humans. However, if we consider that '[a]n important feature of a system's stability (or instability) is that it is a property of the whole system and can be assigned to no part of it' (54), then the separating line becomes unnecessary and we need to look at Ashby's solution in terms of definition to cope with the problem. He simply states: 'Given a system, a variable not included in it will be described as a parameter' (72) – therefore, within a system, a factor which is external to the system's stability is a parameter. We can think of digital attacks as parameters; lack of electricity; connection faults; unethical use of the technology; and so on.

[13] Note by the editor/Dr Galanos: It is a happy thing that purely as a function of the order of receipt of contributions to this paper, that the preceding appendix, by Sir Dermot Turing, also refers to Ross W. Ashby's cybernetic thinking. Dr Galanos' text provides a descriptive explanation of the formula in Appendix A.

It becomes very difficult for experimenters (psychologists or physicists) to be constantly aware and distinguish between the parameters that they control and the variables they are observing; in our case, the confusion appears to have an extra layer as we explore the ways we can control already given variables and observe potential parameters-as-threats, since, in a given system, 'a change of stability can only be due to change of value of a parameter, and change of value of a parameter causes a change in stability' (79). Therefore, much work needs to be done to define as best as possible within given periods of time the variables and the parameters, that is, separate between wanted and unwanted behaviour (I am adding here a relatively moral value to the differentiation as a resilient, homeostatic system is the one which preserves its variables – hence we need to render them as 'wanted'; parameters tend to change the stability of the system, and to the extent that we wish to preserve its resilience, we might need to render them as 'unwanted'; however, these sentences can be easily contested if we think about control: parameters are more easily controlled than variables which are observed and as long as we are within the broader system in question what we might encounter, and the difficulty of the present question, lies in the fact that there might be an oscillation between what we treat as variable and as parameter. I will return to this later).

Is there a balance to be struck – and if so how – between the resilience of people in, or affected by, the system and the estate(s) or technology(ies) that comprises the rest of the system?

This is perhaps the most difficult question to answer from this survey. It is a question of feedback and we have briefly encountered it earlier. Feedback happens to occur within two interacting systems and so far, it has not been clear whether humans using digital systems are treated as a system themselves extraneous to the digital systems, or whether humans plus their systems are forming a single whole. Given that we have already examined the possibility of treating systems as different kinds of systems according to points of view, let us now examine few further principles on the junction of systems and this junction's relation to stability.

> '(a) Two systems may be joined so that they act and interact on one another to form a single system: to know that the two systems when separate when both stable is to know nothing about the stability of the system formed by their junction: it be stable or unstable.
>
> (b) Two systems, both unstable, may join to form a whole which is stable.
>
> (c) Two systems may form a stable whole if joined in one way, and may form an unstable whole if joined in another way.
>
> (d) In a stable system the effect of fixing a variable may be to render the remainder unstable' (55)

It is important that this question is placed here, as it becomes easier for me now to articulate the precise difficulty: a single resilience separated into sub-resiliences may cause unprecedented trouble; and the next question will qualify this claim.

Should we treat an attack on our digital lives like the harm of a physical attack? When we lose data or access to 'digital' are we bereaved? Do we suffer grief? Will we react to the next bleep from our devices with symptoms of PTSD?

Yes. To separate the feelings would mean that we treat digital lives as something extraneous to non-digital ones. But given that digital lives are a subgroup within the general group of non-digital lives, values, virtues, and vices such as dignity, anger, love, hate, deprivation, and so on, have every right to be considered equally in the sphere of the digital as well as the non-digital; most importantly, in the only realistic sphere which is the degrees between the two as there is no digital life without the non-digital one and there is increasingly smaller and smaller probability of non-digital lives being somehow related to digital ones. The question becomes how many and how much of these potential attacks and their effects are accessible to measurement and observation. Again, we are in need of defining what is the irritation, when it occurs, why, and then apply trial-and-error processes until the irritation stops. The generation(s) considered to have grown up as digital natives are now adults. This allows relatively easy access to empirical investigations (interviews, surveys, focus groups) of younger age-groups and their experiences of such technologies. This generation, according to scholars from the late 20th century, should be the one who grew up without the dichotomies of the analogue/digital, online/offline, here/there, and so on, and given that work has been done on the measurement (or at least fine description) of such harms among people who 'lived' the 'transition' (and of course pre-digital societies), it would be useful to compare these younger generations tendencies towards such feelings.

