

Decision Support

Nature plays with dice – terrorists do not: Allocating resources to counter strategic versus probabilistic risks

Boaz Golany^{a,*,1}, Edward H. Kaplan^b, Abraham Marmur^c, Uriel G. Rothblum^{a,1}

^a Faculty of Industrial Engineering and Management, Technion, Israel Institute of Technology, Haifa 32000, Israel

^b School of Management, School of Medicine, Faculty of Engineering, Yale University, New Haven, CT 06520, USA

^c Center for Security Science and Technology, Department of Chemical Engineering, Technion, Israel Institute of Technology, Haifa 32000, Israel

Received 1 December 2006; accepted 3 September 2007

Available online 11 September 2007

Abstract

Probabilistic uncertainty is caused by “chance”, whereas strategic uncertainty is caused by an adverse interested party. Using linear impact functions, the problems of allocating a limited resource to defend sites that face either probabilistic risk or strategic risk are formulated as optimization problems that are solved explicitly. The resulting optimal policies differ – under probabilistic risk, the optimal policy is to focus the investment of resources on priority sites where they yield the highest impact, while under strategic risk, the best policy is to spread the resources so as to decrease the potential damage level of the most vulnerable site(s). Neither solution coincides with the commonly practiced proportionality allocation scheme.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Risk analysis; Game theory; Resource allocation

1. Introduction

September 11, 2001 and December 24, 2004 will be remembered for a very long time – on both days a shocked world found itself facing unexpected disasters of horrendous magnitude. Though property damage and loss of life occurred on both of these days, the causes were drastically different. On September 11, a focused and determined group of terrorists executed a vicious attack on the US, whereas on December 24, the destructive tsunami that hit Southeast Asia was an “act of nature.”

Functioning in environments in which unexpected events occur and taking action without having full control

over the resulting outcome is a fundamental problem that every living creature faces. In fact, an important characteristic of human beings is the ability to exercise reason and rationality when making decisions in situations where outcomes are uncertain, that is, not fully predictable. Such uncertainty reflects the influence of factors that the decision-maker does not control. In this regard, one has to distinguish between uncertainty that is the result of coincidence (or chance) and uncertainty that is the result of actions taken by interested parties; we refer to the former as *probabilistic uncertainty* and to the latter as *strategic uncertainty*.² In both cases, we refer to risk as the joint influence of the uncertainty and the associated impact on the decision-maker.

Once the notions of probabilistic and strategic risks are formalized, it is evident that decision-making differs in

* Corresponding author. Tel.: +972 4 829 4512; fax: +972 4 829 5688.

E-mail addresses: golany@ie.technion.ac.il (B. Golany), edward.kaplan@yale.edu (E.H. Kaplan), marmur@technion.ac.il (A. Marmur), rothblum@ie.technion.ac.il (U.G. Rothblum).

¹ Research of these authors was partially supported by Avnet 37, The Magnet Program, Chief Scientist, Israel Ministry of Industry, Commerce and Employment.

² Our use of the expression “strategic” refers to environments in which intentional actions (of interested parties) are taken, and not in the sense (used frequently in the Operations Management literature) of a higher level planning – above the operational and tactical levels.

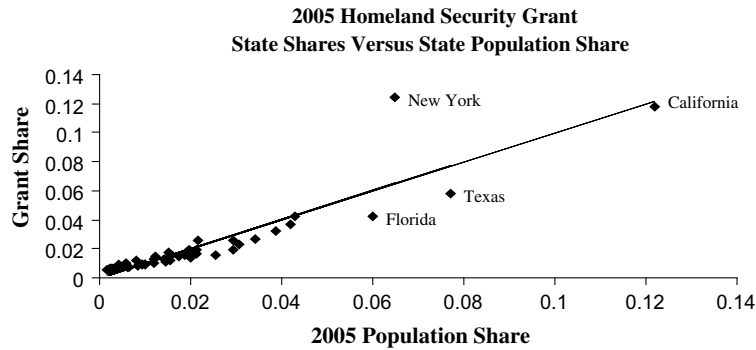


Fig. 1. 2005 state shares in homeland security grants vs. state population.

these environments. Probability theory, utility theory, decision analysis and their extensions have developed tools for understanding and assisting decision making in environments that exhibit probabilistic risk, whereas game theory is the field that explores strategic environments. Studies of similarities and differences between environments that exhibit probabilistic and strategic risks have emerged recently (e.g., [3,10,19]). Still, no comprehensive quantitative comparative analysis of these environments has been conducted to date.

The distinction between probabilistic and strategic risks has been at the center of recent discussions in the US Senate when debating ways to appropriate funds for homeland security purposes. Fig. 1 is based on data for the 2005 allocation of homeland security grants to states in the US. It demonstrates that this funding was, by and large, proportional to the population size in each state.

The proportional allocation came under severe criticism when the senate debated the appropriations for 2006 as many lawmakers argued that areas containing national symbols or strategic assets are vulnerable to terrorists' attacks regardless of their population size. The debate led to a new bill that favored "risk-based" rather than "formula-based" allocations.³ However, a quick look at the allocations that were actually made in 2006 (see Fig. 2), reveals that with the exception of a few states (notably, Washington DC), the allocation has once again resulted in funds being allocated in rough proportion to state populations.

From the dollar allocations alone, one cannot determine if these allocations were constrained to return proportional

results, or if homeland security officials "discovered" that newly-constructed risk measures yielded allocations proportional to population. For example, if officials decided to allocate in proportion to "risk" (however measured), and measured risk turned out to be proportional to population, then the resulting allocations would not be surprising. Still, the question remains – if you were a decision-maker facing a similar resource allocation problem, how would you define risk and consequently, what model would you use to determine the best possible allocation of resources in the face of that risk?

The current paper conducts a comparative study of probabilistic versus strategic risks with respect to one specific aspect. We consider the problem of a decision-maker who is allocating a resource to several sites so as to reduce the expected damage of uncertain mishaps, in the presence of either probabilistic or strategic uncertainty. The allocation problem in the probabilistic environment is that of a defender who attempts to protect several sites against potential mishaps that might be caused by acts of nature like floods and earthquakes. The corresponding problem in a strategic environment is that a decision-maker in the Homeland Security (HS) arena has to decide on the best way to protect certain sites from potential terrorist attacks. We formalize the two problems and identify optimal solutions that yield distinctly different optimal policies, demonstrating substantial differences between the two environments.

