Int. J. Emergency Management, Vol. 2, No. 4, 2005 263

Protection from extreme events: using a socio-technological approach to evaluate policy options

Richard G. Little*

University of Southern California, USA
The Keston Institute for Infrastructure
331 D. Lewis Hall, Los Angeles, CA 90089, USA
Fax: 01 203 740–6170
E-mail: rglittle@usc.edu
*Corresponding author

Elise A. Weaver

Worcester Polytechnic Institute, USA
Department of Social Science and Policy Studies
100 Institute Road, Worcester, MA 01609, USA
Fax: 01 508 831–5896
E-mail: eweaver@wpi.edu

Abstract: Effective ways to address the vulnerabilities of urban areas to terrorism and other hazards have been subject to considerable discussion, debate and reflexive defensive measures. However, a coherent strategy for protecting public spaces while maintaining access to them has yet to emerge. Current approaches have pitted 'security' against 'openness,' neglecting critical issues, such as what constitutes publicly-valued levels of protection, a prudent government response, and sustainable public expenditures. Although a desire for drastic measures is certainly understandable, it is not based on a true assessment of risk, nor will it lead to an effective, let alone cost effective, approach to the threats of hazards and urban terrorism. This paper will discuss three tools from the social and policy sciences that can be used to develop balanced and prioritised approaches to physical security. It will demonstrate that these tools also can be applied to planning for extreme natural events.

Keywords: physical security; extreme events; social judgment theory; risk management.

Reference to this paper should be made as follows: Little, R.G. and Weaver, E.A. (2005) 'Protection from extreme events: using a socio-technological approach to evaluate policy options', *Int. J. Emergency Management*, Vol. 2, No. 4, pp.263–274.

Biographical notes: Richard G. Little is Director of the Keston Institute for Infrastructure, a policy research centre at the University of Southern California. He is the immediate past Director of the Board on Infrastructure and the Constructed Environment of the National Research Council and has directed studies and lectured and written extensively on urban security strategy and decision making for extreme events. He is Member of the American Planning

Association, the American Institute of Certified Planners, and the Society for Risk Analysis. He holds a BS in Geology and an MS in Urban-Environmental Studies, both from Rensselaer Polytechnic Institute.

Elise A. Weaver is Assistant Professor of Psychology and System Dynamics at Worcester Polytechnic Institute. She is affiliated with the Center for Risk and Security at Clark University and the Center for Policy Research at the University at Albany. Dr. Weaver's unique training in mathematics, physics, and psychology informs her modeling of multi-stakeholder situations characterised by differing perspectives. Her theoretical projects address how policy makers change thresholds for decisions over time and the effects of persuasion and individual differences on judgement. Her applied research is a study of insider threats in cyber-security, working with technical experts from Carnegie Mellon University.

1 Introduction

Since 11 September 2001, the vulnerabilities to terrorism of our urban areas, and how best to address them, have been subject to considerable discussion, debate and reflexive defensive measures. Physical access control measures have ranged from 'temporary' concrete barriers, to planters and street furniture, to permanent bollards. These are usually supplemented by armed guards aided by closed-circuit TV cameras and other surveillance measures. Buildings have received retrofit window treatments and structural enhancement of columns and slabs to mitigate blast effects, and some buildings are being considered for systems of sensors and filters to guard against chemical and biological agents. Although these direct physical responses to the frightening events of 11th September are certainly understandable, they are driven more by the desire to protect people and assets from what could happen rather than what is likely to happen. In other words, these measures address the *vulnerability* of people, buildings and other public spaces to certain types of attack and not a true assessment of the *risk* of that attack. This approach essentially removes the public from the decision process and eliminates from consideration any willingness on the part of the public to accept some portion of that risk. It is not surprising, then, that the debate over appropriate security for public places has often degenerated to a simple binary set – secure but unaesthetic on one side and attractive but vulnerable on the other. In this article, three tools will be introduced for the development of an indicator for level of safety, a threshold of that indicator for policy, and a regulatory environment that responds to changes in that threshold. This toolset will be discussed in the context of the appropriate protection of public spaces against terrorist attacks and the appropriate time to order a mass evacuation in the event of a natural disaster.