It seems that while one thinks of a 'cyberattack' to a 'cybersystem' with regard to the competition between variables and parameters, it is really a question of non-essential variables within the system that might cause such harms: for example, heated debates over the web, enough to ruin one's day, propaganda (fake news), the internet's fantastic yet very problematic tendency to remember everything in a very mechanical way used by very human manners, revenge pornography, are only but few examples of attacks within the system by the system. It appears as if certain essential variables (the ability to store information and learn, data processing, connectedness) transform into critical parameters that the experimenter has no control upon (not to mention the difficulty in defining the 'experimenter' as we will see later). Given that much work needs to be done with respect to commonalities and differences in the perception of harm, not between digital and non-digital lives, but between people unaccustomed to the digital, accustomed through transition, and without the need or choice of being accustomed, the changes within the digital system which is meant to be resilient much come as a very long series of very small changes to allow for corrections, alternative pathways, resistance within the (presumably democratic) system, and scholastic study of the available lessons. I will return to these small changes in the last question.

DOES 'CYBER RESILIENCE' EXIST AS A THING OR – TAKING A SYSTEMIC APPROACH – DOES IT BECOME THE WAY OF ARTICULATING RESILIENCE?

Once again, a simple Google Scholar search returns many relevant hits to answer the first part of the question; whether the usage of the prefix 'cyber-' is correct or not, this is left to historians of citations and etymologists. Regarding the second part of the question, I absolutely agree; and yet I think that with the present text I manage to show very briefly how occasioned by what is called 'cyber-resilience,' we may take up the opportunity and revisit cybernetic principles in order to articulate resilience and, why not, contribute towards the resilience of digital systems as well.

CAN YOU ADJUST A STATE OF RESILIENCE TO ENCOURAGE MORE FAVOURABLE OUTCOMES? HOW MIGHT WE COMPENSATE FOR EXPECTATIONS OF RECOVERY TO THE OLD STATE WHEN WE'LL BE DEALING WITH A SYSTEM THAT HAS COME THROUGH A CHANGE IN ITSELF?

Favourability is quite immeasurable. Perhaps, the first question should be placed in a different way: is resilience always favourable? What if the current state of resilience is not that favourable and the variables rendering a system homeostatic need to be changed? Certain variables within the system (perhaps not the mechanic parts as much as humans) might need to be trained. A system is corrected by external observers according to feedback signals and hence, the typical way to provide feedback back to a trained system is through positive and negative signals: 'All training involves some use of "punishment" or "reward"' (112) – Reward 'usually involves the supplying of some substance (e.g. food) or condition (e.g. escape) whose absence would act as "punishment"' (113) – the problem here, is that Ashby refers to experimenters external to the studied systems, being able to exchange feedback signals within each other, whereas it is difficult for humans (designers, users, legislators, journalists, etc.) to use such controlled methods of reward as parts of a broader system including members of their own species as feedback might be perceived in a distorted manner (or not perceived at all). Given the difficulties we have already encountered about the inability to clearly separate between variables and parameters, yet another problem occurs as to the identity of 'we': who are we who have the right to manipulate this system which contains us? We might see us then as some sort of internal variables, aware of the system we aim to make resilient and homeostatic working towards its refinement. Reward/punishment in that sense, would mean something similar to the very small trial-and-error changes in sensorimotor perception when somebody tries to balance using a bicycle, adjusts one's attire in order to cause less discomfort, calculates the best framing for a photograph, and so on. Such processes are already taking form in the digital realms as minor regulations which co-ordinate collective behaviour: restrictions about vernacular speech or nudity on various social media are good recent examples (not necessarily successful ones in terms of algorithm training); previous debates about the virtues of netiquette pointed towards a similar direction, and contemporary future-orientated philosophical debates about the ethics of (imagined versions of what some people call) artificial intelligence are the system's homeostatic behaviour allowing it to protect its resilience by its own intrinsic threatening parameters/variables. The advantage of treating such processes using the cybernetic nomenclature is that we are allowed, at least for the argument's sake, to define temporary variables and parameters and work in order to control them by asking 'how much should we allow connectedness between this and this parts of the system?', 'is such a behaviour harmful?', 'did a relatively harmful behaviour lead to a more resilient version of the system?', 'which variables were responsible for this?' and so on.

