A bottom-up risk-based resource allocation methodology to counter terrorism

Luca Quadrifoglio

Zachry Department of Civil Engineering,
Texas A&M University,
College Station, TX 77843-3136, USA
Fax: +1-979-845-6481
E-mail: lquadrifoglio@civil.tamu.edu

Abstract: After September 11th, government officials have begun questioning the appropriateness of the current federal allocation formula to counter terrorism. Newly appointed Secretary Chertoff demanded the development of more rigorous risk-based mechanisms. This paper is one of the first addressing this request by developing the corresponding resource allocation optimisation model based on the current largely accepted rigorous definition of risk and suggesting a practical implementation of a decomposition methodology to solve it. It is shown how the proposed procedure works and can be carried out in practice within the current geographical hierarchical partitioning of the nation in a 'bottom-up' fashion: 'child' subsets provide their 'parents' with piecewise linear relationships indicating the cost-effectiveness of the allocated budget in reducing the overall risk. This information will be sufficient to solve the overall allocation problem. A successful implementation of the methodology would result in reducing the 'barrier' between theory and application, performing a better resource allocation to counter terrorism, ultimately using fewer resources more efficiently to make the nation safer.

Keywords: terrorism; resource allocation; decomposition; risk; game theory.

Reference to this paper should be made as follows: Quadrifoglio L. (2008) 'A bottom-up risk-based resource allocation methodology to counter terrorism', *Int. J. Society Systems Science*, Vol. 1, No. 1, pp.4–25.

Biographical notes: Dr. Luca Quadrifoglio earned his PhD (2005) from the Epstein Department of Industrial and Systems Engineering at USC (Los Angeles, CA) and worked as a Postdoctoral Research Associate for the Center for Risk and Economic Analysis of Terrorism Events (CREATE). In 2006, he joined the Faculty of the Zachry Department of Civil Engineering at Texas A&M University in College Station, TX. Dr. Quadrifoglio won the 2004 CUTC National Best Paper Award and the 2006 Pritsker Doctoral Dissertation Award (3rd place). His research interests are within the broad area of operations research and performance/risk/decision analysis applied to a variety of fields.

1 Introduction

1.1 Background

The events on September 11th have been a wake up call for the US. Besides realising how much damage the terrorist organisations are capable of causing even within the US territories, government officials have begun questioning whether the current allocation formula of federal funds to prevent, respond and recover from terrorist attacks is appropriate. While the anti-terrorism federal grant mechanism has gone through several changes in past ten years and is still rather complicated, the distribution has been more the result of needed political compromises rather than the outcome of a proper and rational allocation methodology aiming to provide the funds where they are needed.

From 2001 to 2005 the largest fraction of funds was distributed according to a two-part formula (2001 Patriot Act), using a base amount of 0.75% of the total allocation for each state and 0.25% of the total allocation for each US Territory with the balance of funds being distributed on a population-share basis. The result has been an allocation of 40% of the total evenly divided among the States, with the remaining 60% distributed to states based on share of population (the '40/60' formula). State funding has been then sub-granted to local governments.

Things have changed to some extent in 2006, after the newly appointed Secretary of Homeland Security Michael Chertoff officially requested that a more rigorous 'risk-based' allocation methodology be developed ("We have to put the resources where the highest threats are"). The US has been divided in Urban Areas (UAs) not necessarily corresponding to the states, but of comparable size. While the 'fixed' 40% portion is still being allocated in the same way, the 'variable' 60% portion will be distributed across the US based on risk, taking in consideration several other factors in addition to population. However, the new methodology is still in its infancy and further details are currently classified.

1.2 Current allocation critiques

While it is clear that the federal government is making an effort to increase the effectiveness of the funding mechanism, there are still several issues that need further clarification: how is the current '40/60' formula justified? Is the 'fixed' 40% portion reasonable or even needed? Some UA could have a very low 'risk score' to begin with not justifying at all the fix amount of their share of funds, which would be much more effective if allocated somewhere else. Experts in the field have reached different and contradicting conclusions, but they all seem to agree that the '40/60' formula used now still disconnects the funding from a requested risk-based approach. As noted by Hall (2003):

> "Widespread criticism continues among stakeholders regarding how the federal anti-terrorism funding is being distributed to states…Critics argue that the funding allocation fails to take into account the heightened needs of some areas of the Country, while it provides funds to other states, cities and localities that have a low-level of need for such funding".