2 Providing physical protection

Until modern explosives and aerial bombardment rendered them moot, most physical protection strategies for cities and towns were aimed at keeping an attacker at bay by means of moats, walls and other physical obstacles. Even today, *stand-off* (the distance between a bomb and its intended target) is still considered the most effective defense against a terrorist vehicle bomb because blast energy falls off with the cube of the distance and dissipates very quickly. It is for this reason that the initial reaction of those faced with ‘doing something’ to enhance security is often to move in concrete barriers and checkpoints in the hope of maintaining an adequate stand-off distance. However, when effective stand-off distance is not possible or cannot be enforced by these measures, other steps can be taken to protect targeted buildings from bomb damage.

The numerous terrorist bombing attacks experienced in the past 25 years have generated considerable research into the effects of bomb blasts on buildings and people (NRC, 2000). As a result, the vulnerabilities of buildings to deliberately placed bombs are reasonably well understood, as are the relative effectiveness of various countermeasures (Little, 2002). Blast-resistance in buildings is generally provided by passive features, such as additional reinforcement and connections in the structural frame for increased ductility, composite fiber wraps to prevent the shattering of columns and slabs, and high-performance glazing materials that resist blast pressures (AMPTIAC, 2003). When such structurally enhanced buildings are attacked, these measures have been shown to be effective in reducing damage and casualties (Mlakar *et al.*, 2003).

3 Security and risk

Although governments and other stewards of public welfare have a clear responsibility to provide for the safety of those entrusted to their care, the government must also consider the cost of providing that level of safety to buildings and infrastructure. Ultimately, a choice must be made whether an investment to reduce risk to those directly affected is of greater benefit to society than expending the funds for some other purpose (NRC, 1985). Given the high cost of providing security to all public buildings and spaces, this is a question that reasonably should be considered by the public at large. In any event, as has already been determined in the provision of seismic resistance, questions of this type are not for engineers to answer alone (NIST, 1994).

Risk assessment has classically been defined by three questions (Kaplan and Garrick, 1981):

1. What can go wrong?
2. What is the likelihood that it would go wrong?
3. What are the consequences of failure?

Although these questions are relatively straightforward, in practice they often prove difficult to define precisely. Therefore, risk management seeks answers to a second set of questions (Haimes, 2002):

1. What can be done and what options are available? (What is the mix of site selection and configuration, building features and management practices that will provide the desired level of protection?)
2. What are the associated trade offs in terms of all costs, benefits and risks? (For example, reduced risk and improved confidence in security normally would be traded off with increased cost.)

3 What are the impacts of current management decisions on future options? (Policy options that seem cost-effective at present must be evaluated under plausible future changing conditions. For example, providing certain physical protective features may preclude building modifications to increase functionality in the future.)

These questions are particularly relevant to the current discussion because experience has shown that all too often 'temporary' security measures become *de facto* permanent solutions (NRC, 2003). In effect, these measures represent a precautionary approach to the possibility of future attacks without further discussion or assessment of risk, costs or benefits. Such physical protective features typically target generic vulnerabilities and are not generally selected based on a quantified (even if somewhat subjective) risk calculation.

Given the high cost of implementing an effective physical security strategy for public buildings and spaces, the participation and knowledge of all affected parties, including citizens, policy makers, law-enforcement officials, building owners and occupants, planners, architects, engineers and security specialists, should be elicited. Much of the current debate on security in an open society is unproductive because it fails to recognise the distinct difference between the technical elements in the risk calculation (*e.g.*, terrorist threat levels, tactics, bomb sizes and delivery methods, building construction) on the one hand, and community value judgments (*e.g.*, architectural aesthetics, freedom of movement, *etc.*) that must be incorporated on the other. The social and policy sciences provide some interesting and useful tools to frame this complex and often emotionally charged discussion, and three of these tools will be discussed.