Questioning what is the desired outcome, and applying such training methods to achieve it, another sub-question occurs: is there one or many such outcomes? And more crucially, if a desired adaptation state is reached, could it be the case that a more desired adaptation will be needed in the future, and if this is the case, and if the system reaches it by training, will it be able to return in its previous state of adaptation if needed, or will previous adaptations be destroyed? This question is important when it comes to education of new technologies, as well as digital security when similar technologies are considered in different regions of the world, with different infrastructures and cultures. What may seem appropriate as a technical variable (or parameter!) in terms of software, hardware, internet protocol and so on, in one place in 1995 for the given local system's resilience, may be appropriate for another local system in 2010; or vice versa. Revolutionary acts provide with good examples: Occupy Wall Street protesters made good use of human voice as a means of communication rejecting everything digital, whereas most current file sharing activists make use of relatively 'outdated' technologies. The overall system's ability to return to previous states of adaptation, in other words, knowing one's history, allows for easier adaptation to novel circumstances and control over more possible variables/parameters within the system and recognition of potential harm.

To describe resilience (adaptation, stability, homeostasis, and everything in between) means to find measurable variables/parameters and actually get to measure and control them. Generally speaking, the current state of the human plus digital technologies system can be described as relatively resilient as some apparently trivial tasks are conducted in a habitual manner with relatively little difficulty; however, if one thinks about how to compute the average correct choices among the many available ones online (in communicating, 'liking', marketing, bargaining, creating, and so on), the system appears to function well even without as being thoroughly able to describe all these processes in fine detail. This problem (or fact) is often stressed by roboticists who verify the difference between a human and a machine learning device: humans tend to need very few examples in order to learn something, whereas machine learning devices (or algorithms) rely on vast databases. The fact that some processes cannot be fully quantified and measurable renders them an open-ended task, and perhaps the quest for resilience is as complex (even futile) as the search for intelligence. Humberto Maturana and Gloria D. Guiloff's short paper The Quest for the Intelligence of Intelligence should perhaps be revisited as a quest for the resilience of resilience, as part of the novel Ratio Club's future endeavours. More broadly, it might be interesting to trace the social history of the concept of resilience. What socio-historical circumstances brought the concept of resilience into account? Were there commonalities or differences, say, in the 1970s when the concept was mainly introduced and the last decade when it received much recent attention?

Note: a very important concept advanced in Ashby's book (and work in general) is the notion of step-functions as a key to understand systems' mechanics. For reasons of space, and due to the very experimental nature of the notion, I have not elaborated on them. Similarly, the same applies for the notions of ultrastable and multistable systems (while they are very interesting, as theories they depend a lot upon the acceptance of step-functions; thus are relatively unreliable in contrast to the rest of the cybernetic notions examined here).

I should acknowledge the fact that much of the thinking presented here has been heavily shaped by the discussions of our Cybernetics Today group conversations in Edinburgh last year, although the opinions expressed here are clearly my own. I know that colleagues participating in the group might shed further light through their expertise in areas I do not feel so comfortable speaking of, such as the Viable Systems Model or Niklas Luhmann's systems approach.

FURTHER READING

Ashby, W. R. (1954) *Design for a Brain* 2nd ed. London: Chapman and Hall.

Beer, S. (1984) *The viable system model: Its provenance, development, methodology and pathology* Journal of the operational research society, 35(1), 7-25.

Dijkstra, A. (2007, July) *Cybernetics and resilience engineering: Can cybernetics and the viable system model advance resilience engineering?* Proceedings of the Resilience Engineering Workshop; 25-27 June; 2007; Vadstena; Sweden (No. 023, pp. 23-29). Linköping University Electronic Press.

Luhmann, N. (1986) *The autopoiesis of social systems* Sociocybernetic paradoxes, 6(2), 172-192.

Raab, C. D., Jones, R., and Székely, I. (2015) *Surveillance and resilience in theory and practice* Media and Communication, 3(2).

Walker, J., and Cooper, M. (2011) *Genealogies of resilience: From systems ecology to the political economy of crisis adaptation* Security dialogue, 42(2), 143-160.

Appendix C. Nigel Jones on the theme of resilience

WHAT CHARACTERISES RESILIENCE?

As a concept, resilience offers a view of outcomes that security does not. Security for security's sake is meaningless. 'Resilience' communicates a determination to continue to operate, even whilst under stress. A former US Department of Defense (DoD) Chief Information Officer (CIO) once told a Cyber Security Knowledge Transfer Network (KTN) event in Paris in 2009, that his role was to ensure that the DoD could 'fight through degradation'. This might be akin to the notion do 'robustness'.

