Choosing Optimal Ways to Increase Resilience of Critical Infrastructures

Sandra König
Austrian Institute of Technology
Sandra.Koenig@ait.ac.at

ABSTRACT

Increasing resilience is a core interest in critical infrastructure (CI) protection that involves many challenges. It is necessary to agree on a common understanding of resilience and identify potential strategies to improve it. Once this is done, the question arises how to choose among these strategies. We propose to decide based on a game-theoretic framework that allows identification of optimal actions under various scenarios. This framework considers different threat scenarios as attacks to the CI and the identified strategies to improve resilience as defense strategies for the CI. Since the payoff of the game, namely the resilience of the CI, can hardly be measured with certainty we choose an extension of classical game theory that allows taking uncertainty into account and still finds provably optimal solutions. This approach is especially useful in a situation where we aim to optimize a quantity that is difficult to measure (such as resilience). The result of this analysis is two-fold: it identifies an optimal defense but also provides information about the resilience in the worst case. The approach is illustrated with a small example using a publicly available implementation.

Keywords

Resilience, critical infrastructure, optimization, game theory.

INTRODUCTION

In the context of critical infrastructures (CIs) resilience is especially important as the impact of a failure or even limited availability of a CI has a significant impact on society and economy (Boumphrey and Bruno 2015; National Infrastructure Commission 2017; Royal Academy of Engineering 2018). Several measures for resilience exist for specific sectors of CIs, e.g., for the energy domain (Panteli et al. 2017) or the water supply (Cuisong and Hao 2008), or for specific scenarios, e.g., for hurricanes (Tokgoz and Gheorghe 2013). More advanced approaches to resilience take into account viewpoints from organizations and people (Gouglidis et al. 2016) but are more time-consuming to evaluate. In this work, we use a more general formulation of resilience in the context of CIs to develop a general methodology that may be refined for specific sectors. We choose a more general measure as proposed in (König et al. 2019) that is simple but in line with a risk management process according to ISO 31000. This measure builds on the notion of risk as the product of likelihood and impact by defining the resilience against a specific scenario as the product of likelihood and a factor depending on the expected impact due to the risk scenario and the preparedness against it. It is constructed in such a way that positive values indicate that preparedness is on average higher than the expected damage and negative values indicate that preparedness is not sufficient on average.

Once agreed on a measure for resilience, the aim is to identify ways to improve it as much as possible. This involves several challenges. First, it is important to identify relevant threats and potential countermeasures. Then, the question is which of these countermeasures should be chosen to optimize resilience. Answers to the first question are manifold and may include consultation of vulnerability databases but also adversary modelling. Still, such a list of threats will always be incomplete due to the risk of unexpected events (i.e., black swans). The second question on the other hand may be threated rigorously by applying optimization techniques. As the resulting resilience depends on both the threat and the countermeasure we here apply a game-theoretic framework where an attacker causes problems that reduce the CI's resilience while the CI operator tries to maximize the resilience in the light of all possible attack. In this somewhat abstract view the attacker is not necessarily a real person but anything that threatens the CI, including nature. In this setting it is impossible to model

the adversary explicitly since we do not have any information on his intention. However, considering the worst case that the adversary's intentions are exactly opposite to the defenders ones (i.e., playing a zero-sum game) provides a lower bound to the resilience if the CI. This game theoretic optimization of resilience may benefit from existing risk management processes where for example risks have been identified and evaluated. Such information may be incorporated in the described approach.
This paper is structured as follows. After a discussion on how to measure resilience in CIs we focus on ways to improve it based on the proposed game theoretic framework. This consists of the identification of threat scenarios and strategies to improve resilience as well as estimation of the resilience for each scenario to build a game. The choice on an optimal improvement strategy is found by solving this game. An example is provided to illustrate how to apply the theoretic model. Open issues are discussed in the conclusion.

MEASURING RESILIENCE IN CRITICAL INFRASTRUCTURES

When aiming for improvement of resilience, it is necessary to have a measure of resilience so that different situations can be compared. While a general definition is hard to find, formalizations are possible for specific fields. Concerning CIs, the term resilience can be understood (analogously to the term risk) as

$$R = \sum_{s=1}^{N} \omega_s \cdot (P_s - E\,[I_s]), \qquad (*)$$

where the sum is over all scenarios s, ω_s is the likelihood of occurrence of scenario s, P_s is the preparedness against scenario s, and $E\,[I_s]$ is the expected impact on the CI in case scenario s happens (König et al. 2019). In order to compute this value, it is first necessary to agree on a set of scenarios that threaten the CI. These include both natural disasters and intentional attacks. For each of these scenarios it is then necessary to estimate the likelihood of occurrence, which is often based on historical data about similar events. Then, every single scenario needs to be investigated to answer two questions:

- What is the expected impact on the CI for this scenario?
- How well is the CI prepared for this scenario in the sense that it is able to keep its functionality to a certain level despite a realization of the scenario?

