Industrial and Manufacturing Systems Engineering
Publications

Industrial and Manufacturing Systems Engineering

2015

Allocating Resources to Enhance Resilience, with Application to Superstorm Sandy and an Electric Utility

Cameron A. MacKenzie
Iowa State University, camacken@iastate.edu

Christopher W. Zobel
Virginia Polytechnic Institute and State University

Allocating Resources to Enhance Resilience, with Application to Superstorm Sandy and an Electric Utility

Abstract

This article constructs a framework to help a decisionmaker allocate resources to increase his or her organization's resilience to a system disruption, where resilience is measured as a function of the average loss per unit time and the time needed to recover full functionality. Enhancing resilience prior to a disruption involves allocating resources from a fixed budget to reduce the value of one or both of these characteristics. We first look at characterizing the optimal resource allocations associated with several standard allocation functions. Because the resources are being allocated before the disruption, however, the initial loss and recovery time may not be known with certainty. We thus also apply the optimal resource allocation model for resilience to three models of uncertain disruptions: (1) independent probabilities, (2) dependent probabilities, and (3) unknown probabilities. The optimization model is applied to an example of increasing the resilience of an electric power network following Superstorm Sandy.

Keywords
resilience, resource allocation, Superstorm Sandy

Disciplines
Industrial Engineering | Systems Engineering

Allocating Resources to Enhance Resilience, with Application to Superstorm Sandy and an Electric Utility

Cameron A. MacKenzie

Department of Industrial and Manufacturing Systems Engineering

Iowa State University

Ames, IA 5011

camacken@iastate.edu

Christopher W. Zobel

Department of Business Information Technology

Virginia Tech

Blacksburg, VA 24061

czobel@vt.edu

ABSTRACT

This paper constructs a framework to help a decision maker allocate resources to increase his or her organization's resilience to a system disruption, where resilience is measured as a function of the average loss per unit time and the time needed to recover full functionality. Enhancing resilience prior to a disruption involves allocating resources from a fixed budget to reduce the value of one or both of these characteristics. We first look at characterizing the optimal resource allocations associated with several standard allocation functions. Because the resources are being allocated before the disruption, however, the initial loss and recovery time may not be known

with certainty. We thus also apply the optimal resource allocation model for resilience to three models of uncertain disruptions: (1) independent probabilities, (2) dependent probabilities, and (3) unknown probabilities. The optimization model is applied to an example of increasing the resilience of an electric power network following Superstorm Sandy.

KEYWORDS: resilience, resource allocation, Superstorm Sandy

1. INTRODUCTION

The capacity for resilience is an important characteristic of any real-world system that is subject to the possibility of disruptions. Given the increasing number of natural and man-made disasters, and their growing impact on our globally interconnected infrastructure systems, improving system resilience and assessing the level and type of resources that must be committed in order to do so effectively is of growing importance. This requires quantifying a system's resilience and the relative cost effectiveness of different approaches for increasing that resilience, subject to constraints on resource availability.

Although the term resilience has a variety of different definitions in the literature, depending on the discipline, a resilient system is frequently considered to be one which is able to withstand, absorb, and successfully recover from the impact of some sort of disruption to normal operations.[1,2] Quantifying system or organizational resilience thus requires measuring both the amount of loss and the length of recovery time.[3,4] From the standpoint of making investments to improve resilience, some activities will have a larger impact on reducing loss while others will more directly and significantly impact the recovery time. This paper considers the trade-offs

between these two types of investments from a resource allocation perspective: given a fixed budget, which investment should a decision maker choose to make to improve the overall resilience of a given system.

This paper makes the following contributions to the growing literature on resilience. Mathematical models determine the optimal allocation of resources toward reducing loss and improving recovery time in order to maximize system resilience prior to a disruption. We prove the conditions under which a decision maker should allocate the entire budget either to reducing loss or to reducing the time to recover for four different allocation functions. Since the optimization model assumes that resources are being allocated prior to a disruption, we also explore the impact of uncertainty on model parameters for (1) independent probabilities, (2) dependent probabilities, and (3) unknown probabilities. We apply the optimization model to an example in which an electric utility company seeks to increase the resilience of its network following Superstorm Sandy. This example demonstrates how an organization can use the resilience model to determine how resources should be divided between reducing the impacts (e.g., hardening) and improving the time to recovery.

We begin our discussion with a brief literature review to establish the context for the resilience-enhancing, resource allocation framework. Following a discussion of the modeling formulation, we examine four characteristic allocation functions under certainty and under uncertainty and explore the optimal solutions given each allocation function. Next, the illustrative example explores how the Consolidated Edison Power Company could use the resource allocation model to optimally spend a $1 billion budget to improve the resilience of its electric power network.

The paper concludes by expanding on the mathematical results and the example to provide general guidelines for decision makers while recognizing the importance and limitations of modeling assumptions.

2. LITERATURE REVIEW

The capacity for a system to withstand, absorb, and recover from a disruption can depend on a number of different physical, social, and economic factors. This multi-dimensional nature of the concept of resilience lends itself to different types of analysis by researchers in different disciplines. In the social sciences, for example, indicator variables are typically used to capture the breadth of societal or economic characteristics that support resilient behavior,[5,6] whereas in engineering it is more easily quantifiable physical measures of changing system performance that are typically analyzed.[7,8]

In this paper, we adopt the engineering-based view that resilience is a process that can be measured and monitored with respect to changes in system performance over time. Much of the this research is related to measuring resilience as a function of the calculated area beneath the performance curve (see Fig. 1), with the amount of loss suffered and the length of time until recovery serving as important related parameters.[9,10] Under this approach, resilience, R, may be represented as follows, where T represents the length of time that the system is in a state of loss and $\bar{X}$ represents the average loss per unit time during the disruption:[11]

$$R_*(\bar{X}, T) = 1 - \frac{\bar{X}T}{T^*}.$$

(1)

The parameter T^* in Equation (1) represents the maximum possible recovery time. (If $\bar{X}$ or T

equals 0, then there are no losses and the organization's resilience is 1.) This particular

formulation will apply in general even if the system's recovery trajectory is non-monotonic

because the average loss per unit time depends on the recovery trajectory. As the average loss in

system performance, $\bar{X}$ is the fraction of performance lost from the non-disrupted system, and

$0 \leq \bar{X} \leq 1$. From an operational standpoint, a decision maker can improve resilience by

allocating resources to lessen the average impact and/or to decrease the time to recovery.

[Insert Fig. 1]

This basic idea has been extended in a number of ways, including explicitly modeling the trade-

offs between loss and recovery time,[11,12] modeling interconnected infrastructure resilience,[13,14]

capturing multi-dimensionality,[4,15] and tracking the disaster recovery process.[16] Most of these

research efforts ultimately tend to concentrate on the characterization of resilience so as to better

understand system behavior. Simulation, for example, is often used to compare a system's

response to a disruptive event under alternative control policies,[11,17] and techniques like

regression are used to try to uncover the theoretical relationship between decisions and

outcomes.[13]

A related viewpoint depicts resilience as a time-dependent metric defined as the ratio of recovery

at a given time to the loss in performance.[18] Such a metric can be used to identify the most

important components in a network, such as a transportation or power network, with the goal of

prioritizing investments to protect the most influential components.[19] Incorporating uncertainty

with this metric produces a stochastic measure of resilience, and the optimal recovery activity

can be identified via simulation[20] or by comparing the percentiles of different probability

distributions of resilience generated by different recovery activities.[21]

Although the motivation behind measuring resilience typically focuses on helping decision

makers prepare and respond to disruptions, less research has concentrated on how resources

should be allocated to improve an organization or system's resilience. The literature studying

how to optimize system resilience usually defines activities to pursue or components to protect

and then uses a numerical method such as simulation to select the best alternative.[19,20,21] In a

real-world context, however, it is important to keep in mind that there are limits not only on how

quickly a system can recover from a disruption, but also on the physical and financial resources

available to help offset the disaster's impacts. A process manager will thus need to focus on

determining the most effective way to allocate his or her resources in order to improve system

resilience. This paper seeks to help a decision maker determine the most effective alternative by

mapping resources to the resilience metric in Equation (1) via mathematical functions. Non-

linear programming is used to analytically determine the optimal allocation of resources in order

to provide a broader application and lead to more general results than a numerical method tied to

a specific application.

3. ALLOCATION MODEL TO ENHANCE RESILIENCE

Given a system that will be directly impacted by a disruption, let us assume that functions are

known which describe the effect of allocating resources to (1) reducing that system's average

loss per unit time: $\bar{X}(z_{\bar{X}})$, where $z_{\bar{X}}$ is the amount allocated to lessen the impact, and (2) reducing

its recovery time: $T(z_T)$, where z_T is the amount allocated to improve the recovery time. The total budget available is Z. The decision maker may then maximize system resilience by solving the following problem:

$$\text{maximize} \quad R_*\big(\bar{X}(z_{\bar{X}}), T(z_T)\big)$$

$$\text{subject to} \quad z_{\bar{X}} + z_T \leq Z \tag{2}$$

$$z_{\bar{X}}, z_T \geq 0$$

where R_* is defined as in Equation (1). In general, the functional forms for $\bar{X}(z_{\bar{X}})$ and $T(z_T)$ should obey certain properties. First of all, the first derivatives, $\frac{d\bar{X}}{dz_{\bar{X}}}$ and $\frac{dT}{dz_T}$, should be less than 0, which means that the functions are decreasing functions and as more resources are allocated to $\bar{X}$ or T the average loss and recovery time should continue to decrease. Furthermore, the second derivatives $\frac{d^2\bar{X}}{dz_{\bar{X}}^2}$ and $\frac{d^2T}{dz_T^2}$ should be greater than or equal to 0, which signifies constant returns or marginal decreasing improvements as more resources are allocated. The first dollar allocated to reduce $\bar{X}$ or T thus should be at least as effective as the next dollar allocated.