4 Three useful tools for goal setting and decision making

Three useful tools for goal setting and decision making within the context of protecting public spaces are judgement analysis, the Taylor Russell diagram and the system dynamics model. Each of these tools was developed for a different purpose and adapted for policy formulation over time. The first, judgement analysis, can be used to design a safety indicator for a given public space that is based on a consensus judgment of a group of technical experts. The method of judgement analysis and the theory surrounding it, Social Judgment Theory (SJT), is about 50 years old (Cooksey, 1996; Stewart, 1988; Hammond, 1955; Brunswik, 1955). SJT was recently reviewed by Meacham (2004) as a potentially useful approach to fire risk problems. The second, the Taylor-Russell diagram, can be used to engage public values in a decision about the appropriate threshold for that safety indicator, that is, how much safety is 'enough'. This tool dates back to signal detection theory (Swets *et al.*, 1991; Green and Swets, 1966) and has been effectively applied to a variety of policy formulation questions (Swets, 1992; Hammond, 1996). The final tool, the system dynamics model, can be used to investigate a management or regulatory structure to allow for changes to the indicator threshold over time and across contexts. The system dynamics model is a computer simulation tool developed by an electrical engineer who addressed business strategy problems (Forrester, 1961). Like the Taylor-Russell diagram, this tool has been applied outside its original domain to great effect; these applications have been reviewed in depth by Sterman (2000). Each tool is presented in turn, together with the scenario for its application, the method, and the outcome or deliverable that would result.

4.1 Judgement analysis

In the development of a physical protection strategy for public spaces and buildings, it is important to have an index or indicator of safety. One way to construct such a safety indicator would be to gather historical data about how spaces with particular attributes performed in the past under a range of actual bombing attacks. Next, a regression analysis of building performance could be conducted with respect to the various features of the buildings and attacks (*e.g.*, bomb size, stand-off distance, building type). Once the regression parameters were established, other buildings could be scored on how secure they would be against a given type of attack. This approach requires abundant and accurate data that is often difficult or impossible to obtain. But even without such data, it is still possible to develop an indicator by analysing the judgement of experts.

The development of an indicator using judgement analysis involves the following steps (Cooksey, 1996). First, experts could be interviewed about the attributes of the problem that constitute the cues to the judgement of safety, including active security measures, stand-off and building features. While experts might differ on the importance of each of these to overall safety, they may be able to agree on a list.

Second, real buildings could be sampled to assemble a representative set of cases with attributes that occur together in plausible ways. (For more on representative design, consult Hammond and Stewart, 2001; Cooksey, 1996; Stewart, 1988; Brunswik, 1955.) Third, once a set of representative cases has been gathered, each expert could rate each case for safety. Finally, regression analysis could be used to model each expert's 'policy' for making judgements of safety. The judgement policy would represent a weighted combination of important safety factors, with the importance weights determined from expert judgement as influenced by past experience, research results, and so on.

Note that the reason for using regression analysis is that people (including experts) are not always accurate about the judgement policies they hold or that they assume others hold. The use of statistics greatly increases the likelihood that the exercise will produce actual expert judgements rather than the experts' guesses about how they make judgements. The policy of a given expert could then be applied to any building to determine a safety score.

If all the experts' judgement policies matched, the result would be a consensus safety indicator. However, consider a scenario in which there was no clear expert consensus. Perhaps one expert thinks that active physical security is the most important factor in safety and another thinks that passive features such as stand-off and physical hardening are paramount. Where experts show differing judgement policies, one could cluster expert approaches. In this case, indices might be developed to represent the different clusters.