If available, simulation methods may be used to mimic a scenario and estimate its impact. Otherwise, experts may be able to provide a qualitative estimate or historical data may be available. The second question is a bit more difficult to answer as the term preparedness is not clearly defined, but needs to be taken into account (Martin and Ludek 2013). Practically, it is necessary to measure expected impact and preparedness on the same scale to be able to evaluate formula (*) above. A qualitative scale, e.g., a 5-tier scale, is recommended for experts' opinions since precise estimations are difficult to get. In that case, the term *preparedness level* is used. In the context of CI resilience, the term preparedness is understood as the maximal degree of impact that still allows smooth operation (similar to the breaking point of an estimator in statistics), i.e., a damage higher than the preparedness yields to limited functionality of the CI and reduces the overall resilience of the CI (as this scenario contributes a negative value to the sum).

When investigating resilience of a CI it is not always enough to have a single value that provides an average resilience. A deeper analysis may be concerned about the various scenarios and their severity. In particular, we aim to identify scenarios with the low resilience as these threaten the system the most. Thus, we focus on the resilience R_s for a specific scenario s,

$$R_s = \omega_s \cdot (P_s - E\,[I_s]),$$

to be able to compare the different situations. The value is also influenced by changes in the CI, in particular if different strategies for improvement are implemented. For a bunch of improvement strategies we identify those that are optimal, i.e., that reduce resilience the most.

INCREASING RESILIENCE OF A CRITICAL INFRASTRUCTURE

Evaluation of the resilience for the current situation is only the first step in an analysis that aims at improving resilience of a CI. Consider a fixed set of scenarios with known probabilities of occurrences. When applying the formulation of resilience as described above, resilience can be improved by either increasing the preparedness level or reducing the expected impact (or even both at the same time).

Potential strategies for improvement range from implementation of new protection measures (e.g., honeypots to detect malware), improving existing controls (e.g., check pumps more frequently) or training of employees to increase awareness of new threats. While a list of potential actions can be found in discussion with experts, it is

typically not possible to apply all potential actions due to limited resources. The challenge is then to decide which strategies should be chosen. In this work, we apply a game-theoretic framework that describes a threat to a CI as an attack (intentional or not) while the CI operator is in the position to defend, i.e., to protect his system. As we do not have any knowledge on the attackers intentions (especially in the case of natural disasters), we model the situation with a zero-sum game that assumes an adversary who tries to cause as much damage as possible while the defender tries to minimize the suffered damage. Analysing the situation with this approach consists of three steps:

(1) identification of threat scenarios
(2) identification of countermeasures
(3) estimating resilience for each combination of an attack and defence (a scenario)

Threat scenarios are interpreted as attack strategies where the attacker does not need to be a person but also nature, basically anything that causes harm to the CI. Similarly, countermeasures are interpreted as defence strategies since these help reducing the expected damage. Finally, the estimated resilience is the payoff that the defender cares about. He aims to maximize it while the attacker's actions aim at reducing it. In the remainder of this section we elaborate on these three steps.

Identification of Threats and Countermeasures

The identification of relevant threats to the CI and potential countermeasures to protect the system is the most difficult part of the game-theoretic analysis. A game-theoretic equilibrium identifies the worst-case attack and the best way to protect against it, but it can only be chosen from available information. The optimality of the found solution is no longer valid if attacks have not been considered or defence strategies are incomplete or cannot be applied in practice. A list of threats contains information on past incidents as well as issues identified with experts familiar with the system. Recent publications such as (Ghafir et al. 2018; Zimba et al. 2018) describe various scenarios that may harm the system.
A list of countermeasures has at least two sources of information: international standards such as ISO27001 and expert knowledge. Additional actions that may protect the system can be found when analysing past incidents or penetration testing, but also from publications such as (Walker-Roberts et al. 2018). Further, a list of threats may be available form a classical risk analysis that includes risk identification and evaluation. Despite all these different sources, the list of threats will always be incomplete (as it misses for example unexpected events). However, this is a general issues and not a weakness of the game theoretic approach. Another issue is the growth of the strategy space when combinations of attacks are considered. While this is a problem in the theoretical setting it is not such a big issue in practical applications as combination of threats (e.g., a contamination followed by an earthquake) are rather rare and arbitrary combination of protection measures may not be affordable for the CI operator.