Given these constraints, several functional forms could describe the effectiveness of allocating resources to $\bar{X}$ or to T. Although one may not invest directly in reducing the average system loss or the system recovery time, the projects or activities to which one allocates those resources would typically impact one or both of these characteristics. The model assumes that a resource only reduces one of the factors, and future research can examine the allocation if a resource can simultaneously benefit both factors. Despite this limitation, if the optimization model in (2)

reveals that 70% of the resources should be allocated to reduce $\bar{X}$ and 30% to reduce T, a good heuristic may be to select a project in which 70% of the benefits reduce the average losses and 30% reduce the recovery time.

The following discussion explores four such forms for $\bar{X}(z_{\bar{X}})$ and $T(z_T)$ and studies the conditions under which all of the resources should be allocated to reduce either $\bar{X}$ or T, or the budget should be divided between the two factors. Rules for allocating resources are developed first for models with certainty and then for models that incorporate uncertainty in one or more of the parameters.

4. MODEL UNDER CERTAINTY

Let us first assume that all parameters are known with certainty. The functional forms that we will consider are linear, exponential, quadratic, and logarithmic. The parameters $\hat{X}$ and $\hat{T}$ represent the baseline average loss and time to recovery if no resources are allocated.

4.1. Linear Allocation Function

A linear allocation function rests on the assumption that the decrease in one of the resilience factors is constant, no matter how small that factor becomes. It thus assumes constant returns to scale. Mathematically, this gives $\bar{X}(z_{\bar{X}}) = \hat{X} - a_{\bar{X}} z_{\bar{X}}$ and $T(z_T) = \hat{T} - a_T z_T$, where $a_{\bar{X}} > 0$ and $a_T > 0$ are parameters describing the effectiveness of allocating resources. After substituting the linear allocation functions into Equation (1), the objective function in (2) becomes

$$\text{maximize } 1 - \frac{\left(\hat{X} - a_{\bar{X}}z_{\bar{X}}\right)\left(\hat{T} - a_{T}z_{T}\right)}{T^*}. \tag{3}$$

In order to simplify the following proposition, we assume that $\hat{X} \geq a_{\bar{X}}Z$ and $\hat{T} \geq a_{T}Z$. If that assumption proves to be incorrect, then it is trivial to see that perfect resilience, $R_* = 1$, can be achieved by reducing one of the factors to zero.

Proposition 1. The decision maker should allocate resources to reduce only one factor. He or she should choose $z_{\bar{X}} = Z$ if $a_{\bar{X}}/\hat{X} \geq a_{T}/\hat{T}$ and choose $z_{T} = Z$ otherwise.

Proof. Because of the symmetry in the optimization problem, it suffices to just show for one parameter. After expanding the product, optimizing Equation (3) is equivalent to:

$$\text{minimize } \hat{X}\hat{T} - a_{\bar{X}}\hat{T}z_{\bar{X}} - a_{T}\hat{X}z_{T} + a_{\bar{X}}a_{T}z_{\bar{X}}z_{T}. \tag{4}$$

Assume that $a_{\bar{X}}/\hat{X} \geq a_{T}/\hat{T}$, which implies $a_{\bar{X}}\hat{T} \geq a_{T}\hat{X}$. Because $a_{\bar{X}}\hat{T}$ and $a_{T}\hat{X}$ are the two negative components, it is clear that that $z_{\bar{X}} = Z$ and $z_{T} = 0$ is the optimal allocation. If $a_{T}/\hat{T} > a_{\bar{X}}/\hat{X}$, the opposite is true, and all the resources should be used to reduce T. ■

If a linear function is an appropriate model for representing the effectiveness of allocating resources, then the decision maker should focus his or her resources to reduce the factor whose initial parameter is the smallest, assuming the effectiveness parameters are relatively equal. For

example, if the average loss $\hat{X}$ is very small but the time to recovery $\hat{T}$ is large, then the decision maker should allocate resources to reducing the average loss.

4.2. Exponential Allocation Function

In contrast to linear allocation, an exponential allocation function assumes that each resilience factor is always reduced by same fractional amount for a given investment. In other words, if a given investment reduces an average loss of $\hat{X}_1$ by one half, then the same investment would also reduce an average loss of $\hat{X}_2$ by one half, regardless of the actual values for $\hat{X}_1$ and $\hat{X}_2$. Unlike the linear allocation function, an exponential allocation function assumes marginally decreasing improvements and that the first dollar allocated is more effective than the second dollar.

The exponential allocation functions for the two factors may be represented by: $\bar{X}(z_{\bar{X}}) = \hat{X}\exp(-a_{\bar{X}}z_{\bar{X}})$, and $T(z_T) = \hat{T}\exp(-a_T z_T)$, where $a_{\bar{X}} > 0$ and $a_T > 0$ describe the effectiveness of allocating resources.

Proposition 2. The decision maker should choose $\boldsymbol{z_{\bar{X}}} = \boldsymbol{Z}$ if $\boldsymbol{a_{\bar{X}}} \geq \boldsymbol{a_T}$ and $\boldsymbol{z_T} = \boldsymbol{Z}$ otherwise.

Proof. Without loss of generality, assume $a_{\bar{X}} \geq a_T$. After substituting the exponential functions into Equation (1), the objective function in (2) becomes

$$\text{maximize } 1 - \frac{\hat{X}\hat{T}\exp(-a_{\bar{X}}z_{\bar{X}} - a_T z_T)}{T^*}. \tag{5}$$

Equation (5) is equivalent to minimizing $\exp(-a_{\bar{X}}z_{\bar{X}} - a_T z_T)$, which leads to setting $z_{\bar{X}} = Z$.

The same argument applies for the other case. ∎

A decision maker who relies on either the exponential or linear functions to guide allocation decisions should therefore allocate resources to reduce only one of the two factors. Whereas the linear allocation function focuses attention on reducing the smallest initial factor $\hat{X}$ or $\hat{T}$, a decision maker using the exponential allocation function should focus resources on the factor with the largest effectiveness parameter, regardless of the initial values.

4.3. Quadratic Allocation Function

Similar to an exponential function, a quadratic allocation function assumes marginal decreasing returns. The marginal return on allocating resource decreases at a rate linear to the amount of resources allocated. The quadratic allocation functions for the two factors are: $\bar{X}(z_{\bar{X}}) = \hat{X} - b_{\bar{X}}z_{\bar{X}} + a_{\bar{X}}z_{\bar{X}}^2$ and $T(z_T) = \hat{T} - b_T z_T + a_T z_T^2$ where $a_{\bar{X}}, a_T > 0$ and $b_{\bar{X}}, b_T \geq 0$. In order to ensure decreasing functions, it must be true that $z_{\bar{X}} \leq b_{\bar{X}}/(2a_{\bar{X}})$ and $z_T \leq b_T/(2a_T)$.

After taking the product of $\bar{X}(z_{\bar{X}})$ and $T(z_T)$ and replacing z_T with $Z - z_{\bar{X}}$, we minimize

$$\bar{R}(z_{\bar{X}}) = A z_{\bar{X}} + B z_{\bar{X}}^2 + C z_{\bar{X}}^3 + D z_{\bar{X}}^4 \tag{6}$$

where

$$A = \hat{X}b_T - \hat{T}b_{\bar{X}} - 2\hat{X}a_T Z + b_{\bar{X}}b_T Z - b_{\bar{X}}a_T Z^2$$

$$B = \hat{X}a_T - b_{\bar{X}}b_T + 2b_{\bar{X}}a_T Z + \hat{T}a_{\bar{X}} - a_{\bar{X}}b_T Z + a_{\bar{X}}a_T Z^2$$

$$C = -b_{\bar{X}}a_T + a_{\bar{X}}b_T - 2a_{\bar{X}}a_T Z$$

$$D = a_{\bar{X}}a_T.$$

Minimizing $\bar{R}(z_{\bar{X}})$ is equivalent to maximizing $R_*(\bar{X}, T)$ in the optimization problem in (2).

Equation (6) is a quartic equation. As $z_{\bar{X}}$ approaches negative and positive infinity, Equation (6) approaches positive infinity because $D > 0$. Thus, $\bar{R}(z_{\bar{X}})$ has at most two local minima and at most one local maximum. We label the first local minimum $z_{\bar{X}} = z_1$ and the second local minimum $z_{\bar{X}} = z_2$.

Unlike the linear and exponential allocation functions, the quadratic allocation function may result in an optimal allocation of resources to reduce both factors, which occurs at either z_1 or z_2. We describe conditions when it is optimal to allocate to reduce either one factor ($z_{\bar{X}} = 0$ or Z) and when it is optimal to allocate to reduce both factors ($z_{\bar{X}} = z_1$ or z_2).

Proposition 3. The following conditions should guide the decision in determining how to optimally allocate resources:

1. If $A > 0$ and $\bar{R}'(Z) > 0$

$$z_{\bar{X}} = \begin{cases} z_2, & 0 < z_2 < Z \text{ and } \bar{R}(z_2) \leq 0 \\ 0, & \text{otherwise} \end{cases}$$

2. If $A > 0$ and $\bar{R}'(Z) < 0$

$$z_{\bar{X}} = \begin{cases} 0, & 0 \leq \bar{R}(Z) \\ Z, & \text{otherwise} \end{cases}$$

3. If $A < 0$ and $\bar{R}'(Z) > 0$

$$z_{\bar{X}} = \begin{cases} z_1, & z_1 > 0, z_2 < Z, \text{and } \bar{R}(z_1) \leq \bar{R}(z_2) \\ z_1, & z_2 > Z \\ z_2, & \text{otherwise} \end{cases}$$

4. If $A < 0$ and $\bar{R}'(Z) < 0$

$$z_{\bar{X}} = \begin{cases} z_1, & 0 < z_1 < Z \text{ and } \bar{R}(z_1) \leq \bar{R}(Z) \\ Z, & \text{otherwise} \end{cases}$$

where $\bar{R}'(z_{\bar{X}}) = A + 2Bz_{\bar{X}} + 3Cz_{\bar{X}}^2 + 4Dz_{\bar{X}}^3$.