A judgement analysis could provide insight into how diverse experts rate buildings on safety, as well as models of the judgement policies of clusters of experts. With time, one could conduct research to assess the relative effectiveness of the differing judgement policies. Meanwhile, a policy maker, or advisory group charged with developing policy recommendations, could craft a compromise among the judgement policies to create an acceptable indicator of safety. The result would be an experience-based safety index for security features.

4.2 The Taylor-Russell diagram

If it is assumed that there is an indicator with a known success rate at predicting building safety, whether developed by analysing past data or by analysing expert judgment, the next step is to select a safety threshold or cut off point, such that buildings above the threshold would be considered 'safe', and those below it would be considered 'unsafe'. Unless the indicator is perfect, any threshold for a safe/not-safe decision will result in some buildings rated as safe when they are not (false positives) or some buildings being rated as unsafe when they are safe (false negatives). A Taylor-Russell diagram can be used to clarify the components of this situation (Hammond, 1996; Green and Swets, 1966).

A Taylor-Russell diagram is presented in Figure 1. Along the horizontal axis are building safety indicator scores. Along the vertical axis are building security performance scores. Each point represents a particular building. The quality of the indicator is shown by the spread of the points around a line angled at 45 degrees.

Figure 1 The Taylor-Russell diagram

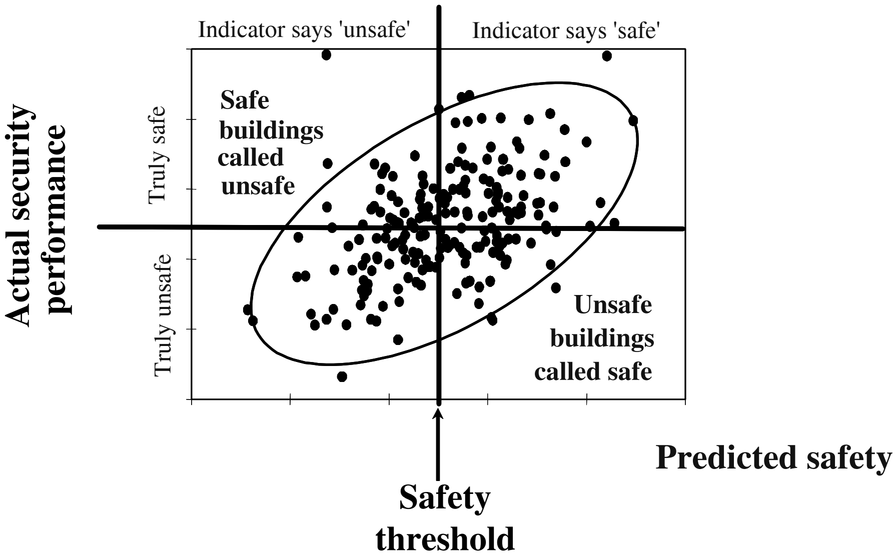

A lower or 'lenient' threshold for an acceptable level of building security may reduce costs, but there is a risk of constructing buildings that will turn out to be unsafe if attacked (false positives). On the other hand, as the threshold for the acceptable level of security is raised or made more 'strict', unnecessary costs on the builder may be imposed as unnecessary security measures are implemented (false negatives). Hammond (1996) describes this trade off between false positives and false negatives for any choice of threshold as the duality of error, which is made instantly visible in a Taylor-Russell diagram.

A Taylor-Russell diagram is constructed by plotting individual buildings as points, with their predictive indicator score along the horizontal axis and their actual performance score along the vertical axis. Next, a threshold for true safety is set.

Finally, a value-based threshold on the predictive indicator is set that results in consequences (false negatives and false positives) that are tolerable to all stakeholders in the process. By splitting the diagram into quadrants, one discovers the number of true positives (*i.e.*, 'safe' according to the indicator and safe in reality), true negatives (*i.e.*, 'unsafe' according to the indicator and unsafe in reality), false positives (*i.e.*, positive according to the indicator for safety, but unsafe in reality), and false negatives (*i.e.*, negative according to the indicator for safety, but quite safe in reality) that result from the threshold selection, given the uncertainty or quality of the indicator.

