Balancing Protection and Resilience

St. George Protection-Resilience Study

Part 4 of 4: Comparison Study

David N. Bristow
Alexander H. Hay

Centre for Resilience of Critical Infrastructure

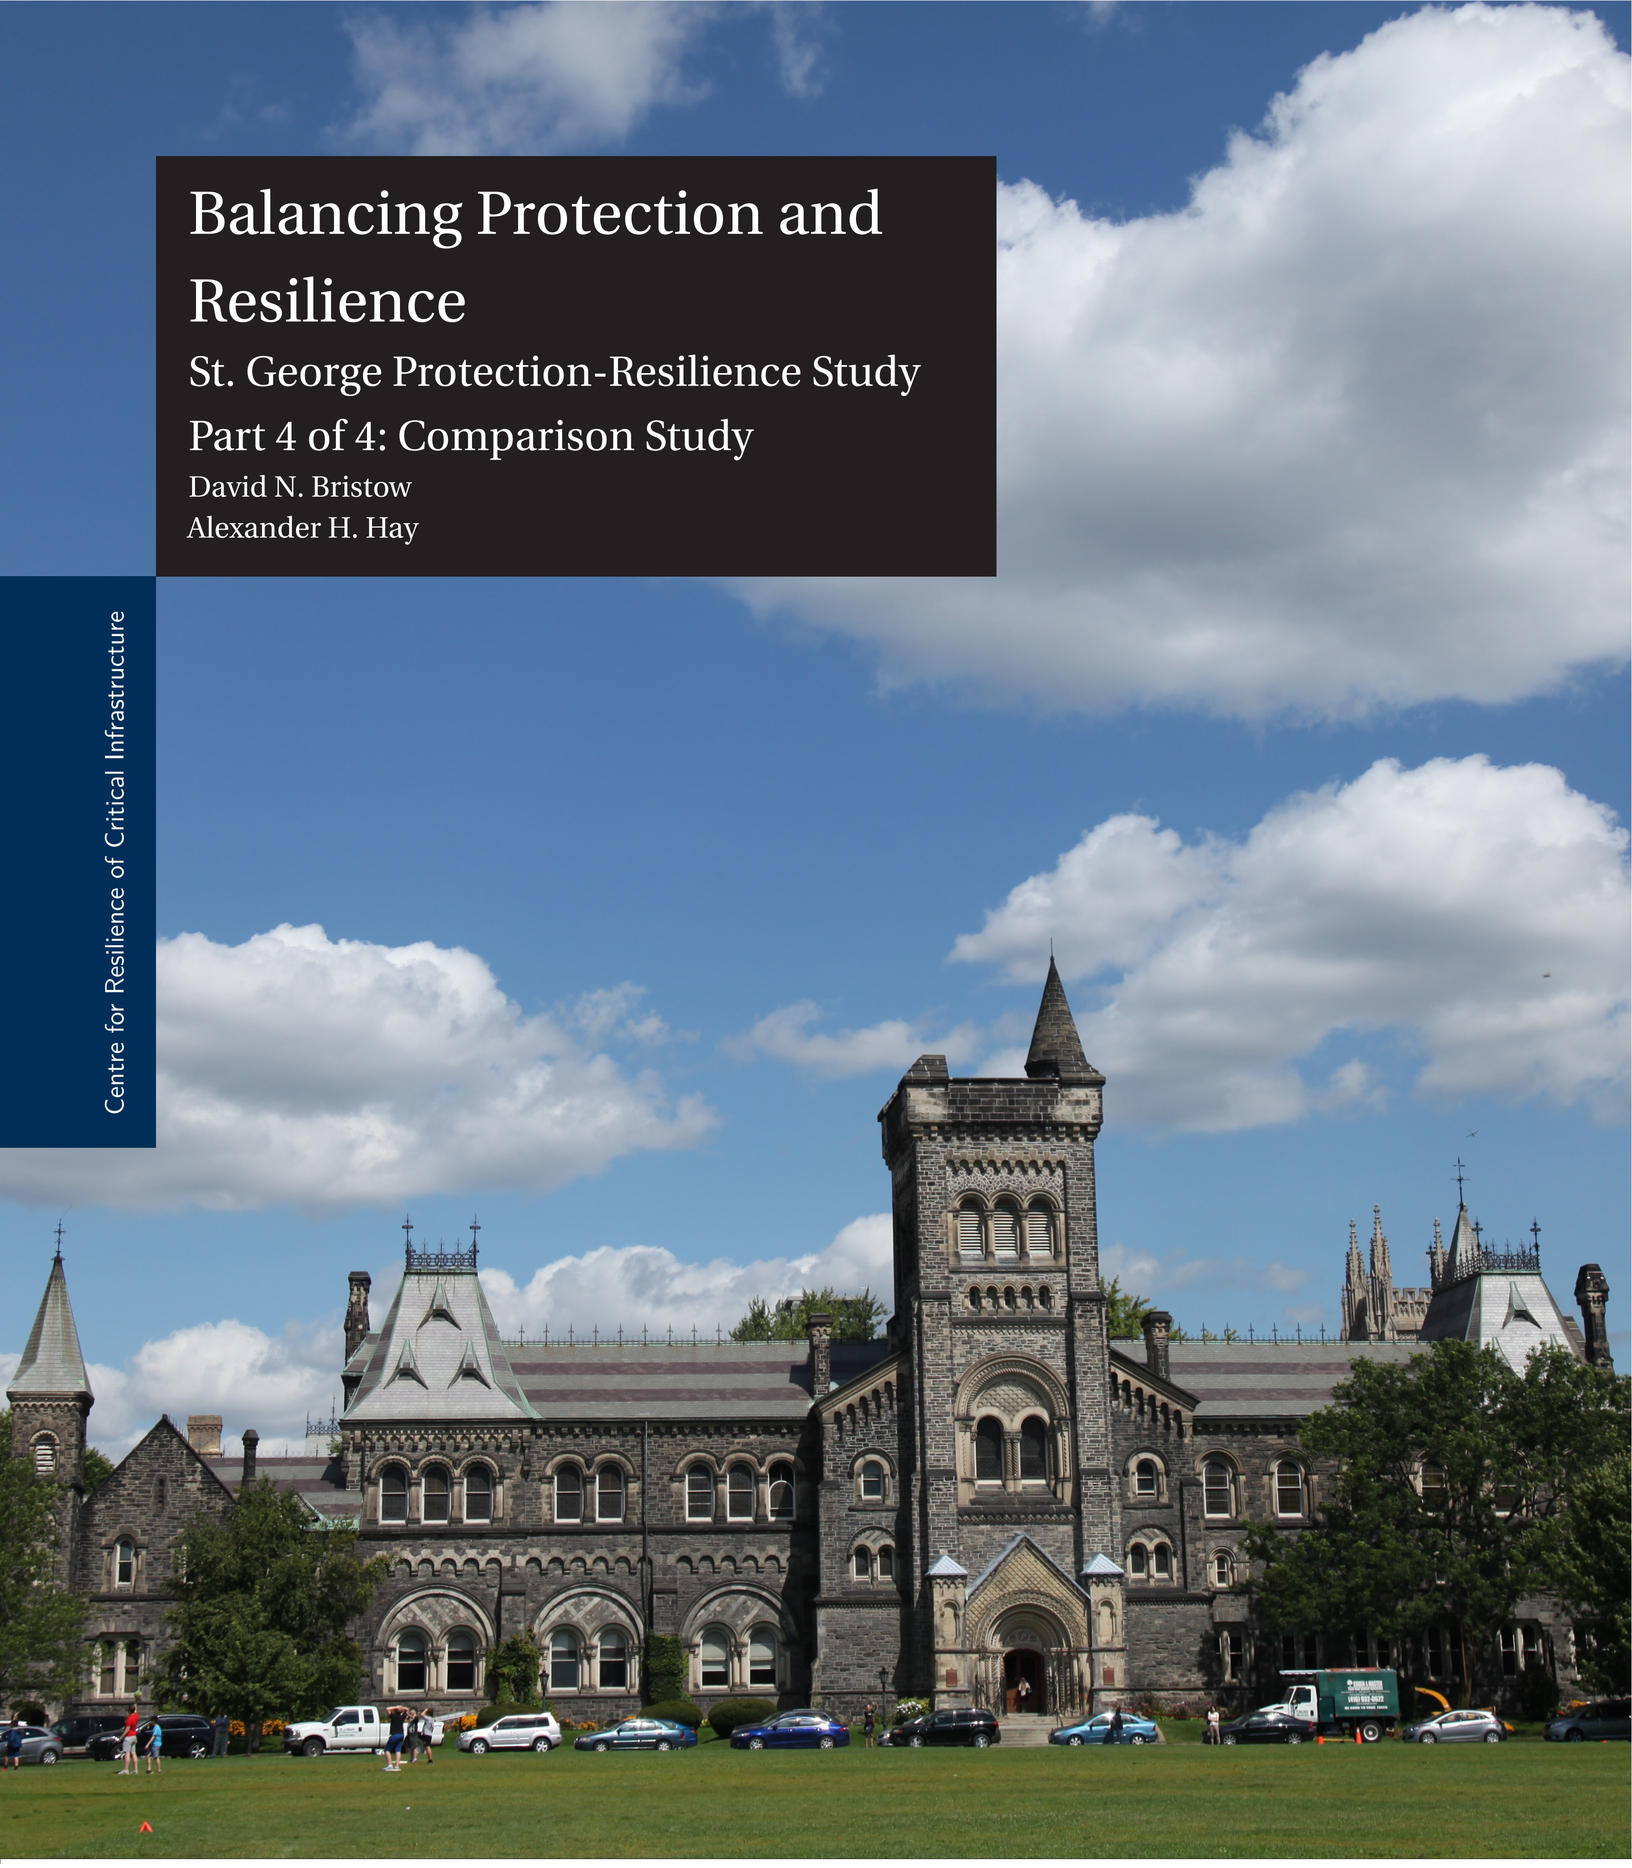

UNIVERSITY OF TORONTO

BALANCING PROTECTION AND RESILIENCE

ST. GEORGE

PROTECTION-RESILIENCE STUDY
PART 4 OF 4: COMPARISON STUDY

by

David N. Bristow
Alexander H. Hay

Centre for Resilience of Critical Infrastructure

This project was graciously funded by the University of Toronto.

An electronic version of this report is available at
http://www.crci.utoronto.ca/.

This report was prepared by David Bristow and Alexander Hay of the Centre for Resilience of Critical Infrastructure at the University of Toronto and is informed by the study completed with guidance from David Black, Coordinator, Emergency Response Planning at the University of Toronto; and Duresameen Ashraf, Coordinator of Public Safety at George Brown College. Michael Munroe, Gary Borges and Ryan Dow provided their survey services. The study included interviews with several campus and utility staff. The material herein reflects our best judgement based on the information available. Outcomes resulting from decisions made based on this report are the responsibility of the decision maker.

UNIVERSITY OF TORONTO

Executive Summary

The Centre for Resilience of Critical Infrastructure was tasked to conduct a comparison of a full protection assessment and a full resilience assessment of the same operation, the St. George campus of the University of Toronto, in order to determine the investment tradeoffs between protection and resilience. Using these two studies the general parameters of operations relevant to achieving a balance of protection and resilience are extracted and formulated into a framework methodology. This report, constituting the detailing of this novel methodology is the final stage of the study.

The study was conducted by The Centre for Resilience of Critical Infrastructure in collaboration with Campus Community Police. The study area comprised of 110 buildings of the St. George campus, excluding federated colleges. The effort was conducted over the course of Fall 2013 through Spring 2014 and included background research, protection and resilience surveys of the buildings and operation, and the full assessments and analysis. The study phases are summarized in the following table.

Phase	Phase Steps		
1	Project definition	Client engagement	Project plan buy-off
2	Mission analysis	All-Hazards analysis	Spatial data collection
3	Survey definition	Fall protection surveys	Fall resilience assessment
4	Fall review	Winter protection surveys	Winter resilience surveys
5	Winter review	Analysis and comparison	Reporting

The reporting is comprised of four parts. Parts 1 to 3, submitted in summer 2014 to the University's coordinator for emergency preparedness, comprising of the contextual overview, the protection assessment, and the resilience assessment, respectively. These parts are confidential and are the sole and express property of University of Toronto Emergency Response Planning due to the exposed vulnerabilities outlined therein. Part 4 (this report) is the comparison methodology and is the subject of public release and is therefore generic and unattributable in nature. A trimmed version of this report is being submitted for publication in a peer-reviewed journal.

Documented herein is a solution to the challenge of balancing protection and resilience. The difficulty to solving this challenge lies in capturing both the hazard risk and the operational risk of interdependent infrastructure systems. This paper provides a generic framework for determining the optimal balance of protection and resilience as a multi-objective, multi-criteria decision problem based on the risk criteria of an operation through the specification of the operation's design requirements for resilience. The method developed here expands on a graph theoretic approach to assessing interdependency risk to allow for the assessment of protection and resilience measures and provides an heuristic, operation-wide method to reduce the search space of possible combinations of measures. This approach is then extended into a new campus infrastructure planning framework that optimizes the protection-resilience balance with sustainability. Successive applications of this planning approach can generate positive impacts for the campus and hence strengthen the cash flow and is hereby recommended as a campus infrastructure planning tool.