The key conclusions of our analysis show that in the face of probabilistic risk, of the form caused by natural disasters, it is optimal to focus the effort by implementing a priority rule that determines which site(s) should receive the maximal possible protection. In contrast, when facing a strategic risk, of the form encountered in terrorist attacks, it is optimal to set a risk-threshold and spread the effort in parallel. Spreading the effort means that only partial protection will materialize at each site. (In an operational environment where the allocation of the resource is used to facilitate repeated actions of similar nature, this effect can be achieved by randomizing over each of the repeated decisions; see the example described in Section 6.) While the optimization techniques we apply are rather standard, the

³ The debate that preceded the approval of the new bill is reflected in the words of Sen. Diane Feinstein: "The bottom line is that if Federal funds are going to be distributed to improve first responders' ability to "prevent, prepare for, respond to, or mitigate threatened or actual terrorist attacks," those funds should be distributed in accordance with a risk-based analysis. Al-Qaida and its allies do not attack based on a formula. *This bill rejects the formula approach in favor of a framework that is flexible and risk focused.*" (Source: Congressional Record, Senate, July 12, 2005, p. S8093, available online at <http://tinyurl.com/yje8cz> or <http://tinyurl.com/ykrowx>).




Fig. 2. 2005–2006 state shares in homeland security grants vs. state population.

most interesting aspect of our findings is the contrast between the similarity in the formulation of the problems and the difference in the structure of the corresponding optimal policies.

The optimal policies for the two allocation problems that we consider are compared with the frequently used rule of allocating resources in proportion to “need” – in our context, “need” means potential risk. Also, an interesting derivative of our results is a marginal analysis which determines the effect of modifying the available amount of the resource.

The solution we obtain for the allocation problem in the strategic environment is shown to be a unique Nash equilibrium of a corresponding game whose players are the attacker and the defender. It should be noted that game theory has already been used to solve resource allocation problems in strategic setups where multiple sites are to be defended from potential attacks by an adversary. Shubik and Weber [21,22] and Shubik [20] have developed and analyzed “system defence games” while Coughlin [6] extended their work. However, these articles modelled such situations as cooperative games while the approach we adopt here is that of non-cooperative games. Harris [8] pointed out the role of game-theoretic models in addressing terrorism risks but he did not develop specific models as is done in the current paper. Finally, Sandler and Arce [18] reviewed the literature on the use of game theory to explore measures of countering terrorist acts and offered some novel insights on anti-terrorism policies.

The problem we investigate under strategic risk has two characteristics that simplify its analysis. First, the setup is one of *complete conflict* – the terrorists’ gain is the damage and loss their attack causes to the defender (see, e.g. [6, p. 197]). In fact, this is the ultimate zero-sum game, a term that is commonly used to capture conflicting interests of (two) parties in systems where the final outcome is a redistribution of the resources held originally by the participants. Second, while in conventional warfare armies are capable of hiding (or masking) some of their important assets (e.g., aircraft in underground bunkers, missiles in

protected silos), the assets in the HS theater are basically exposed. Eventual attacks take place in areas that are open to the public and would be terrorists can move about just like ordinary citizens. In particular, they can collect information from open sources (see [4]), not only about the devices that are deployed, but also about operational routines, security patrols etc., until they believe they are satisfied with the data they have accumulated. They can also assess with high accuracy the potential damage attacks might inflict and the likelihood (probabilities) that such attacks will be successful. Indeed, (unfortunately) evidence from recent years indicates that collecting quality intelligence has become the prevailing practice for terrorists (see [13]). So, a reasonable first approximation is a model with the basic assumption that terrorists have complete information. We incorporate such an assumption in our basic model. In the last section of this paper, we suggest ways in which this assumption might be relaxed.

We assume that the impact of the allocated resource is linear. While this is a crude approximation of reality, it helps us gain insights into the two contexts we explore (probabilistic versus strategic risks), and crystallizes the differences between the nature of the optimal policies that are obtained for them.

The paper is organized as follows. Two resource allocation problems are presented and solved in Sections 2 and 3 – one reflects probabilistic risk while the other reflects strategic risk. Section 4 and its associated appendix prove that the solution of the resource allocation problem under strategic risk is a unique Nash equilibrium of a corresponding game (whose players are the attacker and the defender). Section 5 considers the commonly used proportionality rules and Section 6 describes an example which analyzes potential policies of the Immigration and Naturalization Service (INS) in administering the fingerprinting of visitors and immigrants when entering the US; the example discusses allocation problems in both probabilistic and strategic environments. Finally, Section 7 contains some concluding remarks and offers directions for future research.

2. Allocating resources to counter probabilistic risks

In this section we consider the problem of allocating a single continuous resource to counter the probabilistic risks of an undesired event (disaster) hitting any of n distinct sites, where the goal is to minimize the expected loss.

The characteristics (data) for each site i are assumed to be given by⁴

1. π_i – the probability that site i will be hit.
2. $0 \leq p_i < 1$ – the (conditional) probability that site i will be compromised when hit, given that no investment was made to defend it.
3. $C_i > 0$ – the incurred cost.
4. $\alpha_i > 0$ – the proportion-coefficient of the effectiveness associated with the investment of the resource to defend site i , specifically, investing x_i units of the resource in site i reduces the (conditional) probability that this site will be compromised when hit to $p_i - \alpha_i x_i$ in the event that site i is compromised.

Without any action, the expected damage at site i is the product $\pi_i p_i C_i$.^{5,6} The expected damage at the various sites can be reduced through the allocation of the resource. Consider the decision variables which indicate the (non-negative) amounts x_1, \dots, x_n that are to be invested, respectively, in the protection of n sites. These amounts are subject to individual constraints $p_i - \alpha_i x_i \geq 0$ for $i = 1, \dots, n$, indicating that no investment is to be made beyond full protection. In addition, there is a (joint) capacity/budget constraint of the resource; with B representing the total amount of the resource that is available, one has the constraint $\sum_{i=1}^n x_i \leq B$. Now, if x_1, \dots, x_n units of the resource are allocated to the defence of the corresponding sites, an attack at site i will result in damage C_i with probability $p_i - \alpha_i x_i$, yielding expected loss of $C_i(p_i - \alpha_i x_i)$; the goal of the decision-maker is to minimize the expected global loss which is given by $\sum_{i=1}^n \pi_i C_i(p_i - \alpha_i x_i)$. The optimization problem that the decision-maker faces under the above assumptions is given by the following Probabilistic Risk Allocation Problem (PRAP):

$$\begin{aligned} \text{PRAP:} \quad & \min_x \quad \psi = \sum_{i=1}^n \pi_i C_i (p_i - \alpha_i x_i) \\ \text{subject to:} \quad & \sum_{i=1}^n x_i \leq B, \\ & 0 \leq x_i \leq \frac{p_i}{\alpha_i} \quad \text{for } i = 1, \dots, n. \end{aligned} \tag{1}$$