Estimation of Resilience

Having a set of relevant attacks and feasible defence strategeis, resilience must be estimated for each scenario. For the resilience measure used here this requires estimation of the expected impact and the preparedness for a specific attack and a fixed countermeasure. The likelihood of occurrence remains the same (just as the chances of a lightning when computing the risk), but countermeasures influence the impact of a scenario (i.e., better protection prevents the house from burning down). Different strategies typically do not change the dynamics of the cascading effects but rather cause a faster or slower spreading. In case the expected impact is estimated based on simulations, different scenarios may be analysed by choosing different parameters for the simulation. In case we are not able to mimic the consequences of an attack, we need to ask experts to rate the expected damage. In both cases, we get several estimates for the same situation which yields distribution-valued payoffs. As described above, preparedness is measured through expert assessments on a qualitative scale. Also, this estimate is distributions-valued if we ask several experts for their assessment and do not aggregate their assessments.
In a classical game-theoretic setting, the estimated payoffs need to be real numbers but extensions allow payoffs to be random variables (Rass et al. 2015). This allows in particular dealing with disagreement among experts since it is not necessary to agree on a single estimate. Rather, we collect all available opinions and work with histograms instead of numbers. While this may yield pessimistic estimates (e.g., if one of the experts tends to be anxious) it is in line with a qualitative risk assessment as recommended by the German Federal office of Information Security (BSI). Further, this general approach helps to keep track of single assessments and thus increases the understanding of the decision while such information is lost if different opinions are aggregated.

Choosing Optimal Protection Strategy

In order to find an optimal protection strategy the game needs to be solved. Since our estimate of resilience contains assessments from several experts we apply the more general setting. It allows taking into account every single assessment rather than aggregating all these different opinions. Even in this generalized setting a Nash equilibrium can be computed (Rass et al. 2016) with the (generalized) fictitious play algorithm. It mimics a fixed number of game plays and records how often each strategy has been played. It can be shown that this yields an estimate of a Nash equilibrium (Rass et al. 2016) where no player has an incentive to deviate since he will not be able to improve. An implementation of the methods applied here is publicly available on CRAN (Rass and König 2018). Its use is illustrated with a small example below.
The main difference to existing approaches lies in the way how uncertainty is taken into account. Stochastic models like Bayesian games assume that a player is unsure about what type of adversary he is playing with but is certain about the payoffs in each possible case. Our point of view is different: we think it is generally very changing to exactly predict the payoffs for a given scenario which is why we replace the crisp (real-valued) payoff by a random variable.

ILLUSTRATIVE EXAMPLE

Consider a fictitious water provider that uses an industrial control system such as a SCADA (Supervisory Control and Data Acquisition) system to control its processes. The SCADA server as the central component of the SCADA network is located in a control room to allow monitoring processes in real time. Programmable Logic Controllers (PLCs) supervise processes at the treatment plant and Remote Terminal Units (RTUs) at remote sites such as pump stations.

The game-theoretic model considers the three attacks

- s_1 contamination
- s_2 earthquake
- s_3 attack on SCADA system

and the defence strategies

- d_1 more water reservoirs
- d_2 more frequent controls

Attacks on SCADA systems are manifold. We here consider the situation where data in the SCADA system are manipulated in such a way that a change at the remote site is not recognised or such that fake data provokes an act that causes damage to the system.
Increasing the amount of stored water in reservoirs aims at reducing the damage in case of a contamination. Regular controls help to detect problems with the water quality (e.g., due to a contamination) as well as a potential mismatch between the real status of a site and the data reported by the SCADA system in case of an attack that involves forged data. Countermeasures against earthquakes are complex and expensive (e.g., reconstruction of a building) and are thus not taken into account in a mid-term analysis.
In order to recognise improvements over the current state we also need to include the current situation of the CI as a defence strategy

- d_0 no change (status quo)