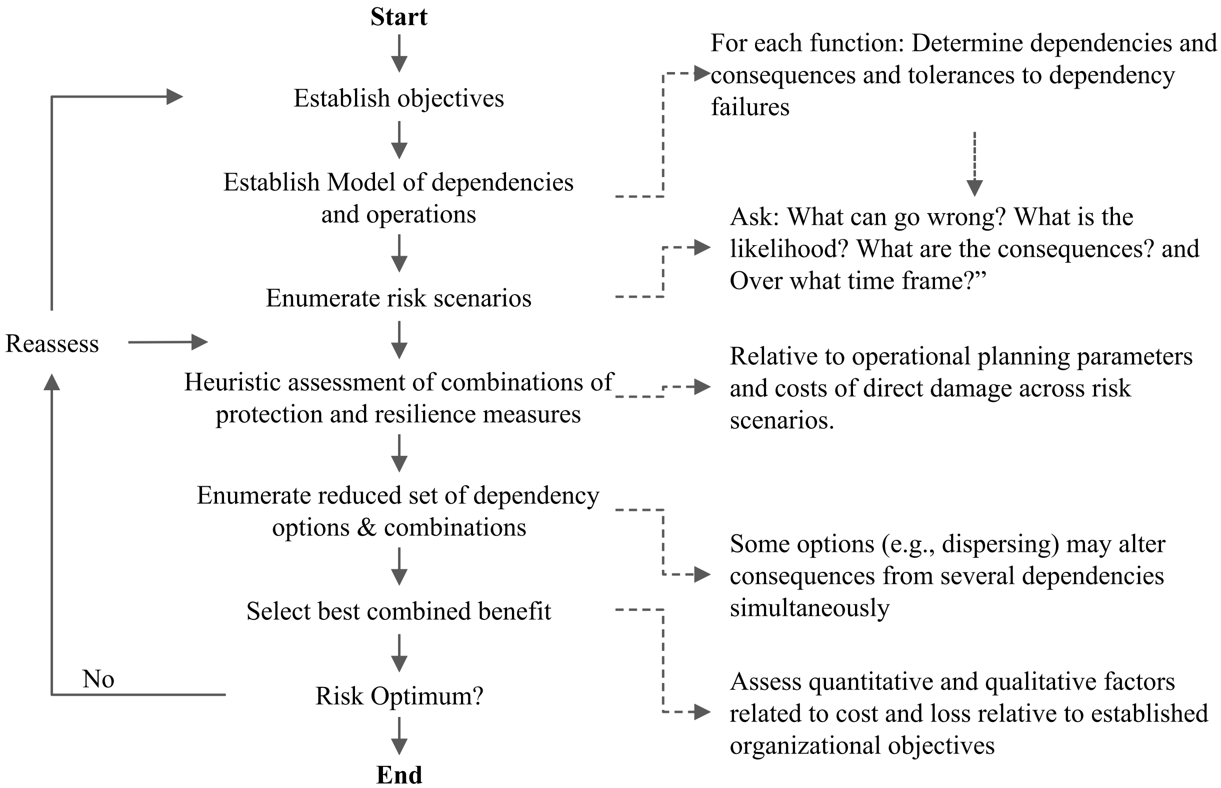

Protection-resilience balancing framework

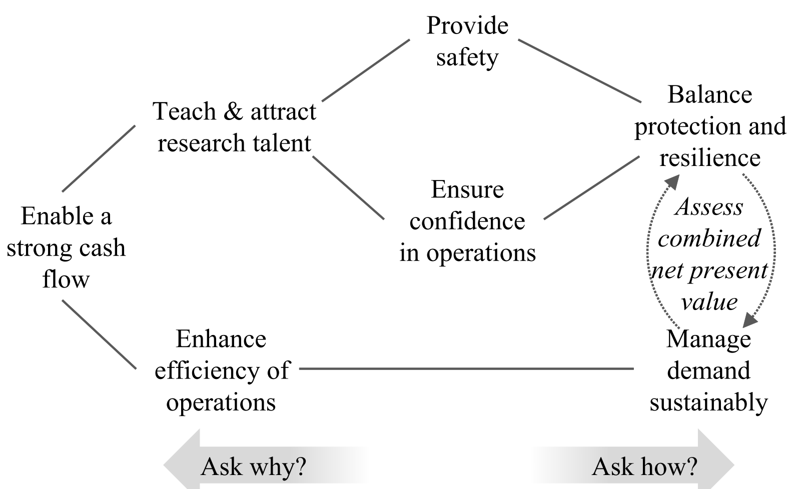

Campus infrastructure development planning framework

Contents

List of Figures

List of Tables

1

INTRODUCTION

Notwithstanding that assessing the value of infrastructure is nuanced and challenging, requiring consideration of different measures (Abouchar 1977), the tight coupling between infrastructure investment and economic activity is commonly observable, if also quite an involved process (Aschauer 1989; Gillen 1996; Kennedy 2011). Without the value delivered by infrastructure economies would not be able to perform as they do. It is therefore of little surprise to learn that the global expenditure on critical infrastructure protection values in the billions and is expected to swell from a current value of $64B to $106B by 2018 (MarketsAndMarkets 2014). There are many important reasons why we wish to protect our infrastructure: to protect life, assets and investments, organizations and communities, even the nation as a whole. Implicit in these reasons is the securing of the resources that enable our future and ensuring our means to sustain and modify our operations to achieve a desired future. We protect a facility that we perceive as enabling our operations in the future. Resilience planning, on the other hand, is more directly concerned with returning operations to desired performance levels following a shock or stress. The evolution of operations that enable the future will have changing demands upon those same facilities. The resilience parameters used, therefore, focus on the operation itself as the enabler of future capability. The application of protection and resilience thinking to infrastructure, it turns out, have different histories, and vastly different outcomes.

Protection risk assessments, especially those where the methodological lineage stems from security risk assessments of predominately malicious threats, commonly concern vulnerabilities and probability distributions of damages to infrastructure, facilities, and life. This sort of assessment usually falls under a fail-safe mindset (Ahern 2011; Ezell et al. 2000; Louis Lebel et al. 2010). Under the protection fail-safe mindset there is negligible consequence for common, low severity scenarios. Subsequently, at a threshold ideally based on a full risk assessment, and not on the funds available – as can often be the case, there is a large increase in consequence. The growing literature on infrastructure resilience assessment on the other-hand commonly examines performance over time (Cimellaro et al. 2010; Cimellaro et al. 2009; Ouyang et al. 2012) to assess both impact to operations and the form of the recovery. This type of assessment takes a safe-to-fail perspective, aiming for operational continuity. The safe-to-fail resilience mindset is partly motivated by situations where history proves a bad predictor of the future. In such cases there is great sensitivity of a selected protection level to the accuracy of the probability estimates of the high severity events (Taleb 2012). Both protection and resilience assessment benefit from standard risk evaluation methodologies. The concept of dependency or interdependency mapping and analysis of interconnected infrastructure systems is increasingly used to not only assess hazards themselves, but the cascade of hazard effect (Casalicchio & Galli 2008; Haimes & Jiang 2001; Macaulay 2008; O'Neill 2013; Ouyang 2014). Partly due to the recognition of the role of dependency and infrastructure risk, the interplay between protection and resilience thinking is rapidly evolving.

In practical terms, protection and resilience influence one another and typically present in some balance in both protection and resilience assessments and plans. Protection, such as a flood wall, can lessen the drop in performance when a risk is realized, while quicker recovery through resilience planning can improve the rate of return of operations to pre-shock levels, or better, by augmenting the existing risk management to also consider the temporal aspect of functionality (Linkov et al. 2014). Balancing protection and resilience can be more efficient than focusing solely on protection (Haimes et al. 2008). The US Government has recognized this tight coupling in their official critical infrastructure policy (DHS 2013). Indeed, despite the varying traditions of thoughts on the matter, protection and resilience assessment, along with dependency risk, are merging into a unified understanding of infrastructure risk (Hay 2013). From a decision making and investment point of view however, there is little guidance on how to optimize the balance of protection and resilience measures and it has been previous argued that new research is required to further explore how this balance can be achieved (Haimes et al. 2008).

The objective of this paper is to provide a means to locate this optimal balance in order to deliver maximum value to an operation, comprised of organization, people and infrastructure, given a set of operational parameters. The proposed approach is first properly embedded within the general language of risk management to establish the context of the generalized formulation to follow. Understanding of this step is critical, as application of the method is not through a template fashion but through a first principles approach to risk management. This is followed by a review and then an extension of dependency risk methodologies. Hereafter a mathematical formulation of the approach is provided for straightforward, single function cases. Building on this foundation the application to multi-function operation-wide protection and resilience planning are described. At this point the operation-wide value of protection and resilience are discussed and a means to determine their optimal balancing point revealed in terms of a quantitative and qualitative multi-objective and multi-criteria decision problem.

The formulation presented herein is based on a comparison study of the St. George campus at the University of Toronto, though all attributable details are omitted. The study consisted of separate protection and resilience assessments of 110 buildings (see Figure 1) and supporting infrastructure against the same risk context as established through an all-hazards approach. The phases of the project are listed in Table 1. With this project structure it was possible to assess how a change in operations alone, encompassing the infrastructure, functions and risk criteria associated with each building, altered the cost of each of a protection and a resilience approach and the relative value to the operation as a whole. The approach herein proposed is hence generic to other operations, yet concerns specifically hazard and operational risk. The implications to other types of risk, such as speculative and strategic are briefly discussed.

Table 1: Study phases and tasks

Phase	Phase Steps		
1	Project definition	Client engagement	Project plan buy-off
2	Mission analysis	All-Hazards analysis	Spatial data collection
3	Survey definition	Fall protection surveys	Fall resilience assessment
4	Fall review	Winter protection surveys	Winter resilience surveys
5	Winter review	Analysis and comparison	Reporting

This document represents the final report of the project and hence concludes the project. The previous three report parts were delivered to the coordinator of emergency preparedness at the University of Toronto in Summer 2014 and detailed the contextual overview, the protection study and the resilience study. The sensitive nature of these parts necessitated confidential filing. This document in detailing the resultant generic method of protection and resiliency comparison is non attributable and released publicly. A trimmed version of this report is being submitted for publication in a peer-reviewed journal.