Proof. See Appendix A. ■

Calculating z_1 and z_2 involves finding the roots of a cubic equation $\bar{R}'(z_{\bar{X}})$. Although a closed formula exists for solving a cubic equation, expressing the solution in terms of A, B, C, and D provides little insight into z_1 and z_2.

4.4. Logarithmic Allocation Function

The logarithmic allocation functions are: $\bar{X} = \hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} z_{\bar{X}})$ and $T = \hat{T} - a_T \log(1 + b_T z_T)$ where $a_{\bar{X}}, b_{\bar{X}}, a_T, b_T > 0$. The function is strictly decreasing if the parameters are greater than 0. The logarithmic allocation function also assumes marginal decreasing returns, but the marginal reduction in a factor decreases at a quicker rate than that of the quadratic allocation function.

The logarithmic allocation function may also result in situations where it is optimal to allocate resources to reduce both factors. Maximizing the resilience function is equivalent to minimizing

$$\tilde{R}(z_{\bar{X}}) = \log\left(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} z_{\bar{X}})\right) + \log\left(\hat{T} - a_T \log(1 + b_T[Z - z_{\bar{X}}])\right) \quad (7)$$

where z_T is replaced by $Z - z_{\bar{X}}$.

As with the quadratic allocation function, the logarithmic allocation function may result in an optimal allocation of resources to reduce both factors. Calculating a solution that reduces both factors involves finding the roots to the numerator of the first derivative of $\tilde{R}(z_{\bar{X}})$:

$$\tilde{R}'^{(z_{\bar{X}})} \quad (8)$$

$$= \frac{a_T b_T\left(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} z_{\bar{X}})\right)(1 + b_{\bar{X}} z_{\bar{X}}) - a_{\bar{X}} b_{\bar{X}}\left(\hat{T} - a_T \log(1 + b_T[Z - z_{\bar{X}}])\right)(1 + b_T[Z - z_{\bar{X}}])}{\left(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} z_{\bar{X}})\right)\left(\hat{T} - a_T \log(1 + b_T[Z - z_{\bar{X}}])\right)(1 + b_{\bar{X}} z_{\bar{X}})(1 + b_T[Z - z_{\bar{X}}])}.$$

Before describing the conditions under which it is optimal to allocate resources to reduce both factors, we determine the maximum number of local minima and maxima for $\tilde{R}(z_{\bar{X}})$. We assume that $\hat{X} > a_{\bar{X}} \log(1 + b_{\bar{X}} Z)$ and $\hat{T} > \log(1 + b_T Z)$ to ensure that $\tilde{R}(z_{\bar{X}})$ is continuous and differentiable over the domain of $z_{\bar{X}}$. If this assumption is violated, the decision maker should allocate resources until either $\bar{X}(z_{\bar{X}})$ or $T(z_T)$ equals 0.

Lemma 1. $\tilde{R}(z_{\bar{X}})$ has at most two local interior minima.

Proof. The number of local extrema for $\tilde{R}(z_{\bar{X}})$ equals the number of solutions to $z_{\bar{X}}$ when the numerator of $\tilde{R}'(z_{\bar{X}})$ in Equation (8) equals 0. The first derivative of the numerator of $\tilde{R}'(z_{\bar{X}})$ is

$$\tilde{R}''(z_{\bar{X}}) = b_{\bar{X}}b_T\big(\hat{X}a_T + \hat{T}a_{\bar{X}} - a_{\bar{X}}a_T \log([1 + b_{\bar{X}}z_{\bar{X}}][1 + b_T(Z - z_{\bar{X}})]) - 2a_{\bar{X}}a_T\big). \tag{9}$$

After setting $\tilde{R}''(z_{\bar{X}}) = 0$, we rearrange Equation (9) so that the equation is quadratic in $z_{\bar{X}}$. Thus, $\tilde{R}''(z_{\bar{X}}) = 0$ has at most two solutions for $z_{\bar{X}}$, which implies that $\tilde{R}'(z_{\bar{X}})$ has at most one local maximum and one local minimum. This implies that $\tilde{R}'(z_{\bar{X}}) = 0$ has at most three solutions for $z_{\bar{X}}$, and $\tilde{R}(z_{\bar{X}})$ has at most three local extrema. At least one of these local extrema must be a local maximum. Consequently, at most two local interior minima exist for $\tilde{R}(z_{\bar{X}})$. ∎

Lemma 2. $\tilde{R}(z_{\bar{X}})$ has at most one local interior maximum.

Proof. We can conclude that $\tilde{R}(z_{\bar{X}})$ has at most two local interior maxima by the same reasoning as in Lemma 1. To show that $\tilde{R}(z_{\bar{X}})$ has at most one local interior maximum, we multiply terms in Equation (9):

$$\tilde{R}''(z_{\bar{X}}) = b_{\bar{X}}b_T\big(\hat{X}a_T + \hat{T}a_{\bar{X}} - a_{\bar{X}}a_T \log\big(1 + b_T Z + [b_{\bar{X}} + b_T Z - b_T]z_{\bar{X}} - b_{\bar{X}}b_T z_{\bar{X}}^2\big) \tag{10}$$
$$- 2a_{\bar{X}}a_T\big).$$

Since $-\log(\cdot)$ is convex and nonincreasing and $1 + b_T Z + [b_{\bar{X}} + b_T Z - b_T]z_{\bar{X}} - b_{\bar{X}}b_T z_{\bar{X}}^2$ is concave, the expression $-a_{\bar{X}}a_T \log\big(1 + b_T Z + [b_{\bar{X}} + b_T Z - b_T]z_{\bar{X}} - b_{\bar{X}}b_T z_{\bar{X}}^2\big)$ is convex.[22] Thus $\tilde{R}''(z_{\bar{X}})$ is convex in $z_{\bar{X}}$.

If $\tilde{R}(z_{\bar{X}})$ has three local extrema, then $\tilde{R}''(z_{\bar{X}}) = 0$ must have two solutions for z_X or two roots. In this case, the convexity of $\tilde{R}''(z_{\bar{X}})$ implies that $\tilde{R}''(z_{\bar{X}}) > 0$ when $z_{\bar{X}}$ is less than the smaller root, $\tilde{R}''(z_{\bar{X}}) < 0$ when $z_{\bar{X}}$ is greater than the smaller root but less than the larger root, and $\tilde{R}''(z_{\bar{X}}) > 0$ when $z_{\bar{X}}$ is greater than the larger root. Thus, $\tilde{R}(z_{\bar{X}})$ is initially convex, then concave, and then convex. Consequently, the first extreme point must be a local minimum, the second extreme point must be a local maximum, and the third extreme point must be a local minimum.

If $\tilde{R}(z_{\bar{X}})$ only has two local extrema, one of them must be a local minimum because of the continuity of $\tilde{R}(z_{\bar{X}})$ over the domain of $z_{\bar{X}}$. Therefore, $\tilde{R}(z_{\bar{X}})$ has at most one local interior maximum. ∎

Since $\tilde{R}(z_{\bar{X}})$ has at most two local minima and one local maximum, we label the local minimum that occurs before the local maximum as z_1 and the local minimum that occurs after the local maximum as z_2. The conditions with the logarithmic allocation function that should guide the decision about the optimal allocation are identical to the conditions with the quadratic allocation function, where $\tilde{R}(z_{\bar{X}})$ replaces $\bar{R}(z_{\bar{X}})$ and $\tilde{R}'(z_{\bar{X}})$ replaces $\bar{R}'(z_{\bar{X}})$.

Proposition 4. The following conditions should guide the decision in determining how to optimally allocate resources:

1. If $\tilde{R}'(0) > 0$ and $\tilde{R}'(Z) > 0$

$$z_{\bar{X}} = \begin{cases} z_2, & z_2 < Z \text{ and } \tilde{R}(z_2) \leq \log(\hat{X}) + \log(\hat{T} - a_T \log(1 + b_T Z)) \\ 0, & \text{otherwise} \end{cases}$$

2. If $\tilde{R}'(0) > 0$ and $\tilde{R}'(Z) < 0$

$$z_{\bar{X}} = \begin{cases} 0, & \hat{X}(\hat{T} - a_T \log(1 + b_T Z)) \leq \hat{T}(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} Z)) \\ Z, & \text{otherwise} \end{cases}$$

3. If $\tilde{R}'(0) < 0$ and $\tilde{R}'(Z) > 0$

$$z_{\bar{X}} = \begin{cases} z_1, & z_1 > 0, z_2 < Z, \text{ and } \tilde{R}(z_1) \leq \tilde{R}(z_2) \\ z_1, & z_2 > Z \\ z_2, & \text{otherwise} \end{cases}$$

4. If $\tilde{R}'(0) < 0$ and $\tilde{R}'(Z) < 0$

$$z_{\bar{X}} = \begin{cases} z_1, & z_1 > 0 \text{ and } \tilde{R}(z_1) \leq \log(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} Z)) + \log(\hat{T}) \\ Z, & \text{otherwise} \end{cases}$$

Proof. See Appendix B. ∎

Calculating the values for z_1 and z_2, if they both exist, is a little more complicated with the logarithmic allocation function than with the quadratic allocation function. The conditions in Proposition 4 can be used to determine whether z_1, z_2, neither, or both should be calculated. Many algorithms exist to find the zeroes of an equation, and these algorithms can be applied to Equation (8) to find the solutions when $\tilde{R}'(z_{\bar{X}}) = 0$.

To summarize Section 4, both linear and exponential allocation functions should result in allocating resources to reduce only one factor. Which factor should be reduced is determined by the initial values and effectiveness parameters for the linear allocation function but is governed by only the effectiveness parameter for the exponential allocation function. Quadratic or logarithmic allocation functions may result in an optimal allocation to reduce both factors, and

the conditions for how to allocate resources under those allocation schemes are outlined in the propositions.

Preferences of the decision maker could also change the resource allocation decisions. As described in Zobel[12], the decision maker may have preferences that alter the shape of the resilience curves and change the objective function beyond minimizing the product of average loss and recovery time. The resource allocation decisions could change depending on those preferences.