The likelihood of occurrence for the considered attacks are set to $\omega_1 = 0.2$, $\omega_2 = 0.2$ and $\omega_3 = 0.3$ based on reports on contamination (Brown and Darby 1988), earthquakes (Field 2015) and ICS attacks (Kovacs 2016).
We assume that estimates of the expected impact as well as the preparedness level are available either from expert knowledge or simulation. For illustration we work with the data given in Table 1, Table 2 and Table 3 where a 5-tier scale is used such that 1 corresponds to smooth operation and 5 to total failure (intermediate values correspond to limited availability, accordingly).[i] These values are artificial and solely meant to illustrate the approach. Still they should reflect the intuition that more water reservoirs (d_1) increase the preparedness against a contamination (as more pure water is available). More frequent controls (d_2) are assumed to increase preparedness against contamination and a SCADA attack (as it is discovered earlier) and to reduce the expected impact due to a contamination or a SCADA attack (as the spreading gets slower due to the controls). In practice, these values come from interviews with experts in the field that provide estimates of the preparedness and the impact. In case experts may only provide a vague prediction such as the most likely value and an assurance level, a distribution over all possible values can be estimated (König and Rass 2018). Further, simulation methods may support the impact estimation.

Table 1 Estimates for current status d_0

Scenario	Preparedness	Estimated Impact
s_1 contamination	2,3,3,3,4	3,3,4,4,4
s_2 earthquake	3,3,4,4,4	3,3,3,4,4
s_3 SCADA server attack	2,3,3,4,4	3,4,4,4,5

Table 2 Estimates for more water reservoirs d_1

Scenario	Preparedness	Estimated Impact
s_1 contamination	3,3,3,3,4	3,3,3,4,4
s_2 earthquake	3,3,4,4,4	3,3,3,4,4
s_3 SCADA server attack	2,3,3,4,4	3,4,4,4,5

Table 3 Estimates for more frequent controls d_2

Scenario	Preparedness	Estimated Impact
s_1 contamination	3,3,3,3,4	3,3,3,3,4
s_2 earthquake	3,3,4,4,4	3,3,3,4,4
s_3 SCADA server attack	3,3,3,4,4	3,3,4,4,5

These data allow computation of the resilience for each combination of a threat scenario and a defense strategy. Since we have several estimates, we use the R package HyRiM (Rass and König 2018) that is able to optimize distribution-valued payoffs. In order to apply that framework, the maximization of the resilience R needs to be translated into minimization of $-R$ and the resulting values need to be transformed such that the values are positive integers and potential gaps need to be filled by smoothing (see description of the package (Rass and König 2017) for details on how to use it). The resulting payoff matrix is shown in Figure 4.