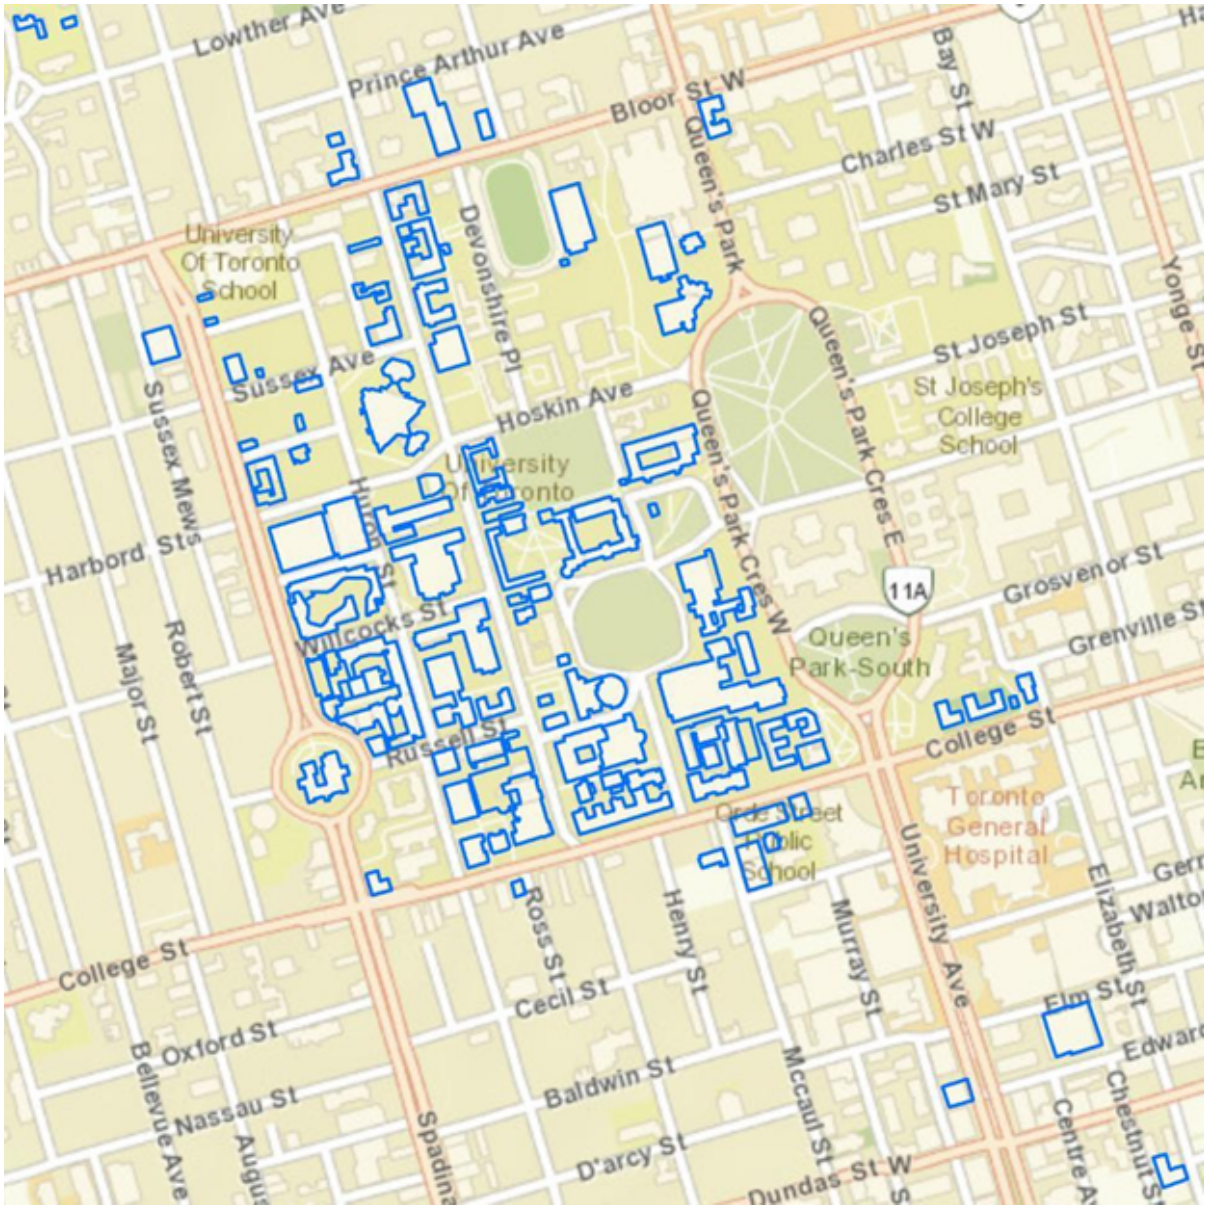

Figure 1: Study area buildings (bold outline – 110 altogether)

2

PROTECTION AND RESILIENCE PLANNING IN RISK MANAGEMENT

Quantitatively, risk is often assessed in the form of triplets such that for a set of scenarios ($i \in 1 \dots N$) the risk is $R = \langle s_i, p_i, \mathbf{x}_i \rangle$, where s is a scenario describing what can go wrong, p is the probability and $\mathbf{x}$ is the consequence (Kaplan & Garrick 1981) taking the form of a scalar (x) or a vector ($\mathbf{x}$). The vector form could for example include both hazard and operational consequence as separate values to allow for a multi-criteria assessment. For convenience the scenario subscript, i, is dropped at times when it is clear that a specific scenario is the subject matter. More generally, the Kaplan and Garrick expression can be extended to include time; doing so extends the questions of risk to: "What can go wrong? What is the likelihood? What are the consequences? and Over what time frame?" (Haimes 2009). This reformulated definition of risk is formally stated here as $R = \langle s_i, p_i, \mathbf{x}_i; \mathbf{t}_i \rangle$, where $\mathbf{t}_i = t_{i,0}, t_{i,1} \dots t_{i,n}$, and $t_{i,0}$ represents the onset of a realized hazard, and $t_{i,n}$ generically represents the end point of an analysis of the realized hazard in an $n+1$ stage evolution of the event that ideally culminates by this point or sooner with a complete recovery and possibly an improved overall performance. When assessing what can be done about risk from a resilience planning standpoint the aspect of time is particularly crucial as it is imperative to understand when things fail, and how long a failure is acceptable (Bristow et al. 2014; Hay 2013).

Assessing risk options can proceed from the time expanded definition of risk. The residual risk of a potential risk avoidance, mitigation or transfer option can be assessed using the time expanded formulation of risk consequence as $\Delta \mathbf{x}_i = \mathbf{x}_i^* - \mathbf{x}_i$, where * represents the consequence with the application of an option. A reduction in risk represents a benefit. Together with costs of implementation, including capital and operation and maintenance implications alternatives can be compared. It is however, more generally appropriate to disaggregate risk into a multi-criteria problem with quantitative and qualitative aspects for the purposes of comparison and for different ends of the severity spectrum as per the partitioned multi-objective risk method (Haimes 2011, Chapter 8).

When incorporating the time aspect, event graphs (also dubbed incident profiles or performance curves) are central to understanding the effect to operations of a shock upon a system and are quite central to the discussion of resilience of societal or engineered systems (see for example Bruneau et al. 2003; Linkov et al. 2014; Ouyang et al. 2012; Vugrin et al. 2010). For the present purpose it is necessary to specify several key planning measures on these graphs as illustrated generically in Figure 2. Subsequent to an event causing a loss of performance of an operation at a time to, there is a reaction to return operations to a minimum operating performance (PM) at which point a more situation-specific response engages to bring the performance up to a minimum sustainable level (P_s) in a pre-determined time ($\Delta t_s = t_2 - t_0$). In private sector terms, PS at t_2 is the point at which you will not go out of business, whereas P_M at t_1 is the point at which the essential functions of the operation can survive and are at a significantly lowered risk of fail-

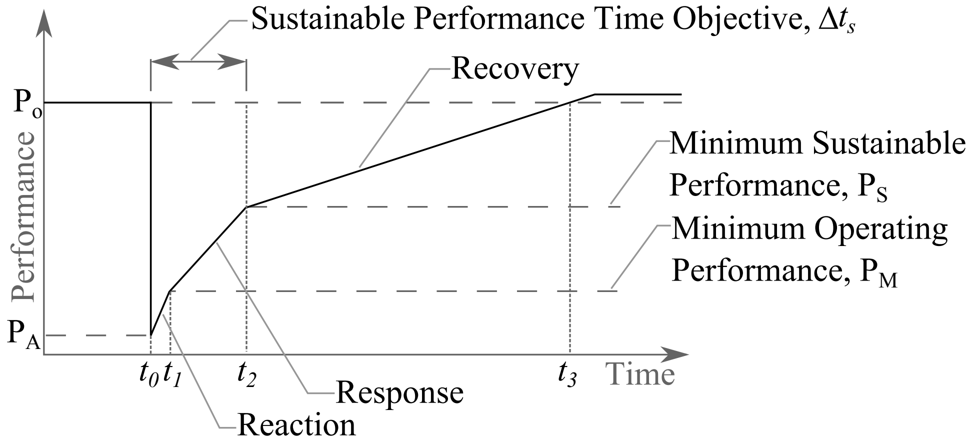

Figure 2: Generic event graph. Performance is often measured in percentage terms)

ure. Once P_S is achieved, the recovery proceeds to return the operation to the pre-incident levels, or above, as is often the case since pre-existing inefficiencies are discovered and subsequently addressed. The event graph is useful for constructing plans to achieve such a strong recovery.

Indeed, determining an operation's minimum operating performance, minimum sustainable performance and sustainable performance time objective establishes the planning objectives. The minimum operating performance is the level at which the operation can minimally function without causing additional damage or system failure. This can be akin to re-opening a facility or plant, or restarting a process at the lowest acceptable level of output. The minimum sustainable performance and sustainable time objective are established concurrently and represent the organization or holistic tolerance of loss or interruption. This proceeds by identifying the functions critical to a sustainable level of performance and by identifying how long the operation can tolerate their operation below the sustainable level. The tolerance can be seconds in the case of a critical infrastructure or health care centre, and up to a day or more for general service firms. Notably, failure to succeed in establishing an appropriate plan based on these objectives can result in a far more protracted recovery, or none at all. This can be the case when protection-only planning is utilized and the threshold of the protection is exceeded by a shock, as occurred in the case of Hurricane Katrina, Superstorm Sandy and the Fukushima Tsunami (Gibbs & Holloway 2013, p.13; Seed et al. 2006, pp.15–1; World Nuclear Association 2014).

The functions across a campus vary in terms of their desirable time to achievement of minimum operating and minimum sustainable performance. The sustainable time performance objectives, however, can lack some specificity if, as is the case here, they are based on the business continuity plan which, at the time of writing, only groups the most time critical functions into those that must be active within 30 days of a disaster. Hence, these are not plotted in Figure 3, which does however depict the variation in the desirable time to achieve minimum operating performance. Recall this measure is based on the level of performance at which the risk of further damage is unlikely. It is hence possible to extract the time objective for this variable from the tolerances of the various functions on campus to different dependency failures, such as utilities. These tolerances vary seasonally, and depending on the existing infrastructure layout, can be overridden by a tighter tolerance in the case where attending to the tighter tolerance resolves an issue with a looser tolerance. In general the trend of Figure 3 depicts how a varied reaction is desirable and hence planning around this reaction and any necessary infrastructure changes can place emphasis on the tolerances that are most difficult to achieve. Difficulties arise often, though not always, in the cases where the tolerance is tighter. Cases where this may not be true is where there is a loose tolerance, but addressing the failure is unique or challenging enough that achieving this loose tolerance remains relatively difficult.

At this stage it is necessary to review specific types of changes that are made when these planning approaches

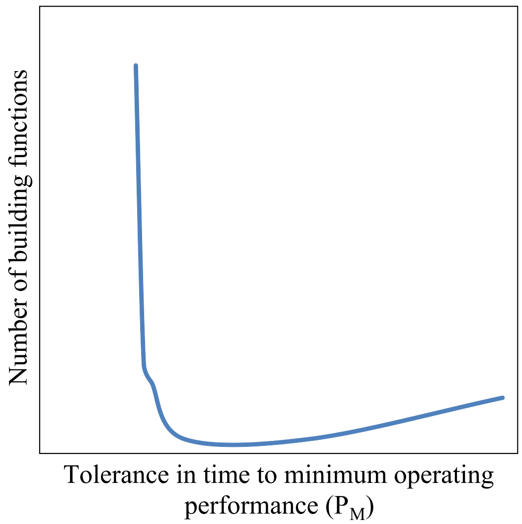

Figure 3: Variation in time to minimum operating performance across functions in campus buildings. The names and location of functions, along with axis values are omitted to ensure the trends are non-attributable.

are applied. These specific measures include (at a minimum): dispersing, hardening, adding flexibility, adding diversity and adding redundancy. Hardening is a typical result of protection assessment, whereby infrastructure is physically strengthened against an identified hazard. Hardening can dampen the operation's susceptibility to shocks and can limit the reduction in operating performance when the protection threshold is exceeded. The remaining techniques are more often associated with resilience planning and are summarized alongside hardening in Table 2.

In protection planning the upper threshold for hardening is ideally based on a thorough risk assessment, but due to high cost is often set lower. Generally, the goal from a resilience planning perspective, conversely, is to identify how a series of measures can be used in tandem to enhance confidence of operational continuity regardless of the magnitude of event. Alternatively an upper bound set to the lesser of either the decision makers' maximum acceptable risk

Table 2: Overview of generic WHEDR measures

Treatment	Overview
W. Disperse	Spreading out functions to a wider area means each of the activities of the operations in those areas have different dependency maps, each of which must be assed each in regards to WHEDR *e.g., Creating satellite campuses.*
H. Harden	Reduce likelihood of failure of the dependency. Note, a hardening measure, on its own, can be a dependency *e.g., Add bollards or street furniture.*
E. Add Flexibility	Provide elasticity by adjusting dependencies to serve multiple and overlapping functions. *e.g., Enable internet to serve telephony needs and enable telephony to serve critical internet needs.*
D. Diversify	If a dependency a is duplicated with b, then b and its unique dependencies are added to the risk context. The consequence of failure of a is reduced subject to the benefits of b. *e.g., Add backup power from a different source.*
R. Add Redundancy	If a dependency a is made redundant by adding dependency b then there is no consequence upon failure of any one of the redundant dependencies, but they could both fail simultaneously as they themselves share dependencies. *e.g., Add a second power or water line from a required resource. If the second line is from a sufficiently differentiated portion of the upstream infrastructure then this strategy becomes a diversification (D) solution.*

(MAR), the level at which the infrastructure or facility would be relocated or closed, or the maximum realizable threat (MRT). The result of this approach is that risk partitioned into this upper region of severity is considered so unlikely that no measures are applied in advance and anything below this threshold receives resilience planning measures designed to ensure continuity of operations (Hay 2013). Of course, appropriate sensitivity analysis is required on this threshold to ensure risk exposure is acceptable (Taleb 2012). The key to this approach is also in understanding the duration of severe events. Event graphs are presented in more detail in Figure 4 against loss curves to demonstrate the impact of protection and resilience planning with this heuristic approach. The wider graphs in the middle column are for a given event severity and duration, while the graphs towards the margins illustrate how the total loss from an event would vary with duration and severity given the type of risk management (none, protection, resilience, or a combination). In the base case (a), the absence of any meaningful protection or resilience means the performance of the operation drops to nothing and requires a very long time to return to the pre-event performance (P_o). As the event severity graph shows, for an event of a given duration the total loss is theoretically fixed since events of all severities cause complete failure in performance. As the duration of the event is increased, however, the total loss increases from ∘ to • since the longer event inhibits the start of any recovery. Now, this base case abstraction is never realized due to building codes, compliance requirements, etc. that ensure a basic level of protection. The rules governing these basic levels, however, can be applied inconsistently across all components, achieving more risk reduction and value in some areas compared to others, where the protection can even interfere with the operation. This basic level of protection varies across infrastructure and circumstance.