Ransdell (2004), while comparing the DHS funding vs. other federal programs, noted that "A small-state minimum of 0.75% is unusually large". He also added that:

> "Most formula grant minimum percentages are applied after the administering agency has already made an initial allocation of funds, and only if needed…In contrast, DHS begins by allocating to each state the minimum amount, after which it distributes remaining funds, including those that already received considerably more than their share because of the small-state minimum".

De Rugy (2005) says: "the underlying theory behind this all-state-minimum formula" – the 40% portion – "is that terrorists could strike anywhere, but the theory that money should be spent smoothly across states has not been supported by reasoned analysis during the public policy debate". In the Congressional Quarterly (2004), Cox points out: "It is not the case that American Samoa should receive proportionately less or should receive more or less than anywhere else, except if security needs require it". Of course, smaller and at-lower-risk states pushes to maintain the status quo and bigger at-higher-risk states do the opposite. Furthermore, the 9/11 Commission Report (2004) recommended that "Federal homeland security assistance be distributed to state and local government based on risk and vulnerability only". It also added that "Federal homeland security assistance should not remain a program for general revenue sharing. It should supplement state and local resources based on the risks or vulnerabilities that merit additional support". Finally, Secretary Chertoff explains in a statement before the House Appropriations Committee's Homeland Security Subcommittee (Congressional Quarterly, 2005): "I want to emphasise that our analysis of threats and risks posed to the USA by terrorists will drive the structure, operations, policies and missions of the department, and not the other way around". In other words, grant money should only be distributed based on an evaluation of risk and security need and nothing else.

1.3 Risk definition and assessment

There are still ongoing debates on how to clearly and uniquely define risk. For a comprehensive discussion on the general definition of risk the reader might refer to the work of Kaplan and Garrick (1981). In the contest of terrorism, Willis et al. (2005) proposed an expected value definition that seems to be today the most widely accepted and adopted by the federal government and Secretary Michael Chertoff. Terrorism risk is viewed as the product of three main components:

1 the threat to a target

2 the target's vulnerability to the threat

3 the consequences should the target be successfully attacked.

This definition is easy to grasp, but it may not always be the most accurate, primarily because it ignores the distribution tales and variances of the variables and most importantly overlooks risk tolerance issues, which can be crucial in the decision making when dealing with catastrophic events (such as nuclear blasts).

No matter how risk is defined, a comprehensive risk assessment of all the US assets is needed for any rigorous methodology and this is not an easy task. Thus, easily measurable indicators, such as the current population, population density, location quotients, etc. have been used by other models as proxy variables for risk to drive the allocation, but are not necessarily correct. An allocation proportional to the actual risk distribution (as proposed by Willis et al., 2005) would definitely be a better choice;

however, it would still not necessarily be accurate, since the largest risk reductions can be achieved not necessarily where the risk is higher.

1.4 Research objectives

The large uncertainty related to risk assessments and the needed benefit/cost analyses of alternative countermeasures would lead to developing complex decision trees models, such as, for instance, the analytic hierarchy process proposed by Saaty (1986, 1994) and Saaty and Vargas (2001) or the entropy method approach in Zeleny (1982) and Hwang and Yoon (1981). However, as noted by Triantaphyllou (2000), the perfect decision making framework may never be found. Everything is further complicated by game-theoretic issues; see Bier (2005) and Bier et al. (2005) for related work. In fact, terrorists are intelligent agents that will strategically and dynamically act in response to our defences to maximise their objective that is to inflict the largest amount of damage to the western world and the US in particular. They can basically strike anywhere at anytime and most likely where and when they are less expected to do so. Thus, the complexity and dimension of the problem and the intrinsic uncertainty of the data itself do not seem to justify the development and implementation of complex models, which would not be able to eliminate the underlying stochastic nature. The development of a simpler model, with its practicality, adaptability and ease of use would instead be more appropriate and also more likely to encourage decision makers to take advantage of the tool, since they would be reluctant to utilise any overly complicated even if more accurate model, especially because the allocation procedure needs to be done yearly in the best possible and practical manner.