Note that the number of false positives and false negatives depends not only on the threshold chosen, but also on the degree of association between the indicator and the true safety rating. The correlation between these values represents the quality of the indicator. The indicator will show as much uncertainty as is currently present in the predictive science regarding security measures and actual safety.

Working with a Taylor-Russell diagram, individuals can discuss the number of false positives and/or false negatives they willing to accept and, given the quality of the indicator, they can select an appropriate threshold. The Taylor-Russell diagram provides a means of envisioning simultaneously the connection between:

- the choice of threshold
- the effectiveness of an indicator
- the resulting consequences.

4.3 The system dynamics model

Imagine that an indicator has been chosen, data obtained on its performance, and a threshold selected. Ideally, the community at large would accept this as a stable threshold. Unfortunately, it cannot be guaranteed that the selected trade off of false positives and false negatives will be shared by others or that the particular threshold selected will always be appropriate. In fact, if a highly salient event occurs for which there is a false positive, say, a building turns out unexpectedly to be unsafe (*i.e.*, it performs less well than expected) those constituents concerned with safety will pressure policy makers to move the threshold higher. This debate is currently raging about the performance of the twin towers of the World Trade Center in New York, with implications for both structural design and fire protection. As additional safety measures are implemented for buildings that were already safe enough, there will be diminishing returns on investment. The community concerned with the cost of buildings might then pressure the policy makers to lower the index. Swets (1992) has described the potential for a threshold to oscillate in light of recent salient events. Hammond (1996) suggested that the oscillation would occur as stakeholders respond to recent events and pressure policy makers to change the threshold. Weaver and Richardson (2002) designed a system dynamics simulation to analyse the systemic requirements for such an oscillation to occur.

As the selection of a threshold represents a value-based decision, it may become outmoded as societal values respond to recent events and as indicators improve. In order to have a responsive policy context that is protected from too rapid and vigorous an overreaction to recent events, it may be necessary to build in legal structures that manage or regulate the threshold in an appropriately responsive manner. These phenomena are particularly amenable to computer simulation that can model a much broader range of scenarios than would be found in practice.

A model that includes the legal and political regulatory structures that affect the policy threshold can be created that will permit the testing of different regulatory environments, including penalties for non-compliance, avenues for complaint, community values about outcomes, and the predictive quality of the safety indicator. The simulation might include stakeholder pressures to change the threshold, the quality of the index, the resulting false positives and false negatives and how the entire stakeholder community would react to an unacceptable number of either of these. It could include not only the information for which clear data are available, but also could embed the rich intuitions of experts, so that a broad range of possible scenarios could be tested. System dynamics has been used to simulate similar conditions in other contexts (Sterman, 2000).

The outcome of such a modelling effort would be a simulation that would allow policy makers to test the consequences of various threshold choices and a safety index that improves its predictive quality over time or with resource investment. In addition, it would allow them to set up and test a management and regulatory environment that would build in constraints against too sensitive a response to recent events, while guaranteeing the flexibility to update the model.

5 Application to extreme natural events

We will now turn to recent natural disasters such as the Indian Ocean tsunami of 26 December 2004 and hurricane Katrina in August 2005 to see if these tools could add value to the decision processes for extreme natural events. In the case of the tsunami, we will generalise the decision problem to whether or not to issue an evacuation order to coastal areas if undersea seismic activity has been detected. Although there was no detection system in place in the Indian Ocean at the time of the tsunami, it is expected that one will be deployed shortly (IOC, 2005). Ocean-based sensors called tsunameters can determine the occurrence of a tsunami with reasonable accuracy by measuring changes in ocean levels (as small as 1 cm in 6000 m of water) and the rate of propagation and direction of these changes. Once a warning has been received on land, alarms of various types can be sounded and people instructed to reach pre-determined evacuation routes or vertically evacuate within identified places of refuge. Vertical evacuation (Ruch *et al.*, 1991) assumes that buildings are constructed in such a way that they withstand tsunami waves, while affording shelter from the highest waves. Many tall (more than three storeys) resort buildings constructed to modern standards withstood the tsunami waves very well (Dalrymple and Kriebel, 2005). Had people on the beach known of the impending tsunami, many might have had time to reach the buildings from the beach and ascend to a place of safety.