At its most basic, one might view resilience as the ability to bounce back from a crisis. This is in line with the word's Latin origins.

We therefore have two views so far – one of robustness and the other of bouncing back. However, today, neither this nor the concept of robust 'continuity' is sufficient as a stretch target for organisations. These concepts are hampered by communicating a stable and unchanging notion of the organisation's function, structure and future. For organisations to succeed, stay relevant, and resilient in changing risks, learning and adapting must form part of the way an organisation (unit) works, and must be reflected in how they learn from stress tests.

I am therefore of the view that resilience should be promoted as a combination of robustness, bounce-back and adaptability and should be defined in some combination of these terms.

HOW DO YOU MEASURE RESILIENCE BEFORE IT'S TESTED BY REAL WORLD EVENTS?

The answer to this is in my view lies in a combination of modelling, simulation and exercise. It is essential that we develop tools that help planners, leaders and procurement staff ask 'what if' questions. I do not believe we are at the stage yet of being able to quantify how much is enough in terms of investment in resilience in all contexts. We should be able however to model an environment and to apply scenarios to it, in order to illuminate our judgement...

IS THERE A BALANCE TO BE STRUCK – AND IF SO HOW – BETWEEN THE RESILIENCE OF PEOPLE IN, OR AFFECTED BY, THE SYSTEM AND THE ESTATE(S) OR TECHNOLOGY(IES) THAT COMPRISES THE REST OF THE SYSTEM?

As we are in the realm of judgment, this reflects the reality of trade-offs – and the skill of the analyst. If we model a business environment or organisation as if it were a system, we can look at our assumptions and judgements reflecting the role of people, technology, and organisations. The balance to be struck should be based on leadership's ability to assess desired outcomes, interventions and the culture in which proposals for change sit. However, the notion of a 'struck' balance may give the impression that it's a one-time solution. Rather the balance to be struck is more like a balancing board on a cylinder, requiring ongoing adjustments, even as the length of the board and diameter of the cylinder change (and wind strength and direction…and someone is trying to knock you off (pardon the expression).

The general issue of how individuals relate to communities/ organisations and the national level is one we are picking up in IAAC and was reflected in our Symposium September 2018. How these levels relate to one another is of acute interest to me and IAAC.

Should we treat an attack on our digital lives like the harm of a physical attack? When we lose data or access to 'digital' are we bereaved? Do we suffer grief? Will we react to the next bleep from our devices with symptoms of PTSD?

I feel that analogy of grief and trauma should perhaps be left for instances of grief and trauma as we might commonly understand them. However, that does not mean that we cannot think about harm or indeed anxiety associated with our digital lives, which for some may be very traumatic indeed. I spoke to the man who had experienced this[14] and it was (is) indeed a traumatic experience.

However, at an organisational level, we can do much more to help people think about their role in resisting crime and it affects. Work begins before it happens and not just afterwards. In some ways we can inoculate people for living in the digital world, but I don't know if much has been happening in this regard. I don't think it has. Rather digital resilience in terms of school aged people, has I think tended to focus on how to stay safe rather than on how to respond emotionally when things go wrong.

On the whole, on the issue of harm, a scale of harm is more appropriate than thinking about it in analogous terms of absolute trauma. This thinking may however help define the scale.

Does 'cyber resilience' exist as a thing or – taking a systemic approach – does it become the way of articulating resilience?

Yes, the method can be a thing or entity in its own right – *and* the 'thing'/way in which it is communicated. In current organisational structure this is constructed in different ways. For example, one major bank has a head of resilience and a head of cyber security. The head of cyber security does the cyber part of resilience and the two have regular meetings – and sit close to one another. However, it is increasingly hard to disentangle cyber from other security functions such as physical and personnel security. So a whole organisational view of the cyber enabled company is also appropriate – aligned with how the business generates value. The company should be viewed as if it were a system, even if it isn't. Much more work on resilience security and value need to be done – and we are doing that in IAAC too.

Can you adjust a state of resilience to encourage more favourable outcomes? How might we compensate for expectations of recovery to the old state when we'll be dealing with a system that has come through a change in itself?

See above regarding modelling, balancing and adaptability.

[14] https://metro.co.uk/2018/03/04/married-man-discovers-photos-used-scam-women-dating-sites-7359889/