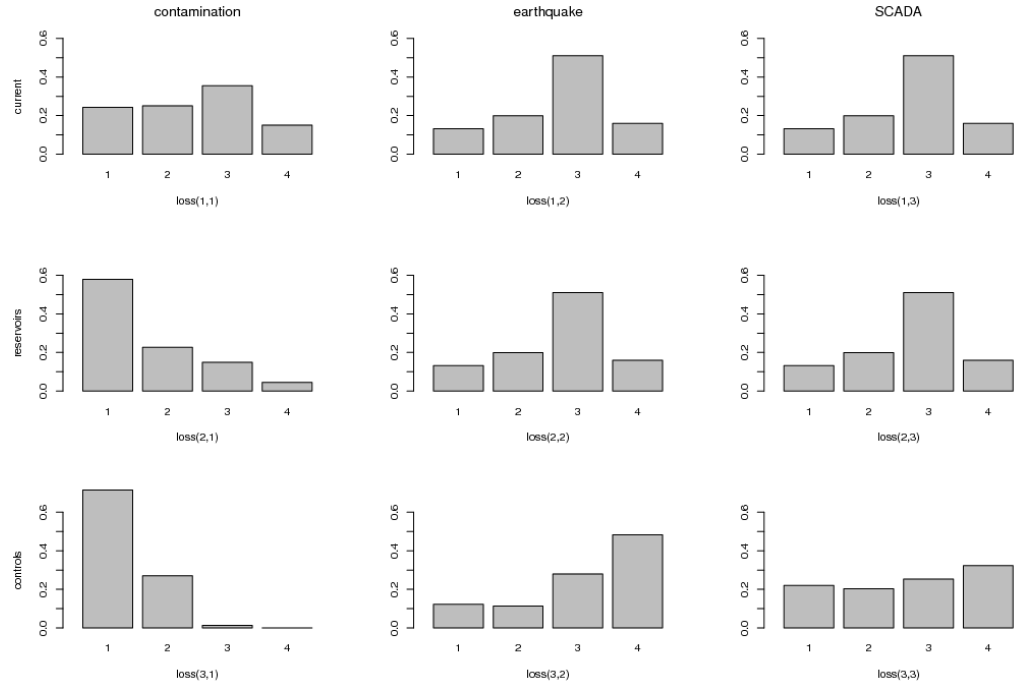

Figure 4 Payoff Matrix

Applying the generalized fictitious play algorithm to solve this game yields the Nash equilibrium (d_0, s_2), that is, the scenario that reduced resilience the most is an earthquake (s_2) and the best way to protect against it is to keep the current state (d_0). In other words, the two defense strategies d_1 and d_2 do not yield to an improvement in this setting and it is necessary to look for other strategies to improve resilience under these threat scenarios. The solution is illustrated in **Figure 5** and reads as follows: from the potential defense strategies, choose "current" $(d_0)$with probability 1 and similarly choose attack strategy "earthquake" (s_2). In case these two strategies are played the resulting distribution of the resilience is shown in the lower right corner.

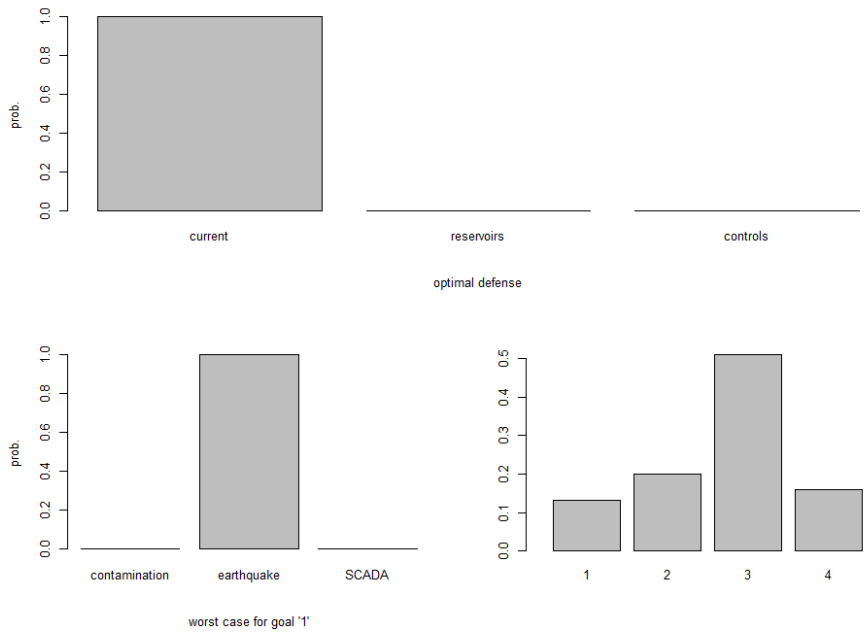

Figure 5 Optimal Strategies and Worst Case Resilience

In a more elaborated example where the list of potential defenses and threats is longer, the optimal solution may not only involve a single defense action but a combination of several actions (a mixed equilibrium). Then, the strategies with a positive likelihood need to be played with relative frequencies that match these likelihoods. For the risk scenarios that constitute the attack strategies this means that several scenarios may occur successively (with relative frequencies corresponding to the likelihoods from the resulting equilibrium). The framework allows optimizing several goals simultaneously which makes it possible to minimize implementation costs of countermeasures at the same time.

CONCLUSION

In the context of critical infrastructures several measures for resilience exist for different domains. Some of these measures allow comparison between different CIs but also for different situations a single CI may face. If it is possible to estimate the resilience for different situations, consequences of protection measures can be compared. Game-theoretic models can be applied to determine the optimal way to protect a system if a list of potential improvements is available. If various threat scenarios are considered it is further possible to determine a worst-case scenario that yields to minimal resilience and estimate the resulting resilience in this case.

While the introduced approach is general in the sense that it is applicable to a wide range of CIs, many issues are still to be investigated in more detail. A big challenge is the estimation of the preparedness level. An estimation of this value depends on many factors such as the weather conditions, so that a single value may not be enough to describe this parameter. Methods for impact estimation as well as the identification of strategies need to be examined in more depth, with different approaches for different sectors. Another issue ignored so far time that plays an important role when talking about resilience. The resilience measure used is time-independent and the game is only played once so that it optimizes the current situation.

ACKNOWLEDGMENT

The author thanks Julian Magin from the Austrian Institute of Technology for interesting discussions.