Part of the difficulty in understanding the protection-resilience interplay is in the connected or interdependent nature of infrastructure. Defined broadly, infrastructure includes physically designed systems, organizations and human resources, all which interact in an emergent fashion (Haimes et al. 2008). Within the complete set of risk scenarios are those resulting from all of the possible cascades of dependency failures. On the surface this seems a sufficiently large set of scenarios so as to be unmanageable, however, O'Neill (2013) provides a powerful method based on graph theory to capture the suite of possibilities. The approach proceeds as follows: a graph comprised of nodes and edges, where the edges represent the existence of dependency relationship between nodes, and the nodes represent components or entities of an infrastructure system,

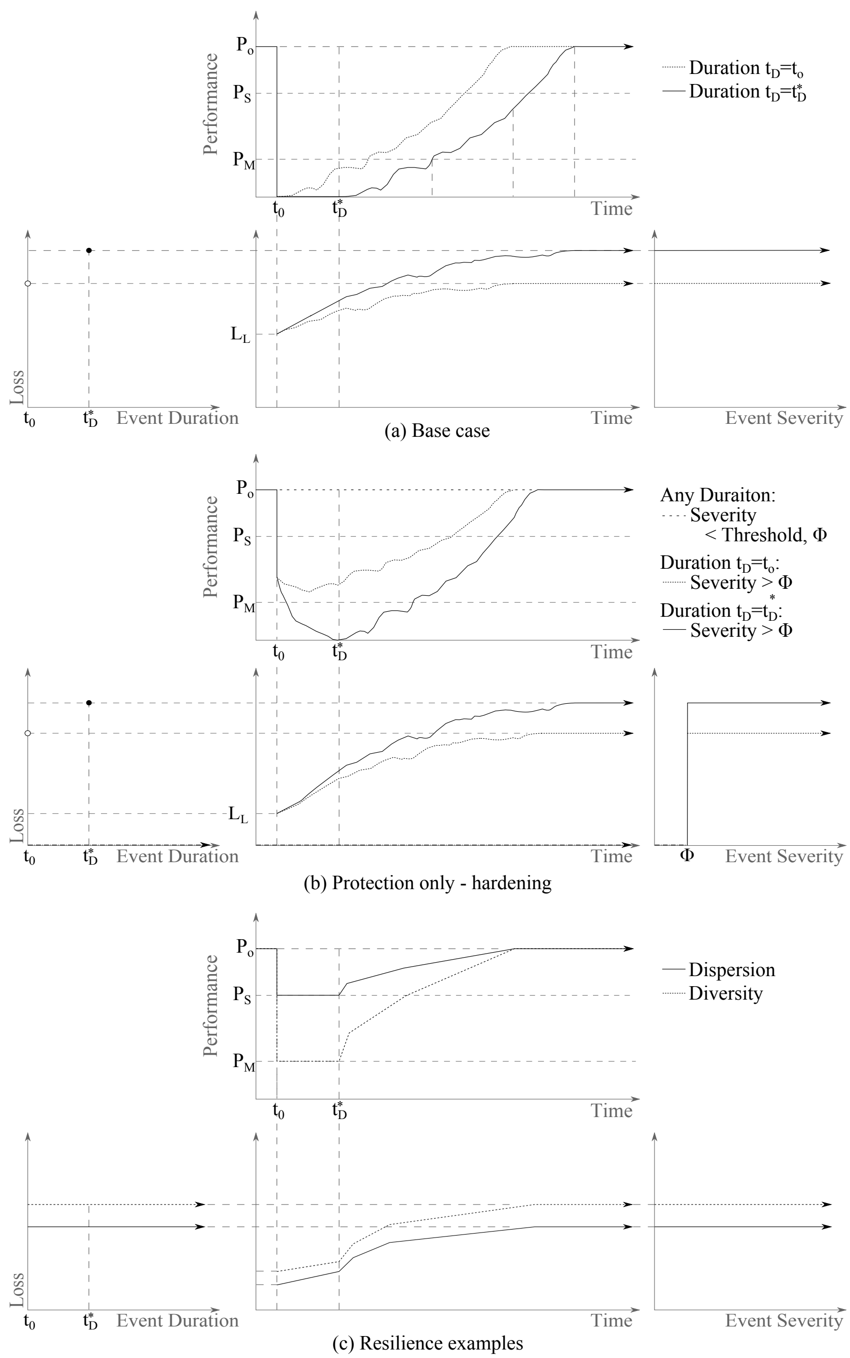

Figure 4: Performance and cumulative loss for generic dependency protection and resilience measures compared to how changing event duration and severity impact loss under the MAR/MRT threshold. The noise in the base case and protection only performance and loss curves illustrates the anecdotal evidence of the tenuous and difficult nature of recovery in these cases. ∘ and • represent losses due to different durations

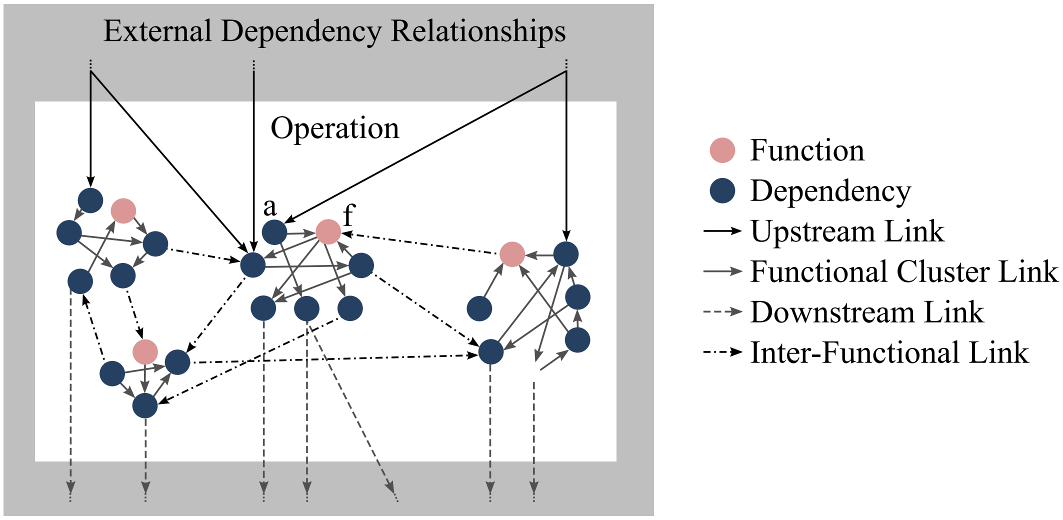

Figure 5: Schematic of functional dependency maps of an operation within a broader external set of dependencies. The arrowhead indicates that that node is dependent on the originating node (In the above, j depends on i).

is compiled for a system under study. Dependencies are assigned likelihoods of failure and degree of direct consequence on the downstream node upon this failure. O'Neill then provides an aggregation methodology for assessing system-wide global risk and system-wide risk associated with a given node whereby the propagated consequence is the lowest consequence along a path between two nodes. Subsequently the outcome of a change to the network can be tested non-destructively in simulation and the risk benefit of the change assessed. For practical purposes this would be accomplished within a sub-system of a larger global dependency network as illustrated in Figure 5. Since O'Neill's publication the approach has been extended to incorporate indirect consequences (such as political impacts); both qualitative and quantitative factors; and the specification of event severity threshold triggers (personal communication); though even further changes, as discussed next, are required for the assessment of the protection-resilience tradeoff.

3

Protection-resilience assessment of dependencies

The generic representation of risk and dependencies presented above is formulated here into a means to assess the balance of protection and resilience by first extending it to include the specifics of the WHEDR measures, such that their pairing can then be assessed to determine the optimal protection-resilience balance. The extension of the dependency methodology is presented schematically in Figure 6. In each case there is an additional cost (ΔC) associated with the measure to remove or lower losses within a range of risk magnitudes and probabilities. Schematically, with the arbitrary axes, the difference in terms of loss between hardening and the other measures is that the hardening achieves zero losses within the threshold design range.

The losses associated with each WHEDR measure are mathematically represented here in purely quantitative terms, but as is the case with O'Neill's original formulations the methodology can easily be extended to include qualitative factors through nominal or ordinal representation. We commence by restructuring how consequence is recorded in the time expanded definition of risk introduced in the previous section. Recall that at each point in time following an event the consequence can be measured in terms of $j \in 1...l$ variables. These vectors of consequence can be compiled into a matrix for each scenario (X_i) such that the first column is direct losses resulting from an event (at t_0), and the remaining columns include the operational losses from reduced performance at the stages following the event (for $t_k, k] \in 1...n$). Hence, for the i_{th} scenario, and the j_{th} measure of consequence, the consequence at the k_{th} time interval is recorded as:

$$x_{i,j,k}^{o,a\to f} = \begin{cases} l_{i,j,0}^{o,a\to f} & k = 0 \\ l_{i,j,k}^{o,a\to f} = \int_{t_k^{o,a\to f}}^{t_{k-1}^{o,a\to f}} P_o - P_{i,j,k}^{o,a\to f}(t)\mathrm{d}t & k > 0 \end{cases}$$

The matrix $\mathbf{X}_i$ for the consequence on a function f relative to an effect on dependency a when no additional measure is applied ($M = o$, the base case), and for the $i_t h$ scenario then takes the form:

$$\mathbf{X}_i^{M=o,a\to f} = \begin{bmatrix} x_{i,1,0}^{o,a\to f} & \cdots & x_{i,1,n_i^{o,a\to f}}^{o,a\to f} \\ \vdots & \ddots & \vdots \\ x_{i,l,0}^{o,a\to f} & \cdots & x_{i,l,n_i^{o,a\to f}}^{o,a\to f} \end{bmatrix}$$

and, for those interested in reformulating the risk in terms of triplets, the risk becomes $R = \langle s_i, p_i, \mathbf{X}_i \rangle$. For a generic application of dispersion, $M = W$, the format is similar, though, the value of the losses change. The direct Losses, $l_{i,.,0}^{W,a\to f}$, are generally less than in the base case due to the geographically separated nature of a dispersed approach. If however, the i_{th} scenario represents an event that impacts both f and f' then the direct losses can be even greater than the base case since there is an additional cost of implementation (ΔC from Figure 6) that can be lost through damage caused by the event. Heuristically, dispersion should be designed such that performance never drops below the min-