PRAP is a standard Continuous Knapsack Problem (e.g. [11]), which is known to have an optimal solution that follows a simple priority rule. Specifically, the solution is obtained by first ranking the sites by the products $\pi_i C_i \alpha_i$, resulting in an enumeration i_1, \dots, i_n of the sites so that

$$\pi_{i_1} C_{i_1} \alpha_{i_1} \geq \dots \geq \pi_{i_n} C_{i_n} \alpha_{i_n}. \tag{2}$$

The optimal allocation is then to direct as much as is possible of the resource to the top-ranked sites, subject to the constraints. Specifically, this policy can be determined as follows: allocate the amount $x_{i_1}^* = \min \left\{ \frac{p_{i_1}}{\alpha_{i_1}}, B \right\}$ of the resource to site i_1 ; if there is still some resource left, allocate $x_{i_2}^* = \min \left\{ \frac{p_{i_2}}{\alpha_{i_2}}, B - x_{i_1}^* \right\}$ to site i_2 ; continue in this manner until the resource capacity is exhausted or all sites are fully protected. When ties are present in (2), the allocation of the resource can be directed to either of the tied sites, or it can be split arbitrarily among them.

We observe that the above solution is parametric, i.e., it provides optimal solutions of PRAP for all possible capacity levels. Further, it is robust in the sense that an optimal plan made under any assumed capacity level can be easily adjusted to any modification of the capacity level; specifically, use of the resource that has been decided on before the change of the available amount remains part of the optimal solution as long as it does not violate the new availability constraint. In particular, the analysis provides a marginal (sensitivity) analysis to the problem showing that when an extra unit of the resource becomes available, it will be allocated in its entirety to the site with the highest priority (according to (2)) of those that are not yet fully protected. Finally, if one or more new sites are added, the optimal solution is available through the adjustment of (2) and the targeting of the resource accordingly. In particular, the internal ranking of the previous sites will not be changed, and no “rank-reversal” will take place.

3. Allocating resources to counter strategic risks

We next consider the problem of allocating a single continuous resource to counter strategic risks in n distinct sites that are subject to a potential attack by a strategic adversary.

The characteristics (data) for each site i are similar to those of PRAP and are given by

1. $0 \leq p_i < 1$ – the probability that site i will be compromised when attacked, given that no investment was made to defend it.
2. $C_i > 0$ – the incurred cost in the event that site i is compromised.

⁴ The calculation of these characteristics is not a trivial task. Examples involving multi-attribute utility theory, multi-objective programming, value trees, hierarchical holographical models, and more appear in [1,12].

⁵ For FY 2006 The Department of Homeland Security of the US Government defines “risk” in the following way (see [7]): “For FY2006, risk is defined as the product of three principle variables: the consequences of a specified attack to a particular asset, the vulnerability of that asset to the particular threat, and the threat to that asset” – the mentioned three

⁶ An alternative interpretation for C_i and p_i is for C_i to be the maximal potential damage that site i might suffer, i.e., “total loss”, and for p_i to be the (expected) proportion of the maximal potential damage C_i that site i will suffer when attacked.

- $\alpha_i > 0$ – The proportion-coefficient of the effectiveness associated with the investment of the resource to defend site i , that is, investing x_i units of the resource in site i reduces the probability that this site will be compromised when attacked to $p_i - \alpha_i x_i$.

As in the PRAP, the decision variables are the amounts x_1, \dots, x_n that are to be invested, respectively, in the protection of the n sites. These amounts are subject to the same constraints as under PRAP, with the available amount of the resource denoted by B . Also, if x_1, \dots, x_n units of the resource are allocated to the defence of the corresponding sites, an attack at site i will result in damage C_i with probability $p_i - \alpha_i x_i$, yielding an expected loss of $C_i(p_i - \alpha_i x_i)$. But, here we assume that attacks are strategic, carried out by an opponent that has full information and is determined to inflict the heaviest possible damage.⁷ Thus, an attack will take place at a most vulnerable site, resulting in an expected damage of $\max_{i=1, \dots, n} C_i(p_i - \alpha_i x_i)$. Here again, the goal of the decision-maker is to minimize expected loss which is an optimization problem given by the following Strategic Risk Allocation Problem (SRAP):

$$\begin{aligned} \text{SRAP:} \quad & \min_x \theta = \max_{i=1, \dots, n} C_i(p_i - \alpha_i x_i) \\ \text{subject to:} \quad & \sum_{i=1}^n x_i = B, \\ & 0 \leq x_i \leq \frac{p_i}{\alpha_i} \quad \text{for } i = 1, \dots, n. \end{aligned} \tag{3}$$

The optimization problem that expresses SRAP has a unique optimal solution which has a different structure than that found for PRAP. Specifically, denote the optimal objective value of SRAP by θ^* (its explicit expression depends on the capacity B and is about to be determined below, in (8) and (9)). The value θ^* serves as a threshold level for the expected loss in the various sites; this means that the expected damage under optimal allocation levels x_1^*, \dots, x_n^* satisfies

$$C_i(p_i - \alpha_i x_i^*) = \begin{cases} \theta^* & \text{if } C_i p_i \geq \theta^*, \\ C_i p_i & \text{if } C_i p_i < \theta^*. \end{cases} \tag{4}$$

The condition about x_1^*, \dots, x_n^* expressed in (4) is equivalent to

$$x_i^* = \begin{cases} \frac{C_i p_i - \theta^*}{C_i \alpha_i} & \text{if } C_i p_i \geq \theta^*, \\ 0 & \text{if } C_i p_i < \theta^*. \end{cases} \tag{5}$$

A formal proof of the optimality of x_1^*, \dots, x_n^* as given by (5) and the derivation of an explicit expression for θ^* , provided below in (9), is given in the Appendix. Of course, (4) implies that

$$C_i(p_i - \alpha_i x_i^*) \leq \theta^* \quad \forall i = 1, \dots, n. \tag{6}$$

⁷ The assumption about full information can be relaxed by assuming that the terrorists have probabilistic information about the damage that their actions will cause and hence they apply the criterion of expected damage.