The ultimate goal of this research project is to perform a more efficient resource allocation, reducing the risk for terrorism and making the nation safer. Two recent works address the same issue by proposing alternative ways to allocate funds to counter terrorism with respect to the current federal practice. Weinlein (2004) proposes to apply the analytical hierarchical process for the allocation. Brunet (2005) strongly critiques the fixed 40% portion of the allocation, proposing a 100% allocation based on risk; the author develops a non-linear optimisation model and proposes to solve it globally. In this paper, we attempt to reduce the barrier between theory and applications by formulating the resource allocation optimisation model with a rigorous risk-based approach to counter terrorism and by proposing a decomposition procedure to solve it, embedding and integrating this methodological approach with a practical procedure in a bottom-up fashion, as a possible alternative to the current federal allocation formula. The proposed procedure could be the base for solving other similar allocation problems, such as preventing and responding to natural disasters, broadening the contribution of the paper for the research community and eventually for society in general. In Section 2, we describe the modelling framework. In Section 3, we discuss the proposed bottom-up decomposition technique. In Section 4, we suggest two possible decomposition alternatives. Section 5 outlines an example and finally the conclusions are presented in Section 6.

2 Modelling framework

In this section, we review and discuss the ideal rigorous methodology needed to build the optimisation model in the context of providing the best resource allocation to minimise the overall terrorism risk.

2.1 Definition of risk

2.1.1 Scenarios

A rigorous development of any methodology to counter terrorism will need to begin by enumerating all the targets and the means of attack that might be used against them, to identify the possible terrorist attack scenarios in the area considered (the world, the US, a state, etc.). This is a very complex task itself, but it is also crucial. In fact, as the events of September 11th 2001 clearly showed, the creativity of terrorists often eludes any common rational prediction. Forgetting or disregarding possible scenarios would unavoidably increase the exposure to future attacks, especially if terrorists suspect the oversight.

Formally, we have the set of all possible terrorist targets $t \in \mathrm{T}$ (i.e.: a stadium, a building, a plant, a city, an aircraft, a computer network, etc.) and the set of all possible means of attack $a \in \mathrm{A}$ (dirty bomb, manpads, nuclear device, cyber attack, etc.). A scenario $s(a, t) \in \mathrm{S}$ (manpads attack on a commercial aircraft, dirty bomb in downtown Los Angeles, etc.) is the combination of an attack *a* to a target *t*. It is possible to represent the collection of scenarios in a matrix form: the columns and the rows being respectively the set of targets and the set of means of attack (or vice versa). Each cell of the matrix identifies a possible scenario. This two-dimensional 'check list' representation would help to exhaustively take into account all scenarios and to avoid ignoring feasible but not immediately obvious ones. Of course, several cells of the matrix would correspond to unrealistic scenarios, because they would be the result of improbable matches between an *a* and a *t* (such as, for example, a conventional bomb attack to a nuclear plant, which would clearly be completely ineffective in producing any significant damage).

2.1.2 Threat

For each scenario $s(a, t)$, the *threat* is defined as the probability p_s that an attack *a* will be attempted on *t* in a given time interval (i.e.: a year). That is:

$$p_s = \mathrm{P}[a \text{ is attempted on } t]$$

For example, the annual probability that a city's football stadium will be subject to an attack with a radiological weapon is the threat of this specific scenario.

These probabilities, which can also be thought as the combination of the terrorists' intent and capability to attempt an attack *a* on target *t*, are by far the toughest to estimate and they might significantly vary depending on the defensive plan; in fact, terrorists will strategically and dynamically act in response to our defences seeking weaker targets to achieve their goals (game theory).

2.1.3 Vulnerability

For each scenario $s(a, t)$, the *vulnerability* is defined as the probability v_s that an attack a will be successful on t given that the attack is attempted. That is:

$$v_s = \mathrm{P}[a \text{ is successful on } t|\ a \text{ is attempted on } t]$$

For example, the probability that a radiological attack would be successful, causing some sort of damage, on a city's football stadium (given that the attack has been attempted), is the vulnerability of this specific scenario.