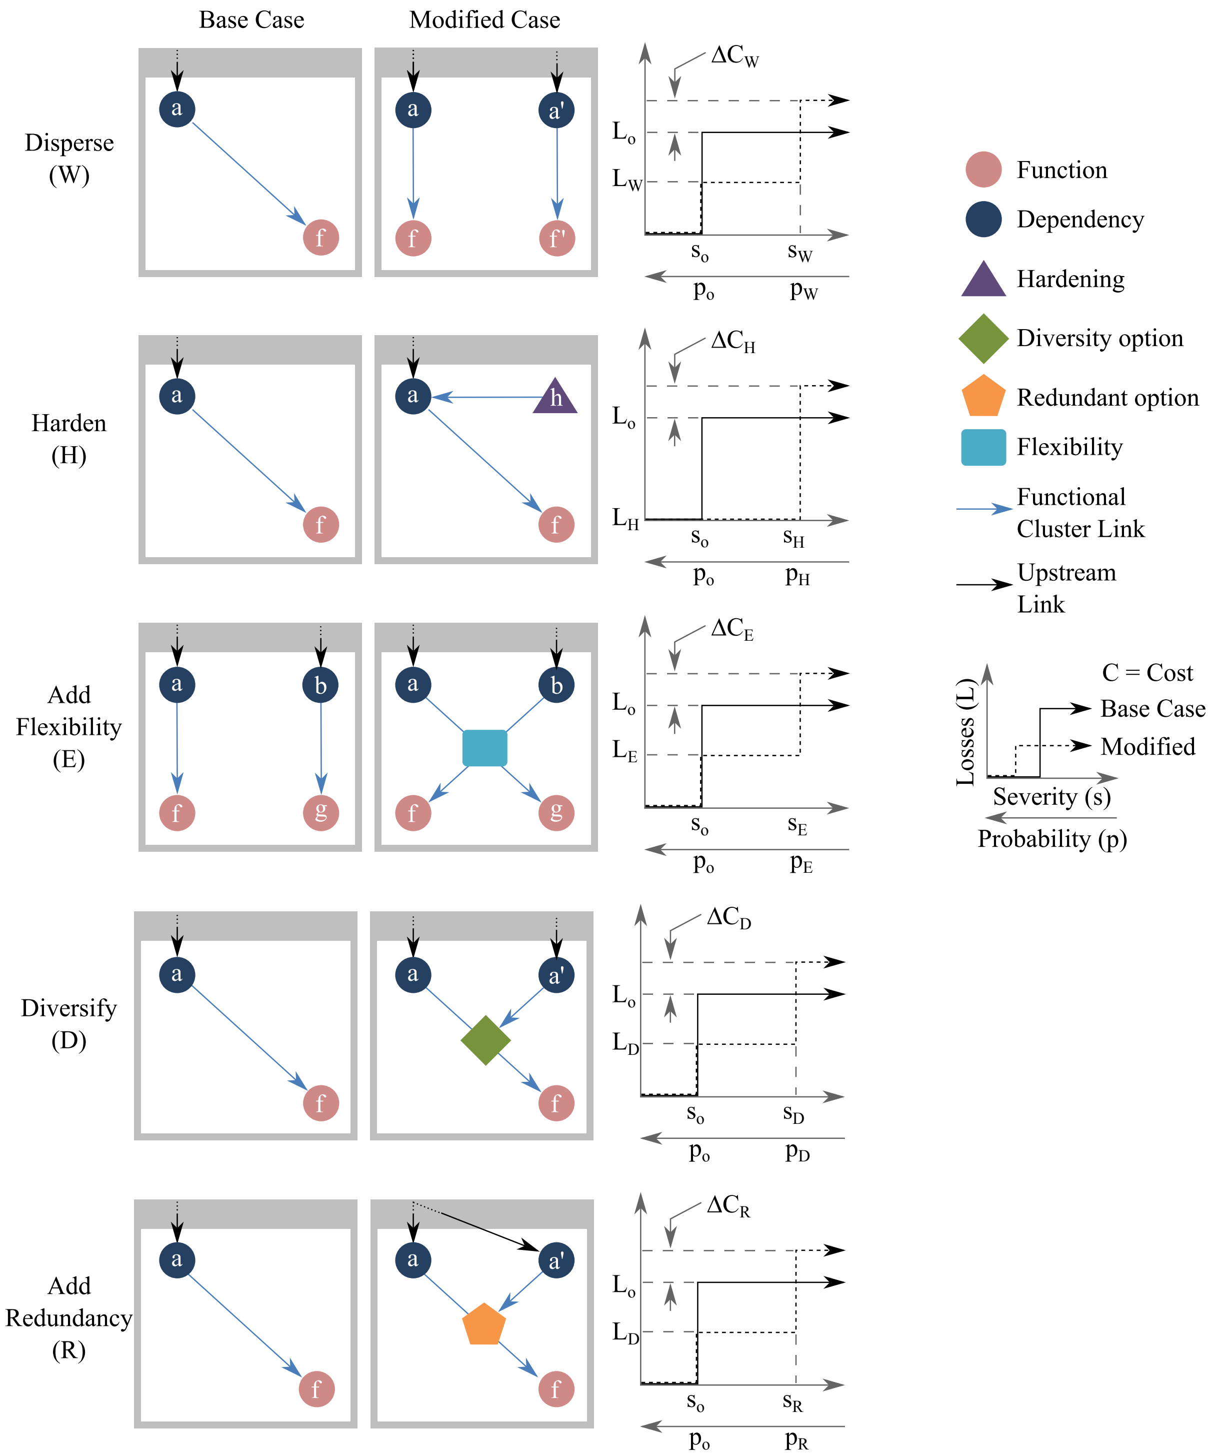

Figure 6: Dependency map reformulation for WHEDR treatments with partitioned loss curves (probabilities, scenarios, losses and costs not to scale as they vary with specific risk criteria and investment). C is the extra cost associated with the measure, a cost that is lost if the threshold of the measure is exceeded.

imal sustainable level (P_S), although P_M can alternatively be used as the planning limit due to the usually high cost of this type of measure, or if setting the limit higher impedes the efficiency of the operation. In any event the operational losses, $l_{i,.,k>0}^{W,a\to f}$ will normally be smaller than in the base case. For the addition of flexibility ($M = E$), redundancy ($M = R$) and duplication ($M = D$), similar loss trends follow as with dispersion depending on the severity and scope of the scenarios, except that the losses are to the entirety of function f since it is not distributed in these measures. Finally, for hardening, the direct losses are usually smaller than the base case, which can reduce the operational losses if sufficient reduction in direct losses is achieved.

Aggregating losses across the entire dependency network for a combination of applications of measures provides a system-wide global assessment. The global assessment of different combinations of measures in terms of losses and costs of implementation constitutes the means by which an appropriate balance of measures, and indeed an appropriate balance of hardening (H) and resilience (WEDR), can be determined as illustrated in Figure 7. This figure is for of a singular function, or multiple unrelated functions. When an entire operation is under consideration, then the step involving establishing the dependency model also includes establishing a model of the operation. This model can take several forms, such as in input-output model (Haimes & Jiang 2001; Macaulay 2008); a control system model (Bristow et al. 2014); a system dynamics or agent based model (Ouyang 2014); or an integration of the dependencies of the functions of the operations on one another within the dependency map itself (City of Toronto 2012). Enumerating all of the possible WHEDR options for a full dependency network, however, is difficult at best for a large dependency network. The Ontario Emergency Management dependency map for example, with 1,300 entities would produce 21,300 different possible new dependency maps assuming only one new WHEDR option per entity (that number is on the

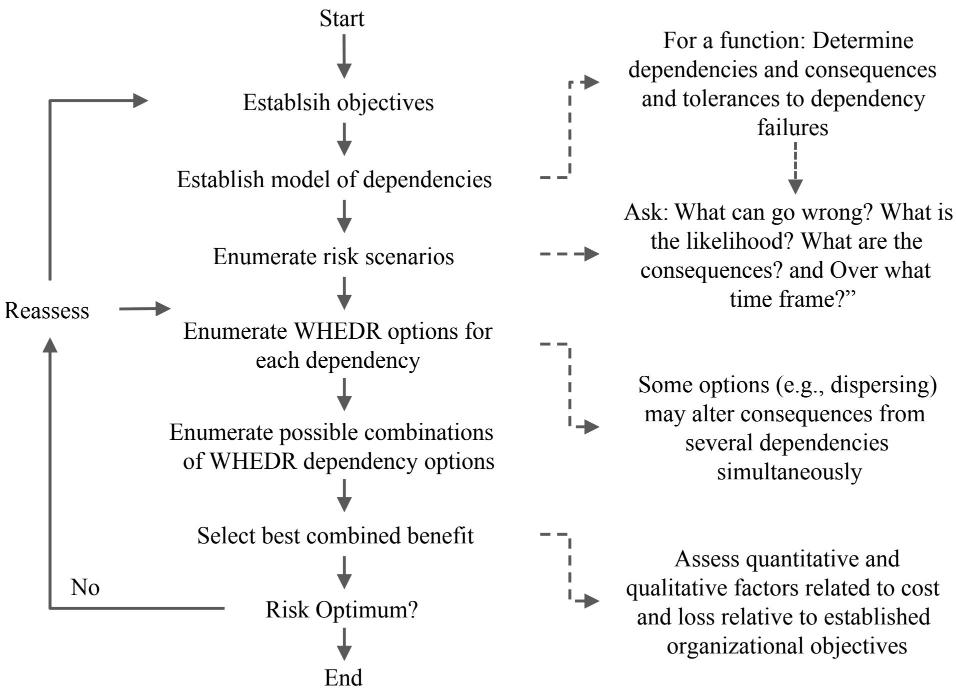

Figure 7: Framework overview for a singular function or independent functions. Evaluation of benefit is subject to risk criteria (in so far as for this study the other factors are fixed).

order of 10 followed by almost 400 zeroes. To provide some meaning to this number, the number of atoms in the observable universe is estimated at something like 10 followed by 80 zeroes). Now, it is possible using O'Neill's method to determine the entities that bring the largest overall risk to a network and hence concentrate efforts on mitigating risk pathways from this much smaller set of entities. However, the mitigation measures may still be applied at the other entities and hence the number of possible maps again scales up quickly in a full assessment. By considering the operation holistically it is possible to greatly reduce the number of combinations that must be considered while trying to achieve a balance of protection and resilience. This is the subject of the following section.

4

Operation-wide protection-resilience balancing

Taking an operation-wide perspective on protection and resilience requires a broader view of event curves. Figure 8 illustrates, generically, how risks can affect the performance of an operation depending on the measures at work, namely protection, resilience or hybrid of the two, and the scale of the realized risk. This scale is divided into two portions: routine and those beyond a critical threshold. In practice this is a useful division, though, more divisions and even probabilities of failures at different thresholds can be employed if necessary (Garcia 2005; Haimes 2011) for increased resolution of analysis. For our purposes in simply differentiating between protection and resilience planning the detail in Figure 8 is sufficient. The area under these curves relative to the initial performance, P_o, is the operational loss, and an area above this level is a benefit. This benefit, occurs in-year and is a form of speculative risk. It is mentioned here and demonstrated in the figure for completeness but is dropped from our formalism for ease and clarity.

The first two sub-graphs illustrate protection only planning. Part (a) illustrates the performance curve when critical asset thresholds are not exceeded, while part (b) illustrates the result when these thresholds are exceeded. The latter indicates the typical initial drop in performance due to direct damage, P_1, followed by a consequential (indirect) drop, P_1' across the dependency network. The consequential loss is loss that resilience planning is concerned with removing. Comparing part (a) to (c) illustrates the differences in routine between a purely protection and a purely resilience based approach. The resilience routine is punctuated by a series of performance dips followed each time by a recovery. To contrast, part (d) illustrates the impact on a resilience only approach when a more extreme event occurs. This figure more clearly indicates the stages that occur in recovery operations. After the initiating event causes a hurried reaction to re-establish operations to the minimum possible level of performance (P_M), at this point a more measured response occurs until a minimum sustainable level of performance is achieved (P_S). From here forward a predetermined sequence of actions are carried out to recover the operation to the initial performance (P_o), or as often occurs to an even better level of performance due to the shedding of inefficiencies. In the hybrid approach illustrated in parts (e) and (f) the minimum operating and sustainable performance levels are used in concert with protection to achieve a modified resilience curve to that seen in either of the individual forms. Part (e), however, may actually look like part (a) depending on the protection threshold selected.

Notably in the hybrid approach the application of protection measures reduces the size of the initial losses compared to the resilience only approach. Due to the partitioning employed in the figure it is convenient to discuss the differences between the application of protection in parts (e) and (f) in terms of two levels of hardening, call them

Φ_α and Φ_β. The benefits of protection in general is the reduced time before response $(t_1 - t_0)$ and the reduced operational losses (area under P_o) in the cases where no dip in performance occurs due to the protection measures. However, the former (Φ_α) type concerns specifically protection of critical assets against residual risk (the consequential risk to an asset due to the catastrophe subject to treatment and is the combination of mitigation of effect and protection of the critical asset, referred to as dampening and hardening respectively). At a minimum Φ_α is comprised of those protections required to meet code and compliance, and can be the result of things such as building codes, industry standards, conventions, or liability concerns. Φ_β is protection above and beyond Φ_α and can achieve the protection of the operation against routine hazards (background risk) or to whatever level deemed desirable. A result of the approach herein for balancing protection and resilience is the locating of this level of additional protection. The cost of this additional protection is the total minus that for code and compliance. Hence, if code and compliance protection is used intelligently, the overall cost of protection, and hence the additional cost of protection up to the desired Φ_β can be lessened.

Compliance protection (Φ_α) can be blunt, achieving excessive protection in some parts of an operation, potentially inhibiting the efficiency of the operation, and achieving insufficient protection elsewhere. The effect of protection of these two varieties is illustrated in Figure 9. In part (a) the total protection is sufficient to protect the entirety of the operation, while in (b) and (c) only a portion is protected, the difference being in (c) that the compliance protection is also exceeded, meaning there was not enough other protection to dampen the effect on these compliance needs.