We next describe a process for computing the optimal objective value θ^* . First, rank the sites by the products $C_i p_i$, yielding an enumeration i_1, \dots, i_n of the sites so that

$$C_{i_1} p_{i_1} \geq \dots \geq C_{i_n} p_{i_n} \tag{7}$$

(notice the difference between this enumeration and the one used in (2) for the solution of PRAP). Next, allocate the resource to the weakest (top ranked) site i_1 until its expected loss is equated with that of i_2 . The amount of the resource allocated to site i_1 at this stage will be $x_{i_1,1}^* = \min \left\{ \frac{C_{i_1} p_{i_1} - C_{i_2} p_{i_2}}{C_{i_1} \alpha_{i_1}}, B \right\}$. If the resource is not depleted in the first stage, one enters the second stage. In this stage, the resource is allocated in parallel to sites i_1 and i_2 , at rates that are inverse-proportional to the products $C_i \alpha_i$; the allocation at these proportions maintains the same protection levels at sites i_1 and i_2 . The second stage continues until the joint expected loss of i_1 and i_2 is equated with that of i_3 , or until the resource is depleted. The process continues until either one runs out of the resource, or one achieves 100% protection at all sites (which means that the expected loss at all sites equals zero). Augmenting the list of sites with a fictitious site $n+1$ with $p_{n+1} = 0$, $C_{n+1} = \alpha_{n+1} = 1$, an explicit expression for θ^* is available by determining

$$t^* \equiv \max \left\{ t = 1, \dots, n+1 : \sum_{s=1}^t \frac{C_{i_s} p_{i_s} - C_{i_t} p_{i_t}}{C_{i_s} \alpha_{i_s}} \leq B \right\}, \tag{8}$$

and observing that

$$\theta^* = \begin{cases} C_{i_{t^*}} p_{i_{t^*}} + \frac{B - \sum_{s=1}^{t^*} \frac{C_{i_s} p_{i_s} - C_{i_{t^*}} p_{i_{t^*}}}{C_{i_s} \alpha_{i_s}}}{\sum_{s=1}^{t^*} \frac{1}{C_{i_s} \alpha_{i_s}}} & \text{if } t^* \leq n \\ 0 & \text{if } t^* = n+1. \end{cases} \tag{9}$$

The above solution of SRAP enjoys the same features as the solution of PRAP. First, it is *parametric* – that is, it provides a solution of SRAP for all possible capacity levels. Second, it is *robust* in the sense that an optimal plan made under any assumed capacity level can be easily adjusted to any modification of the capacity level; decisions that had been implemented before the change remain part of the optimal solution (as long as they do not violate the new capacity constraint). In particular, the analysis provides a *marginal (sensitivity) analysis* to the problem showing that when an extra unit of the resource becomes available, it will be distributed in parallel to the sites that are most vulnerable under the current allocation; the proportions of the parallel allocation are to be determined by the reciprocals of $C_i \alpha_i$ of these (most vulnerable) sites. Also, if a site is added, the optimal solution can be adjusted by modifying (7), without “rank-reversal.”

4. Optimal strategic allocation and economic Nash equilibrium

In this section, we view SRAP as part of a two-person zero-sum game and show that the optimal solution of SRAP is part of a Nash equilibrium for that game. In fact,

we show that under a simple non-degeneracy assumption, this Nash equilibrium is unique.

Consider the two-person zero-sum game, where the two players are the defender and the attacker. In this game, the attacker’s gain is the defender’s loss – the unfortunate reality in terrorist attacks. We note that this game is an ultimate zero-sum game, expressing conflicting interests of the players; this is the case not because payments are transferred, but from the pure perspective of preferences.

We will restrict the defender to pure strategies of allocating the resource (continuously) to the different sites, while the attacker is assumed to apply randomized strategies, meaning that it is allowed to use a lottery to determine a single target. So, the attacker’s decision variables are probabilities w_1, \dots, w_n of attacking the different sites. The loss to the defender (the attacker’s gain) when the attacker selects such a policy, while the defender allocates x_1, \dots, x_n to defend these sites, is $\sum_{i=1}^n w_i C_i (p_i - \alpha_i x_i)$. The probabilities w_1, \dots, w_n are subject to the constraints $w_i \geq 0$ for $i = 1, \dots, n$ and $\sum_{i=1}^n w_i = 1$.

The set of strategies of the defender is given by

$$\mathcal{X} \equiv \left\{ x \in R^n : 0 \leq x_i \leq \frac{p_i}{\alpha_i} \text{ for } i = 1, \dots, n, \sum_{i=1}^n x_i \leq B \right\}, \tag{10}$$

while the set of randomized strategies of the attacker is given by

$$\mathcal{W} \equiv \left\{ w \in R^n : w \geq 0, \sum_{i=1}^n w_i = 1 \right\}. \tag{11}$$

The problem that the defender faces is a (formally) modified, but equivalent, variant of SRAP that we denote Rand-SRAP

$$\text{Rand-SRAP} : \min_{x \in \mathcal{X}} \max_{w \in \mathcal{W}} \sum_{i=1}^n w_i C_i (p_i - \alpha_i x_i). \tag{12}$$

When the attacker takes the pessimistic point of view (from his/her perspective) that the defender allocates his resource optimally, the problem that the attacker faces is

$$\text{Dual-SRAP} : \max_{w \in \mathcal{W}} \min_{x \in \mathcal{X}} \sum_{i=1}^n w_i C_i (p_i - \alpha_i x_i). \tag{13}$$

A variant of the well-known von Neumann–Morgenstern Min–Max theorem (e.g. [23,16],) assures that the Rand-SRAP (and its equivalent version SRAP) and the Dual-SRAP have the same optimal objective function. Further (e.g. [14]), a pair $\bar{x} \in \mathcal{X}$ and $\bar{w} \in \mathcal{W}$ are, respectively, optimal for these problems if and only if they form a Nash equilibrium, meaning that

$$\begin{aligned} \sum_{i=1}^n w_i C_i (p_i - \alpha_i \bar{x}_i) &\leq \sum_{i=1}^n \bar{w}_i C_i (p_i - \alpha_i \bar{x}_i) \\ &\leq \sum_{i=1}^n \bar{w}_i C_i (p_i - \alpha_i x_i) \end{aligned} \tag{14}$$

for all $x \in \mathcal{X}$ and $w \in \mathcal{W}$,

and with (\bar{x}, \bar{w}) satisfying either of the above conditions,

$$\bar{\theta} \equiv \sum_{i=1}^n \bar{w}_i C_i (p_i - \alpha_i \bar{x}_i) \tag{15}$$

is the common optimal objective function on SCRП, Rand-SRAP and Dual-SRAP.

We shall use the above equivalence in the Appendix to formally prove the optimality of the solution of SRAP that is presented in Section 3 along with the optimality of the solution of Dual-SCRП which we present below in (16) and (17); this is done by verifying that this pair of vectors forms a Nash equilibrium.