5. MODEL UNDER UNCERTAINTY

Because the decision maker needs to allocate resources prior to a disruption, the nature and possible consequences of the disruption will not be known for certain. The average loss and time to recovery if no resources are allocated, and the parameters describing the effectiveness of allocating resources, may all be uncertain. This section assumes that the parameters $\hat{X}$, $\hat{T}$, $a_{\bar{X}}$, $b_{\bar{X}}$, a_T, and b_T are uncertain. We explore three situations with uncertainty: (1) the uncertain variables have known probability distributions and are independent of each other; (2) the random variables are probabilistically dependent; and (3) a robust optimization method if probability distributions are unknown.

5.1. Allocation Assuming Independence

If the decision maker can assign probability distributions to all of the uncertain parameters, the decision maker's objective could be to maximize expected resilience, $E[R_*(\bar{X}(z_{\bar{X}}), T(z_T))]$. This subsection assumes that the random variables are independent of each other and explores how

the allocation decisions may change compared to the previous cases with certainty. The effectiveness of allocating resources as represented by model parameters $a_{\bar{X}}, b_{\bar{X}}, a_T$, and b_T may be independent of the average impacts, $\hat{X}$, or the time to recovery, $\hat{T}$, especially if the same type of tasks need to be accomplished whether the disruption is large or not. Even if the parameters are probabilistically dependent, assuming independence may be justified as a way to simplify the problem or because the correlation is unknown.

If the allocation functions are linear, quadratic, or logarithmic, it is mathematically possible that $\bar{X}(z_{\bar{X}}) < 0$ or $T(z_T) < 0$ for a given realization of the random parameters even if $E[\bar{X}(z_{\bar{X}})] \geq 0$ and $E[T(z_T)] \geq 0$ for all $0 \leq z_{\bar{X}}, z_T \leq Z$. Because it is physically impossible to have a negative impact or a negative time to recovery, the expected resilience with independence equals $1 - E[\bar{X}(z_{\bar{X}})^+]E[T(z_T)^+]/T^*$ where $\bar{X}(z_{\bar{X}})^+ = \max\{\bar{X}(z_{\bar{X}}), 0\}$ and $T(z_T)^+ = \max\{T(z_T), 0\}$. Consequently, it may be optimal to allocate resources to reduce both factors even if the allocation functions are linear.

If the allocation functions are exponential, $\bar{X}(z_{\bar{X}})$ and $T(z_T)$ will always be positive so the max function is unnecessary. Although the decision maker should allocate to reduce only one factor under certainty, he or she may allocate to reduce both factors under uncertainty.

Proposition 5. If the allocation functions are exponential, the decision maker should allocate to reduce both factors if and only if

$$\max\{E[-a_{\bar{X}}], E[-a_T]\} < \min\left\{\frac{E[-a_{\bar{X}}\exp(-a_X Z)]}{E[\exp(-a_{\bar{X}} Z)]}, \frac{E[-a_T\exp(-a_T Z)]}{E[\exp(-a_T Z)]}\right\}. \tag{11}$$

Proof. See Appendix C. ∎

If the conditions of Proposition 5 are satisfied, the decision maker should choose $z_{\bar{X}}$ such that

$$E[(a_T - a_{\bar{X}})\exp([a_T - a_{\bar{X}}]z_{\bar{X}} - a_T Z)] = 0, \tag{12}$$

which is derived from the conditions for optimality in Equation (16) in Appendix C and by replacing z_T with $Z - z_{\bar{X}}$.

5.2. Allocation Assuming Dependence

The allocation decisions may change if the random variables are probabilistically dependent. This subsection assumes that the decision maker still desires to maximize expected resilience, but the uncertain parameters $\hat{X}, \hat{T}, a_{\bar{X}}, b_{\bar{X}}, a_T$, and b_T are dependent. A disaster that results in large average impacts, $\hat{X}$, will likely also result in a longer recovery time, $\hat{T}$, and the effectiveness of allocating resources could decline as the severity of the disruption increases. Future research can try to assess more explicitly the correlation, if any, among the model parameters. As with independence, calculating expected resilience must account for the mathematical possibility that $\bar{X}(z_{\bar{X}}) < 0$ or $T(z_T) < 0$ when the allocation functions are linear, quadratic, or logarithmic. Expected resilience equals $1 - E[\bar{X}(z_{\bar{X}})^+ T(z_T)^+]/T^*$, and it may be optimal to allocate to reduce both factors for these allocation functions.

If $P(\hat{X} < a_{\bar{X}}Z) = 0$ and $P(\hat{T} < a_TZ) = 0$ when the allocation functions are linear, it is always

optimal to allocate resources to only one factor. This can be seen by expanding the objective

function in Equation (3) where the decision maker should minimize

$$E[\hat{X}\hat{T}] - E[a_{\bar{X}}\hat{T}]z_{\bar{X}} - E[a_T\hat{X}]z_T + E[a_{\bar{X}}a_T]z_{\bar{X}}z_T. \tag{13}$$

Because $E[a_{\bar{X}}\hat{T}]z_{\bar{X}}$ and $E[a_T\hat{X}]z_T$ are the two negative components, it is clear that $z_{\bar{X}} = Z$

minimizes the above equation if $E[a_{\bar{X}}\hat{T}] \geq E[a_T\hat{X}]$ and $z_T = Z$ minimizes Equation (13) if

$E[a_T\hat{X}] \geq E[a_{\bar{X}}\hat{T}]$.

If the allocation functions are exponential, the optimization problem is convex in the two

decision variables, but the allocation rules described in Proposition 5 do not necessarily hold

because of the dependent random variables. Unlike the situation when assuming independence,

the random variables $\hat{X}$ and $\hat{T}$ may influence the optimal allocation if these variables are

probabilistically dependent with $a_{\bar{X}}$ or a_T.

5.3. Robust Allocation

If probability distributions cannot be determined for the uncertain parameters, the decision maker

may choose a "robust" solution where he or she maximizes the worst-case resilience,

i.e. maximize min $R_*(\bar{X}(z_{\bar{X}}), T(z_T))$. This subsection assumes that bounds on each parameter are

known such that

$$\underline{\hat{X}} \leq \hat{X} \leq \bar{\hat{X}}$$

$$\underline{\hat{T}} \leq \hat{T} \leq \bar{\hat{T}}$$

$$\underline{a}_{\bar{X}} \leq a_{\bar{X}} \leq \bar{a}_{\bar{X}} \tag{14}$$

$$\underline{a}_T \leq a_T \leq \bar{a}_T$$

$$\underline{b}_{\bar{X}} \leq b_{\bar{X}} \leq \bar{b}_{\bar{X}}$$

$$\underline{b}_T \leq b_T \leq \bar{b}_T$$

The robust allocation follows the same principles as the allocation under certainty but replaces the certain parameters with the worst-case values of the parameters. If the allocation functions are linear, the decision maker should follow Proposition 1 and allocate the budget Z to the factor that corresponds to $\max\{\underline{a}_{\bar{X}}/\bar{\hat{X}}, \underline{a}_T/\bar{\hat{T}}\}$. If the allocation functions are exponential, allocate Z to the factor that corresponds to $\max\{\underline{a}_{\bar{X}}, \underline{a}_T\}$. If the allocation functions are quadratic, Proposition 3 applies after assigning $\hat{X} = \bar{\hat{X}}, a_{\bar{X}} = \bar{a}_{\bar{X}}, b_{\bar{X}} = \underline{b}_{\bar{X}}, \hat{T} = \bar{\hat{T}}, a_T = \bar{a}_T$, and $b_T = \underline{b}_T$. If the allocation functions are logarithmic, the decision maker should follow the rules in Proposition 4 after assigning $\hat{X} = \bar{\hat{X}}, a_{\bar{X}} = \underline{a}_{\bar{X}}, b_{\bar{X}} = \underline{b}_{\bar{X}}, \hat{T} = \bar{\hat{T}}, a_T = \underline{a}_T$, and $b_T = \underline{b}_T$.

6. ILLUSTRATIVE EXAMPLE: RESILIENCE OF ELECTRIC POWER NETWORK TO SUPERSTORM SANDY

Superstorm Sandy struck the east coast of the United States during the final days of October 2012 and was the second costliest hurricane in U.S. history. Millions of homes were left without electricity, trains and subways were cancelled due to flooding and other impacts, and gasoline

shortages lasted for several weeks in New York and New Jersey. Measuring the resilience of industries or the economic region in the face of Superstorm Sandy provides insight into the preparedness of industries for this type of natural disaster, and the allocation models for enhancing resilience can help these industries optimize their preparedness for a future disruption.

We concentrate on the resilience of the Consolidated Edison Power Company in the New York metropolitan area where about one-fifth of its customers lost power due to Superstorm Sandy. One of us previously collected data (Table I) on the number of households and the proportion of ConEdison's households who lost power for five New York areas (Manhattan, Queens, Brooklyn, Bronx, and Westchester) over a span of 13 days.[16]

[Insert Table I]

After Superstorm Sandy, ConEdison[23] announced a post-Sandy enhancement plan to increase the resilience of its electric power network to severe storms and other natural disasters. The utility's initiatives total $1 billion over four years. Some activities focus on hardening the network to storms, which reduces the impacts or $\bar{X}$ in the resilience model. Hardening measures include trimming trees around power lines, building higher flood barriers, and having backup power for substations. Other activities would help ConEdison restore power more quickly, which reduces the recovery time T. Restoration measures include installing smart-grid technologies, preemptively shutting down steam plants to be able to restart them more quickly after a disruption, and deploying advance teams before a storm.

Table I can be used to estimate that 21.2% of ConEdison's New York customers lost power because of Superstorm Sandy and the average proportion of customers without power per day was 0.073. We use the average daily proportion of customers without power as the most likely value for $\hat{X}$. The minimum value for $\hat{X}$ is 0.030, or the average daily proportion of customers in Brooklyn who lost power, and the maximum value for $\hat{X}$ is 0.216, the average daily proportion of customers in Westchester without power. Because it took ConEdison 13 days to restore electricity after Sandy, the most likely value for $\hat{T}$ is 13 days. Johnson[24] reviews electric outage data for 10 years and records that the minimum duration for an outage from a hurricane was 3 days and the maximum duration was 26 days. We use these values as the minimum and maximum values for $\hat{T}$, and set $T^* = 26$ days to correspond to the maximum value of $\hat{T}$.

