

Human-centered strategies for cyber-physical systems security

E.N. Ceesay^{*}, K. Myers² and P.A. Watters³

¹Johns Hopkins University, Baltimore MD, USA

²Dalferes Enterprises, Alexandria VA, USA

³La Trobe University, Melbourne VIC, AUSTRALIA

Abstract

Human error contributes to information system losses. Exposure to significant risk will continue and is not effectively addressed with conventional training. Broader strategy that addresses the social system is recommended. Such strategies have been successfully developed in industrial settings to deal with workplace hazards that are functionally similar to cyber loss. Four of these strategies are reviewed and found to be relevant to the needs of the IT-enabled organization in mitigating cyber security risks. These strategies are not consistent with each other or uniformly applicable, however, and would need to be adapted to contemporary knowledge work settings and used cautiously. Long-term institutionalization and development of organizational practices pose further challenges. While a holistic, sociotechnical systems (STS) approach to cyber security requires significant effort, IT-enabled organizations, as industrial organizations before them, will realize the effort is justified.

Keywords: Sociotechnical systems; human factors; cyber security.

Received on 20 April 2018, accepted on 12 April 2018, published on 15 May 2018

Copyright © 2018 Ebrima Ceesay *et al.*, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi: 10.4108/eai.15-5-2018.154773

^{*}Corresponding author. Email:enceesay@ucdavis.edu

1. Introduction

Insecure Information Technology (IT) systems lead to large losses, and human action is often implicated along the path leading to a loss. As Greg Shannon, Assistant Director for Cybersecurity Strategy in the Obama Administration, Office of Science and Technology Policy (OSTP) has said, "...it will always be true that malicious insiders and human error can create problems..." [43]. The possibilities for human-implicated problems are rife and can range from an operator error that overlooks an open port, to an insider who steals, to a user who inadvertently hits a malicious link, to a fraudster who spoofs filtering procedures. Yet, for an IT system to be of any value in most organizations, it needs to be open and flexible. This requirement sets a limit on technological remedies and shifts attention back to human action. An approach is needed to reduce loss by increasing safe

behavior. Where does one begin? Heavy industry has experienced essentially the same phenomena over decades and has developed a refined understanding of how to deal with human errors that lead to loss. The cyber world does not need to recapitulate the painful and protracted learning process already conducted in industrial settings. We suggest that IT-enabled organizations adapt what industry has learned to the unique conditions of cyber security.

Industry's effort is generically described as workplace safety. Computer technicians and operators may scoff that workplace safety is far from their concern. Early industrial technicians agreed. Machines were fragile, and the technician's only concern was to make the machines work and keep them working. If a human operator was implicated in a mechanical breakdown, it was a nuisance best solved by training the operator to perform exactly as the machine required. Better yet, when technical advances allowed, eliminate the role of the operator altogether. This was the

default strategy in industrial operations a generation ago, and in many IT shops it is the default ‘best practice’ today. Yet industry has moved on. The IT security field should move on as well, by recognizing the parallel between workplace safety and cyber security and learning from it. The American Society for Industrial Security (ASIS) makes the point well: “Because of its many similarities between preventing accidents and preventing security breaches, not only will security practitioners improve their understanding of myriad safety issues, but they will also be able to apply many of these concepts to their security duties” [1].

The authors argue the following: First, it should be recognized that the IT system is a component within the sociotechnical system, and specifically that people are ‘in the loop’ in most cyber losses. Doing something about cyber loss requires addressing the social system directly. The argument takes a brief excursion to examine how approaches to human error in the industrial setting have evolved and reviews what is different about the cyber setting and why human errors in that setting are particularly challenging. The paper then critiques popular approaches to the human side of IT security that do not take advantage of sociotechnical insights. Finally, the paper sketches elements of a sociotechnical strategy and points to some options that flow from that strategy

2. Viewing the IT-Enabled Organization as a Sociotechnical System

The technical system and the human system need to be considered separately for how each can succeed. This was the initial insight of sociotechnical systems theory. The social system at work needs its own coherence, to afford meaning and motivation for workers. The technical system will also have its own inescapable requirements that constrain the human system. Though today, there are fewer such constraints, due to the flexibility of digital technology. Sociotechnical systems theory would have us consider each system for their independent needs to remain coherent, and to then jointly optimize the two. This will often appear to compromise the technical system. Yet the aim is to make the overall, integrated system more effective than would be the case if the performance of the technical system were optimized independently.

The object of our concern is the IT-enabled sociotechnical system. This system generates many functions, among them the following basic cyber security functions:

- Maintain protections and security capabilities
- Monitor for anomalies, threats, and losses and describe sources, sequences, and effects
- Coordinate security efforts with other functions and organizations
- Respond to threats and losses

Technical systems have a major role in supporting these functions, of course, but they will be incomplete; even near-autonomous technical systems require human supervision. Typically, there is a two-part social system that is

responsible for these security functions. The normal production staff accepts a security role as a secondary responsibility, and the IT staff takes security roles as a primary function among many. In addition, there are specialists and outside services or advisers.

