FEDERAL DECISION MAKING FOR HOMELAND SECURITY

Mapping the Normative/Descriptive Divide

L. VALVERDE, JR.

AI Systems, Inc.
Kirkland, WA, USA
drljva@hotmail.com

S. FARROW

University of Maryland
Baltimore County, MD, USA
farrow@umbc.edu

Abstract: The events of 9/11 have dramatically shifted public and private sector priorities aimed at addressing the threat of transnational terrorism. An important issue facing public decision makers is how best to allocate scarce resources in the face of significant uncertainty concerning potential threats and hazards, together with uncertainty concerning the potential costs and benefits associated with possible prevention and mitigation strategies. Viewing this problem from the vantage point of modern economic theory, normative theories of choice provide guidance on how agents should make decisions if they wish to act in accordance with certain logical principles. Often, however, there is a discord between normative theory and how people behave in real-world decision contexts. In this paper we explore several aspects of current homeland security resource allocation practices within the federal government. We begin with an examination of two normative investment models, and we explore the linkages that exist between actual practice and the insights that economic theory lends to these problems. We then present the rudiments of a prescriptive approach to homeland security decision making and risk management that seeks to guide decision makers toward consistent, rational choices, while recognizing their real-world limitations and constraints.

1. Introduction

The events of 9/11 have brought about dramatic shifts in government and private sector investments to address the threat of transnational terrorism.

I. Linkov et al. (eds.), Real-Time and Deliberative Decision Making.

An important issue facing federal agencies and public decision makers charged with managing the security of the homeland is how to best allocate scarce resources in the face of large uncertainties concerning the evolving nature of the threat, together with uncertainty concerning the potential costs and benefits associated with possible prevention and mitigation strategies. In particular, federal agencies within the homeland security domain face a number of challenges in deciding how best to allocate scarce resources in the pursuit of a broad range of strategic goals and objectives—program effectiveness and economic efficiency, to name just two. In this decision context, the allocation of resources is made difficult by:

1. The existence of multiple decision makers and stakeholders
2. The presence of multiple and often conflicting objectives
3. The prevalence of significant uncertainty surrounding key facets of the terrorism problem

In a complex, dynamic, and uncertain context like this, decision makers can avail themselves of guidance and decision aids from a variety of sources, ranging from informal, qualitative methods to the most formal, quantitative methods. In this regard, it is natural to distinguish between two types of theories: *normative* theories of choice on the one hand, which seek to provide guidance on how agents should make decisions based on logical principles; and alternatively, *descriptive* theories of choice, which seek to provide empirical explanations for actual decision-making behavior in these environments.

In this paper we explore these two decision-making perspectives, with a view towards ultimately informing a *prescriptive* view of how homeland security decision making might best be improved, given all of the attendant constraints and uncertainties. As several decades of empirical psychological research have shown, there is often a discord between normative theories of choice and observed behavior in real-world decision contexts characterized by risk and uncertainty. Our pursuit of this line of inquiry is motivated, in the first instance, by our witnessing a plurality of viewpoints and methodologies currently being applied in the homeland security domain. There are, we feel, a number of lessons to be gleaned from the current state of affairs. How issues are framed in these complex environments, how rational or cognitive decision rules are utilized, how key uncertainties are characterized and evaluated, how values are aggregated—all of these factors influence both the decision-making process itself and, ultimately, the likely ensuing outcome.

Our discussion is organized along the following lines. First, we present an illustrative pair of canonical normative investment models under uncertainty that attempt to capture and represent several salient features of the homeland security problem. In this discussion, our point of departure is a normative model for allocating security expenditures across multiple sites,

given a specified security budget. A generalization of this model then allows us to capture two central and related problems in terrorism risk management; namely, how to allocate resources across *probability*- and *damage-reducing* activities. With this as background, in Section 3 we discuss current general practices within the federal government for allocating homeland security resources. This, in turn, motivates a discussion in Section 4 on the rudiments of a prescriptive framework for approaching these problems. Ultimately, the framework seeks to guide decision makers toward consistent, rational choices, while recognizing multiple limitations and constraints (e.g., cognitive, organizational, and other). We conclude with some closing remarks and a brief discussion of possible future research directions.

2. Normative Investment Models Under Uncertainty

Normative investment models under uncertainty span a wide conceptual range—from individual, utility-maximization models to market-based welfare models of rational choice. In this section, we take the rational actor model as a point of departure for highlighting several normative bases for choice in the homeland security domain.

We begin by looking, first, at a utility maximizing model for an individual decision maker who—in the context of our discussion here—considers numerous possible outcome dimensions as being important (e.g., national welfare, agency mission, government costs, political support). Given the decision maker's preferences across these dimensions, the decision maker chooses the option with the greatest expected utility. To illustrate key issues we use the simplest expected-value approach, expected-value maximization, which assumes that the decision maker values increases and decreases in risk equivalently.

The rational expected utility model provides a useful starting point for the issues under consideration here. For the purposes of our discussion, we ignore debates as to whether decision makers actually make decisions according to the classical model [8]; the position we take here is that models grounded on the maximization principle may be useful as benchmarks for evaluating the *quality* of actual decisions made in these environments. The models we consider here are intended to integrate decision, probability, and outcome information in ways that seek to inform decisions on government expenditures directed at managing homeland security.

As we discuss below, different analytical models, in effect, pose different questions. The simplest normative model involves expenditures to reduce the probability of attack at independent sites. A key result of the basic model we present is that some sites are left unprotected if, after the updating of probabilities for investment, the marginal social costs of an attack

on the site are less than a threshold that is exogenously constrained by the available funds.

Presented below are two short variations based on independent sites, and consideration of both prevention and mitigation investments.

2.1. ALLOCATING DEFENSIVE EXPENDITURES ACROSS MULTIPLE, INDEPENDENT SITES

We begin with an expected cost minimization model for optimally allocating defensive expenditures across multiple, independent sites.[1] The model presented here is easily extended to allow for the treatment of complexities such as dependency between sites and other variations (see, e.g., [4]).