Since we have identified minimum, maximum, and most likely values for $\hat{X}$ and $\hat{T}$, we assume that each of these parameters follows a triangle distribution. If these parameters are probabilistically dependent, we assume the correlation coefficient between $\hat{X}$ and $\hat{T}$ is 0.8, which signifies that recovery is likely to take longer when more customers lose power. Based on these values and assumptions, $R_*(\bar{X}, T) = 0.963$ if no additional resources are allocated and the most likely values for $\hat{X}$ and $\hat{T}$ are used. If uncertainty is introduced, $E[R_*(\bar{X}, T)] = 0.943$ assuming independence and $E[R_*(\bar{X}, T)] = 0.937$ assuming dependence. If $\hat{X} = 0.216$ and $\hat{T} = 26$, then the worst-case resilience is 0.784.

We derive the costs and benefits of hardening (see Table II) and recovery (see Table III) activities from a cost-benefit analysis of possible utility upgrades to protect against severe storms,[25] ConEdison's post-Sandy enhancement plan, [23] and a news story on restoring

electricity post-Sandy.[26] First, we separate the dozen activities from these studies into activities

that reduce $\bar{X}$ and those that improve T, and we scale the costs of the activities to reflect the 3.3

million customers that ConEdison serves in the New York area. The benefits from each activity

are expressed in terms of the percentage of damages avoided or the percentage of restoration

time reduction. Second, after sorting each group of activities from the greatest to the smallest

benefit-cost ratio, we calculate the cumulative benefits of multiple activities by multiplying

probabilities. For example, vegetation removal reduces impacts by 7.6% and backup power for

substations reduces impacts by 30%. Choosing both activities reduces impacts by $1 -$

$(1 - .076) * (1 - .3) = 0.353 = 35.2\%$. Finally, $\bar{X}$ and T are estimated for each combination of

activities by multiplying the most likely initial values $\hat{X}$ and $\hat{T}$ by $1 - p$ where p is the

cumulative benefit.

[Insert Table II]

[Insert Table III]

Although ConEdison could use the information in Tables II and III to select the most cost-

effective activities, it is not clear how much ConEdison should allocate towards hardening versus

recovery activities. Also, the tables depict discrete activities, and modeling the resource

allocation as a continuous optimization problem could provide insight into the marginal benefit

of increasing the budget by a small amount. Since much of the data is derived from Brown,[25]

the activities do not map perfectly to ConEdison's plans for enhancing its network. In general, an

organization might not have information on the benefits and costs of each activity or project, but

it still might have some knowledge or make assumptions about the benefits and costs of reducing

impacts and improving recovery time. This example illustrates how the resource allocation model for resilience could help ConEdison determine how best to make its electric power network more resilient to severe weather and account for uncertainty in the severity of the storm and effectiveness of its activities.

Tables II and III are used to estimate the effectiveness parameters for each of the four allocation functions discussed earlier. The most likely effectiveness parameter is estimated so that the corresponding allocation function represents a best-fit line (see Fig. 2) to the cumulative costs and benefits from Tables II and III. The minimum and maximum values for each parameter are calculated to encompass a 99% confidence interval for the best-fit line. As with $\hat{X}$ and $\hat{T}$, we assume each of the effectiveness parameters follows a triangle distribution when uncertainty is included in the model. Table IV depicts the most likely, minimum, and maximum value for each of the parameters used in the optimization model. We assume the effectiveness parameters are negatively correlated with $\hat{X}$ and $\hat{T}$ when the model accounts for dependent probabilities. Negative correlation means that resources are less likely to be effective if more customers lose power or restoring power takes longer. Specifically, the correlation coefficient between $\hat{X}$ and $a_{\bar{X}}$ and $b_{\bar{X}}$ is -0.8 and the correlation coefficient between $\hat{T}$ and a_T and b_T is -0.8. The budget Z equals $1 billion, which is the amount that ConEdison has planned to spend to improve its electric power network over the next four years.

[Insert Fig. 2]

[Insert Table IV]

The optimal allocations and the resulting resilience are displayed in Table V for the different

allocation functions under certainty and under uncertainty. The linear and exponential allocation

functions recommend opposite allocation amounts. According to Proposition 1, ConEdison

should allocate the entire budget to reducing recovery time if it uses a linear allocation function

because $a_T/\hat{T} = 0.00061 > a_{\bar{X}}/\hat{X} = 0.00049$. According to Proposition 2, ConEdison should

allocate the entire budget to reducing the number of customers who lose power if it uses an

exponential allocation function because $a_{\bar{X}} = 0.0088 > a_T = 0.0008$. These allocation amounts

should not change if the parameters follow the triangle distribution, and Proposition 5 is not

satisfied. Because $\underline{a_{\bar{X}}}$ is so small for both the linear and exponential allocation functions,

planning for the worst case means that ConEdison should allocate the entire budget to reduce T.

[Insert Table V]

If a quadratic or logarithmic allocation function is used, ConEdison should spend money on both

hardening and improving recovery. $762 million should be spent on hardening and $238 million

should be spent on improving recovery with the quadratic allocation function and the most likely

parameters, but introducing uncertainty with independent probabilities changes the optimal

allocation to $z_{\bar{X}} = \$556$ and $z_T = \$444$ million. If the uncertain parameters are correlated, $z_{\bar{X}} =$

$850 and $z_T = \$160$ million, which is an example of how assuming dependence among

uncertain parameters can significantly change the optimal allocation.

According to the logarithmic allocation function with certainty, ConEdison should spend $648

million on reducing the number of customers who lose power and $352 million on reducing the

restoration time. ConEdison should distribute the allocation approximately evenly if the parameters are uncertain: $z_{\bar{X}} = \$494$ and $z_T = \$506$ million assuming independence and $z_{\bar{X}} = \$470$ and $z_T = \$530$ million assuming dependence. A robust allocation means that ConEdison should allocate \$286 million on hardening and \$714 million on recovery because the worst-case effectiveness parameters for reducing the impacts are small.

This example demonstrates the variety of solutions that may be optimal depending upon which allocation function is used and whether or not the parameters are uncertain. The linear model recommends allocating the entire budget to recovery, the exponential model recommends allocating the entire budget to hardening, and the quadratic and logarithmic models recommend dividing the budget between the two activities. ConEdison should follow the recommendation from the model that it believes is the most accurate. Fig. 2 shows how well each allocation function matches the data using the most likely parameters. The logarithmic allocation function appears to fit the data the best, which is reasonable considering the cumulative benefits are the product of multiple percentages.

Another way to compare the models is to examine the resilience of ConEdison's network if it allocates resources based on the wrong model. Table VI depicts the resilience if ConEdison allocates resources based on one model and another model is correct. If ConEdison follows the logarithmic function with certainty and chooses $z_{\bar{X}} = \$648$ and $z_T = \$352$ million but there really is randomness with the parameters, ConEdison's expected resilience is 0.977 (assuming independence) or 0.969 (assuming dependence). These resilience metrics are identical to three decimal places to the expected resilience if ConEdison allocates $z_{\bar{X}}$ and z_T according to the

optimal solution from the logarithmic model with uncertainty. As long as the logarithmic

allocation function is correct, the resilience of the network hardly changes based on whether

parameters are certain or uncertain.

[Insert Table VI]

If ConEdison allocates $z_{\bar{X}} = \$648$ and $z_T = \$352$ million based on the logarithmic model but

the linear model is correct, ConEdison's resilience is 0.980 compared to a resilience of 0.986 if

the entire budget is spent on recovery as the linear model recommends. Since resilience without

any resources is 0.963, $R_*(\bar{X}, T) = 0.980$ represents almost 75% improvement to the best

resilience of 0.986. If ConEdison allocates $648 million to hardening and $352 million to

recovery, and the linear model is correct, a daily average of 159,000 customers will be without

power for 10.2 days. Allocating $1 billion to recovery means a daily average of 232,000

customers without power for 5.1 days. According to the resilience function, the latter scenario is

preferable because this scenario has, on average, fewer customers without power per day.

However, the former scenario does not seem much worse than the latter scenario.

If ConEdison spends $1 billion on restoration according to the linear model, but the logarithmic

model is correct, $R_*(\bar{X}, T) = 0.977$ as opposed to 0.989 if $z_{\bar{X}} = \$648$ and $z_T = \$352$ million. If

there is uncertainty, $E[R_*(\bar{X}, T)] = 0.963$ if $z_T = \$1$ billion but $E[R_*(\bar{X}, T)] = 0.977$ if $z_{\bar{X}} =$

$648 and $z_T = \$352$ million. Given that $R_*(\bar{X}, T) = 0.963$ and $E[R_*(\bar{X}, T)] = 0.943$ if no

resources are allocated, using the allocation according to the linear model only improves

resilience a little more than half as much as the optimal allocation if the logarithmic model is

correct. Allocating the entire budget to restoration as recommended by the linear model means

that a daily average of 232,000 customers lose power for 8.0 days according to the logarithmic model. Allocating $648 million to hardening and $352 million to restoration means that a daily average of 101,000 customers lose power for 8.7 days. Misallocating resources because the incorrect model is used can cause 131,000 additional customers to lose power with only a gain of 0.7 days in restoration time. Given the parameters in this example, following the optimal allocation from the logarithmic still provides good resilience even if another model is correct, but following the optimal allocation from another model such as the linear model can be very sub-optimal if the logarithmic model is correct. Consequently, ConEdison should spend $500 to $650 million to reduce the number of customers that lose power and spend $350 to $500 million to improve the recovery time.