Since cyber security functions are secondary for the production staff, they may discount its importance, an attitude that may be reinforced by compensation schemes or performance feedback. Staff members may also assume that the IT staff has the problem handled. Even if the actions, thinking, and habits of the production staff are not properly aligned with what the technical system requires, one often assumes that this can be fixed by increased conformity to rules. But a social system has its own ways and constraints. Simple orders to conform, even if they are followed, may damage the social system. Orders to conform may contradict what a person and group understands to be a correct or fair behavior and can even cause resentment, backlash, or disregard for the rules, counter to what was intended. Often a solution requires a balance or trade off, wrapped into an overall joint optimization where each system works well independently as well as together.

3. Viewing the IT-Enabled Organization as a Sociotechnical System

The source of loss that we are addressing is human error mediated by the IT system. The standard typology of human error follows [34]:

- *Execution failures*
 - Slip: An action contrary to intention.
 - Lapse: Omissions of expected behavior, or acting with suspended or mixed intention.
- *Planning failures*
 - Rule-based mistake: Rules that are followed may be incorrect or misapplied.
 - Knowledge-based mistake: Acting based on incorrect assumptions or knowledge.

Across many industrial settings, research has found that 70% of execution errors are detected and corrected, while 50% or less of rule-based and knowledge bases errors are detected and corrected [11]. These are sobering figures. Many errors are of little consequence, but in IT systems, some of these errors can be quite damaging and unrecoverable. The social system and the technical system both have to be shaped to reduce errors.

Safety engineers recognize that errors continue after technological solutions are applied. In the past, this drove the development of human factors engineering, whereby the characteristics of human operators were more fully incorporated in technological design. In recent years, however, the analysis of complex system failures has led to additional focus on underlying factors that influence how people think and behave, especially in more complex work situations that involve judgment and incorporate roles beyond direct operation of machines, such as administrators,

inspectors, and programmers. Additional shaping factors include policies, organizations, and work cultures. Before exploring how these newer strategies apply, however, it is necessary to consider what is different about the IT-enabled organization, compared to traditional industrial organizations.

4. How the IT Setting Differs from the Industrial Setting

Five features of the IT-enabled sociotechnical system make risk reduction especially challenging. First, adversaries of the organization are actively attempting to induce human error through deception. This is different from, say, a factory where workers are not commonly tempted into error by saboteurs. Authentic-looking phishing messages can easily induce a slip that releases a virus, or a worker may commit a lapse by answering seemingly legitimate queries with personal information [23]. The scope of deception is so large that the Federal Bureau of Investigation (FBI) now says that cybercrime has become a top law enforcement activity [15]. The organization needs to take into account not only friendly cognition and behavior, but the cognition and behavior of adversaries (including the insider threat). These adversaries are often determined, well-funded and highly skilled [27].

Second, the environment is open. In order to be at all helpful in business, the IT system needs to support communication and information exchange broadly. This sets a limit on technical security restrictions, and allows malware infections to spread easily [26].

Third, the environment is both complex and tightly coupled. Perrow has argued that high complexity and high coupling, as exists in many IT-enabled organizations, make accidents inevitable or normal [32]. There is disagreement on this point, but even critics agree that protections in such systems often add to complexity and create additional layers of error [24]. Once errors occur in such systems, the effects propagate and the consequences are often difficult to control. Complexity also makes dealing with security incidents a challenge: a single networked application, such as Skype, can utilize numerous protocols and contain a mixture of plaintext and encrypted traffic, making forensic analysis challenging [4].

Fourth, the IT system is opaque. While an error-inducing condition will often seem simple once it is identified, the inner workings of the IT system are relatively invisible. Those who wrote the code may understand more, but they are long gone. Operators often just trust the system and are therefore liable to fall into hidden traps. A final feature is power and speed. A simple keystroke can cause enormous damage, and the consequences can roll out in milliseconds. Introducing data flow controls remains a challenge [45].

Risks can be enhanced by an operator's negative attitudes, even if he is not an outright adversary. The user may be ignorant of rules and procedures, pig-headed about

doing it the way he wants to rather than the way it should be done, hostile to instruction or safe practices, or simply heedless of procedure and preferring to guess. Such users exist, and we must deal with them, but there are many others who, with the best intentions remain distracted, forgetful, or unclear, and they can make just as many errors. Finally, it is not just what is wrong with humans that lead to errors, but what is right about them. Humans are social, they trust, and they are curious. The IT system is like a free candy store for the garrulous and curious. This situation calls for risk mitigation directed to the social system. The immediate response has been awareness training, which we will now examine.