Whether viewed from a national perspective, or from the vantage point of a decision maker charged with infrastructure protection, we assume a unitary decision maker with two or more independent sites for which defensive resources must be allocated. The decision maker ultimately wishes to select those defensive options that minimize the expected costs associated with a terrorist attack. We begin by defining

e_i ≡ Level of defensive expenditure on site i, for $i = 1, \ldots, n$;

Z ≡ Aggregate expenditure level over all sites and vulnerability pathways;

$\Pr(e_i)$ ≡ Probability of a successful terrorist attack, with $\Pr'(e_i) < 0$ and $\Pr''(e_i) > 0$;

$S(e_i)$ ≡ Non-governmental costs of the investment expenditures, with $S'(e_i) > 0$;

$C(e_i)$ ≡ Social cost, given that an attack occurs, with $C'(e_i) < 0$ and $C''(e_i) > 0$.

The government's decision problem is to choose an optimal level of expenditure, $e_i^* \geq 0$, for each site i, minimizing expected costs

$$\min \sum_{i=1}^{n} \{\Pr(e_i)[e_i + C(e_i) + S(e_i)] + [1 - \Pr(e_i)][e_i + S(e_i)]\},$$

subject to the constraints

$$\sum_{i=1}^{n} e_i = Z \text{ and } e_i \geq 0.$$

[1] Interdependencies—both positive and negative—are a central concern in evaluating homeland security investments. Positive interdependencies among sites have a possible public good component, in that expenditures at one site may have beneficial effects at other sites. Border security is an obvious example: if potential attackers are stopped at the border, the probability of an event at a number of sites is reduced. Alternatively, should an attack occur, improvements in response capabilities may mitigate or reduce damages at multiple sites.

Looking first at those sites where positive expenditures occur, we formulate the Lagrangian expression for this problem, yielding the following necessary conditions for optimization with exhaustion of the budget:

$$\Pr'(e_i)C(e_i) + C'(e_i)\Pr(e_i) + S'(e_i) = \lambda - 1, \tag{1}$$

The left-hand side of this equation is simply the *marginal expected social cost avoided* (MESCA) through each additional unit of expenditure, while being net of the non-governmental cost associated with each expenditure, $S(e_i)$.

All sites $i \neq j$ with positive expenditures are equated to the common shadow price of funds $(\lambda - 1)$:

$$\Pr'(e_i)C(e_i) + C'(e_i)\Pr(e_i) + S'(e_i) = \Pr'(e_j)C(e_j) + C'(e_j)\Pr(e_j) + S'(e_j),$$

such that the MESCA is equal across all sites. In this formulation it is important to note that some defensive expenditures, e_i, can be zero. Sites without expenditures are those where the MESCA is less in absolute value than the cutoff level of the shadow price of funds. Prescriptively, the model stipulates that some sites are sufficiently "small"—taking both the probability of success and the potential ensuing damages into account—that it is optimal to do nothing to protect them. Of course, all sites are characterized by some level of risk exposure, regardless of whether defensive expenditures occur. The asymmetric nature of the attacker and the intended victim(s) precludes the possibility of reducing the risk to zero.[2]

Prescriptively, then, in allocating defensive funds across independent sites, *for sites that exceed a threshold of potential impact, equate the marginal expected social cost avoided for all sites and vulnerabilities*. In this way, for any given site, there is a cutoff marginal social cost avoided where it is optimal not to expend anything on that site.

2.2. ALLOCATION OF EXPENDITURES ACROSS DAMAGE AND PROBABILITY REDUCING ACTIVITIES

A crucially important policy question in the homeland security domain is the optimal balance between actions and processes that *prevent* attacks and those that *mitigate* (partially or fully) the potential adverse consequences associated with these attacks. In practical settings, the problem may be one of deciding how best to allocate budgets between intelligence-related activities (that are, by their very nature, directed towards preventing attacks) and

[2] A lucid argument for this line of reasoning is provided by Posner [11].

the hardening of vulnerable physical infrastructure (aimed at minimizing the adverse effects associated with an attack). For this particular model, let

e_i ≡ Probability-reducing expenditures at site i;

h_i ≡ Damage-reducing expenditures at site i;

Z ≡ Aggregate level of probability- and damage-reducing expenditures;

$\Pr(e_i)$ ≡ Probability of a successful attack given defensive expenditure e_i;

$C(h_i)$ ≡ Social cost given that an attack occurs.

Consistent with our earlier discussion, we assume that $\Pr'(e_i) < 0$ and $\Pr''(e_i) > 0$, and that $C'(h_i) < 0$ and $C''(h_i) > 0$.

As before, we assume a unitary decision maker who is charged with maintaining a finite number of sites, labeled $i = 1, 2, \ldots, n$. The decision maker wishes to choose an optimal level of expenditures

$$\mathbf{e}^* = (e_1^*, e_2^*, \ldots, e_n^*) \geq 0 \quad \text{and} \quad \mathbf{h}^* = (h_1^*, h_2^*, \ldots, h_n^*) \geq 0$$

that minimize the total expected cost

$$\min \sum_{i=1}^{n} \{\Pr(e_i)[e_i + h_i + C(h_i)] + [1 - \Pr(e_i)][e_i + h_i]\},$$

subject to the constraints

$$\sum_{i=1}^{n} (e_i + h_i) = Z \text{ and } e_i, h_i \geq 0.$$

As before, Lagrangian methods are used to solve this constrained optimization problem, yielding the following necessary conditions for optimality:

$$\Pr'(e_i)C(h_i) = \lambda - 1 \ \forall i; \tag{2}$$

$$\Pr(e_i)C'(h_i) = \lambda - 1 \ \forall i. \tag{3}$$

Equations (2) and (3) imply the equality of the marginal expected social cost at each site, with positive expenditures for each individual type of expenditure and across both types of expenditure (the latter when Eqs. (2) and (3) are set equal to each other).