7. CONCLUSIONS

Based on a widespread definition and mathematical formula for resilience, allocating resources to enhance resilience can either lessen the impacts from a disruption or improve recovery time. How much a decision maker should allocate to one of those two factors depends to a large extent on how he or she believes resources affects those factors and on the level and type of uncertainty. If resources reduce a factor in either a linear or exponential manner with no uncertainty, the decision maker should allocate all the resources to reduce only one factor. The marginal benefits for quadratic and logarithmic allocation functions with certainty decrease more rapidly than linear or exponential functions, and it is often optimal to allocate resources to lessen the impacts and improve recovery with quadratic and logarithmic functions.

Practically all situations that call for enhancing resilience will involve some uncertainty, and maximizing expected resilience means that a decision maker should allocate resources to reduce both factors more often than if he or she is maximizing resilience with certainty. If probabilities cannot be determined and the decision maker wishes to maximize the worst-case resilience, he or she should follow the same rules as for allocation under certainty but the optimal resilience in the robust case will be less than the optimal resilience in the certain case.

Because of the assumption that resilience is calculated as the product of impact and recovery time, a good heuristic is to focus resources on the factor with a small initial value and high effectiveness. That heuristic is complicated when uncertainty is considered. Assuming that the effectiveness and initial value parameters are uncertain but independent may mean that resources are allocated to reduce both factors, especially when it is believed that the marginal benefits decrease rapidly as in the case of the quadratic or logarithmic functions. Dividing the resources in an approximately equal manner is often a good strategy under this assumption.

Allocating resources to lessen the impacts may include system hardening, building redundancy, and better emergency response systems that can reduce the impacts. Allocating resources to improve recovery time may focus on activities such as improving capabilities for repairing and rebuilding and prepositioning supplies that will be needed for recovery. This paper has applied the resilience model to ConEdison's plan for building a more resilient electric power network after Superstorm Sandy. We estimate the effectiveness parameters for the different allocation functions based on ConEdison's own analysis and another cost-benefit analysis of electric power network upgrades. The logarithmic allocation function appears to be the best fit to the data

derived from the studies, and ConEdison should allocate between 50 and 65% of its budget to reduce the number of customers who lose power and the rest of the budget to improve recovery time.

The ConEdison example demonstrates how the resilience allocation model can guide decision makers in preparing for disruptions. The parameters for allocation functions can be estimated either by matching the functions to data or by asking experts about the benefits and costs of activities. After choosing the allocation function that best fits the data or meets the assumptions about how the benefits change as more resources are allocated, the decision maker can divide his or her budget according to optimal allocation as recommended by the optimization model. The recommended solution can be tested for robustness to model inaccuracy—the optimal allocation according to the logarithmic model performed fairly well in the ConEdison example if the linear model is correct—and can be analyzed for how uncertainty impacts the optimal allocation.

After determining the proper division between reducing the average impacts and improving recovery time, the decision maker needs to determine the specific projects to pursue based on the cost-effectiveness of these projects. For example, Brown's[25] analysis finds that removing hazardous trees, having backup power generators, and hardening existing distribution components as the most cost-effective projects to reduce the number of customers who lose power. If ConEdison follows this analysis, it should allocate much of the $650 million designated to reduce $\bar{X}$ to these projects. Future research can analyze how the resource allocation model for resilience can be expanded to decide which individual projects to fund.

APPENDIX

Appendix A: Proof of Proposition 3.

Condition 1. If $A > 0$, $\bar{R}(0)$ is increasing because $\bar{R}'(0) = A$. Since $\bar{R}(z_{\bar{X}})$ approaches positive infinity as $z_{\bar{X}}$ approaches negative infinity, $\bar{R}'(z_{\bar{X}}) < 0$ for a very negative $z_{\bar{X}}$. $\bar{R}'(0) > 0$, which means that the first local minimum z_1 occurs when $z_{\bar{X}} < 0$. If $z_2 < Z$, then $z_{\bar{X}} = z_2$ is a global minimum if $\bar{R}(z_2) \leq 0$ because $\bar{R}(0) = 0$. Otherwise, $z_{\bar{X}} = 0$ is the global minimum because $\bar{R}(z_{\bar{X}})$ is increasing when $z_{\bar{X}} = Z$ if $\bar{R}'(Z) > 0$.

Condition 2. As with the first condition, if $A > 0$, the first local minimum z_1 occurs when $z_{\bar{X}} < 0$. Since $\bar{R}(z_{\bar{X}})$ approaches positive infinity as $z_{\bar{X}}$ approaches positive infinity, $\bar{R}'(z_{\bar{X}}) > 0$ for a very positive $z_{\bar{X}}$. If $\bar{R}'(Z) < 0$, the second local minimum z_2 occurs when $z_{\bar{X}} > Z$. Thus, only two possible minima occur when $z_{\bar{X}} = 0$ or Z. If $\bar{R}(Z) \geq 0$, a global minimum occurs when $z_{\bar{X}} = 0$. Otherwise, the global minimum occurs when $z_{\bar{X}} = Z$.

Condition 3. If $A < 0$, then $\bar{R}(0)$ is decreasing and $z_{\bar{X}} = 0$ cannot be a minimum. If $\bar{R}'(Z) > 0$, then $\bar{R}(Z)$ is increasing and $z_{\bar{X}} = Z$ cannot be a minimum. Thus, at least one local minimum exists between 0 and Z. If $0 < z_1, z_2 < Z$ and $\bar{R}(z_1) \leq \bar{R}(z_2)$, $z_{\bar{X}} = z_1$ is a global minimum. If $z_2 > Z$, then $0 < z_1 < Z$ and the global minimum occurs at $z_{\bar{X}} = z_1$. If neither of those conditions are met, then $z_1 < 0$, which implies that $z_2 < Z$, and $z_{\bar{X}} = z_2$ is the global minimum.

Condition 4. As with the third condition, if $A < 0$ then $z_{\bar{X}} = 0$ cannot be a minimum. As with the second condition, if $\bar{R}'(Z) < 0$, the $z_2 > Z$. If $z_1 > 0$ and $\bar{R}(z_1) \leq \bar{R}(Z)$, then $z_{\bar{X}} = z_1$ is a

global minimum. If $z_1 < 0$ or $\bar{R}(z_1) > \bar{R}(Z)$, then $z_{\bar{X}} = Z$ is the global minimum because $\bar{R}(Z)$

is decreasing. ∎

Appendix B: Proof of Proposition 4

The proof is similar to that for Proposition 3. Condition 1. If $\tilde{R}'(0) > 0$, $\tilde{R}(z_{\bar{X}})$ is increasing

when $z_{\bar{X}} = 0$. Because $\tilde{R}(z_{\bar{X}})$ is increasing, we know the first extreme point (if one exists) is a

local maximum. If two local minima exist, $z_1 < 0$. If $z_2 < Z$, then $z_{\bar{X}} = z_2$ is a global minimum

if $\tilde{R}(z_2) \leq \log(\hat{X}) + \log(\hat{T} - a_T \log(1 + b_T Z)) = \tilde{R}(0)$. Otherwise, $z_{\bar{X}} = 0$ is the global

minimum because $\tilde{R}(z_{\bar{X}})$ is increasing when $z_{\bar{X}} = Z$ if $\tilde{R}'(Z) > 0$.

Condition 2. As with the first condition, if $\tilde{R}'(0) > 0$, the first local minimum z_1 is less than 0 if

it exists. $\tilde{R}'(Z) < 0$ implies that $\tilde{R}(Z)$ is increasing. If z_2 were less than Z, this would mean that

two local maxima exist because the first local maximum is less than z_2. From Lemma 2, we

know that at most one local maximum exists. Thus, $z_2 > Z$ and the only two possible minima

occur when $z_{\bar{X}} = 0$ or Z. If $\hat{X}(\hat{T} - a_T \log(1 + b_T Z)) \leq \hat{T}(\hat{X} - a_{\bar{X}} \log(1 + b_{\bar{X}} Z))$, $\tilde{R}(0) \leq \tilde{R}(Z)$,

and a global minimum occurs when $z_X = 0$. Otherwise, the global minimum occurs when $z_{\bar{X}} =$

Z.

Condition 3. If $\tilde{R}'(0) < 0$, then $\tilde{R}(0)$ is decreasing and $z_{\bar{X}} = 0$ cannot be a minimum. If $\tilde{R}'(Z) >$

0, then $\bar{R}(Z)$ is increasing and $z_{\bar{X}} = Z$ cannot be a minimum. Thus, at least one local minimum

exists between 0 and Z. If $0 < z_1, z_2 < Z$ and $\tilde{R}(z_1) \leq \tilde{R}(z_2)$, $z_{\bar{X}} = z_1$ is a global minimum. If

$z_2 > Z$, then $0 < z_1 < Z$ and the global minimum occurs at $z_{\bar{X}} = z_1$. If neither of those

conditions are met, then $z_1 < 0$, which implies that $z_2 < Z$, and $z_{\bar{X}} = z_2$ is the global minimum.

Condition 4. As with the third condition, if $\tilde{R}'(0) < 0$ then $z_{\bar{X}} = 0$ cannot be a minimum. As

with the second condition, if $\tilde{R}'(Z) < 0$, then $z_2 > Z$. If $z_1 > 0$ and $\tilde{R}(z_1) \leq \log\big(\hat{X} -$

$a_{\bar{X}} \log(1 + b_{\bar{X}}Z)\big) + \log(\hat{T}) = \tilde{R}(Z)$, then $z_{\bar{X}} = z_1$ is a global minimum. If $z_1 < 0$ or $\tilde{R}(z_1) >$

$\log\big(\hat{X} - a_X \log(1 + b_{\bar{X}}Z)\big) + \log(\hat{T}) = \tilde{R}(Z)$, then $z_{\bar{X}} = Z$ is the global minimum since $\tilde{R}(Z)$ is

decreasing. ∎

Appendix C: Proof of Proposition 5

Because the random variables are independent, maximizing expected resilience is equivalent to

$$\text{minimize } \log E[\exp(-a_{\bar{X}}z_{\bar{X}})] + \log E[\exp(-a_T z_T)]. \tag{15}$$

Each component of Equation (15) is identical to a cumulant-generating function, which is a

convex function.[22] A solution that satisfies the Karush-Kuhn-Tucker conditions is the unique

minimum to Equation (15).