To present the optimal solution of Dual-SRAP, let θ^* be the scalar given by (7)–(9) and let x^* be the vector given by (5). We recall that it was argued (and will be proven in the Appendix) that θ^* is the optimal objective value of SRAP and that x^* is the unique optimal solution for the SRAP (which is also optimal for Rand-SRAP). Let

$$J \equiv \{i = 1, \dots, n : x_i^* = 0\} = \{i = 1, \dots, n : C_i p_i \leq \theta^*\} \tag{16}$$

(the equality following from (5)) and let w^* be the vector given by

$$w_i^* = \begin{cases} \left(\frac{1}{C_i \alpha_i}\right) \left(\frac{1}{\sum_{j \in J^c} C_j p_j}\right) & \text{if } i \notin J, \\ 0 & \text{if } i \in J. \end{cases} \tag{17}$$

In addition to showing in the Appendix that w^* is an optimal solution of Dual-SRAP, it is shown that w^* is the unique optimal solution under the following (non-degeneracy) assumption:

$$\{i = 1, \dots, n : C_i p_i = \theta^*\} = \emptyset. \tag{18}$$

Here, the proof (in the Appendix) is based on verifying that (under the above non-degeneracy assumption) w^* is the only vector that forms a Nash equilibrium together with x^* .

We end this section with an informal motivation of (17). Under the non-degeneracy assumption, the equilibrium attack probabilities at site i are positive if and only if site i is among the most vulnerable sites under the optimal allocation. These optimal attack probabilities are then calculated so as to equate the expected damage at the most vulnerable sites, denying the defender’s possibility of reducing the expected damage by shifting the resource allocation. In particular, the solution is such that the expected damage at any target attacked never exceeds a maximum level θ^* ; the only way to reduce this maximal expected damage level is to increase the available resource (the defence budget).

We finally note (but without a proof here or in the Appendix) that without the non-degeneracy assumption, the set of all optimal solutions for Dual-SRAP is obtained by the attacker dividing the attack-probabilities among the sites of $J^c = \{i = 1, \dots, n : C_i (p_i - \alpha_i x_i^*) = \theta^*\}$, assigning equal probabilities to the sites in $J^c = \{i = 1, \dots, n : x_i^* > 0\}$ and lower probability levels to the sites in $J^c \setminus J^c$.

5. Proportional allocations

A common practice in resource allocation aimed at countering negative effects is to allocate the resource in proportion to the potential needs. The most basic way to implement this scheme is to allocate in proportion to the C_i 's, regardless of the p_i 's. Indeed, billions of HS dollars have been allocated to US states and territories in proportion to population (see [5, pp. 7–8]). Only recently, the HS Department started to consider the application of a so-called “risk-effects” policy which means that the allocation will be done in proportion to $\pi_i p_i C_i$ as discussed below, a so-called “risk-based” policy [15].

A critical view of the approach of proportional allocation, in the context of allocating funds to HIV prevention, can be found in [9] (which considers a resource allocation problem in the context of a probabilistic environment, as discussed in Section 2). Using our notation, the potential risk in the context of the allocation of a resource to counter probabilistic risk at site i is given by $\pi_i C_i p_i$. Allocation levels under a proportional scheme are then given by $\gamma \pi_i C_i p_i$ with $\gamma = B / \sum_i \pi_i C_i p_i$. Of course, an underlying assumption is that $B < \sum_i \pi_i C_i p_i$, under which $\gamma < 1$. The expected loss at site i under this policy will be $\pi_i C_i p_i - \alpha_i \gamma \pi_i C_i p_i = \pi_i C_i p_i (1 - \gamma \alpha_i)$. It is easy to verify that this allocation rule is optimal for a decision-maker who considers the following Log-Probabilistic Risk Allocation Problem (Log-PRAP):

$$\begin{aligned} \text{Log-PRAP: } \quad & \min_x \quad \rho = \sum_{i=1}^n \pi_i C_i p_i \log x_i \\ \text{subject to: } \quad & \sum_{i=1}^n x_i \leq B, \\ & 0 \leq x_i \leq \frac{p_i}{\alpha_i} \quad \text{for } i = 1, \dots, n. \end{aligned} \quad (19)$$

There is nothing intrinsically improper in the logarithmic objective function of the form appearing in Log-PRAP. But, the corresponding proportional rules are used without any documentation (or reference) that the logarithmic utility function is knowingly employed, let alone, with a justification for its use. In particular, the proportional solution is suboptimal with respect to (the more natural) optimization problems which we have formulated and solved in Sections 2 and 3.

Marginal analysis under proportional allocation rules is very simple – it implements the “cut/spend across the board” rule. Specifically, under this rule, proportional- (that is, percentage-wise-) cuts and increases in the available resource are passed uniformly to all budget-items. But, here again, such marginal analysis is optimal only under a very narrow class of objective functions – see, e.g. [17].

6. Demonstrating the three allocation models

To demonstrate the potential differences in the outcome of the PRAP, SRAP and Proportional allocation models,

we return to the US Senate example which was discussed in Section 1. Let us assume for sake of argument that the actual allocations in 2005 were proportional to C_i . Thus, if a_i denote the actual observed allocations, then we set $C_i = k \cdot a_i$ for some constant k . To simplify the demonstration, we assumed that $\alpha_i = \alpha$ for all i and set $p_i = 1$, where p_i is the baseline conditional probability of damage given an attack as defined in Section 2. We then made the optimistic assumption that the total damage averted by the observed allocations exactly offsets the total maximum damage, that is: $\sum_i C_i \cdot \alpha \cdot a_i = \sum_i C_i = k \cdot B$ where B is the total budget ($B = \sum_i a_i$). Substituting $k \cdot a_i$ for C_i on the left-hand side yields an estimate for α ⁸

$$\alpha = \frac{B}{\sum_i a_i^2}. \quad (20)$$

Finally, we took π_i (which is needed for solving PRAP) as proportional to population. Then, we solved for the allocations expected from the proportional (to population, not to C_i), PRAP and SRAP models using the linear damage functions $C_i \cdot (1 - \alpha \cdot x_i)$ (where x_i is determined from the models, as opposed to the actual observed allocations a_i which were used earlier to set the parameters). The outcome of this exercise is given in Fig. 3; note that our interest is in the resulting allocation patterns, as opposed to the specific damage assessments, which depend upon the scale parameter k that is set to unity in the figure.