5. Critique of Current Approaches to Human Cyber Security

Providing information on cyber risks makes sense to the logically minded: tell people what the problem is, and they will avoid it. Yet the results from informational training are unimpressive.

The U.S. federal government and defense contractors provide cyber security training, particularly phishing awareness to improve end-users' ability to recognize the signs of cyber attacks. In a recent survey conducted by ISACA and RSA Conference, 87% of respondents reported having a security awareness program in place; of those, 72% believe it to be effective [20]. However, the survey results reported that enterprises that do not have awareness training are doing 12% better than those that have training programs. Attackers are more frequently penetrating enterprise security among enterprises that have an awareness program in effect. These data show that awareness training is unlikely to be sufficient to ensure safe behavior, and it may even be counterproductive.

Many observations have been offered to explain such weakness in training:

- The learning technique is flawed. Greitzer argues that awareness training is often conducted using linear training paradigms that place the learner in a rigid, passive training environment [16].
- Incentives are misaligned. Herley argues that users' rejection of the security advice they receive is entirely rational from an economic perspective [17].
- The technical system is excessively lax. Schneier argues that by building systems that are vulnerable to the worst case raises risks for the average case. It would be better if we designed systems that conform to their user's security habits, rather than forcing them to learn new habits. [Schneier, 2013].

Awareness training can work in relatively simple and unchanging systems where there is immediate feedback and clear incentives. In CyberCIEGE, for instance, the learner is placed in a meaningful context where decisions have clear consequences that strengthens learning experience and thus help improve the potential for internalization of the acquired knowledge. In addition, use of personally meaningful

projects will enhance learning [6, 31].

Knowledge can also be viewed as schemas representing relationships among facts and concepts; knowledge structures contain schemas that may vary in their degree of automaticity [22, 49]. Schemas allow many elements of knowledge to be treated as a single element in working memory [2, 29], which reduces demands on working memory compared to controlled, conscious processing that requires higher cognitive loads [3, 40, 41]. This redirection from items to schemas is exactly what has occurred in industrial sociotechnical approaches to safety, often coupled with a more holistic approach in more meaningful circumstances. In the next section we review and compare some of the industrial strategies for reducing human error and loss.

6. A Sociotechnical Strategy for Cyber Security

A roundtable on cyber security at the National Academy of Science (NAS) emphasized a cautious approach, concluding that:

- Definitions of the problem need further exploration.
- Cyber security has unique characteristics, requiring rethinking of advice developed in industrial settings.
- Much of the available advice is relevant, but it is also inconsistent and sometimes contradictory. [30].

Reason, speaking from the viewpoint of industrial safety, makes a similar point. He warns against risk management consultants promising simple solutions and providing overconfident answers. He writes that, "Safety is a goal that has to be constantly striven for rather than achieved. Safety is not a state of grace but a guerilla war..." [35]. This caution is embedded in his "Swiss cheese model" for developing safety. He uses this imagery to communicate that there is no straightforward course toward complete protection, that holes are everywhere, caused by latent threats and errors such as company pressure, tight schedules, awkward rules that are hard to follow, long working hours, inadequate rest, lack of on-going training, lack of safety awareness practices, impractical policy, etc. By "latent" he means the holes are built into the system. They are not always the cause of accidents but can contribute when coupled with "active failures" such as human errors. Defenses can reduce latent threats but negligence can seep into this process as well. What he recommends is a safety culture, which is much more than organization charts or process diagrams. It is an organizational strategy that guides everyone's thoughts and actions as they perform functions, involving philosophy, policy and enforcing procedures, and habituated practices to guard against latent and active failures. Such a strategy aims to create a successful social system effectively matched to a technical system. Effective matching will sometimes require accommodations in either system to meet requirements of the other.

Beginning with the whole organization in this way (rather than with the particulars of error events) may seem too roundabout, but a sociotechnical approach sees it differently. Everything is eventually addressed, but starting with strategy is important because it provides the consistency and reinforcement for whatever changes are made at the individual level, changes that might not persist without explicit, broader principles and organizational support [13].

7. Organization-level Strategies for Error Reduction

Selecting from recent literature [9, 14, 36, 38, 39, 44, 50], four organizational strategies for risk mitigation stand out as relevant for cyber security. We summarize the advantages of each, then consider their interactions and how they may be reconciled. These are: high reliability, safety culture, sensemaking, and anti-fragile.

7.1. High Reliability

There are work situations where complex tasks must be performed with great precision and where mistakes can be catastrophic. Such tasks can sometimes be delegated to highly reliably technical systems, but sometimes this is not possible. The flight deck of an aircraft carrier is one such situation where delegation is not possible and a human crew is necessary. Crews in this and similar situations have been studied for what makes them successful, as reported in the literature on the high-reliability organization (HRO) [14, 18].