It is important to note that this model does not distinguish between expenditures that are earmarked for "homeland security" and those that are directed at other types of risks or hazards. In the homeland security domain, it may, for example, be useful to distinguish between manmade hazards (like acts of terrorism) and natural hazards (like extreme weather events).

The above model can, of course, be generalized to allow for this kind of "all-hazards" conception of how best to allocate prevention and response investments.

3. What to Protect: A Descriptive View

Risk management is, in many ways, an endemic feature of public decision making in the 21st century. As a matter of course, the federal government manages a panoply of risks, ranging from employment, environment, finance, and public health to national security [1]. Managing this last component—the national security interests of the country—is, to be sure, a multifaceted task that is fraught with risk and complexity. The specter of transnational terrorism exists throughout the world, in a number of guises.[3] As a practical necessity, managing this evolving threat requires the ability to trace out the expected consequences—economic and otherwise—associated with potential acts of terrorism.

In this light, risk management in a homeland security context is seen to entail various attempts to:

1. Characterize the nature of the threat environment
2. Characterize the vulnerability of people and systems to these threats
3. Value the potential monetary and non-monetary impacts associated with these threats and vulnerabilities

In a management and planning context, decision makers utilize this information to prioritize capitol investment decisions geared at the *prevention* of undesirable events or at the *mitigation* of adverse consequences. Ultimately, the goal is to arrive at adequate levels of protection against these risks and hazards, within specified constraints.

Of course, in the wake of 9/11, all of these considerations sit in an organizational setting and context that is vastly more complex than the one that preceded it. The U.S. Department of Homeland Security (DHS) consists of 23 separate agencies with more than 183,000 employees. Given both the scale and urgency of this undertaking, the challenges that federal decision makers face are, in the first instance, *organizational*. How an organization of this size and complexity takes its congressionally legislated mandate and drives it programmatically through the entire organization is, of course, a key challenge.[4]

[3] For a discussion of recent trends, see, e.g., Chalk et al. [2].

[4] For one DHS insider's perspective on these organizational challenges, see Ervin [3].

At the heart of DHS's mandate is a fundamental desire to protect people and property against a broad range of potential extreme events—both manmade and natural. How, in this context, strategic intent is construed and executed rests, in large measure, on the ability to create and foster a *risk-based* culture that takes as its point of departure a coherent and rational appraisal of the threat/hazard environment, together with a flexible and adaptive organizational structure that is able to prepare for, and respond to, these threats.

Any incremental steps to this end must, in the first instance, be informed by a strategic roadmap that lays out how risk management principles should inform a broad range of homeland security decisions. Central in this regard is the ability to provide—at every level of the organization—clear and direct guidance on how risk management principles should be applied in these strategic, tactical, and operational settings. At the present time, there is little in the way of systematic guidance for how risk management principles should be applied, though some progress has been made in certain areas in recent years. In light of this situation, it is not surprising that, in the homeland security domain, there are a broad range of risk assessment models currently in use at the federal level. The diversity of models found in these environments reflects, to a large extent, the domains and mission areas from which they stem, with applications including agro-terrorism, aviation security, cargo security, port security, rail security, and critical infrastructure protection.

In the post-9/11 era, much emphasis has been placed on models that proceed from a threat, vulnerability, and criticality (TVC) mindset, for which the U.S. Government Accountability Office (GAO) provides the following characterization [13, 15]:

- *Threat Assessment:* An attempt to identify relevant threats, and to characterize their potential risk
- *Vulnerability Assessment:* The identification of weakness and susceptibility in a system
- *Criticality Assessment:* An attempt to systematically identify and evaluate an organization's assets and operations by the importance of its mission or function (and perhaps other key attributes, such as national security, public health and safety, etc.) and individuals at risk

Looking, first, at the threat assessment component, much effort currently focuses on identifying and evaluating a number of potential threats and hazards.

Specific steps in this process usually include:

1. The identification of *threat categories*, together with potential adversaries
2. The characterization of adversary *motivations*, *intentions*, and *capabilities*

3. The estimation of frequencies or likelihoods for specific threat scenarios

At the conclusion of this type of analysis, decision makers often rank threats along various dimensions; e.g., greatest likelihood or potential impact. A vulnerability assessment then takes this threat information and assesses the manner and degree to which a system's integrity and viability are compromised by specific threats. Finally, criticality assessment entails the prioritization of assets, as determined by how a particular asset compares with other valued assets, given specified threats and vulnerabilities. Often this will take the form of a prioritized list of risks (asset, threat, and vulnerability combinations) that inform resource allocation decisions. In this regard, various countermeasures can be considered in order to reduce specific vulnerabilities linked to risks that are deemed unacceptable.

The constellation of models currently in development and use represent an important first step in the government's efforts to assess and manage terrorism risk. As we discuss below, however, these models place a myopic focus on risk assessment *per se*, to the exclusion of other factors and considerations that are central to a more fully realized conception of risk management.

Current analytical approaches are characterized by several notable features. First, as mentioned above, is the focus on TVC-based approaches [10]. Second is the use of multicriteria analysis (MCA) methods [6]. Increasingly, MCA-type methods are used in homeland security applications, largely because costs and benefits are not always easily monetized. In general, these methods provide decision makers with

- A way of looking at complex problems that are characterized by a mixture of monetary and non-monetary objectives
- A set of analytical techniques for breaking complex problems into manageable pieces, allowing for data and expert judgments to be brought to bear on individual elements of the problem
- Analytically tractable ways to reassemble the pieces, and to present a coherent overall picture to decision makers

The U.S. Coast Guard's Port Security Risk Assessment Tool (PS-RAT) provides a useful case in point. This risk assessment tool is used by the Coast Guard leadership to help prioritize the allocation of scarce resources to key mission areas and activities.[5] On the threat side, the methodology is scenario-driven, with emphasis on the combination of *target* and *means of attack*.