The figure exhibits drastic differences in the allocation schemes that were generated by the three models. First, we observe that the PRAP solution (which focuses each time on the state that yields the largest marginal reduction in risk) would bring the expected damage of several states down to zero while leaving other states exposed to significant expected damage. It is interesting to note that even if one believes that the likelihood of attack is proportional to population (recall that the π_i we used were taken as proportional to the population), PRAP does not result in a proportional allocation of the budget! The proportional allocation leaves most mid-size states and one very large state (Texas) exposed to considerable expected damage. But, the most important observation is that the SRAP establishes a damage threshold that is significantly smaller than the maximal values of damage expected under the other models, and allocates the funds so that no state goes above that threshold. In a world of perfect information, it would be easy for the terrorists to identify and then attack the states with the largest expected damage if the PRAP or the proportional allocation models were implemented. In contrast, the SRAP solution would force the terrorists to choose one of the states that share the upper limit threshold of expected damage, knowing that an attack there would yield smaller expected damage than could be inflicted if one of the other two models was used.

⁸ One could more realistically assume that the observed allocations would only offset a fraction $f \leq 1$ of the total maximum damage; this would rescale α to equal $fB / \sum_i a_i^2$.

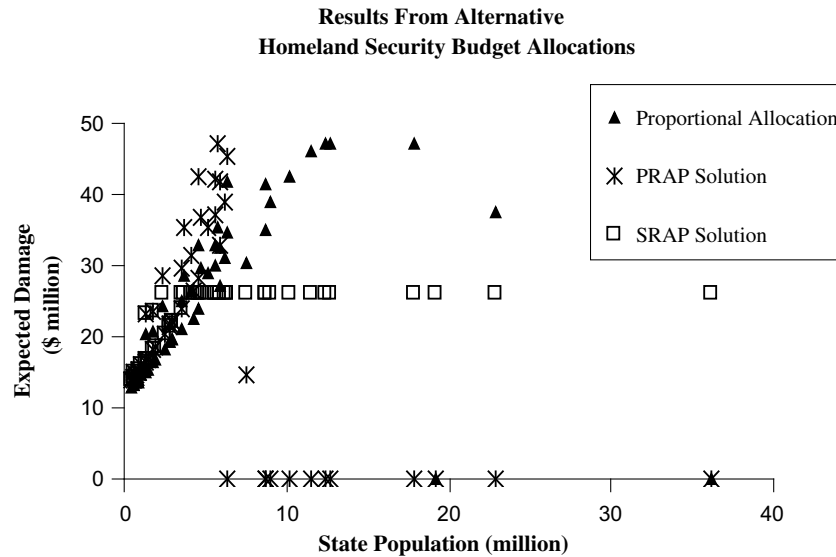


Fig. 3. Results from alternative homeland security budget allocations.

7. A real-world example

In this section, we describe a real-world example that illustrates the potential use of the models considered in Sections 2 and 3, while highlighting the different circumstances that apply to these models. The formulation of our example is a simplification of a more detailed model developed in [24]. Further analysis and model formulations of the real system has recently appeared in [2].

Following the attacks of September 11, 2001, there has been increased concern about the potential security-threat of visitors and immigrants who enter the US one of the actions taken by the Department of HS was the introduction of a multi-billion dollar US Visitor and Immigrant Status Indicator Technology (US-VISIT) Program. Among the actions adopted under this program, visitors and immigrants are fingerprinted when they enter the US and their fingerprints are compared and scored against a watchlist that contains several million sets of fingerprints that belong to known criminals and suspected terrorists. When the similarity between the fingerprints of a visitor and someone on the watchlist scores sufficiently high, the visitor is further investigated. Under the standard test, prints of the two index fingers are taken and analyzed. An additional test of 10 fingers is also available and is administered occasionally. The 10-finger test is more reliable, in particular when the quality of the two-finger prints is low. The alternative test is also more time-consuming (in execution and in the analysis of the results).

The main outcomes of the research reported in [24] were that while the average false-negative and false-positive probabilities that were obtained by the US-VISIT program were 4% and 0.3%, respectively, the actual chances of matching a person with the watch list changed dramatically with the quality of his or her fingerprints. Since terrorists could easily gather information about the performance of these tests, the

study concluded that the strategy that was implemented in the US-VISIT program was far less effective than the average detection level reported by the authorities.⁹

Inspired by [24], we discuss the use of potential policies of administering 10-finger tests versus two-finger tests. We start by characterizing the decision problem of administering the fingerprinting policy. Fingerprints are classified into eight categories of quality, labelled $1, \dots, 8$, with one denoting the best quality and eight denoting the worst. Data on the reliability of tests that compare fingerprints to those on the watchlist are available for each of the eight categories. Specifically, for each $i = 1, \dots, 8$, two numbers are available – p_i , which is the probability of a mistake in identifying an individual on the watchlist whose prints are of category i when the two-finger test is executed, and q_i , which is the corresponding probability under the 10-finger test (administered following the results of the two-finger test). Assuming that the fingerprint-quality labels are arranged in decreasing quality levels, we have that $p_1 \leq p_2 \leq \dots \leq p_8$ and $q_1 \leq q_2 \leq \dots \leq q_8$. Also, for each $i = 1, \dots, 8$, $p_i > q_i$.

Assuming that N visitors are handled daily; a capacity B of 10-finger tests per day; and π_i as the proportion of the population whose fingerprint-quality is i , we consider the decision problem of minimizing the probability of a mistake in identifying an individual on the watchlist. The decision variables for this problem are the numbers x_i of individuals in category $i = 1, \dots, 8$ who will be selected at random from among all those in category i to receive a 10-finger test. Thus, each individual with fingerprint-category i will be required to have a 10-finger test with proba-

⁹ Details can be found in [24] and references therein where considerable specific data are available. The research outcomes were also discussed by a US Congressional subcommittee on homeland security – see <http://faculty-gsb.stanford.edu/wein/Personal/Testimony%20of%20Lawrence%20Wein.htm>.

bility $\frac{x_i}{\pi_i N}$. As a result, the probability of a mistake in identifying an individual on the watchlist whose fingerprint-quality is i will be

$$N \left[\left(\frac{\pi_i N - x_i}{\pi_i N} \right) p_i + \left(\frac{x_i}{\pi_i N} \right) q_i \right].$$

The decision problem of minimizing such a mistake when conducted on a random population is then expressed as

$$\begin{aligned} \min_x \quad & \sum_{i=1}^8 \pi_i \left[\left(\frac{\pi_i N - x_i}{\pi_i N} \right) p_i + \left(\frac{x_i}{\pi_i N} \right) q_i \right] \\ & = \sum_{i=1}^8 \pi_i \left[p_i - x_i \left(\frac{p_i - q_i}{\pi_i N} \right) \right] \end{aligned} \quad (21)$$

$$\text{subject to: } \sum_{i=1}^8 x_i \leq B,$$

$$0 \leq x_i \leq \pi_i N \quad \text{for } i = 1, \dots, 8.$$

Evidently, (21) is a special instance of PRAP with $C_i = 1$ and $\alpha_i = \frac{p_i - q_i}{\pi_i N}$ for each $i = 1, \dots, 8$. Its optimal solution is then available from the results of Section 2; in particular, the ranking in (2) is based on $p_i - q_i$, that is, the absolute improvement in the performance of the 10-finger test over the two-finger test. The optimal policy is then to apply all 10-finger tests to the categories with the “best improvement”.