HRO training is best applied where there are well-understood procedures that must be repeated without variation or error. The emphasis in HRO is on controlling the worker's environment to eliminate distractions and to have tasks that fall within a narrow band of variability. The work system and environment are well-bounded, controllable, and perfectible. The group develops a practice of observing every deviation, encouraging each other to achieve exact standards, and isolating themselves in order to achieve reliable performance.

There is some danger in going down this path, since the HRO is not a perfect fit for the demands of cyber security, and it can even inadvertently reduce capabilities that are needed. There are clearly some security tasks that can be treated in this manner. A strict routine for testing and installing updates and patches and performing analytic sweeps would qualify. But some functions involve inherently non-routine responsibilities that require a different mindset and a different kind of observation, testing, and learning. This is a valuable approach for some tasks, and perhaps for some selected staff, but not for global application in an IT-enabled organization that must deal with a turbulent, rapidly changing environment.

7.2. Safety Culture

Cyber security work can be messy with a lot of activity and data that may or may not be relevant. This situation is similar to industrial plants where there are many small and variable activities that may or may not interact or be of concern. For this situation, DuPont pioneered the development of a culture of safety along the lines recommended by Reason [34]. Practitioners developed a superior safety record, not just by finding ways to avoid disasters, but by relentless identification of anomalous conditions. The workers do not restrict their focus to levers out of place or equipment out of date, for which there are standards. They notice conditions that might be a precursor to trouble, even if there is no applicable standard. A simple water spill becomes reportable, and workers are rewarded for identifying and dealing with such conditions, regardless of their importance. This ethic is somewhat similar to Toyota's where any worker on the assembly line may stop the line if something appears out of place, even if it is not obviously an important matter [12]. The workers are not punished, but rather, they are recognized for taking responsibility. Chemical process safety inspectors are keenly aware of the value of a safety culture. One chemical process safety inspector explained that, in his experience, a very telling indicator of whether there is trouble in a plant is whether the floors are swept [personal communication]. He was not able to write up dusty floors as a violation of standards, and he also never told his clients that he was looking for such indicators, but he nevertheless used these observations as evidence of awareness and an attitude of care that is conducive to safety.

A cyber organization could benefit from developing a version of safety culture. The equivalent of a dirty floor might be a backlog of documentation or phone messages not returned. Is the condition a potential starting point for trouble? Can we change things (technology or attitudes) to avoid such conditions? The organization needs to develop an ethic whereby workers are encouraged to be mindful in this way, and to speak up. It becomes a skill and a point of pride. The organization, including all production staff as well as IT staff develop and celebrate a living culture of awareness and initiative.

7.3. Sensemaking

Organizations need an ability to frame novel cyber security situations provisionally, act experimentally, and learn quickly from interactions. The rulebook is not entirely ignored, but in emergency situations there may be no rules that immediately fit because the situation is unique. Rules do not always apply. Keen observation, framing, and judgment are called upon, and this collection of skills has been named "sensemaking" [21, 50]. Sensemaking in ambiguous, complex, and pressure-filled situations is difficult to teach because it is not a procedure but requires a rare cognitive capability that only comes from practice. Yet many such

situations are rare, thus learning from experience may be insufficient. Simulation helps, but not if it amounts to drilling for the correct answers. What is needed is not the right answer so much as the right thinking to arrive at actions that are likely to be better.

There are many complex technical systems that get into trouble where sensemaking is required. Two famous cases of sensemaking failure in complex sociotechnical systems are the Three-Mile Island nuclear plant accident and the accidental shooting of the Iranian Airbus [8, 33]. At Three-Mile Island, the operators took action that made the situation worse. This occurred because they were using an incorrect mental model of how the plant worked, an implicit model that had never been identified or tested. Feedback from actions failed to invalidate or improve the faulty mental model. This all occurred as the warning horns and flashing lights were creating irrelevant signals as well as distractions. Because many cyber security situations are ambiguous and require fast action without a proven script, a sensemaking strategy that prepares staff to deal with surprises and uncertainty would appear to be wise.

7.4. Anti-fragile

For the mechanical portions of our organization, we seek efficiencies and less risk. Continuous effort in this regard is needed because machines are fragile and they only get worse with use or as environments change and render them unfit. But the human portion of sociotechnical systems does not really work that way. At the gym, people become stronger after experiencing stress. For that reason, people are not fragile, but rather anti-fragile, in Taleb's term [44]. Taleb suggests if we do not allow the stressors that make anti-fragile elements stronger, we make it more likely that a big stress will lead to loss. A strategy of anti-fragility, applied to the social system, protects the ability to learn and improve behavior based on environment stressors. Workers should be allowed to take risks and fail to allow learning. Learning from failure in this way does not necessarily mean that the organization will be able to identify and survive a truly catastrophic event, but it will be better suited to a changing environment. An organization managed according to a machine metaphor, on the other hand, is focused on lean efficiency. Mistakes and failures may be punished and covered up, creating a fragile organization that is wary of change, and eventually exposed to more risk, not less.