[5] A detailed description and critique of the PS-RAT is provided in [15].

Relative threat frequencies are assigned for each scenario. Potential target vulnerabilities are scored based on perceived susceptibility in four potential dimensions of vulnerability:

1. *Availability*
2. *Accessibility*
3. *Organic security*
4. *Target hardness*

Consequences are similarly valued in a multi-attributed way; specifically, consequences are measured in terms of their impact on five attributes, namely:

1. *Death/injury*
2. *Economic impact*
3. *Impacts on national defense*
4. *Symbolic effect*
5. *Follow-on homeland security threat*

These attributes are combined using a simple additive value function, and a probabilistic event tree is then used to structure the information in a way that gives decision makers a snapshot view of the expected consequences associated with a given threat scenario.

4. A Prescriptive Framework for Homeland Security Decision Making

The centrality of risk management as an organizing principle around which problems of scarce resource allocation are structured and evaluated is an idea that permeates most contemporary efforts within the federal government to assess and manage the potential adverse consequences associated with extreme events—both manmade and natural [16]. To be sure, the panoply of decision-aiding and risk assessment tools currently being developed will continue to evolve and improve as new methodologies and ways of thinking are brought to bear on these complex issues. Still, as our discussion in the previous sections suggests, there is value to be gained in mapping the hinterland that exists between normative theory, on the one hand, and descriptive decision-making reality, on the other, as it relates to managing the security of the homeland. Understanding the conceptual and pragmatic terrain that defines this hinterland helps inform a *prescriptive* view of how homeland security decisions under uncertainty should be construed and evaluated. In what follows, we set out the rudiments of a prescriptive framework for decision making and risk management that encompasses a number of elements

that are important in any reasoned and systematic effort to appraise and manage homeland security risks.

4.1. ELEMENTS OF THE FRAMEWORK

Our approach to risk management begins, in the first instance, with an awareness and understanding of the fact that assessing and evaluating complex risks presents decision makers with a unique set of challenges, especially in situations or contexts where the risks are ill-defined or poorly understood.[6] As we discuss in detail below, any attempt to characterize and evaluate homeland security risks leads, naturally, to a consideration of possible risk mitigation alternatives, whether at the strategic, tactical, or operational level. In the evaluation of strategic alternatives, decision makers will typically integrate and weigh knowledge and information from a variety of sources, including organizational or societal values. In evaluating potential courses of action, decision makers will also look to explore fundamental trade-offs between risk and return, short-term versus long-term gain, and so on. In the management selection process, other issues may be considered, including relevant organizational constraints and risk tolerances. And finally, any selection of risk mitigation options will entail a program for implementation and monitoring.

The prescriptive framework presented here is based on a synthesis of published literature, and is intended as an all-hazards approach, with particular emphasis on homeland security issues. The framework is designed so that the individual components of the approach do not become ends in themselves; rather, the framework entails a full cycle of activities, ranging from strategic planning all the way through to implementation and monitoring. The five elements of the framework are as follows [15]:

- *Strategic goals, objectives, and constraints*
- *Risk assessment*

[6] A large technical and professional literature addresses these issues. The field of risk assessment has a long history, with much attention focused on the analysis of complex systems (e.g., energy, space systems) and the evaluation of environmental problems. Various risk analysis techniques can be used in evaluating risk mitigation strategies. Fault trees, for example, can be used to focus attention and logical analysis on undesirable events. Failure modes and effects analysis is often used to analyze the effects of possible failure modes on system performance. These and other techniques are often used in probabilistic risk analyses, which seek to measure the risks inherent to a particular system's design or operation. For an overview of relevant methods and techniques, see, e.g., Haimes [5], Morgan and Henrion [9], Raiffa [12] and Viscusi [16].

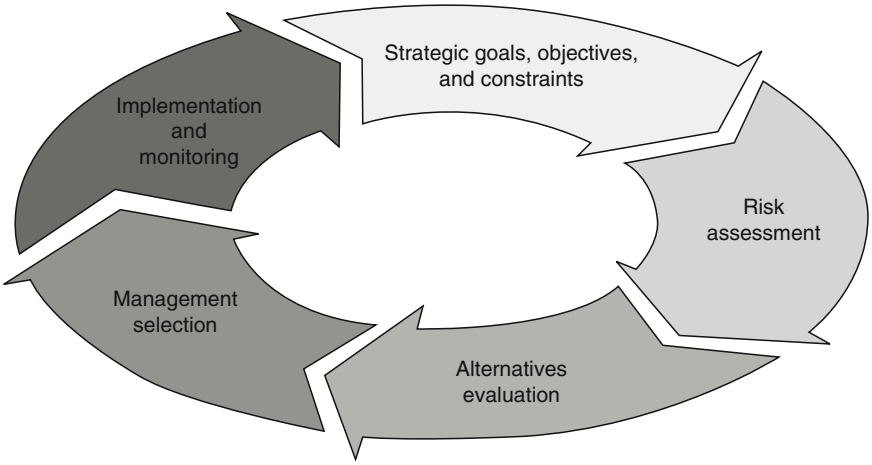

Figure 1. Elements of the Prescriptive Decision making and Risk Management Framework [15].

- *Alternatives evaluation*
- *Management selection*
- *Implementation and monitoring of risk mitigation measures*

Figure 1 illustrates the cyclical nature of the framework. Proceeding through the framework's steps is generally a linear process, though loops may feed back from later to earlier steps in the cycle. Once the process is complete, one or more iterations through various aspects of the framework are possible. The nature of the framework is such that new information can enter any element at any stage in the overall decision making and risk management process.

4.1.1. Strategic Goals, Constraints, and Objectives

The pursuit of *goals* and *objectives* lies at the very foundations of any modern conception of *strategic intent*, and this viewpoint is the conceptual starting point for our prescriptive framework. Modern management practices embed tactical and budgetary decisions in the context of a strategic plan, with clearly articulated goals and objectives that identify resource issues and external threats/hazards.