As it turns out (see [24]), the quality of one’s fingerprints can be deliberately reduced, for example, through the use of sandpaper, chemicals or even surgery. When the decision problem intends to protect against an insurgent who selects his/her fingerprint category, the problem of minimizing the probability of a mistake in identifying an individual on the watchlist has to be modified (in the spirit reported in Section 3) to

$$\begin{aligned} \min_x \quad & \max_{i=1, \dots, 8} \left[\left(\frac{\pi_i N - x_i}{\pi_i N} \right) p_i + \left(\frac{x_i}{\pi_i N} \right) q_i \right] \\ & = \max_{i=1, \dots, 8} \left[p_i - x_i \left(\frac{p_i - q_i}{\pi_i N} \right) \right] \end{aligned} \quad (22)$$

$$\text{subject to: } \sum_{i=1}^8 x_i \leq B,$$

$$0 \leq x_i \leq \pi_i N \quad \text{for } i = 1, \dots, 8.$$

Here, we observe that (22) is a special instance of SRAP, with $C_i = 1$ and $\alpha_i = \frac{p_i - q_i}{\pi_i N}$ for each $i = 1, \dots, 8$. Its optimal solution is available from the results of Section 3; in particular, the ranking in (7) is based on p_i – that is, the performance of the two-finger test. The corresponding optimal policy is then to define a performance-threshold and use the 10-finger tests to get all cumulative tests to or above that threshold.

8. Conclusions and directions for future research

This paper presents optimal resource allocation policies under probabilistic and strategic risks. The main contribu-

tion of the paper is the illustration of the difference between these scenarios through the resource-allocation optimization problems that they generate and the explicit solutions of these problems. Indeed, there is a stark difference between the corresponding optimal policies. When facing probabilistic risk, one has to focus on sites where the investment of the resource has the highest impact and target all the effort to these sites, in sequence, until the resource capacity is depleted. In contrast, when facing strategic risk, one starts by investing in the site which is most vulnerable and continues until its potential damage level equates that of the second most vulnerable site. At that point, one has to invest in parallel in both sites (in different proportions) until the potential damage level of the third most vulnerable site is reached, and continue in that manner.

From a managerial perspective, there are two fundamental differences between the optimal policies that were derived for the probabilistic and strategic risk scenarios. First, the criterion for assigning priority is different. Under probabilistic risk the criterion is effectiveness in reducing expected cost per unit investment of the resource. But, under strategic risk, the criterion is vulnerability. Second, we have a difference between targeting the effort under probabilistic risk and spreading the effort under strategic risk, the latter with the goal of reaching threshold damage levels. In an operational environment, spreading the effort can be achieved by randomization of individual decisions.

The insight that stems from the strategic risk scenario might have immediate consequences for the way decision-makers in charge of HS theaters allocate their resources. To adopt the policy we advocate here, decision-makers should be able to identify the “breakpoints” where the potential damage level of sites becomes equal and then, be able to authorize unequal investments that maintain some balance among the most vulnerable sites.

Understanding the intricacies of the proposed policies and the potential difficulties in implementing them is not an easy task. That is the reason why, in this paper, we considered only the linear versions of PRAP and SRAP. Future research will involve the more realistic nonlinear case where diminishing returns of scale are in effect (that is the impact caused by allocating the amount x_i of the resource to defend a site will be a concave function of x_i) as well as various scenarios of incomplete information.

Appendix

Consider the scalar θ^* and the vectors x^* and w^* defined, respectively, by (7)–(9), (5), (17). In this Appendix, we prove that x^* is an optimal solution for SRAP (and Rand-SRAP), that w^* is an optimal solution for Dual-SRAP and that θ^* is the joint optimal objective function of these optimization problems. Our proof verifies that $\bar{\theta} = \theta^*$, $\bar{x} = x^*$ and $\bar{w} = w^*$ satisfy (14) and (15); therefore, the desired conclusions follow from the results stated in Section 4 about Nash equilibria. We will further show that

under the non-degeneracy assumption (18) of Section 4, w^* is the only vector that satisfies (14) together with x^* , implying that in this case it is the only optimal solution of Dual SRAP.

We first consider the case where $\sum_{i=1}^n \frac{p_i}{\alpha_i} \leq B$. In this case, t^* (defined by (8)) equals $n + 1$, $\theta^* = 0$ and the vector x^* has $x_i^* = \frac{p_i}{\alpha_i}$ and $C_i(p_i - \alpha_i x_i^*) = 0$ for each $i = 1, \dots, n$. It then trivially follows that θ^* , x^* together with any solution $w \in \mathcal{W}$ satisfy (14),(15); in particular, this is the case for θ^* , x^* and w^* .

Next, assume that $\sum_{i=1}^n \frac{p_i}{\alpha_i} > B$. In this case, $t^* \leq n$, $\theta^* > 0$ and x^* depletes all of the resource, that is, $\sum_{i=1}^n x_i^* = B$. With J given by (16) and $J^c \equiv \{1, \dots, n\} \setminus J$, we then have that

$$B = \sum_{i=1}^n x_i^* = \sum_{i \in J^c} x_i^* = \sum_{i \in J^c} \frac{C_i p_i - \theta^*}{C_i \alpha_i}. \tag{23}$$

It follows from (5) that if $i \in J^c$, that is, $x_i^* > 0$, then $C_i(p_i - \alpha_i x_i^*) = \theta^*$. This implication combines with (17) to show that

$$\sum_{i=1}^n w_i^* C_i(p_i - \alpha_i x_i^*) = \sum_{i \in J^c} w_i^* C_i(p_i - \alpha_i x_i^*) = \sum_{i \in J^c} w_i^* \theta^* = \theta^*. \tag{24}$$