There tends to be a wide variation in the degree of code and compliance protection (Φ_α) on campuses, especially older campuses due to the changes in construction mate-

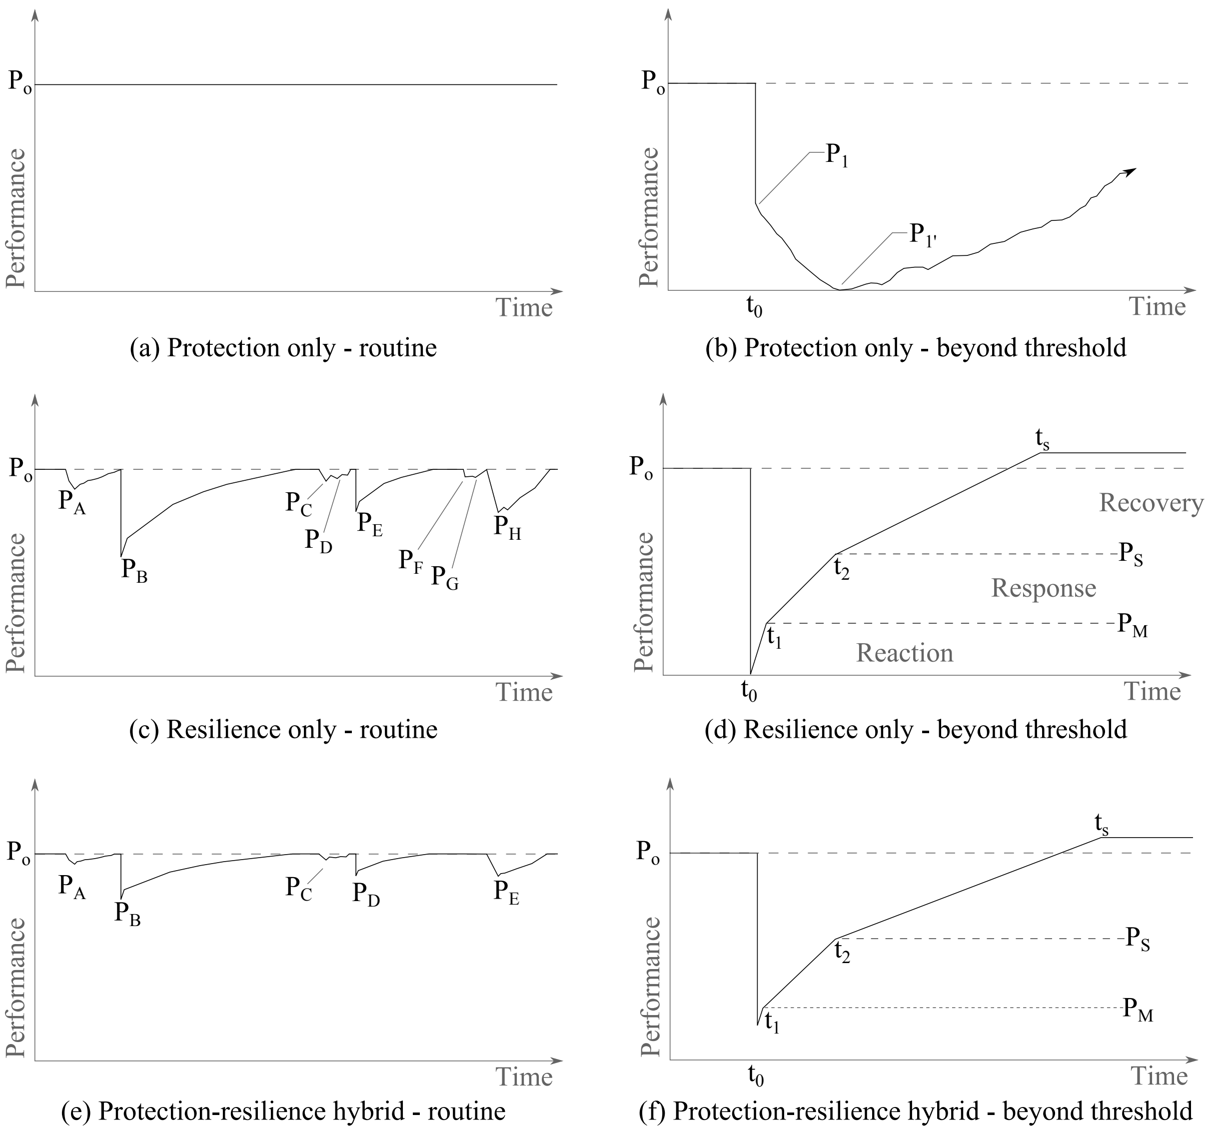

Figure 8: Generic event diagrams for different combination of protection, resilience, an event sizes. Axes are arbitrary. The default level of protection (P_o) represents the point at which there are no losses, areas under this represent cumulative operational losses, areas above are cumulative operational gains. P_M is the minimum operating performance and P_S is the minimum sustainable operating performance.

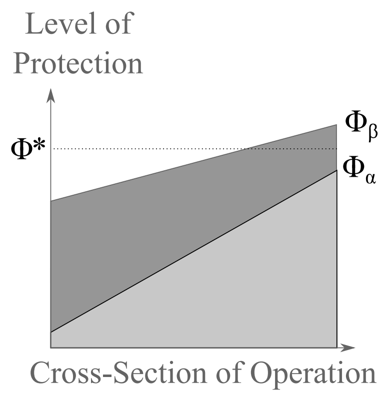

(a) Before event

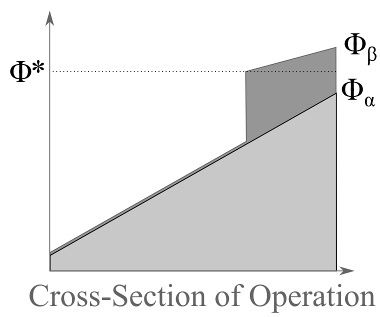

(b) After Event: Additional protection provided enough damping such that code threshold was not exceeded at all

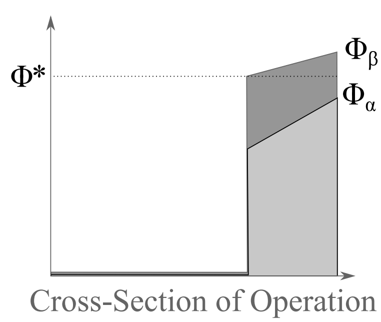

(c) After event: Additional protection did not provide enough damping to stop the code threshold from being exceeded in all areas

Figure 9: Example of code and compliance protection (Φ_α) versus critical asset protection (Φ_β) damage following an hazard event over a cross section of an operation.

rials, design trends, codes and standards of practice over the decades. Newer buildings can provide protection advantages simply due to improvements in codes and standards of practice. Raising electrical and mechanical rooms above grade, for instance, is an effective flood protection measure. The disparity between buildings of different ages is not always such that newer is better, however. Buildings are often retrofitted in various ways, such as the addition of modern fire protection systems, for example. Further, as building layouts change, exposure to shooting hazards, for instance, can be worse in newer, open concept buildings. The application of additional protection (Φ_β) hence must be assessed on a building by building basis, taking into account the full history of the facility and thee risk criteria.

Figure 10 is further instructive on how losses and costs of implementation change as a function of the level of protection for a fixed operation (a given P_S, P_M and Δt_S) and a given event (pictured are a low and high severity event). The level of protection for an event effectively establishes how much direct damage to infrastructure occurs for that event, the remainder being protected, and the initial drop in performance associated with the event. Hence as the level of design protection increases the direct damage to infrastructure decreases, and so too do the losses to the operation (the area under P_o on the event graph) along a path that achieves the requirements of the operation (P_S, P_M and Δt_S) and assuming the resilience costs in the lower graph are spent. Concurrently, the amount of damage to the implemented protection increases initially from zero (since when there is no protection, there can be no damage to it) up to some maximum then returning to zero when the level of design protection ensures no performance loss (P_o). As the level of protection increases, so does the cost of implementation, however, as the level of protection increases the cost to implement resilience that meets the design objectives decreases since less performance drop occurs as protection increases. Unless the threshold is exceeded.

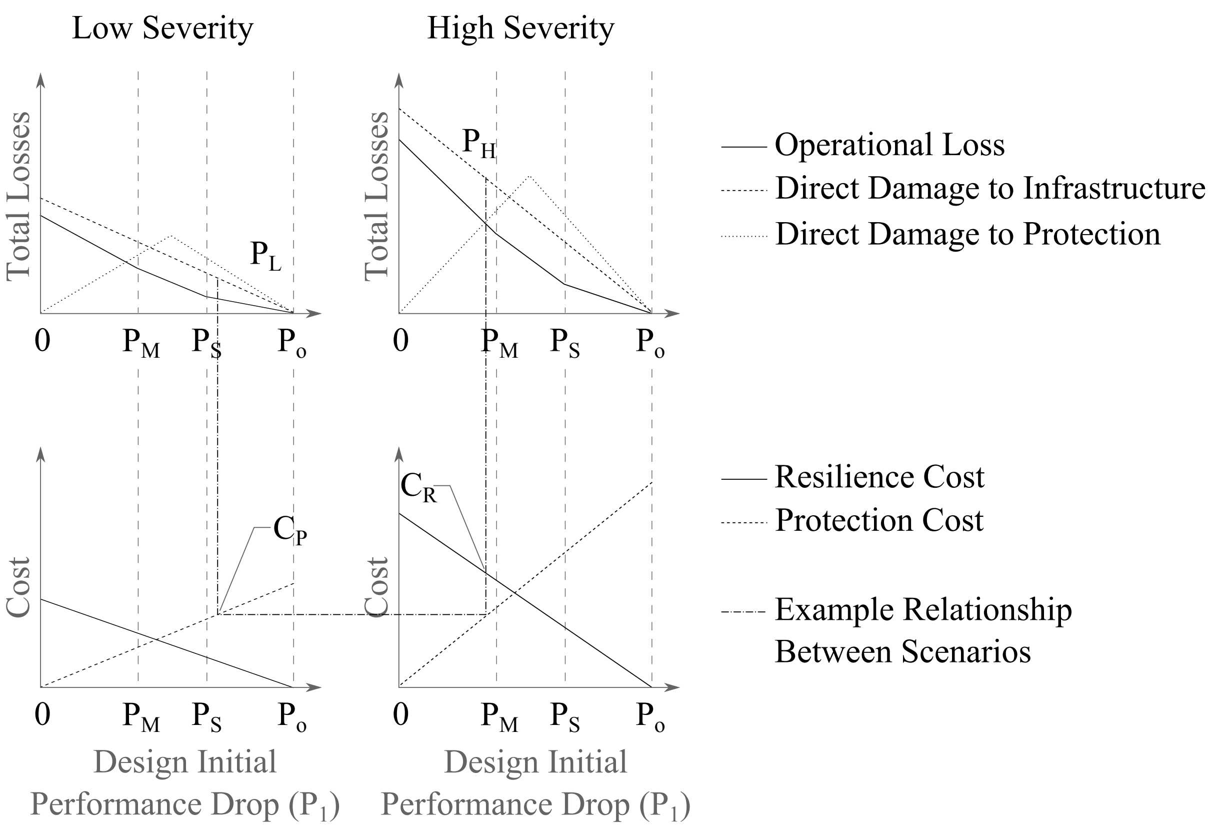

Figure 10: Generic total losses and costs as a function of design protection for an event of a given severity.

With these curves established, based on the design requirements (P_S, P_M and Δt_S) and the associated costs, assessing the balance of protection and resilience proceeds as follows: Selecting a protection level to guard against a low severity event that achieves a drop in performance to P_L corresponds to a level of protection in a high severity case that achieves a drop in performance to P_H and costs C_P. To meet the planning requirements in the high severity case (P_M, P_S, Δt_S) subsequently costs C_R and easily meets the resilience requirements of the low severity case. Trying different levels of P_L across the full spectrum of severity produces a series of options, their costs, and the losses associated with each severity. A comprehensive assessment would include full life-cycle costs (net present value or similar time-value representation) and also assessment of qualitative costs, benefits and losses. Qualitatively, the concern is on the efficiency of operation. Relevant questions in this regard include: ability to negotiate lower insurance premiums or acquire more attractive loans. Other forms of risk, such as strategic risk, can also be included with a suitable analysis thereof. The assessment can then be simplified down by portioning ranges of severity and presenting the overall protection-resilience balance problem as a multi-criteria, multi-objective decision problem with assessments of costs, benefits and losses for each partition for the set of different combinations of protection and resilience. In this fashion, therefore, the balance of protection and resilience can be selected and adjusted as needed based on the risk criteria (P_M, P_S, Δt_S).

Importantly, across a campus such as St. George, the hazard and loss curves of Figure 10 can vary significantly. This trend is well summarized by Figure 11. Each of the points on this graph correspond to a building on campus. Losses to the points nearest the upper right corner are, on a relative basis, of significant concern from both a protection and a resilience perspective, especially given the logarithmic scales of the chart. Points closest to the lower right corner

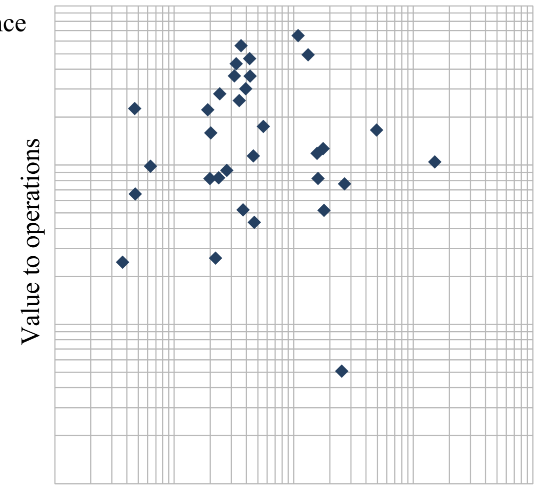

Figure 11: Trends of asset and operational value on St. George campus. Data from the table of values for campus insurance. Note that while the asset value is the potentially maximum direct loss, the value to business interruption is not an upper bound on operational loss but a relative measure used for business continuity insurance purposes. Exact values and building names are omitted to ensure the data remains non-attributable.

are those where protection will play a higher role than resilience. Buildings closest to the upper left corner are those whose risk treatment is most suited, relative to the cost of implementation, to resilience. Buildings nearest the lower left corner are relatively unimportant in terms of risk treatment. The costs associated with treating these buildings for protection and resilience is of course a function of what is already in place.