In our framework, effort is, in the first instance, directed at *structuring* strategic objectives in ways that are meaningful to decision makers, with particular attention paid to the manner in which objectives *relate to*—and

potentially conflict with—one another. Ultimately, this focus on objectives enables decision makers to:

1. Uncover hidden objectives
2. Improve communication and facilitate involvement among stakeholders
3. Enhance the coordination of interconnected strategies and programs

4.1.1.1. Fundamental Objectives, Means Objectives, and Objectives Hierarchies. For our purposes here, it is useful to distinguish between *fundamental objectives* and *means objectives* [7]. As the name implies, fundamental objectives are those objectives that matter most to decision makers. Means objectives, on the other hand, are objectives that provide the instrumental means by which fundamental objectives are achieved.

An examination of national strategies can serve to illustrate these concepts.[7] In particular, we take the National Strategy for Homeland Security (NSHS) as a specific case in point. The overarching objective of the NSHS is, perhaps, best summarized as *maximizing homeland security*. Four fundamental objectives are seen to define this overarching objective:

- *The prevention of terrorist attacks*
- *Reducing vulnerability to attacks*
- *Minimizing damage resulting from attacks*
- *Enhancing recovery*

4.1.1.2. Linking Means and Ends Objectives. Having structured the fundamental objectives hierarchy, the next stage in our process calls for relating means objectives to the fundamental objectives in a manner that conveys the interrelationships between these entities. This linking of means and ends objectives is accomplished via a so-called *means-ends objectives network* [7]. In such a network, the goal is to provide tangible linkages between the decision makers' fundamental objectives and the instrumental means by which these objectives are realized or accomplished. In this regard, it is instructive to pose the question of how the fundamental objectives of the NSHS are achieved via means objectives. These means objectives—and their relation to the fundamental objectives of Figure 2—are depicted in the means-ends objectives network shown in Figure 3.

[7] For an overview of national strategies pertaining to national security and terrorism, see, e.g., [14].

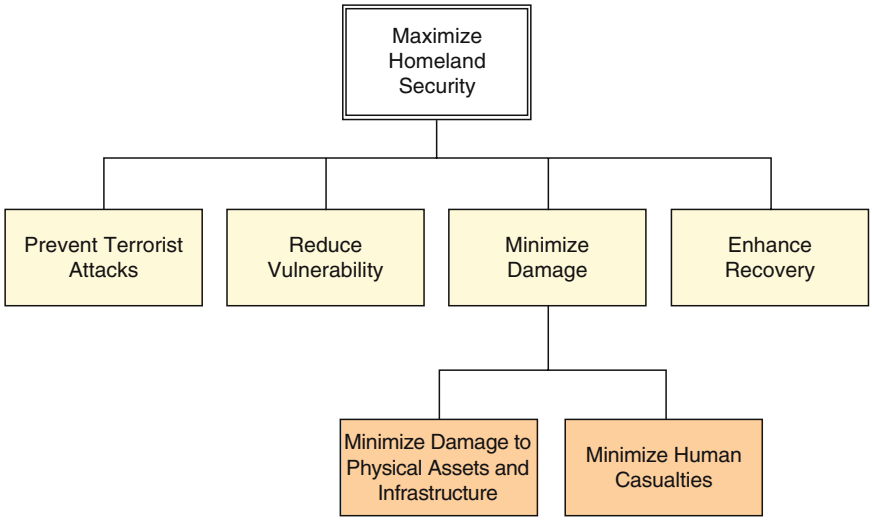

Figure 2. Fundamental Objectives Hierarchy for Homeland Security.

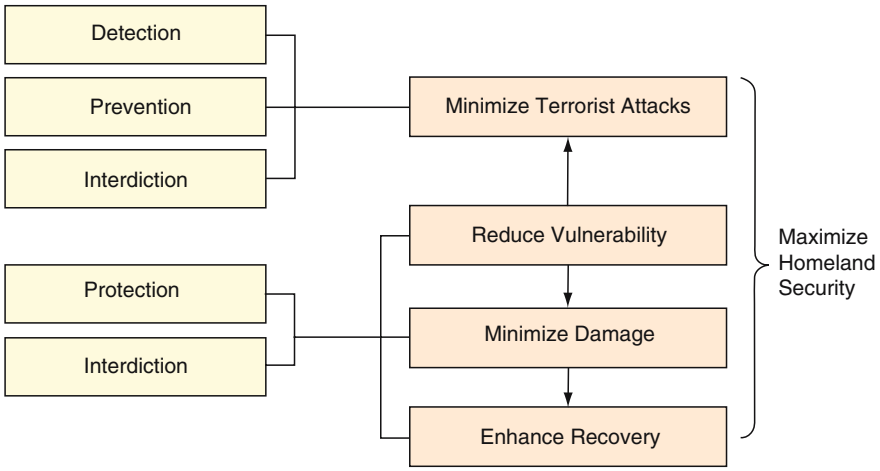

Figure 3. Relation of Key Homeland Security Mission Areas to Means-Ends Objectives.

4.1.2. Risk Assessment

Risk assessment enables decision makers to characterize and evaluate potential adverse consequences under uncertainty. In a typical risk assessment, the following questions are addressed:

- What can go wrong?
- What is the likelihood that something will go wrong?

- What are the consequences associated with these events? There may be multiple dimensions of effects, which may, in turn, be mapped into multi-attribute or benefit-cost analyses.

As a field of professional practice, risk assessment provides a powerful set of analytical tools for assessing the likelihood of events, together with their associated possible consequences. Risks can be evaluated by various methods, depending on the specific application, the available knowledge and information, and management's preferences.