At the same time, an appropriately calibrated decision system could avoid the disruption and cost of unnecessary evacuations, *i.e.*, 'false positives'. For example, on 17 November 2003, the Deep-Ocean Assessment and Reporting of Tsunamis (DART) system in the Pacific Ocean detected a small tsunami generated by an earthquake near Adak, Alaska, but based on data collected from DART buoys, no warning was issued for this event, which saved Hawaii an estimated $68 million. This demonstrated that sometimes the greatest benefit of a warning system is *knowing when* *not to evacuate*. This success story can be contrasted with an event of similar magnitude in 1986 in the same region that resulted in the evacuation of Hawaii's coastal areas. At the time, predictions of the amplitude of tsunamis were difficult to make, and the tsunami that eventually struck the coastline was less than a foot in height. It caused no damage, but the Hawaii Department of Business, Economic Development and Tourism estimated that the evacuation cost the state $40 million in lost productivity and business (Lautenbacher, 2005).

This scenario is quite amenable to analysis using the methods described in this paper. Figure 2 illustrates how a decision table based on a Taylor-Russell diagram could be structured to model the 'evacuate/don't evacuate' decision. The threshold might initially be set based on empirical evidence and expert opinion. However, owing to the often short period of time between the initiation of a tsunami and its landfall, the systems dynamics model would suggest a fairly conservative threshold for signalling the need to evacuate. If this results in an unacceptable number of false positives or Type II errors, with the attendant costs and disruptions, a more liberal threshold will probably evolve. Whether or not this threshold will cycle as suggested by the systems dynamics model will depend to a large degree on the number and severity of future events.

Figure 2 A decision table based on the Taylor-Russell diagram applied to an evacuation decision

		The proper course of action	
		Don't evacuate	Evacuate
Decision	Don't evacuate	Correct decision No disruption or economic loss	**Type I error** Tsunami occurs: large loss of life
	Evacuate	**Type II error** Unnecessary disruption and wasted resources	Correct decision Lives saved

Null hypothesis: evacuation should occur

The modelling of the evacuation decision process for major storms such as hurricane Katrina is similar. Public officials have traditionally been reluctant to order mass evacuations in the face of approaching storms, even large hurricanes, because these storms often fail to make landfall where predicted and the impacts of many storms have been less than expected. This uncertainty must be weighed against the known cost and disruption of a mass evacuation. However, in the aftermath of the enormous damage caused by hurricane Katrina in August 2005, mass evacuations were ordered in September 2005 for hurricanes Ophelia and Rita. In both cases, post-event analysis determined that the evacuations were more extensive than required. The system dynamics model suggests that the evacuation threshold will shift if major storms do not continue on a regular basis, which is borne out by experience that shows that communities become complacent and more at risk with the passage of time. What this suggests is the need to continue to improve predictive capabilities for hurricane landfall so as to minimise Type II errors, and also to improve the capability of structures to resist the forces of extreme winds so as to lessen the need for evacuation.

6 Summary and conclusions

In the case of terrorism, we find ourselves in a time where former contexts of threat, vulnerability and target have all changed and continue to do so. Threats are unpredictable and the full range of threats probably unknowable. Under these circumstances, we will never be able to anticipate all possible threats, and even if we could, there is not enough money to deploy technologies to address them. Security in this situation needs to be flexible and agile and capable of addressing new threats as they emerge while still meeting the demands of the public for attractive architecture and free access to public spaces. This cannot be a one-time investment but rather an effort that will need to be revisited periodically as threats, resources and community values continue to change. At the same time, we will continue to be at risk from natural extreme events such as hurricanes, earthquakes and tsunamis. Addressing these risks will require the application of sound risk management principles and structured decision processes.