This anti-fragile strategy is closely related to the literature on the learning organization [47], and resilience [10, 37] which share the realization that, in complex environments that keep changing, people must keep moving and learn from errors, not only to keep up, but to gain speed and even get ahead, and in this case to get ahead of cyber adversaries.

7.5. Synthesizing Strategies

There are situations where one strategy will be highly appropriate and the others will have much less to contribute. As we mentioned, high reliability is a good strategy for tasks that can be safely reduced to strict procedures, but that rarely covers all that an IT-enabled organization needs to accomplish in its cyber security functions. This suggests that the answer is to not employ any ‘pure’ strategy, but to create a mix. This is not a simple matter because the strategies, at least when considered in their general thrust or pure form, can actually work against each other.

The strongest mutual contradiction is between the high-reliability and anti-fragile strategies. Anti-fragile might be accommodated in a high-reliability organization by creating separate space for safe failure, such as by simulating attacks and failure in a mirrored system, or by keeping a safe fail-over system while learning from failures in the main system [48]. There are additional tensions marked in yellow, but these are abstract considerations and would need to be worked out in detail within specific settings. The main lesson to take is that the human contribution to security cannot be perfected. One can try to get as close as possible to perfection by treating the humans as if they were machines, but this is typically shortsighted for two reasons. The unique capacity of humans to reframe and learn is not being used, and the environment keeps changing such that the target state changes and becomes uncertain. The fuller version of safety culture tends to put up less resistance to all the other perspectives and might be the best starting point. Recalling Reason’s admonition, the safety culture strategy accepts that there is no final or simple conclusion, nor one without compromise and uncertainty

8. Implementing a Holistic Cyber Loss Mitigation Strategy

A sociotechnical strategy provides direction on further programmatic aspects of cyber security. Assuming that one incorporates the concept of safety culture within the sociotechnical strategy, actions will need to be taken to change daily habits and attitudes, and in ways that will be socially validated and reinforced [25]. We will also assume that cultivation of sensemaking is incorporated in the strategy, as a bulwark against surprising and unprecedented events that might occur. These two aspects of the strategy, while not entirely compatible, can be pursued simultaneously. We will not outline a complete program but rather illustrate how such a mixed strategy may be played out in efforts to both train and to institutionalize social aspects of cyber security.

In addition to the general weaknesses of awareness training that we have mentioned above, we can add here that awareness training does not even aspire to changing habits and attitudes, nor does it have much to do with sensemaking, in terms of recognizing and responding to

unprecedented events. Something quite different is required. The individual needs to be engaged as a responsible manager of his own learning. The learner’s attention as well as motivation becomes focused on a practice, beyond assent to a logical argument. This sets us on a path quite different from selecting an approved curriculum and delivering it through an approved instructor.

A good approach is to simulate risky situations and responses. The simulation can occur at several levels, presenting situations during the course of work that challenge personnel. These situations will of course not put the individual or organization at real risk, but the challenges are in a natural context and are similar to real threats. A wide range of simulated threats can be generated such as phishing messages, suspicious insider activity, and a variety of anomalous monitoring results. The point here is to present a wide variety of signs that, to be recognized, require both habituated sensitivity and imagination.

The realism of the challenge is important, but in the normal course of events it is actually difficult to learn anything from such challenges because feedback is sparse and ambiguous. For example, if a worker reports spam, there is no word back from technicians whether they adjusted the spam filter, or whether the spam had harmful coding; and if the worker does not know whether reporting had any positive effect, he might not bother the next time. The answer here is to use psychological principles to make feedback effective in a way that accelerates learning. Changes in even stubborn negative habits have been achieved by “gamifying” the learner’s experience [28]. To gamify means to overlay a minimum set of structures that are proven to make games very engaging and fun. McGonigal uses several structures that will be familiar to gamers such as “quests, bad guys, powerups, epic wins” and so forth. Over a specific gaming period, the player is asked to conduct a minimum set of daily behaviors, such as completing one quest, confronting one bad guy, and using two powerups. The learners do not all have to follow the same path but can work on different problems. The results for each person are not graded centrally. Instead, each person reports on progress to “allies” of their own choosing. This feature is conducive to an organizational culture that reinforces positive habits concerning security. Global performance reports, such as the number of seeded phishing messages that went undetected, or the number of ‘bad guy’ inside players who were not found out, can drive the enterprise-wide game forward. At the same time, the number of real threats detected and countered is also reported. Much broader scenarios can be simulated as well, during which players discuss what they think is happening and what actions they propose. These episodes can later be analyzed in after-action review sessions, using procedures pioneered by the Army [5]. This gaming approach can be used to shift emphasis to build pervasive habits of action, consistent with the safety culture strategy, and also reflection on and reframing of unique situations, consistent with the sensemaking strategy [19]. Whatever specific approach is used, we need to design selection, training, and work processes in a way that neither makes unrealistic

expectations for the human component nor neglects creating a disciplined human capability that makes the most of the technical capability.