Even though estimating vulnerabilities is certainly not an easy task, they will not be affected by game-theoretic issues and they can, therefore, be evaluated with greater confidence than the threats.

2.1.4 Consequences

The consequences of a specific scenario are the type and the magnitude of damage resulting from a successful terrorist attack. The definition adopted today by the federal government is a weighed sum of four different terms: economic losses (including immediate and future), human health (deaths and injuries), strategic and psychological impact. We identify this weighed multi-term consequence function with K. Defining, agreeing upon and estimating these terms and the weights are very difficult tasks themselves. A proper assignment of the weights would allow combining these unrelated terms in a unique category measurable by a single scale (similar to the Richter scale adopted for earthquakes, for instance). See Keeney (1992), Keeney and Raiffa (1993), Chankong and Haimes (1983) and Hammond et al. (1999) for a comprehensive review over multi-objective models. Each successful terrorist attack scenario $s(a\ ,t)$ has therefore an associated expected consequence K_s. Formally:

$$\mathrm{K}_s = E[\mathrm{K}|\ a \text{ is successful on } t]$$

For example, the consequences of the scenario identifying a successful radiological attack on a city's football stadium would be an estimate of the number of fatalities, the economic losses and the strategic and psychological impact on the nation.

2.1.5 Risk

Finally, the risk suffered by a specific terrorism scenario is the product of the three components mentioned above: threat, vulnerability and consequences. In this way, risk for each scenario is ultimately representing the expected annual consequences and is defined as $r_s = p_s v_s \mathrm{K}_s = l_s \mathrm{K}_s$. And, as shown by Willis (2006), it can be represented as the intersection of the three terms (see Figure 1).

Figure 1 Risk is the intersection of threat, vulnerability and consequences

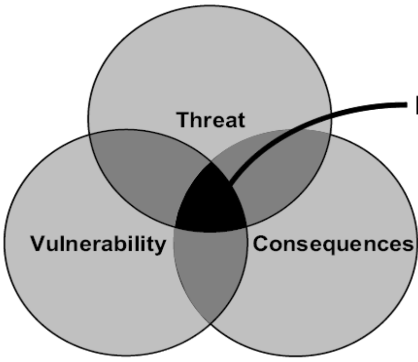

Formally:

r_s = P[a is attempted on t]×P[a is successful on t| a is attempted on t]×E[K| a is successful on t]

The risk associated with a specific area (i.e.: the world, the US, a state, a jurisdiction, a city, etc.) is $\sum_{s \in S} r_s$, with all the possible scenarios affecting the area included in S.

2.1.6 Risk tolerance

The above 'expected value' definition of risk is simplistic and may not be the best; in fact, it does not consider the distribution tales and the variances of the variables, ignoring risk tolerance issues, which are crucial especially when dealing with catastrophic events (Haimes, 2004). This is emphasised by the fact that there is a high degree of uncertainty, because of the very difficult assessments to be made. However, the use of utility functions or proper adjustments to the definition of the consequences would allow incorporating risk tolerance issues by paying special attention to devastating events, endorsing the above developed expected value model. Nevertheless, future developments should carefully consider looking at alternative definitions of risk.

2.2 Countermeasures

The allocation procedure will require decision makers to choose from a set I (i = 1,…, N) of countermeasures available in the area considered. Some countermeasures are exclusively dedicated to have an effect on a specific scenario (such as: installing Anthrax detectors in specific buildings). Some countermeasures may affect a specific target only, even though they could be beneficial for more than one scenario associated with the target (i.e.: building a fence around a plant could prevent several types of attack, affecting more than one scenario). Others affect more than one target and several scenarios (e.g.: increasing the number of policemen throughout the State of California).

A countermeasure i can be viewed as a Y/N decision (e.g.: a fence could be built for a certain cost to protect a building; it would not make sense to build only a part of it). In this case, i is associated with a binary variable y_i.

A countermeasure i can be also viewed as having different levels of effectiveness, depending on how much is invested in it (e.g.: a police force can be increased to protect a certain area; this countermeasure's effectiveness can be considered proportional to the amount of additional policemen needed). In this case, i is associated with a variable y_i (integer or continuous), indicating how much has been spent for it.