Three tools from the social and policy sciences have been introduced for use at different stages of the process of developing a rational and holistic approach to security from terrorist attack, as well as from extreme natural events. Judgement analysis could be used to develop an index of safety; the Taylor-Russell diagram to select an appropriate policy threshold for that index; and the system dynamics model to simulate a policy environment that would respond to unexpected events with appropriate adjustments. Figure 3 illustrates how these three tools could be used together to include all stakeholders, experts and laypersons alike, in the development of consensus levels of security for public buildings and spaces that achieve aesthetics and openness with reasonable security and safety. At the same time, they could be used to model the decision process to be employed when considering evacuation in the face of an impending storm or tsunami event.

Figure 3 A model for applying three useful tools for goal setting and decision making

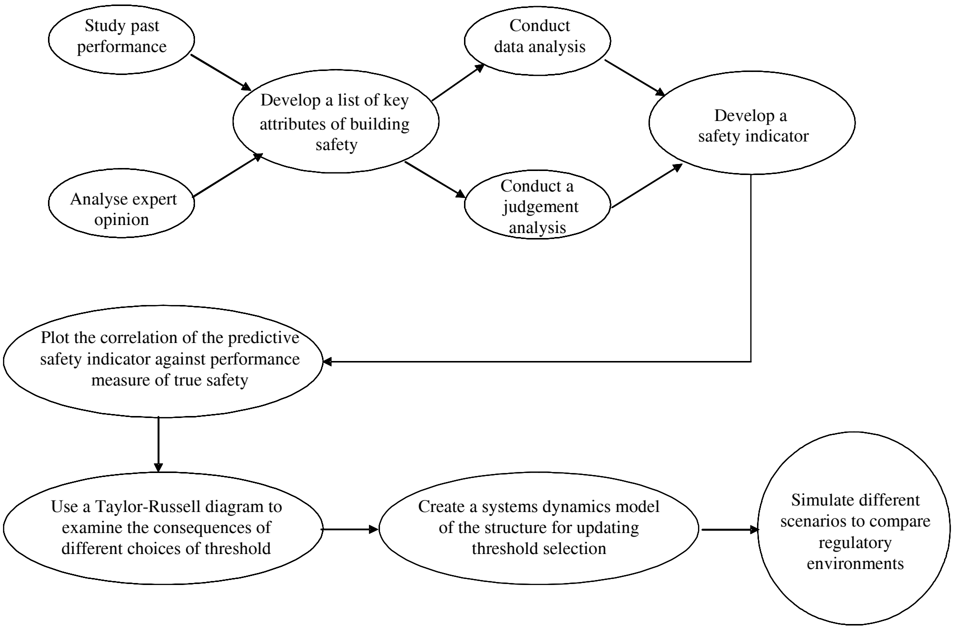

References

Advanced Materials and Processes Technology Information Analysis Center (AMPTIAC) (2003) 'Protecting people at risk: how DoD research reduces the impact of terrorism', *The AMPTIAC Quarterly*, Vol. 6, No. 4.

Brunswik, E. (1955) 'Representative design and probabilistic theory in a functional psychology', *Psychological Review*, Vol. 62, No. 4, pp.193–217.

Cooksey, R. (1996) *Judgment Analysis: Theory, Methods, and Applications*, New York: Academic Press.

Dalrymple, R.A. and Kriebel, D.L. (2005) 'Lessons in engineering from the tsunami in Thailand', *The Bridge*, Vol. 35, pp.4–16.

Forrester, J.W. (1961) *Industrial Dynamics*, Cambridge, MA: Productivity Press.

Green, D.M. and Swets, J.A. (1966) *Signal Detection Theory and Psychophysics*, New York: Wiley.