We have established that there is a human role in cyber security, and that individual thinking practices can be addressed, but any strategy is easily broken if it is not aligned with all other forces in the organization. We can again look to industrial cases for guidance in instituting a holistic strategy. The US Army's efforts in the 1980s are instructive [6]. The Army confronted a disturbing pattern. The design and testing phases of system development were apparently very successful, but when deployed, these systems had serious problems. In one famous case, a hand-held missile performed perfectly on bench tests, but no soldier was every able to hit a target with it. On launch, the recoil was so severe that the soldier lost his grip, and thus, the missile would either shoot skyward or directly into the ground. Many similar acquisitions promoted human error and loss and could not be remedied with late-stage patches. The Army concluded that the development process as a whole was deeply flawed. The conditions under which systems were to operate in the field, and to be incorporated within the whole enterprise, were not adequately understood at any point in the process. In particular, the human contribution to systems was overlooked or misrepresented. While human factors engineering had always been part of Army acquisition, testing occurred far too late in the development process, and the focus on immediate operator controls was far too narrow.

The Army launched an acquisition reform program called Manpower Personnel Integration (MANPRINT) [46], bringing forward all human domains for consideration during design and throughout development. Human factors experts shifted their roles from tester to designer, but the new strategy was far more important. Integration of the broad human domains of manpower, personnel, training, safety, health, and human factors engineering with each other and with technical engineering, constituted the new sociotechnical strategy. Practitioners feared that layers of new procedures would overburden development, but a simultaneous shift to agile development concepts made the new strategy workable. Use of the new strategy in some quite complex programs yielded products that performed well in the field and were readily adopted into service.

Despite a strong effort to institutionalize MANPRINT, the lessons faded under the ever-present pressure to quickly converge on low-cost technical

solutions, to succumb to the temptation to simplify what is not simple, to allow engineers to run ahead to "get the job done" with technology alone and to bypass the human component. The challenge is to recognize that the influence of human environments cannot be eliminated or safely ignored and must be incorporated throughout the lifecycle. Late-stage patches are insufficient. The military has not completely forgotten this legacy, however. A sociotechnical strategy has been reborn in the Defense Department's cyber realm where threats are as intense as in any enterprise [51]. This initiative emphasizes the high reliability strategy and points to nuclear submarine operations as an instructive case, but for reasons mentioned above, we would caution against over-emphasis of that strategy.

We have argued for a more systemic and strategic approach, but of course any strategy that is devised needs further elaboration. We still need synthesis among competing strategies, and guidelines to implement and institutionalize the strategy. There is a place for awareness training, but we have argued that conventional training techniques are not particularly effective in reducing cyber human errors and loss, partly because the problem is inappropriately framed and reinforced. We reviewed comprehensive human error mitigation strategies, ones that have been successfully demonstrated in industrial situations that are parallel to those faced by complex, IT-enabled organizations.

What does a holistic, sociotechnical systems approach cost? In some sense, it costs nothing for everyone to act safely. Or one could argue that the effort to change habits and underlying structures is just an extension of effort that would have been put into organizational development and training anyway, and so is not an extra expense. Yet it is evident that cyber security operations are different in many ways from industrial organizations, ways that make the organizational and cognitive design for cyber security especially complex. Further, any cultural shift entails a great deal of discussion over long periods, no matter how simple the shift. It is never just a message, but is a practice which requires deep-seated change. The social side of cyber security should remain a matter for inquiry whenever significant investments are made to modify the technical system. We are not going to eliminate the human contribution to cyber risk with any model or strategy, but we can make progress by thinking differently, widely, continuously, and with some caution

Acknowledgements.

We would like to acknowledge the support given by Dr. David Honey (the Intelligence Community's Director of Science and technology (DS&T)), colleagues and staff of the Acquisition Technology and Facilities/Science and Technology for many influential discussions and helpful comments. Paul Waters is funded by Australian Research Council grant DP160100601.

References

- [1] AMERICAN SOCIETY FOR INDUSTRIAL SECURITY (ASIS) (1997). Commentary on Managing Risks of Organizational Accidents.
- [2] ANDERSON, J.R., AND BOWER, G.H. (1973). *Human Associative Memory*. V.H. Winston and Sons, New York, NY.
- [3] ATKINSON, R.C., AND SHIFFRIN, R.M. (1968). Human memory: A proposed system and its control processes. In *The Psychology of Learning and Motivation, Volume 2*.