Other countermeasures can be more complicated to be represented mathematically, depending also on the level of details involved. However, they can generally be modelled using combinations of binary/integer/continuous variables (most of the time with acceptable approximations).

2.2.1 Additional logic constraints

In order to better represent the reality and the logic interdependences between different countermeasures, other constraints might and/or should be considered.

For example, contingency relationships stating that pursuing a countermeasure i_1 would be possible only if countermeasure i_2 has already been implemented (e.g.: it is not reasonable to buy an antivirus for a security computer network, if the security computer network system has not yet been purchased) are mathematically represented by $y_1 \le y_2$ (assuming y_1 and y_2 binary). In this way, y_1 can not be 1 if y_2 is 0.

Mutually exclusive relationships between countermeasures might also be needed. If only one countermeasure between i_1 and i_2 could be pursued (e.g.: when having to choose among two models of a particular security system) is mathematically represented by $y_1 + y_2 \le 1$. In this way, either y_1 or y_2 could be 1 (or none), but not simultaneously both.

Other similar constraints can be added to represent even more complicating relationships involving contingency, mutual exclusiveness or other logical interdependences among two or more countermeasures. Also, countermeasures represented by continuous or integer variables can be included in the logic constraints by employing a few modelling 'tricks'.

2.3 Overall cost

Each countermeasure i reduces the risk r_s of one or more scenarios s, either by reducing their threat p_s (prevention), their vulnerability v_s (protection) and/or their consequence K_s (mitigation). Let the vector $\mathbf{y} = y_1, y_2, \ldots, y_N$ represents the collection of variables associated with each countermeasure $i \in I$. The relationship between $\mathbf{y}$ and p_s, v_s or K_s can be of any nature (linear or non-linear) and may include synergistic or antagonistic effects involving several y_i. In fact, a countermeasure can be more or less effective, depending on whether other countermeasures are simultaneously adopted. Also, some scenario could have its threat p_s increased due to countermeasures applied to some other scenarios (game theory). Thus, a considerable amount of assessment is needed. We can express the risk r_s associated with each scenario s as a function of $\mathbf{y}$, since:

$$\left.\begin{aligned} p_s &= p_s(\mathbf{y}) \\ v_s &= v_s(\mathbf{y}) \\ K_s &= K_s(\mathbf{y}) \end{aligned}\right\} \qquad r_s(\mathbf{y}) = p_s(\mathbf{y})\, v_s(\mathbf{y})\, K_s(\mathbf{y}) \qquad \forall s \in S$$

Each countermeasure i has an associated estimated cost c_i (for proportional countermeasures associated with continuous or integer variables, c_i represents the maximum cost possible). Thus, the overall cost function C can be defined as:

$$\mathrm{C}(\mathbf{y}) = \sum_{s \in S} r_s(\mathbf{y}) + w \sum_{i \in I} c_i y_i$$

where w is the parameter needed to properly combine the two terms.

2.4 Optimisation model

The optimal allocation of available resources should aim to minimise the overall cost C in a selected area. The non-linear mixed integer programming (MIP) problem associated with the resources allocation would therefore be the following model (0):

$$\begin{array}{ll} \text{Min} & \mathrm{C}(\mathbf{y}) \\ \text{s.t.} & \sum_{i \in I} c_i y_i \le \mathrm{B} \\ & \left.\begin{array}{l} \text{[possible logic constraints among the } y_i \text{ (Section 2.2)]} \\ y_i \text{binary / integer, whenever needed} \\ \text{upper and / or lower bounds on } y_i, \text{ where appropriate} \end{array}\right\} \forall i \in \mathrm{I} \end{array}$$

where B is the total budget available to decision makers for the allocation. This is a nonlinear Knapsack problem with additional constraints, which is NP-Hard. Provided that the ideal approach would be solving the whole model to optimality, see Bretthauer and Shetty (1995, 2002) for a comprehensive review of the latest techniques, this would be impractical because of its complexity and dimension. Furthermore, the time required to perform an accurate centralised data collection would most likely be more than a year, which is the time span allotted by the federal government to make a decision. In addition, the process of collecting, gathering, sharing and conveying such highly detailed information in a centralised manner about possible terrorist targets could also create additional vulnerabilities itself.