The optimal solution must satisfy the following equations:

$$\frac{E[-a_{\bar{X}}\exp(-a_{\bar{X}}z_{\bar{X}})]}{E[\exp(-a_{\bar{X}}z_{\bar{X}})]} - \lambda_{\bar{X}} = \frac{E[-a_T\exp(-a_T z_T)]}{E[\exp(-a_T z_T)]} - \lambda_T \tag{16}$$

$$z_X + z_T = Z$$

where $\lambda_{\bar{X}}, \lambda_T \geq 0$ are the Lagrange multipliers for the non-negativity constraints for $z_{\bar{X}}$ and z_T, respectively. The Lagrange multiplier $\lambda_{\bar{X}} > 0$ if and only if $z_{\bar{X}} = 0$ and $\lambda_T > 0$ if and only if $z_T = 0$.

We first prove that if Equation (11) is true, it is optimal to allocate resources to reduce both factors. Without loss of generality, assume that $E[-a_{\bar{X}}] = \max\{E[-a_{\bar{X}}], E[-a_T]\}$ and

$$E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)] = \min\left\{\frac{E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]}{E[\exp(-a_{\bar{X}} Z)]}, \frac{E[-a_T \exp(-a_T Z)]}{E[\exp(-a_T Z)]}\right\} \text{ and that}$$

$E[-a_{\bar{X}}] < E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)]$. If $z_{\bar{X}} = 0$ and $z_T = Z$, $\lambda_{\bar{X}} > 0$ and $\lambda_T = 0$. But then Equation (16) would not be true because $E[-a_{\bar{X}}] <$

$E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)]$. Thus, $z_{\bar{X}} = 0$ and $z_T = Z$ is not an optimal allocation.

If $z_{\bar{X}} = Z$ and $z_T = 0$, then $\lambda_{\bar{X}} = 0$ and $\lambda_T > 0$. In order for Equation (16) to be true, it must be true that $E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)] < E[-a_T]$. Due to the convexity of the optimization problem, the gradient is strictly increasing in each variable and $E[-a_X] <$

$E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)]$. Since $E[-a_{\bar{X}}] \geq E[-a_T]$, it must be true that

$E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)] > E[-a_T]$. Thus $z_{\bar{X}} = Z$ and $z_T = 0$ is not an optimal allocation.

If $E[-a_{\bar{X}}] = \max\{E[-a_{\bar{X}}], E[-a_T]\}$ and $E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)] =$

$\min\left\{\frac{E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]}{E[\exp(-\bar{X} Z)]}, \frac{E[-a_T \exp(-a_T Z)]}{E[\exp(-a_T Z)]}\right\}$, the same argument applies. If $z_{\bar{X}} = 0$ and $z_T = Z$,

Equation (16) would not be true because $E[-a_{\bar{X}}] < E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_X Z)] \leq$

$E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)]$. If $z_{\bar{X}} = Z$ and $z_T = 0$, Equation (16) would not be true because $E[-a_T] \leq E[-a_{\bar{X}}] < E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)]$.

Therefore, if Equation (11) is true, the only solution that satisfies Equation (16) is when $z_{\bar{X}}, z_T > 0$, and it is optimal to allocate resources to reduce both factors.

We next prove that if Equation (11) is not true, then either $z_{\bar{X}}$ or z_T equals 0. Assume without loss of generality that $E[-a_{\bar{X}}] = \max\{E[-a_{\bar{X}}], E[-a_T]\}$ and

$$E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)] = \min\left\{\frac{E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]}{E[\exp(-a_{\bar{X}} Z)]}, \frac{E[-a_T \exp(-a_T Z)]}{E[\exp(-a_T Z)]}\right\} \text{ but that}$$

$E[-a_{\bar{X}}] \geq E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)]$. If $z_{\bar{X}} > 0$ and $z_T > 0$, then $\lambda_{\bar{X}} = \lambda_T = 0$. Such a solution does not satisfy Equation (16) because each gradient is strictly increasing and the gradient when $z_{\bar{X}} = 0$, $E[-a_{\bar{X}}]$, is greater than or equal to the gradient when $z_T = Z$,
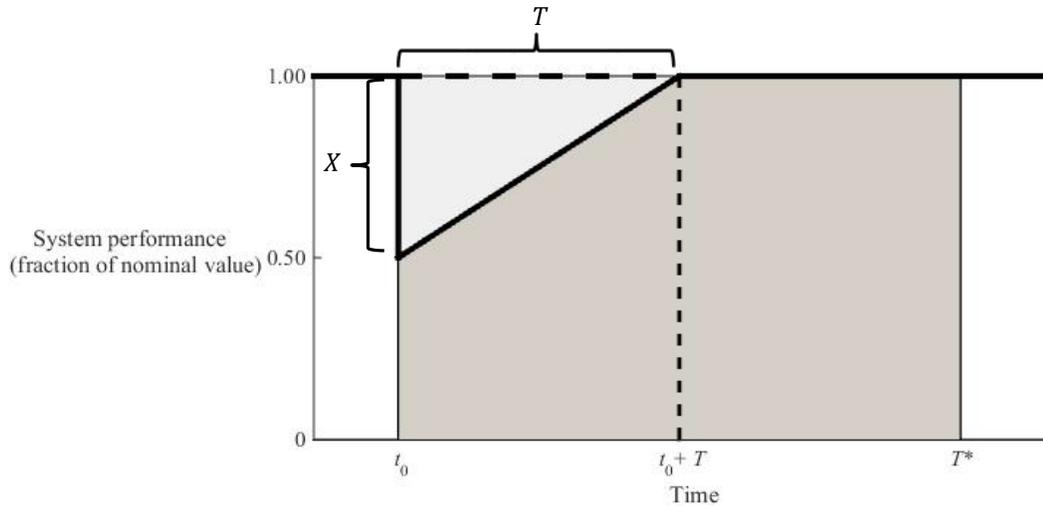
$$E[-a_T \exp(-a_T Z)]/E[\exp(-a_T Z)].$$

Because each gradient is strictly increasing, it is impossible that $E[-a_{\bar{X}}] \geq E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)]$. If Equation (11) is not true, it cannot be true that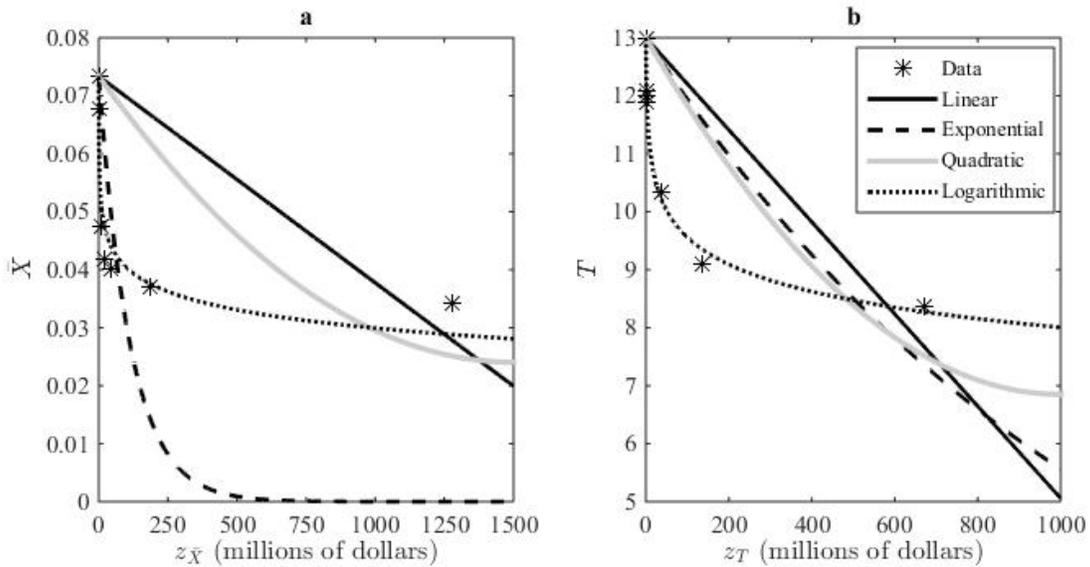 $E[-a_{\bar{X}}] = \max\{E[-a_{\bar{X}}], E[-a_T]\}$ and $E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]/E[\exp(-a_{\bar{X}} Z)] = $

$\min\left\{\frac{E[-a_{\bar{X}} \exp(-a_{\bar{X}} Z)]}{E[\exp(-a_{\bar{X}} Z)]}, \frac{E[-a_T \exp(-a_T Z)]}{E[\exp(-a_T Z)]}\right\}$. Thus, if Equation (11) is not true, $z_{\bar{X}} > 0$ and $z_T > 0$ is not an optimal allocation. ■

Fig. 1. Measuring disaster resilience.

Fig. 2. Best-fit allocation functions for (a) $\bar{X}(z_{\bar{X}})$ and (b) $T(z_T)$.

Table I. Power Outages in New York due to Superstorm Sandy
Data from Zobel[16]

	Manhattan	Queens	Brooklyn	Bronx	Westchester
Number of households without power	250,000	108,000	87,000	45,000	180,000
Initial proportion of customers without power	0.3458	0.1436	0.0954	0.1047	0.5172
Average proportion of customers without power per day	0.0935	0.0624	0.0302	0.0346	0.2163

Table II. Costs and Benefits of Hardening Activities

Activity (ordered from greatest to smallest benefit-cost ratio)	Cost (millions of dollars)	Benefits (percentage of damage reduction)	Cumulative cost (millions of dollars)	Cumulative benefits (percentage of damage reduction)	$\bar{X}$ given the cumulative benefits
1. Vegetation removal	1.3	7.6	1.3	7.6	0.0678
2. Backup power for substations	9.9	30.0	11.2	35.3	0.0475
3. New substations	9.9	11.6	21.1	42.8	0.0420
4. Ground-based inspection	21.8	4.0	42.9	45.1	0.0403
5. Hardened distribution	145.0	8.0	187.9	49.5	0.0371
6. Hardened transmission	1087.7	8.0	1275.7	53.5	0.0341

Note: Cumulative costs and benefits are calculated using the costs and benefits of an individual activity combined with all the activities that have greater benefit-cost ratios.