- K.W. Spence and J.T. Spence, Eds. Academic Press, New York, NY.
- [4] AZAB, A., WATTERS, P.A. and LAYTON, R. (2012). Characterising network traffic for Skype forensics. *Proceedings of the 3rd Workshop on Cybercrime and Trustworthy Computing*.
- [5] BAIRD, L; HOLLAND, P.; DEACON, S. (1999) Learning from action: Imbedding more learning into the performance fast enough to make a difference. *Organizational Dynamics*, 27 (4) 19-32. Spring 1999.
- [6] BOOHER, H. R. (1990). *MANPRINT an approach to systems integration*. New York: Van Nostrand Rheinhold.
- [7] BRUCKMAN, A. (1998) Community support for constructionist learning. *Computer Supported Cooperative Work (CSCW)*, 7 (1-2), 47-86.
- [8] CARROLL J.S. (1995) Incident Reviews in High-Hazard Industries: Sense Making and Learning Under Ambiguity and Accountability, *Organization & Environment*, June 1995 9: 175-197.
- [9] DEKKER, S. (2005). Ten Questions About Human Error: A New View Of Human Factors And System Safety. Mahwah, NJ: Lawrence Erlbaum Associates.
- [10] DEKKER, S., HOLLNAGEL, E., WOODS, D. AND COOK, R. (2009). Resilience Engineering: New directions for measuring and maintaining safety in complex systems. Final Report, November 2008, Lund University School of Aviation.
- [11] EUROCONTROL. (2002). Technical Review of Human Performance Models and Taxonomies of Human Error in ATM (HERA) (Technical Report No. HRS/HSP-002-REP-01). Brussels, BE.
- [12] EVERETT, R.J. & SOHAL A.S. (1991) Individual Involvement and Intervention in Quality Improvement Programmes: Using the Andon System. *International Journal of Quality & Reliability Management*, 8 (2).
- [13] GORMAN, J, COOKE, N., SALAS, E. (2010) Preface to the Special Issue on Collaboration, Coordination, and Adaptation in Complex Sociotechnical Settings. *Human Factors* 52 (2).
- [14] GRABOWSKI, M. & ROBERTS, K.H. (1997) Risk Mitigation in Large-Scale Systems: Lessons from High Reliability Organizations. *California Management Review*, 39 (4), Summer 1997; (pp. 152-162)
- [15] GRANVILLE, K. "Nine Recent Cyberattacks Against Big Businesses." *New York Times*, February 5, 2015. Retrieved from http://www.nytimes.com/interactive/2015/02/05/technology/recent-cyberattacks.html?_r=0
- [16] GREITZER, F.L., KUCHAR O.A. AND HUSTON K. (2007). "Cognitive Science Implications for Enhancing Training Effectiveness in a Serious Gaming Context." *ACM Journal of Educational Resources in Computing*, 7 (3).
- [17] HERLEY, C. (2009). So Long, And No Thanks for the Externalities: The Rational Rejection of Security Advice by Users. Retrieved from <http://research.microsoft.com/en-us/um/people/cormac/papers/2009/SoLongAndNoThanks.pdf>
- [18] HOLLNAGEL E. (1993). Human reliability analysis: Context and control. Academic Press.
- [19] IP G. (2015) Foolproof: Why Safety Can Be Dangerous and How Danger Makes Us Safe. New York: Little, Brown & Company.
- [20] ISACA & RSA CONFERENCE. (2015). State of Cybersecurity: Implications for 2015 An ISACA and RSA Conference Survey. Retrieved from http://www.isaca.org/cyber/Documents/State-of-Cybersecurity_Res_Eng_0415.pdf
- [21] KLEIN G., MOON B., HOFFMAN R. R. (2006). Making sense of sensemaking 2: a macrocognitive model. *IEEE Intell. Syst.* 21, 88–92.
- [22] KOTOVSKY, K., HAYES, J. R., AND SIMON, H. A. (1985). Why are some problems hard? Evidence from the Tower of Hanoi. *Cognitive Psychology*, 17, 248-294.
- [23] LAYTON, R., & WATTERS, P.A. (2009). Determining provenance of phishing websites using automated conceptual analysis. *Proceedings of the APWG E-crime Research Summit*
- [24] LEVENSON, N.G. (2009) Moving Beyond Normal Accidents and High Reliability Organizations: A Systems Approach to Safety in Complex Systems. *Organizational Studies*, 30 (2-3) 227-249
- [25] LEVELSON, N.G. (2011) Engineering a Safer World: Systems Thinking Applied to Safety. MIT Press, Cambridge, MA.
- [26] LOBO, D., WATTERS, P.A., & WU, X. (2010). RBACS: Rootkit Behavioral Analysis and Classification System. *Proceedings of the International Conference on Knowledge Discovery and Data Mining (WKDD 2010)*
- [27] McCOMBIE, S., PIEPRZYK, J., WATTERS, P.A. & LAYTON, R. (2012). Russia, Ukraine and Global Cybercrime: The Australian Perspective. *Proceedings of the 2nd International Conference on Cybercrime, Security & Digital Forensics*
- [28] MCGONIGAL, J. (2015) Superbetter: A Revolutionary Approach to Getting Stronger, Happier, Braver and More Resilient – Powered by the Science of Games. New York City: Penguin.
- [29] MILLER, G. A. (1956). The magical number, seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.
- [30] NATIONAL ACADEMIES OF SCIENCE (NAS). Roundtable on Cyber Resilience, unpublished briefing November 2015.
- [31] PAPERT, S. (1991) "Situating Constructionism." Chapter 1 in *Constructionism*. I. Harel and S. Papert, Eds. Ablex Publishing, Norwood, NJ.
- [32] PERROW, C. (1999) Normal Accidents: Living With High Risk Technologies, Princeton: Princeton University Press.
- [33] PEW, R. W., MILLER, D. C. AND FEEHRER, C. E. (1981). Evaluation of proposed control room improvements through analysis of critical operator decisions. Palo Alto, CA: Electric Power Research Institute NP-1982.
- [34] REASON, J. (1990). *Human error*. Cambridge, UK: Cambridge University Press
- [35] REASON J. (2008). The human contribution: unsafe acts, accidents and heroic recoveries. Aldershot, UK; Burlington, VT: Ashgate.
- [36] ROBERTS, K. H. (1990). Some Characteristics of High-Reliability Organizations. *Organization Science*, 1, 160-177.
- [37] RODIN, J. (2014). *The Resilience Dividend: Being Strong in a World Where Things Go Wrong*, New York, Public Affairs.
- [38] ROE, E., & SCHULMAN, P. R. (2008). *High Reliability Management: Operating on the Edge*. Palo Alto, CA: Stanford University Press.
- [39] SAGAN, S. D. (1993). The Limits of Safety: Organizations, Accidents, and Nuclear Weapons. Princeton, N.J.: Princeton University Press.