Therefore, rather than looking for an exact optimal solution for a complex problem which is intrinsically inaccurate, because of the unavoidable uncertainty of risk assessments, we aim to solve it with reasonable approximations by breaking it down into sub-problems, which would be easier and more practical to tackle and faster to solve, using a decomposition approach, as illustrated in the next section.

3 Model decomposition

In this section, we review the optimisation model decomposition, especially highlighting the needed assumptions and tradeoffs required to carry it out within our application context.

Let's express the set of all the scenarios S, the set of all countermeasures I and the corresponding vector **y** as unions of j = 1 to n subsets, the overall cost C as the summation of the costs of each subset and the budget B as the summation of the budgets to be allocated among the countermeasures within each subset:

$$\mathrm{S} = \mathrm{S}_1 \cup \mathrm{S}_2 \cup \ldots \cup \mathrm{S}_\mathrm{n}$$
$$\mathrm{I} = \mathrm{I}_1 \cup \mathrm{I}_2 \cup \ldots \cup \mathrm{I}_\mathrm{n}$$
$$\mathbf{y} = \mathbf{y}_1 \cup \mathbf{y}_2 \cup \ldots \cup \mathbf{y}_\mathrm{n}$$
$$\mathrm{C} = \mathrm{C}_1 + \mathrm{C}_2 + \ldots + \mathrm{C}_\mathrm{n}$$
$$\mathrm{B} = \mathrm{B}_1 + \mathrm{B}_2 + \ldots + \mathrm{B}_\mathrm{n}$$

Let's also consider the following assumptions:

- *Assumption 1*: each countermeasure $i \in \mathrm{I}_\mathrm{j}$ and its cost c_i 'belong' only to the budget B_j of subset j. Basically, the cost to implement (totally or partially) countermeasure $i \in \mathrm{I}_\mathrm{j}$ is covered by B_j and not by any other budget $\mathrm{B}_{\mathrm{k} \neq \mathrm{j}}$.
- *Assumption 2*: there are no logic constraints (defined in Section 2.2) linking together countermeasures which belong to different subsets.
- *Assumption 3*: the risks r_s of the scenarios s included in each subset S_j depend only on the countermeasures in I_j and on the $\mathbf{y}_\mathrm{j}$ (not on the whole $\mathbf{y}$) and so does the corresponding overall cost $\mathrm{C}_\mathrm{j}(\mathbf{y}_\mathrm{j})$. This assumption implies game-theoretic independence between subsets, saying that whether or not countermeasures belonging to subsets $\mathrm{S}_{\mathrm{k} \neq \mathrm{j}}$ are implemented does not have any effect on the risk level and the overall cost C_j of subset S_j.

A careful and appropriate selection of the subsets, which would reasonably satisfy the above Assumptions 1, 2 and 3, would guarantee their *independence* from each other and the original model (0) can be rewritten as the following model (1):

$$\begin{array}{llllll}
\text{Min} & \mathrm{C}_1(\mathbf{y}_1) & + \mathrm{C}_2(\mathbf{y}_2) & + \ldots & + \mathrm{C}_\mathrm{n}(\mathbf{y}_\mathbf{n}) & \\
\text{s.t.} & \sum_{i \in I_1} c_i y_i \le \mathrm{B}_1 & & & & \\
 & & \sum_{i \in I_2} c_i y_i \le \mathrm{B}_2 & & & \\
 & & & \ldots & & \\
 & & & & \sum_{i \in I_n} c_i y_i \le \mathrm{B}_\mathrm{n} & \\
 & \mathrm{B}_1 & + \mathrm{B}_2 & + \ldots & + \mathrm{B}_\mathrm{n} & = \mathrm{B}
\end{array}$$

$$\left.\begin{array}{l}
\text{[possible logic constraints among the } y_i \text{ (Section 2.2)]} \\
y_i \text{ binary / integer, whenever needed} \\
\text{upper and / or lower bounds on } y_i\text{, where appropriate}
\end{array}\right\} \forall i \in \mathrm{I}_1$$