Further guidance on the balance of protection and resilience can be achieved by considering how the curves in Figure 10 vary with the planning requirements (P_M, P_S, Δt_S). These parameters effectively capture how operations can differ in risk criteria, even for a fixed set of hazards. The key is to consider how large changes in the level of protection compared to resilience can be achieved without having to generate a lot of options and manipulate the dependency graph an excessive number of times. Figure 12 considers how, for a fixed investment in protection and resilience (fixed rates of reaction, response and recovery) and for a given event severity, the losses vary with the planning requirements. In the figure L_o is the loss from the initial damage associated with the event. The remaining losses are associated with the three regions of reduced performance in the generic event graph (Figure2): L_1 is from t_o to t_1 (P_1 to P_M); L_2 is form t_1 to t_2 (P_M to P_S); and L_3 is form t_2 to t_3 (P_S to P_o again). Because the investments are held constant, when one of P_M or P_S is changed in the graphs, so too is Δt_S. This allows for investment options to quickly be removed from contention and hence remove the number of combinations of measures that must be assessed.

For instance, for a varying minimum operating performance (Figure 12a) the sustainable time objective and losses are as shown when there is no protection. Operations with smaller minimum operating performances, all else equal, incur more loss since there is a longer slower climb to the sustainable performance objective. In such situations increased resilience investment for the response phase (L_2) would reduce losses. At the other end of the spectrum resilience investment in the reaction phase would reduce losses. Both of these incur an overall larger cost to the initial case. Alternatively protection can be increased. As shown in Figure 12b, a protection increase and resilience decrease (to maintain cost parity with the initial case) can remove the losses for operations with small minimum operating performance requirements (Losses are zero below P_1^{**}). However, as apparent from Figure12a, the addition of protection at the cost of reducing resilience can negatively impact the time to sustainable performance; hence, if constrained by this time objective resilience as opposed to protection must be improved. As P_M approaches P_S, there is no clear winner between protection and resilience to achieve less loss for less cost, unless the sustainable time objective is a constraint, then resilience is again the winner.

In the case of the minimum sustainable performance

(P_S), Figure 12c, there is initially a large increase in losses for P_S near P_M since the pace of performance improvement is slow beyond P_S for the set investment. In this case where P_S is close to P_M, it can be impractical to invest in protection up to such a high level, except for the smallest severity events. Hence, resilience is the better alternative, focusing on capability for increasing the recovery from P_S to P_o to reduce L_3. Either resilience or protection, or a balance of the two is appropriate as P_S nears P_o, but if the sustainable time objective is a tight constraint, resilience improvements to the pace of response, to lower L_1, is the key.

The quantities of direct loss, L_o, in the graphs of Figure 12 are for a fixed ratio to the initial drop in performance. Changing risk criteria can also alter the value of this direct loss, perhaps due to the complexity of the operation or its historical and cultural value. At sufficiently large values this direct loss can theoretically overshadow operational losses, increasing the protection investment balance, and is hence an importance case to keep in mind.

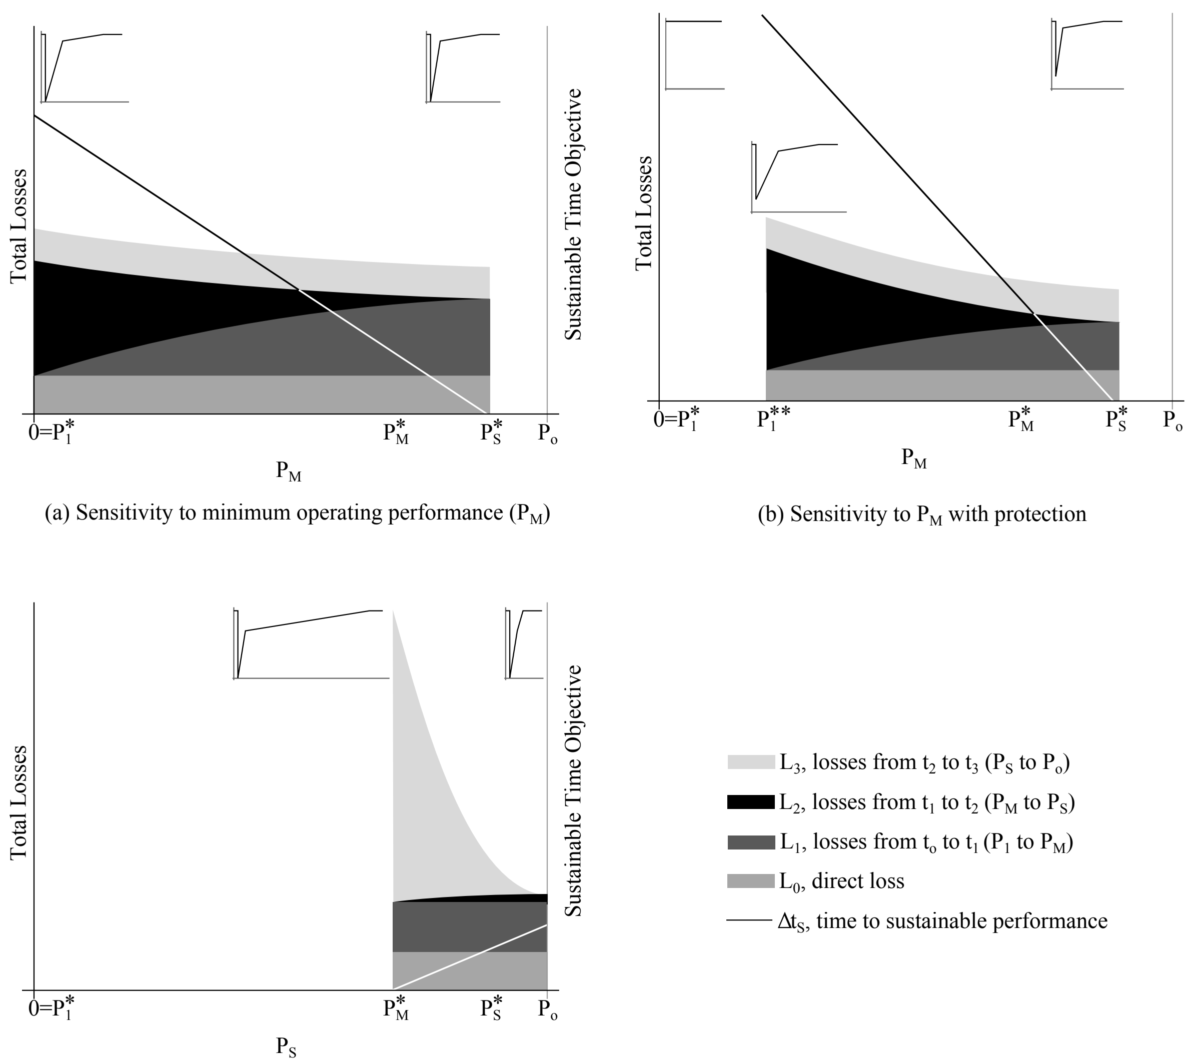

(a) Sensitivity to minimum operating performance (P_M)

(b) Sensitivity to P_M with protection

(c) Sensitivity to minimum sustainable performance (P_S)

Figure 12: Sensitivity of losses to changes in planning parameters for a given hazard. Inset at either end of each graph is the corresponding performance curve at that extreme of the sensitivity range. These graphs are on an operation-wide basis for a singular event and measure of consequence, hence vector notation and indices are dropped. $P_o = 1$, $P_M^* = 0.7$, $P_S^* = 0.9$.

5

Application to practice

The central point of the progression through the past sections is that the dependency map approach to infrastructure interdependency needs the proposed extensions to assess the full impact of protection and resilience investment combinations. As this dependency map becomes sufficiently complex for any real operation, however, the assessment of operational resilience must, at this time, be limited to a few alternatives that are initially deduced through holistic understanding of the operation. The overall integration of these steps follows a modified process to that depicted in Figure 7, as shown below in Figure 13. This generic framework is broadly applicable to public and private infrastructure in any risk context for many types of organizations, such as all levels of government, firms, or campuses.

The above protection-resilience balancing framework methodology serves as a novel, yet integral, component to campus infrastructure planning. Moreover it can be extended into a broader campus infrastructure planning framework. The connection to such a broader framework, to be introduced shortly, is two-fold. First, the basis of the balancing methodology is that it is specifically formulated to consider full-life cycle costs and can also be augmented with

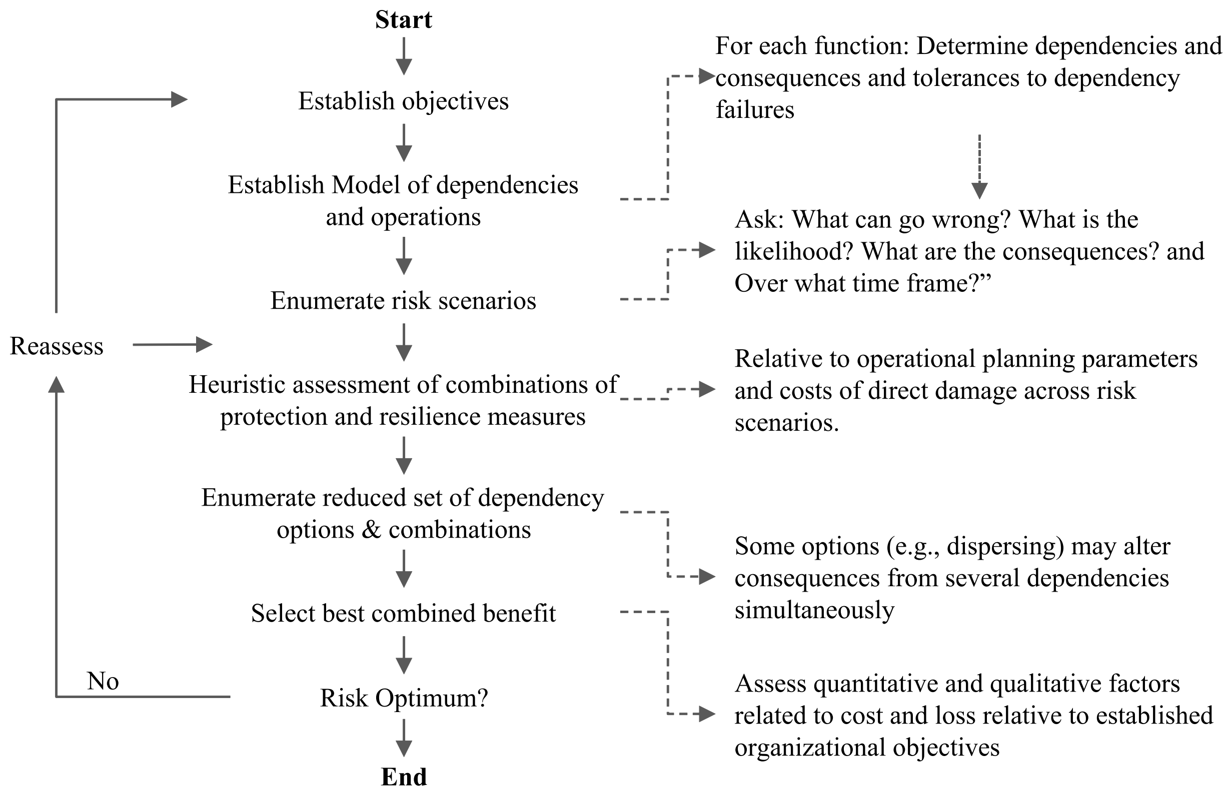

Figure 13: Protection-resilience balancing framework.

other planning criteria such as qualitative reputational or political factors. Secondly, there is tight coupling between the life-cycle cost outcomes of protection-resilience measures and those of sustainability measures – specifically those pertaining to demand reduction. Demand reduction is the practice of removing wastes and finding efficiencies in things like electricity, heat, and even space use and is a common goal for many University campus' as part of their sustainability initiatives. Demand reduction is also a key component of demand management, a strategy of resilience planning. Demand management, also requires a good degree of separation and direction of demand. Together, reduction, separation are best used when in balance with protection and sustainability. Taken too far, for example, demand reduction alone can weaken infrastructure or overly constrain operations, but when planned alongside protection and resilience, these three approaches can achieve overall far more benefit and efficiency.