Table III. Costs and Benefits of Recovery Activities

Activity (ordered from greatest to smallest benefit-cost ratio)	Cost (millions of dollars)	Benefits (percentage of restoration time reduction)	Cumulative cost (millions of dollars)	Cumulative benefits (percentage of restoration time reduction)	T given the cumulative benefits
1. Deployment of restoration teams in advance	1.8	7.1	1.8	7.1	12.1
2. Advanced power network monitoring	0.2	0.5	2.0	7.6	12.0
3. Automatic fault location	0.7	1.0	2.7	8.5	11.9
4. Distribution automation and smart feeders	36.5	13.0	39.1	20.4	10.3
5. Distributed generation penetration	95.9	12.0	135.1	30.0	9.1
6. Advanced metering infrastructure	537.1	8.0	672.2	35.6	8.4

Table IV. Parameters for Optimization Model

		Most likely	Minimum	Maximum
	$\hat{X}$	0.0734	0.0302	0.2163
	$\hat{T}$	13	3	26
Linear	$a_{\bar{X}}$	$3.56 * 10^{-5}$	$1.00 * 10^{-7}$	$9.24 * 10^{-5}$
	a_T	0.00794	$5.53 * 10^{-4}$	0.0153
Exponential	$a_{\bar{X}}$	0.00878	$1.00 * 10^{-7}$	0.0331
	a_T	$8.49 * 10^{-4}$	$1.00 * 10^{-5}$	0.00220
Quadratic	$a_{\bar{X}}$	$2.19 * 10^{-8}$	$5.00 * 10^{-11}$	$1.71 * 10^{-7}$
	$b_{\bar{X}}$	$6.58 * 10^{-5}$	$1.00 * 10^{-7}$	$5.14 * 10^{-4}$
	a_T	$6.15 * 10^{-6}$	$2.76 * 10^{-7}$	$1.47 * 10^{-5}$
	b_T	0.0123	$5.53 * 10^{-4}$	0.0294
Logarithmic	$a_{\bar{X}}$	0.00455	0.00215	0.00698
	$b_{\bar{X}}$	14.1	8.24	68.4
	a_T	0.677	0.589	0.765
	b_T	1.60	1.51	1.73

Table V. Optimal Allocation Amounts and Optimal Resilience for Different Allocation Functions
(allocation amounts in millions of dollars)

	Certainty					Uncertainty with independence					Uncertainty with dependence					Robust allocation				
	None	Lin	Exp	Quad	Log	None	Lin	Exp	Quad	Log	None	Lin	Exp	Quad	Log	None	Lin	Exp	Quad	Log
$z_{\bar{X}}$	0	0	1000	762	648	0	0	1000	556	494	0	0	1000	840	470	0	0	0	21	286
z_T	0	1000	0	238	352	0	1000	0	444	506	0	1000	0	160	530	0	1000	1000	979	714
$R_*(\bar{X}, T)$	0.963	0.986	1.000	0.986	0.989	0.943	0.974	1.000	0.985	0.977	0.937	0.965	0.999	0.976	0.969	0.784	0.788	0.786	0.786	0.832

Note: Lin = linear allocation function; Exp = exponential allocation function; Quad = quadratic allocation function; Log = logarithmic allocation function

Table VI. Resilience Given $z_{\bar{X}}$ and z_T According to Different Allocation Models

	Certainty				Uncertainty with independence				Uncertainty with dependence				Robust allocation			
	Lin	Exp	Quad	Log	Lin	Exp	Quad	Log	Lin	Exp	Quad	Log	Lin	Exp	Quad	Log
$z_{\bar{X}} = 0, z_T = 1000$	0.986	0.984	0.981	0.977	0.974	0.977	0.968	0.963	0.964	0.970	0.961	0.957	0.788	0.786	0.786	0.820
$z_{\bar{X}} = 648, z_T = 352$	0.980	1.000	0.986	0.989	0.966	0.999	0.985	0.977	0.958	0.998	0.975	0.969	0.785	0.785	0.785	0.830
$z_{\bar{X}} = 494, z_T = 506$	0.981	1.000	0.985	0.989	0.967	0.999	0.985	0.977	0.959	0.997	0.974	0.969	0.786	0.785	0.786	0.831

REFERENCES

1. The National Academies (US), Committee on Increasing National Resilience to Hazards and Disasters, Committee on Science, Engineering, and Public Policy. Disaster Resilience: A National Imperative [Internet]. Washington, DC: The National Academies Press; 2012 [cited 2015 Jul 23]. Available from: http://www.nap.edu/catalog.php?record_id=13457.

2. The White House. Presidential Policy Directive: Critical Infrastructure Security and Resilience [Internet]. Office of the Press Secretary, PPD-21, Washington, DC, 2013 Feb 12 [cited 2014 Sep 21]. Available from: http://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil.

3. Bruneau M, Chang SE, Eguchi RT, Lee GC, O'Rourke TD, Reinhorn AM, Shinozuka M, Tierney K, Wallace W, von Winterfeldt D. A framework to quantitatively assess and enhance the seismic resilience of communities. Earthquake Spectra, 2003; 19(4):733-752.

4. Zobel CW. Representing the multi-dimensional nature of disaster resilience. In: Santos MA, Sousa L, Portela E, editors. From Early-Warning Systems to Preparedness and Training. Proceedings of the 8th International Conference on Information Systems for Crisis Response and Management; 2011 May 8-11; Lisbon, Portugal.

5. Chan NH, Wong HY. Data mining of resilience indicators. IIE Transactions, 2007; 39(6):617-627.

6. Cutter SL, Barnes L, Berry M, Burton C, Evans E, Tate E, Webb J. A place-based model for understanding community resilience to natural disasters. Global environmental change, 2008; 18(4):598-606.

7. Ayyub BM. Systems resilience for multi-hazard environments: Definitions, metrics and valuation for decision making. Risk Analysis, 2014; 34(2):340-355.

8. Pant R, Barker K, Zobel CW. Static and dynamic metrics of economic resilience for interdependent infrastructure and industry sectors, Reliability Engineering & System Safety, 2014; 125(1):92-102.

9. Cimellaro G, Reinhorn A, Bruneau M. Seismic resilience of a hospital system. Structure and Infrastructure Engineering, 2010; 6(1):127-144.

10. Zobel CW. Comparative visualization of predicted disaster resilience. In: French S, Tomaszewski B, Zobel C, editors. Defining Crisis Management 3.0. Proceedings of the 7th International Conference on Information Systems for Crisis Response and Management; 2010 May 2-5; Seattle, WA.

11. Zobel CW, Khansa LZ. Quantifying cyberinfrastructure resilience against multi-event attacks. Decision Sciences, 2012; 43(4):687-710.

12. Zobel CW. Representing perceived tradeoffs in defining disaster resilience. Decision Support Systems, 2011; 50(2):394-403.

13. MacKenzie CA, Barker K. Empirical data and regression analysis for estimation of infrastructure resilience with application to electric power outages. Journal of Infrastructure Systems, 2013; 19(1):25-35.

14. Baroud H, Barker K, Ramirez-Marquez JE, Rocco CM. Inherent costs and interdependent impacts of infrastructure network resilience. Risk Analysis, 2015; 35(4):642-662.

15. Bruneau M, Reinhorn A. Exploring the concept of seismic resilience for acute care facilities. Earthquake Spectra, 2007; 23(1):41-62.

16. Zobel CW. Quantitatively representing nonlinear disaster recovery. Decision Sciences, 2014; 45(6):1053-1082.

17. MacKenzie CA, Barker K, Grant FH. Evaluating the consequences of an inland waterway port closure with a dynamic multiregional interdependency model. IEEE Transactions on Systems, Man, Cybernetics, Part A: Systems and Humans, 2012; 42(2):359–370.

18. Henry D, Ramirez-Marquez JE. Generic metrics and quantitative approaches for system resilience as a function of time. Reliability Engineering and System Safety 2012; 99:114-122.

19. Barker K, Ramirez-Marquez JE, Rocco CM. Resilience-based network component importance measures. Reliability Engineering and System Safety, 2013; 117:89-97.

20. Pant R, Barker K, Ramirez-Marquez JE, Rocco CM. Stochastic measures of resilience and their application to container terminals. Computers & Industrial Engineering, 2014; 70:183-194.

21. Baroud H, Ramirez-Marquez JE, Barker K, Rocco CM. Stochastic measures of network resilience: Applications to waterway commodity Flows. Risk Analysis, 2014; 34(7):1317-1335.

22. Boyd S, Vandenberghe L. Convex Optimization. Cambridge, UK: Cambridge University Press; 2006.

23. Consolidated Edison Co. of New York. Post Sandy enhancement plan [Internet]. Orange and Rockland Utilities, 2013 [cited 2015 Mar 23]. Available from: http://www.coned.com/publicissues/PDF/post_sandy_enhancement_plan.pdf.

24. Johnson BW. After the disaster: Utility restoration cost recovery [Internet]. Report prepared for the Edison Electric Institute, 2005 Feb [cited 2015 Apr 15]. Available from: https://www.efis.psc.mo.gov/mpsc/commoncomponents/viewdocument.asp?DocId=404818.

25. Brown R. Cost-benefit analysis of the deployment of utility infrastructure upgrades and storm hardening programs [Internet]. Report prepared for the Public Utility Commission of Texas. Quanta Technology, 2009 [cited 2015 Mar 23]. Available from: https://www.puc.texas.gov/industry/electric/reports/infra/Utlity_Infrastructure_Upgrades_rpt.pdf.

26. Terruso J, Baxter C, Caroom E. Sandy recovery becomes national mission as countless workers come to N.J.'s aid. The Star-Ledger (Newark) [Internet]. 2012 Nov 11 [cited 2015 Mar 23]. Available from: http://www.nj.com/news/index.ssf/2012/11/sandy_relief_becomes_national.html.