Haimes, Y.Y. (2002) 'Risk of terrorism to cyber-physical and organizational-societal infrastructures', *Public Works Management and Policy*, Vol. 6, No. 4, pp.231–240.

Hammond, K.R. (1955) 'Probabilistic functioning and the clinical method', *Psychological Review*, Vol. 62, No. 4, pp.255–262.

Hammond, K.R. (1996) *Human Judgment and Social Policy: Irreducible Uncertainty, Inevitable Error, Unavoidable Injustice*, New York: Oxford University Press.

Hammond, K.R. and Stewart, T.R. (Eds.) (2001) *The Essential Brunswik: Beginnings, Explications and Applications*, New York, NY: Oxford University Press.

International Oceanographic Commission (IOC) (2005) *Toward the Establishment of a Tsunami Warning and Mitigation System for the Indian Ocean*, http://ioc.unesco.org/indotsunami/mauritius05/mauritius05_outcome.htm (accessed 3 August 2005).

Kaplan, S. and Garrick, B.J. (1981) 'On the quantitative assessment of risk', *Risk Analysis*, Vol. 1, No. 1, pp.11–27.

Lautenbacher, Jr., C.C. (2005) 'Tsunami warning systems', *The Bridge*, Vol. 35, pp.21–25.

Little, R.G. (2002) 'A probabilistic approach for protecting people and buildings from terrorist attack and other hazards', *Proceedings of the International Conference on Protecting Structures Against Hazards*, Singapore: C.I. Premiere, pp.49–56.

Meacham, B. (2004) 'Decision making for fire risk problems: a review of challenges and tools', *Journal of Fire Protection Engineering*, Vol. 14, No. 2, pp.149–168.

Mlakar, P.F., Dusenberry, D., Harris, J.R., Haynes, G., Phan, L. and Sozen, M. (2003) *The Pentagon Building Performance Report*, Reston, VA: American Society of Civil Engineers, pp.58–59.

National Institute of Standards and Technology (NIST) (1994) *1994 Northridge Earthquake: Performance of Structures, Lifelines, and Fire Protection Systems*, Gaithersburg, MD: National Institute of Standards and Technology.

National Research Council (NRC) (1985) *Safety of Dams, Flood and Earthquake Criteria*, Washington, DC: National Academy Press.

National Research Council (NRC) (2000) *Blast Mitigation for Structures: 1999 Status Report on the DTRA/TSWG Program*, Washington, DC: National Academy Press.

National Research Council (NRC) (2003) *Working in Olmsted's Shadow: Guidance for Developing a Scope of Services for the Update of the Master Plan for the U.S. Capitol and Grounds*, Washington, DC: National Academy Press.

Ruch, C., Miller, H.C., Haflich, M., Farber, N.M., Berke, P.R. and Stubbs, N. (1991) *The Feasibility of Vertical Evacuation*, Monograph No. 52, Institute of Behavioral Science, University of Colorado.

Sterman, J.D. (2000) *Business Dynamics: Systems Thinking and Modeling for a Complex World*, Boston: Irwin McGraw-Hill.

Stewart, T.R. (1988) 'Judgment analysis: procedures', in B. Brehmer and C.R.B. Joyce (Eds.) *Human Judgment: The SJT View*, Amsterdam: North-Holland.

Swets, J.A. (1992) 'The science of choosing the right decision threshold in high-stakes diagnostics', *American Psychologist*, Vol. 47, No. 4, pp.522–532.

Swets, J.A., Getty, D.J., Pickett, R.M., D'Orsi, C.J., Seltzer, S.E. and McNeil, B.J. (1991) 'Enhancing and evaluating diagnostic accuracy', *Medical Decision Making*, Vol. 11, No. 1, pp.9–18.

Weaver, E.A. and Richardson, G.P. (2002) 'Threshold setting and the cycling of a decision threshold', *Paper Presented at System Dynamics '02*, Palermo, Italy, http://www.wpi.edu/News/Transforamtions/2003Spring/pendulum.pdf.