4.1.2.1. Risk-Ranking Methods. Much current risk assessment practice depends on the qualitative, relative ranking of identified risks. Such rankings may be purely qualitative (using, perhaps, ad-hoc judgments), while others may have a more formal process, using multi-attribute or multi-objective approaches. In some simple cases, direct risk ranking is possible in decision situations where the outcomes are of the same type. In most settings, though, different types or levels of outcomes occur and more complex analyses involving weights or trade-offs are required. In these latter cases, ranking risks typically follows a sequence of steps that include:

1. Identifying consequence attributes (such as exposure or consequence)
2. Defining weights and scales for the attributes
3. Scoring event-consequence scenarios on these attributes
4. Aggregating the weighted scores

From a prescriptive vantage point, the following are some questions useful for evaluating risk-ranking models:

- Is sufficient and reliable information available for the analysis?
- Are attributes that potentially include both government and nongovernment items identified by a reasoned process?
- Is the form of aggregation of the attributes justified? If weights are used in the aggregation process, what justification is given for them?
- Are the upper and lower points of a scale well defined, or at least consistent, across risks in the problem domain?
- If group facilitation or elicitation methods are used to obtain scores or weights, how are the respondents selected? What information is provided to the respondents?
- If ranges or categories (such as 'high,' 'medium,' and 'low') are used, are risks identified as being near analytical boundaries considered in more detail, given the uncertain precision of the responses?

- Is the process formally documented?

4.1.2.2. Quantitative Risk Assessment. Quantitative risk assessments give rise to a wide range of possible outputs (e.g., point estimates, probability distributions). As we discuss below, it is in this step that the discrepancy between normative models of choice such as those sketched in Section 2 and the descriptive practice outlined in Section 3 diverge most markedly. The normative models make a number of unrealistic assumptions concerning the level of precision that is attainable in a complex system such as this (e.g., that the incremental effects of alternative investment options can be distinguished, and that cost information is measurable strictly in dollar terms).

From a prescriptive vantage point, examples of useful quantitative risk assessment questions will include the following:

- Is there a formal, logical model of the risks under consideration?
- What evidence supports the functional forms for the equations that link or functionally relate variables?
- What evidence supports the distributions that are assumed for the uncertain variables?
- What quality control steps are used to assess model validity and calibration?
- Does the analysis conform to accepted practice for the quantitative methods used?

4.1.2.3. Risk Assessments Based on Threat, Vulnerability, and Consequence. As discussed earlier, *threat*, *vulnerability*, and *consequence* are a frequently used decomposition in homeland security risk assessments [10]. In most security settings, all three components are present: a specific threat, a vulnerability in the asset or system that could be exploited by a specific threat, and a damaging outcome associated with specific threat and vulnerability combinations. In the context of our prescriptive framework, questions related to threat, vulnerability, and consequence will include the following:

- Is the threat information credible? How is threat information gathered? Does it come from multiple sources? How is it combined or summarized?
- Are a broad range of threat scenarios used in the risk assessment process?
- Are the threat scenarios generic (oriented toward a general threat environment), or are they particular to specific assets and locations?

- If risk filtering techniques are used to arrive at a manageable set of threat scenarios, how are they implemented? Are 'discarded' scenarios reassessed at some later stage, perhaps in response to new or improved information?
- Are likelihoods (expressed qualitatively or quantitatively) assessed for each identified threat scenario, or are all scenarios assumed to be equally likely? What is the evidence to support the kind of likelihood chosen?
- If likelihood is characterized qualitatively, is it clearly defined?
- Are cognitive biases (such as availability or saliency) managed as part of the threat characterization process?
- How are threat assessments coupled to assessments of vulnerability and consequence?
- What attributes are used to characterize an asset's vulnerability?
- Are weights assigned to each attribute? How are the weights determined?
- How are the consequences associated with specific threats characterized? Is more than one attribute (such as 'lives lost' or 'property damage') used to characterize these outcomes? If so, are the attributes defined clearly and consistently? Are the consequences monetized or used in a benefit-cost analysis?
- If consequences depend on threat, is the threat level clearly specified as part of the consequence valuation process?

4.1.3. Alternatives Evaluation

A risk assessment is likely to identify alternative ways in which decision makers can act to alter either the likelihoods or the outcomes associated with various identified risks. Prevention or damage-reducing actions may also be generated internally or externally through a publicly informed process. The alternatives may include a full range of actions, such as procedural changes, capital investments, regulations, and other actions.

Risks can be reduced appreciably by minimizing their likelihood or by mitigating their impact. In this regard, two concepts are key. The first is that action alternatives should be fed back through the risk assessment process to determine the extent to which risks can be reduced by the alternatives being considered. The initial risk assessment establishes at least part of the structure for evaluating the benefits of alternatives. Consideration should also be given to the possibility that certain actions may simply deflect risk to other assets of the agency, other parts of the government, or to the private

sector, all of which reduce the benefits of the action. The second concept is the role of costs to both government and the public; costs are a key element of alternatives evaluation. Major regulatory actions or capital investments generally require a cost- benefit or cost-effectiveness approach.

Core business and government guidance for evaluating alternatives for budgetary and regulatory purposes focuses on monetized net benefit evaluation. It is here, again, that substantial differences exist between normative best practices and current practice in many homeland security settings, due largely to the lack of accepted methods for quantifying and monetizing the full range of costs and benefits that should be considered as part of the alternatives evaluation process.

4.1.3.1. Structuring Portfolios of Risk Mitigation Strategies. The task of both identifying and structuring the risk mitigation options that will be appraised as part of the resource allocation process is an important aspect of our prescriptive framework. To this end, we are interested in characterizing and evaluating a portfolio of possible risk mitigation strategies. Moreover, we are interested in evaluating this portfolio relative to the kinds of objectives and criteria described earlier.