REFERENCES

Boumphrey, R., and Bruno, M. (2015). “Foresight Review of Resilience Engineering: designing for the expected and unexpected.”

Brown, J. A., and Darby, W. P. (1988). “Predicting the probability of contamination at groundwater based pub-

lic drinking supplies." *Mathematical and Computer Modelling*, 11, 1077–1082.
Cuisong, Y., and Hao, Z. (2008). "Resilience Classification Research of Water Resources System in a Changing Environment." *2008 2nd International Conference on Bioinformatics and Biomedical Engineering*, IEEE, Shanghai, China, 3741–3744.
Field, E. H. (2015). "UCERF3: A new earthquake forecast for California's complex fault system." US Geological Survey.
Ghafir, I., Saleem, J., Hammoudeh, M., Faour, H., Prenosil, V., Jaf, S., Jabbar, S., and Baker, T. (2018). "Security threats to critical infrastructure: the human factor." *The Journal of Supercomputing*, 74(10), 4986–5002.
Gouglidis, A., Shirazi, S. N., Simpson, S., Smith, P., and Hutchison, D. (2016). "A multi-level approach to resilience of critical infrastructures and services." *2016 23rd International Conference on Telecommunications (ICT)*, IEEE, Thessaloniki, Greece, 1–5.
König, S., and Rass, S. (2018). "Investigating Stochastic Dependencies Between Critical Infrastructures." International Journal on Advances in Systems and Measurements, 11(3 & 4), 250–258.
König, S., Schaberreiter, T., Rass, S., and Schauer, S. (2019). "A Measure for Resilience of Critical Infrastructures." *Critical Information Infrastructures Security*, E. Luiijf, I. Žutautaitė, and B. M. Hämmerli, eds., Springer International Publishing, Cham, 57–71.
Kovacs, E. (2016). "IBM Reports Significant Increase in ICS Attacks." *SecurityWeek*.
Martin, H., and Ludek, L. (2013). "The status and importance of robustness in the process of critical infrastructure resilience evaluation." *2013 IEEE International Conference on Technologies for Homeland Security (HST)*, IEEE, Waltham, MA, USA, 589–594.
National Infrastructure Commission. (2017). *Infrastructure and Digital Systems Resilience*. London.
Panteli, M., Mancarella, P., Trakas, D. N., Kyriakides, E., and Hatziargyriou, N. D. (2017). "Metrics and Quantification of Operational and Infrastructure Resilience in Power Systems." *IEEE Transactions on Power Systems*, 32(6), 4732–4742.
Rass, S., and König, S. (2017). "Package 'HyRiM': Multicriteria Risk Management using Zero-Sum Games with vector-valued payoffs that are probability distributions."
Rass, S., and König, S. (2018). "HyRiM: Multicriteria Risk Management using Zero-Sum Games with vector-valued payoffs that are probability distributions." https://cran.r-project.org/package=HyRiM.
Rass, S., König, S., and Schauer, S. (2015). "Uncertainty in Games: Using Probability-Distributions as Payoffs." *Decision and Game Theory for Security: 6th International Conference, GameSec 2015, London, UK, November 4-5, 2015, Proceedings*, M. Khouzani, E. Panaousis, and G. Theodorakopoulos, eds., Springer International Publishing, Cham, 346–357.
Rass, S., König, S., and Schauer, S. (2016). "Decisions with Uncertain Consequences—A Total Ordering on Loss-Distributions." *PLOS ONE*, (S. D. Peddada, ed.), 11(12), e0168583.
Royal Academy of Engineering. (2018). *Cyber safety and resilience - strengthening the digital systems that support the modern economy*. London.
Tokgoz, B. E., and Gheorghe, A. V. (2013). "Resilience quantification and its application to a residential building subject to hurricane winds." *International Journal of Disaster Risk Science*, 4(3), 105–114.
Walker-Roberts, S., Hammoudeh, M., and Dehghantanha, A. (2018). "A Systematic Review of the Availability and Efficacy of Countermeasures to Internal Threats in Healthcare Critical Infrastructure." *IEEE Access*, 6, 25167–25177.
Zimba, A., Wang, Z., and Chen, H. (2018). "Multi-stage crypto ransomware attacks: A new emerging cyber threat to critical infrastructure and industrial control systems." *ICT Express*, 4(1), 14–18.

[i] We chose an odd number of classes to explicitly allow experts to choose a "middle" value to indicate no specific preference to one of the extremes. If this is not desired, a different scale with an even number of classes may be used instead.