- [40] SCHNEIDER, W. AND SHIFFRIN, R. 1977. Controlled and automatic information processing: I. Detection, search and attention. *Psychological Review*, 84, 1-66.
- [41] SCHNEIER, B. Security Awareness Training. Retrieved from https://www.schneier.com/blog/archives/2013/03/security_aware_1.html
- [42] SHIFFRIN, R. AND SCHNEIDER, W. 1977. Controlled and automatic information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- [43] TALBOT, D. Why We're So Vulnerable. MIT Technology Review, Business Report. January 2016. Retrieved from <https://www.technologyreview.com/s/545621/why-were-so-vulnerable/>
- [44] TALEB, N.N. *Antifragile: Things That Gain from Disorder*. New York City: Random House Publishing, 2012. Print
- [45] URECHE, O., LAYTON, R., & WATTERS, P.A. (2012). Towards an implementation of information flow security using semantic web technologies. *Proceedings of the 3rd Workshop on Cybercrime and Trustworthy Computing*
- [46] U.S. ARMY REGULATION (AR) 602-2 (1990). Manpower and Personnel Integration (MANPRINT) in the Material Acquisition Process. Washington, DC: Department of the Army. Retrieved from <http://www.acq.osd.mil/se/docs/Army-FY09-HSI-Plan.pdf>
- [47] VAN DYCK, C., FRESE, M., BAER, M. AND SONENTAG, S. Organizational Error Management Culture and Its Impact on Performance: A Two-Study Replication *Journal of Applied Psychology* 90 (2005), 6, pp. 1228-1240
- [48] VICENTE, K. J., & RASMUSSEN, J. (2002). Ecological interface design: Progress and Challenges. *Human Factors*, 44(1), 62-78.
- [49] WATTERS, P.A. (2013). Modelling the Effect of Deception on Investigations Using Open Source Intelligence (OSINT). *Journal of Money Laundering Control*, 16, 238-248
- [50] WEICK, K.E. (1995) *Sensemaking in Organizations* (Foundations for Organizational Science). Thousand Oaks, California: SAGE Publications.
- [51] WINNEFELD, J.; KIRCHOFF, C.; UPTON, D. (2015) Cybersecurity's Human Factor: Lessons from the Pentagon, *Harvard Business Review*, September 2015. <https://hbr.org/2015/09/cybersecuritys-human-factor-lessons-from-the-pentagon>