To this end, our first task is one of specifying the portfolio of possible risk mitigation strategies. There are numerous methods for accomplishing this task. A useful tool for this purpose is a strategy table, which provides a convenient way of summarizing a sequence of interrelated decisions. To illustrate, take the broadly defined means objectives that we described earlier. Under each of these broad categories, we can specify a set of possible risk mitigation strategies. As Figure 4 illustrates, a strategy table provides a convenient way of summarizing the overall portfolio of decision alternatives. The strategy table lists, in each vertical column, the set of risk reduction strategies identified for each means objective (e.g., 'Detection,' 'Prevention,' etc.). In this way, we are able to specify an entire portfolio of possible risk reduction strategies.

4.1.4. Management Selection

The fourth step in our prescriptive framework, *management selection*, entails choosing among possible alternative courses of action. Management's active participation is important at this stage because risk assessment tools contain various assumptions about preferences that may require value judgments and review at the management level. Management may also have values or information that analysts have not fully assessed. Once decisions have been reached, evidence that they were informed by risk-based information should be documented.

4.1.4.1. Evaluation of Risk Mitigation Strategies. As described earlier, the strategy table shown in Figure 4 represents the portfolio of all possible risk mitigation strategies that are deemed worthy of consideration. In making a

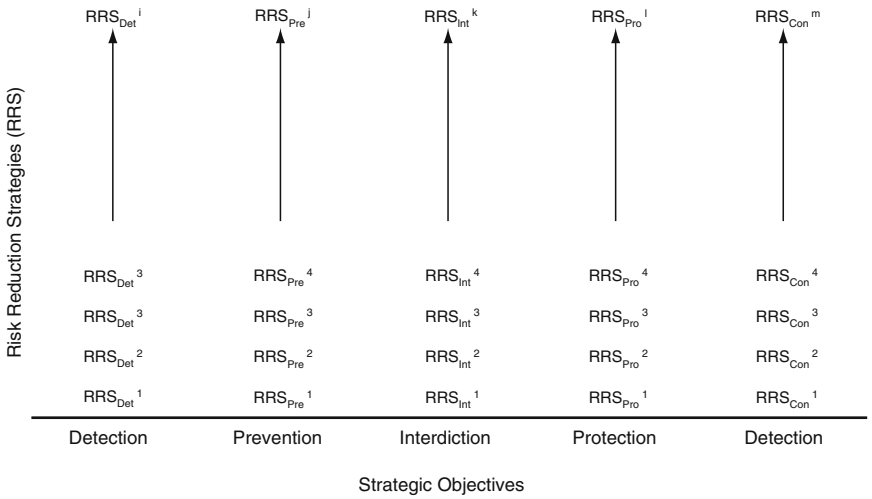

Figure 4. Strategy Table for Risk Reduction Strategies.

strategic resource allocation, our task is one of determining which combination of risk mitigation strategies provides the greatest overall value. In making this determination, decision makers will want to understand and explore key trade-offs, between, say, benefits versus costs or benefits versus risks.

To facilitate this type of analysis, it is possible to utilize objectives hierarchies like those described earlier to make the representation of such trade-offs an explicit feature of the strategic evaluation process. The objectives hierarchy shown in Figure 5 takes elements of our earlier hierarchies and marries them to an explicit consideration of benefit-cost trade-offs. Looking at the leftmost portion of the figure, we begin with the overall objective of maximizing homeland security. To the right of this fundamental objective is the key trade-off to be explored: *Benefits* and *Costs*. In this example, benefits are derived from the pursuit of the fundamental objectives described earlier (e.g., Prevention of Terrorist Attacks, Reduction of Vulnerabilities, etc.). For costs, we distinguish between monetary and non-monetary costs. At the rightmost portion of the diagram are the criteria against which the achievement of each objective is measured. For this illustrative set of criteria, it is, for example, possible to explore the trade-offs that exist between the benefits that might be derived from preventing terrorist attacks and the (social) cost associated with the potential loss of civil liberties.

4.1.5. Implementation and Monitoring

Any conceptual roadmap for how risk management principles can inform homeland security decision making must inevitably confront a number of

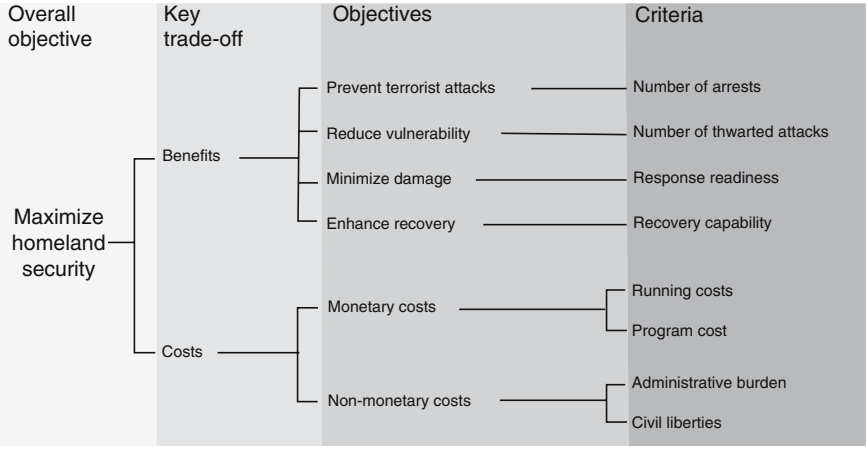

Figure 5. Hierarchical Representation of Objectives and Criteria for Benefit-Cost Trade-Offs.

issues that pertain to *implementation* and *monitoring*. Monitoring is essential to determine whether key objectives and milestones are being met, and whether policies and controls are giving rise to intended outcomes.

Risk management plans should be constructed in ways that ultimately support innovation and improvement, based on a process of continual feedback and learning. Monitoring helps ensure that the entire risk management process remains current and relevant, and that it reflects changes in the effectiveness of the actions and the risk environment in which it operates. Monitoring the risk management plan also involves assessment of the adequacy of strategic objectives and performance measures, as well as ensuring that service delivery and support functions are consistent with design specifications and implemented in accordance with the plan's timeframe.