Table 3 illustrates several examples of where protection, resilience and demand management can be assessed in combination. In each case each of the three measures reduces the requirements on the others. In each situation, hence, it is possible to locate an optimal balance between the different approaches to achieve the best business case. To clarify the interplay, consider the case of grid power outage. In this case the protection measures may already be in place. This fixes a portion of the risk exposure, and hence the relative merits of the business case of the remaining options. If sufficient risk remains, then reducing the electricity demand in key areas to their lowest possible amounts achieves an operational savings while enabling a suite of backup power options such as emergency generators, solar, and wind that would otherwise be too big to fit or afford, thus ensuring the continuity of operations. In this way the symbiosis of protection, resilience and demand management achieve a stronger business case when considered together.

These examples also highlight how the code and compliance protection requirements can be used most effectively on a campus. For instance, there are design guidelines for ventilating cooling equipment and adding redundancy to ensure they do not overheat and fail. Within the context of the planning framework a full assessment relative to resilience and sustainability, and given a warming climate, can expose how small improvements on this required protective cooling can improve operational performance over the life-cycle at next to no cost since only a marginal change on required protection levels is required. Addressing compliance protection requirements within this framework, hence, can save a great deal of expense, then considering protection, resilience and sustainability after the fact and individually.

Table 3: Example applications of the campus infrastructure development planning framework

Case	Protection	Resilience	Demand Management
Grid Power outage	Add buried, flood resistant and redundant power lines	Include a mix of power sources on most critical functions	Aggressively reduce power demand on most critical functions and add separate circuits to these needs
Deluge rain or main burst	Raise equipment in basement	Implement a mobile flood abatement strategy	Prioritize above grade facilities for functions that employ sensitive equipment
Heat outage or extreme cold	Add de-icing systems to, or harden, exposed equipment	A mobile heating strategy, a stationary backup heating solution or off site alternative facilities. Selection depends on size of demands and scope of the heating failure	Aggressively insulate critical heating demand areas
Extreme hot weather	Ensure proper cooling of cooling systems	Isolate critical cooling needs and add backups or mobile cooling solutions.	Employ passive cooling and daylight infiltration management and use efficient appliances to reduce internal heat load.
Water failure or boil water advisory	Lobby for municipal water system upgrades	Store backup water for critical functions. Backup power for boiling water	Aggressively reduce water needs in residence and research labs with essential water needs

Finally, the combination of these three approaches to campus infrastructure planning are integrated into a singular framework in Figure 14. Here the framework is presented, again, predominately in terms of quantifiable costs and benefits for simplification, but qualitative factors can easily be added to augment the analysis and comparison of options. This planning framework in quantitative terms proceeds from the common operation and development drivers of all teaching and research intensive Universities – teaching, attracting the best researchers and ensuring the safety of staff and students on the campus. Protection is necessary to safeguard against injury and resilience enhances

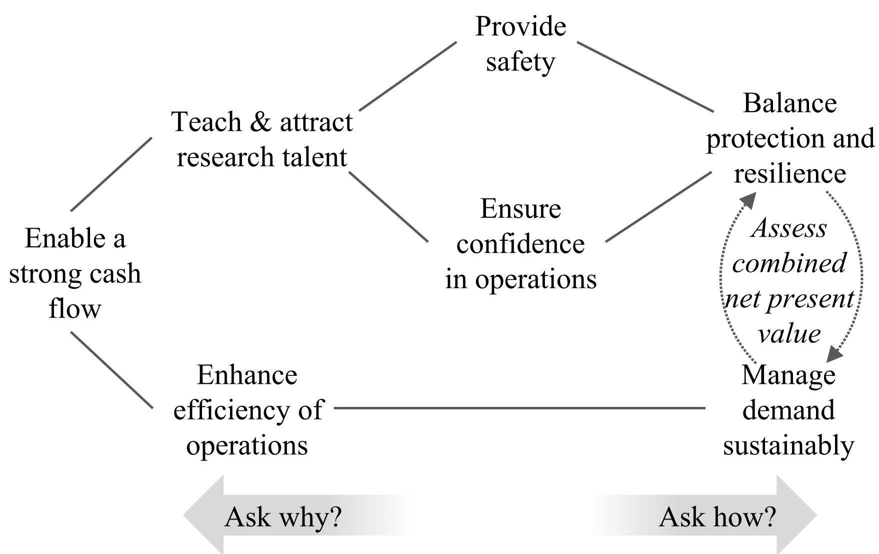

Figure 14: Campus infrastructure development planning framework. Traversing the solid lines from left to right is completed by asking "how" and traversing these lines from right to left is completed by asking "why". The dotted lines correspond to the iterative planning activity between the balancing of protection and resilience and the implementation of sustainably managed demand management measures.

confidence in continuity of functions and of high value research. The demand management efforts of sustainability initiatives form the third pillar of the framework as discussed above. Implementation of or change in protection, resilience or demand management alters the benefits and costs of the others and so must be assessed in combination. Together, when these are on balance, there is a reputational, operational efficiency and a risk reduction improvement that all contribute to a strengthening of the cash flows. It is thus recommended that such a framework be incorporated into campus planning. The combined net savings that can be realised with each successive project can quickly accrue.

6

Epilogue

This work contains several insights and results. First and foremost a generic framework is provided for determining the balance of protection and resilience. This is presented in terms of changes to an operation. Secondly, the graph theoretic approach to assessing risk from dependency relationships of systems or systems of systems of infrastructure is extended to allow for the assessment of the most common measures for reducing risk in such systems, namely hardening, dispersion, diversifying, adding flexibility, and adding redundancy. A process is provided for utilizing the extended dependency map methodology to assess combinations of protection and resilience in order to optimize the protection-resilience balance. Then, an heuristic approach to operation planning is used to guide the generation of combinations of measures to greatly reduce the number of combinations that need be tried to achieve a balance of protection and resilience. The approach provides a decision problem for selecting the balance of protection and resilience based on the operational design parameters (P_M, P_S, Δt_S). Finally, a campus infrastructure development planning framework and example applications are provided.

In general as the planning requirements of operations vary, several general conclusions about the balance of protection and resilience can be drawn. Relative to an initial condition of zero protection, operations with very low minimum operating performance levels can benefit from increased protection, at the cost of resilience. Regardless of the minimum operating performance level or the minimum sustainable performance level, if the sustainable time objective recovery is tightly constrained then resilience, at the cost of protection if necessary, is required. These generalizations are all relative to the cost of direct damage. In some cases this cost can be sufficiently large so as to result in significantly more investment in protection than resilience.

Future work on this subject is both desirable and expected. In future studies application of the methodology on several different operations and risk criteria will be helpful for examining means to enhance the search for optimal balance points. This work shall no doubt reveal the tradeoffs in relaxing some of the constraints related to meeting resilience goals at all ends of the spectrum and allowing for likelihood or probability assessments of treatment success as is common in protection approaches. This methodology can serve as the foundation for the continued improvement of infrastructure risk management in an increasingly dynamic, complex and uncertain world.

7

REFERENCES

Abouchar, A., 1977. Transportation Economics and Public Policy: With Urban Extensions, John Wiley & Sons Inc. Available at: http://trid.trb.org/view.aspx?id=275454 [Accessed September 10, 2014].

Ahern, J., 2011. From fail-safe to safe-to-fail: Sustainability and resilience in the new urban world. Landscape and Urban Planning, 100(4), pp.341–343.

Aschauer, D.A., 1989. Is public expenditure productive? Journal of Monetary Economics, 23(2), pp.177–200.

Bristow, D.N. et al., 2014. Multiple Participant Models of Urban Infrastructure Performance and Decision Support. In GDN 2014. Toulouse, France.

Bruneau, M. et al., 2003. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. Earthquake Spectra, 19(4), pp.733–752.

Casalicchio, E. & Galli, E., 2008. Metrics For Quantifying Interdependencies. In M. Papa & S. Shenoi, eds. Critical Infrastructure Protection II. The International Federation for Information Processing. Springer US, pp. 215–227. Available at: http://link.springer.com/chapter/10.1007/978-0-387-88523-0_16 [Accessed August 11, 2014].

Cimellaro, G.P., Reinhorn, A.M. & Bruneau, M., 2010. Framework for analytical quantification of disaster resilience. Engineering Structures, 32(11), pp.3639–3649.

Cimellaro, G.P., Reinhorn, A.M. & Bruneau, M., 2009. Seismic resilience of a hospital system. Structure and Infrastructure Engineering, 6(1-2), pp.127–144.

City of Toronto, 2012. Critical Infrastructure Resilience Study: Phase 1 Status Report, Environment and Energy Division.

DHS, 2013. National Infrastructure Protection Plan 2013: Partnering for Critical Infrastructure Security and Resilience, Available at: http://www.dhs.gov/national-infrastructure-protection-plan [Accessed September 9, 2014].

Ezell, B., Farr, J. & Wiese, I., 2000. Infrastructure Risk Analysis Model. Journal of Infrastructure Systems, 6(3), pp.114–117.

Garcia, M.L., 2005. Vulnerability Assessment of Physical Protection Systems, 1st edition, Amsterdam; Boston: Butterworth-Heinemann.

Gibbs, L.I. & Holloway, C.F., 2013. Hurricane Sandy After Actioin: Report and Recommendations to Mayor Michael R. Bloomberg, Available at: http://www.nyc.gov/html/recovery/downloads/pdf/sandy_aar_5.2.13.pdf [Accessed September 9, 2014].

Gillen, D.W., 1996. Transportation infrastructure and economic development: A review of recent literature. Logistics and Transportation Review, 32(1), p.39.

Haimes, Y. & Jiang, P., 2001. Leontief-Based Model of Risk in Complex Interconnected Infrastructures. Journal of Infrastructure Systems, 7(1), pp.1–12.

Haimes, Y.Y., 2009. On the Complex Definition of Risk: A Systems-Based Approach. Risk Analysis, 29(12), pp.1647–1654.

Haimes, Y.Y., 2011. Risk Modeling, Assessment, and Management, John Wiley & Sons.

Haimes, Y.Y., Crowther, K. & Horowitz, B.M., 2008. Homeland security preparedness: Balancing protection with resilience in emergent systems. Systems Engineering, 11(4), pp.287–308.

Hay, A., 2013. Operational Survival: Putting resilience at the core of infrastructure planning, London, UK: Explora Research.

Kaplan, S. & Garrick, B.J., 1981. On The Quantitative Definition of Risk. Risk Analysis, 1(1), pp.11–27.

Kennedy, C.A., 2011. The Evolution of Great World Cities: Urban Wealth and Economic Growth, Toronto, Ontario: University of Toronto Press.

Linkov, I. et al., 2014. Changing the resilience paradigm. Nature Climate Change, 4(6), pp.407–409.

Louis Lebel, Bach Tan Sinh & Elena Nikitina, 2010. Chapter 6 Adaptive governance of risks: climate, water, and disasters. In Climate Change Adaptation and Disaster Risk Reduction: Issues and Challenges. Community, Environment and Disaster Risk Management. Emerald Group Publishing Limited, pp. 115–142. Available at: http://www.emeraldinsight.com/doi/full/10.1108/S2040-7262%282010%290000004012 [Accessed September 9, 2014].

Macaulay, T., 2008. Critical Infrastructure: Understanding Its Component Parts, Vulnerabilities, Operating Risks, and Interdependencies, CRC Press.

MarketsAndMarkets, 2014. Critical Infrastructure Protection (CIP) Market: Global Advancements, Emerging Applications, Worldwide Forecasts and Analysis (2013 – 2018).

O'Neill, P., 2013. Protecting Critical Infrastructure by Identifying Pathways of Exposure to Risk. Technology Innovation Management Review, (August 2013: Cybersecurity), pp.34–40.

Ouyang, M., 2014. Review on modeling and simulation of interdependent critical infrastructure systems. Reliability Engineering & System Safety, 121, pp.43–60.

Ouyang, M., Dueñas-Osorio, L. & Min, X., 2012. A three-stage resilience analysis framework for urban infrastructure systems. Structural Safety, 36–37, pp.23–31.

Seed, R.B. et al., 2006. Investigation of the Performance of the New Orleans Flood Protection System in Hurricane Katrina on August 29, 2005, Available at: http://www.ce.berkeley.edu/projects/neworleans/ [Accessed September 9, 2014].

Taleb, N.N., 2012. Antifragile: Things That Gain from Disorder, Random House LLC.

Vugrin, E.D. et al., 2010. A Framework for Assessing the Resilience of Infrastructure and Economic Systems. In K. Gopalakrishnan & S. Peeta, eds. Sustainable and Resilient Critical Infrastructure Systems. Springer Berlin Heidelberg, pp. 77–116.

World Nuclear Association, 2014. Fukushima Accident. Available at: http://www.world-nuclear.org/info/safety-and-security/safety-of-plants/fukushima-accident/ [Accessed September 9, 2014].