Next, let $x \in \mathcal{X}$ and $K \equiv \frac{1}{\sum_{j \in J} \frac{1}{C_j \alpha_j}}$. Combining (17), (23) and the fact that $\sum_{i \in J^c} x_i \leq B$, we conclude that

$$\begin{aligned} \sum_{i=1}^n w_i^* C_i(p_i - \alpha_i x_i) &= \sum_{i \in J^c} \frac{K}{C_i \alpha_i} C_i(p_i - \alpha_i x_i) \\ &= \sum_{i \in J^c} \frac{K p_i}{\alpha_i} - \sum_{i \in J^c} K x_i \\ &\geq \sum_{i \in J^c} \frac{K p_i}{\alpha_i} - K B \\ &= \sum_{i \in J^c} \frac{K p_i}{\alpha_i} - K \sum_{i \in J^c} \frac{C_i p_i - \theta^*}{C_i \alpha_i} \\ &= K \theta^* \left(\sum_{i \in J^c} \frac{1}{C_i \alpha_i} \right) = \theta^*. \end{aligned} \tag{25}$$

Finally, for $w \in \mathcal{W}$, (6) assures that

$$\sum_{i=1}^n w_i C_i(p_i - \alpha_i x_i^*) \leq \sum_{i=1}^n w_i \theta^* = \theta^*. \tag{26}$$

Evidently, (24)–(26) verify that θ^* , x^* and w^* satisfy (14) and (15).

We next verify that under the non-degeneracy assumption (18), w^* is the only vector that forms a Nash equilibrium together with x^* . We first observe that under the non-degeneracy assumption, 5,4,6 imply that

$$\begin{aligned} J &= \{i = 1, \dots, n : x_i^* = 0\} \\ &= \{i = 1, \dots, n : C_i(p_i - \alpha_i x_i^*) < \theta^*\}. \end{aligned} \tag{27}$$

Now, consider a vector \bar{w} that forms a Nash equilibrium together with x^* , that is, $\bar{x} = x^*$ and \bar{w} satisfy (14). We ob-

serve that for such \bar{w} , if $\bar{w}_i > 0$, then $i \notin J$ (for otherwise the attacker will gain from shifting attack probability from site i with $C_i(p_i - \alpha_i x_i^*) < \theta^*$ to a site s with $C_s(p_s - \alpha_s x_s^*) = \theta^*$). Further, $w_i^* C_i \alpha_i$ must be invariant for sites i that are not in J , that is, for sites having $x_i^* > 0$ (for if sites r and s are not in J and $w_r^* C_r \alpha_r < w_s^* C_s \alpha_s$, the defender will gain from shifting the resource from site r to site s where its use would be more effective). Thus, for some constant K , $w_i^* = \frac{K}{C_i \alpha_i}$ for $i \notin J$ and $w_i^* = 0$ for $i \in J$. As $\sum_{i \in J} w_i^* = 1$, $K \left(\sum_{j \in J} \frac{1}{C_j \alpha_j} \right) = 1$ and w^* must satisfy (17).

References

- [1] G.F. Apostolakis, D.M. Lemon, A screening methodology for identification and ranking of infrastructure vulnerabilities due to terrorism, *Risk Analysis* 25 (2005) 361–376.
- [2] M. Baveja, L.M. Wein, An effective two-finger, two-stage biometric strategy for the US-VISIT program, Working paper, Graduate School of Business, Stanford University, 2006.
- [3] V.M. Bier, Should the model for security be game theory rather than reliability theory, *Communications of the Fourth International Conference on Mathematical Methods in Reliability: Methodology and Practice*, Santa Fe, New Mexico, 2004.
- [4] G. Brown, M. Carlyle, J. Salmeron, K. Wood, Defending critical infrastructure, *Interfaces* 36 (6) (2006) 530–544.
- [5] A. Brunet, Grant funding to state and local governments and systematic assessment of vulnerability. CREATE Report 05-018. Center for Risk and Economic Analysis of Terrorism Events, University of Southern California, Los Angeles, California, 2005.
- [6] P.J. Coughlin, Pure strategy equilibria in a class of systems defence games, *International Journal of Game Theory* 20 (1992) 195–210.
- [7] Department of Homeland Security, FY2006 Homeland Security Grant Program, Program Guidelines and Application Kit, December 2005.
- [8] B. Harris, Mathematical methods in combatting terrorism, *Risk Analysis* 24 (2004) 985–991.
- [9] E.H. Kaplan, H. Pollack, Allocating HIV prevention resources, *Socio-Economic Planning Sciences* 32 (1998) 257–263.
- [10] E. Kardes, R. Halls, Survey of the literature on strategic decision making in the presence of adversaries, unpublished manuscript, 2005.
- [11] T. Kuno, H. Konno, E. Zemel, A linear-time algorithm for solving continuous maximum knapsack problems, *Operations Research Letters* 10 (1991) 23–26.
- [12] M. Leung, J.H. Lambert, A. Mosenthal, A risk-based approach to setting priorities in protecting bridges against terrorist attacks, *Risk Analysis* 24 (2004) 963–984.
- [13] National Commission on Terrorist Attacks, The 9/11 Commission Report, Norton, New York, 2004.
- [14] M. Osborne, A. Rubinstein, *A Course in Game Theory*, MIT Press, 1994.
- [15] S. Reese, Risk-Based Funding in Holmeland Security Grant Legislation: Analysis of Issues for the 109th Congress, CRS Report (Order Code: RL33050), 2005.
- [16] T. Rockafellar, *Convex Analysis*, Princeton University Press, 1970.
- [17] M.J. Rosenblatt, U.G. Rothblum, Optimality of cut/expand across the board’ rule for constrained optimization problems with an application to an inventory model, *IIE Transactions* 26 (1994) 102–108.
- [18] T. Sandler, D.G. Arce, Terrorism and game theory, *Simulation and Gaming* 34 (2003) 319–337.
- [19] D. Sarewitz, R. Pielke Jr., M. Keykhan, Vulnerability and risk: some thoughts from a political and policy perspective, *Risk Analysis* 23 (2003) 805–810.
- [20] M. Shubik, *Game Theory in the Social Sciences*, MIT Press, Cambridge, MA, 1982.

- [21] M. Shubik, R. Weber, Competitive valuation of cooperative games, Discussion Paper No. 482, Cowles Foundation For Research in Economics, Yale University, 1978.
- [22] M. Shubik, R. Weber, Systems defense games: Colonel Blotto, command and control, *Naval Research Logistics Quarterly* 28 (1981) 281–287.
- [23] J. Von Neumann, O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, 1944.
- [24] L.M. Wein, M. Baveja, Using fingerprint image quality to improve the identification performance of the U.S. Visitor and Immigrant Status Indication Technology Program, *Proceedings of the National Academy of Sciences* 1020 (21) (2005) 7772–7775.