In assessing, again from a prescriptive vantage point, the implementation of risk mitigation actions, it is useful to pose the following sorts of questions:

- Are objectives and time schedules specified for implementation actions?
- Are mitigation actions implemented as specified?
- Are mitigation actions implemented in a timely manner?
- Do mitigation actions meet cost objectives?
- Are internal controls adequate?
- Are risk communication issues considered?

In addressing monitoring and evaluation activities, critical questions will include the following:

- What types of ongoing monitoring occur as part of the overall risk management process?
- If performance measures exist, what is the outcome of performance measurement protocols and procedures?
- Has the agency previously evaluated the program or does it have a detailed plan for evaluating the program?
- Does the evaluation conform to best practices?
- Are the recommended activities reviewed periodically?
- Are risk scenarios kept up to date and is the system tested periodically?
- How often do decision makers review the entire risk management system?
- What mechanisms identify and deal with risks affected by changing circumstances or new information?
- Do barriers have a significant impact on the agency's ability to achieve its risk management goals?

4.1.5.2. Continual Feedback. Active monitoring is essential to providing feedback to decision makers for continual or periodic improvement of the risk management plan (as dictated by the situation or context), together with information as to whether the plan coordinates effectively with other relevant plans, programs, and agencies. Risk management is a dynamic process and monitoring is a check on whether resources are used effectively and efficiently. Monitoring and evaluation provide information to management and stakeholders about the status of the plan, such as if the plan is in compliance with all current applicable professional standards, and if all memorandums of understanding and mutual aid agreements are in place, and that legal liability concerns have been resolved.

5. Concluding Remarks

In this paper, we have sought to explore a number of issues pertaining to federal decision making for homeland security, looking specifically at the divide—both conceptual and pragmatic—that exists between normative theories of choice and descriptive decision-making practice as it presently exists in the homeland security domain. Our attempts to understand the nature of this divide—and its implications for decision quality, program effectiveness, and economic efficiency, among other things—has motivated a prescriptive framework that seeks, on the one hand, to make the best use of normative insights, while, on the other, candidly confronting the difficulties

(cognitive and otherwise) that decision makers routinely confront in these complex and uncertain realms. If, in our approach, there is a bias, it is in strongly siding with the view that risk management is the *sine qua non* for how extreme events—both manmade and natural—must be construed and managed in the post-9/11 era.

Of course, the tragedy of Hurricane Katrina illustrates what is, perhaps, one of the most vexing challenges in the homeland security domain, namely, how best to allocate scarce resources among the vast panoply of catastrophic risks that can beset mankind in the technological society of the 21st century. Any reasoned risk management approach begins with a cold and dispassionate assessment of the true extent of the nation's vulnerability to a diverse range of threats and hazards. As we have said, at the federal level, the organizational challenges that must be confronted in these domains are significant. In this paper, we have argued for a common set of analytical tools and procedures regarding how the federal government invokes and makes use of risk management concepts and techniques. While current federal approaches to homeland security decision making is evolving towards consistency with the risk management approach articulated here, substantial gaps still exist. The challenge remains one of continued vigilance, flexibility, and resilience in anticipation of, and in response to, an ever-changing threat/hazard environment.

Acknowledgments

We are grateful to a number of colleagues for useful comments and suggestions on many aspects of the work presented here. The prescriptive framework presented in Section 4 draws heavily on work done by both authors at the U.S. GAO, where Neil Asaba, Nancy Briggs, Steve Caldwell, Nancy Kingsbury, Norm Rabkin, and numerous others provided valuable feedback and commentary. Any errors are, of course, our own.

References

1. *Risk Assessment in the Federal Government: Managing the Process,* 1983. National Academy Press, Washington, DC.
2. Chalk, P., Hoffman, B., Reville, R., and Kasupski, A., 2005. *Trends in Terrorism: Threats to the United States and the Future of the Terrorism Risk Insurance Act*. RAND Corporation, Santa Monica, CA.
3. Ervin, C. K., 2006. *Open Target: Where America Is Vulnerable to Attack*. Palgrave Macmillan, New York.
4. Farrow, S., 2007. The economics of homeland security expenditures: foundational expected cost-effectiveness approaches. *Contemporary Economic Policy* 25(1):14–26 (January 2007).
5. Haimes, Y., 2004. *Risk Modeling, Assessment, and Management*. Wiley, Hoboken, NJ.

6. Keeney, R., and Raiffa, H., 1993. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge University Press, Cambridge.
7. Keeney, R. L., 1992. *Value-Focused Thinking: A Path to Creative Decision-Making*. Harvard University Press, Cambridge, MA.
8. Machina, M. J., 1987. Choice under uncertainty: problems solved and unsolved. *Economic Perspectives* 1(1):121–154.
9. Morgan, M. G., and Henrion, M., 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press, New York.
10. Moteff, J., 2004. Risk management and critical infrastructure protection: Assessing, integrating, and managing threats, vulnerabilities, and consequences. Tech. Rep. RL32561, Congressional Research Service, Washington, DC.
11. Posner, R. A., 2005. *Preventing Surprise Attacks: Intelligence Reform in the Wake of 9/11*. Rowman & Littlefield, Lanham, MD.
12. Raiffa, H., 1968. *Decision Analysis*. Addison-Wesley, Reading, MA.
13. United States Government Accountability Office (GAO), 2001. *Key Elements of a Risk Management Approach*. No. GAO–02–150T. Washington, DC.
14. United States GAO, 2004. *Combating Terrorism: Evaluation of Selected Characteristics in National Strategies Related to Terrorism*. No. GAO–04–408T. Washington, DC.
15. United States GAO, 2005. *Risk Management: Further Refinements Needed to Assess Risks and Prioritize Protective Measures at Ports and Other Critical Infrastructure*. No. GAO–0691. Washington, DC.
16. Viscusi, W. K., 1998. *Rational Risk Policy*. Oxford University Press, Oxford.