$$\left.\begin{array}{l}
\text{[possible logic constraints among the } y_i \text{ (Section 2.2)]} \\
y_i \text{ binary / integer, whenever needed} \\
\text{upper and / or lower bounds on } y_i\text{, where appropriate}
\end{array}\right\} \forall i \in \mathrm{I}_2$$

$$\left.\begin{array}{l}\text{[possible logic constraints among the } y_i \text{ (Section 2.2)]} \\ y_i \text{ binary / integer, whenever needed} \\ \text{upper and / or lower bounds on } y_i\text{, where appropriate}\end{array}\right\} \forall i \in \mathrm{I_n}$$

where the B_j are variables and not parameters of the problem and are linked together by the total budget constraint ($B_1 + B_2 + \ldots + B_n = B$). This model can be separated in n 'sub-models'; each of them can be represented by:

$$\begin{array}{ll}\text{Min} & C_j(\mathbf{y}_j) \\ \text{s.t.} & \sum_{i \in I_j} c_i y_i \le B_j\end{array}$$

$$\left.\begin{array}{l}\text{[possible logic constraints among the } y_i \text{ (Section 2.2)]} \\ y_i \text{ binary / integer, whenever needed} \\ \text{upper and / or lower bounds on } y_i\text{, where appropriate}\end{array}\right\} \forall i \in \mathrm{I}_j$$

Each sub-model has exactly the same structure of the overall original model and represents the problem of allocating B_j among the countermeasures in subset I_j within the scenarios S_j in subset j.

Let $C_j(B_j)$ be the relationships indicating how effective the funds B_j are in reducing the corresponding overall cost C_j. The $C_j(B_j)$ are monotonically decreasing with B_j. This is intuitive and easily proven mathematically, since progressively loosening a constraint (by increasing B_j) could never worsen the optimal solution, which would improve or remain the same in the worst case. Deriving each $C_j(B_j)$ would require solving each sub-model for all possible values of B_j and would be impractical. However, each sub-model can be solved for some values of B_j to construct a piecewise linear function $C_j^l(B_j)$, which would approximate the actual relationship (see Figure 2, B_j^{max} represents the maximum possible value that could be allocated to S_j).

Figure 2 Example of $C_j^l(B_j)$ as an approximation of $C_j(B_j)$ (see online version for colours)

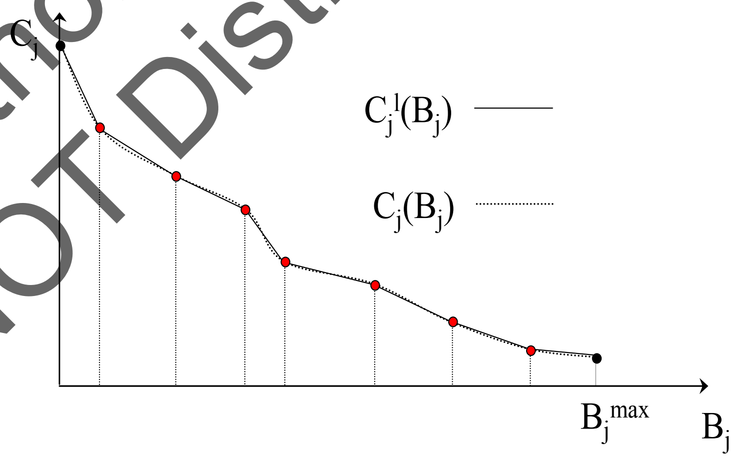

Clearly, the more values of B_j are used, the more accurate $C_j^l(B_j)$ would be in approximating $C_j(B_j)$. Thus, the overall model (1) could be approximated by the following model (2):

$$\begin{aligned} &\text{Min} \quad C_1^l(B_1) + C_2^l(B_2) + \dots + C_n^l(B_n) \\ &\text{s.t.} \quad B_1 + B_2 + \dots + B_n = B \end{aligned}$$

Model (2) is a piecewise linear Knapsack problem with much less variables than model (1) and could be solved fairly quickly by the latest optimisation techniques.

Note that for a given problem with N integer variable (such as the Y/N decision variables for the countermeasures), the number of feasible integer solutions, i.e. the search space, is proportional to 2^N. If decomposed in k sub-problems and assuming that each of them has N/k integer variables, the search space for each of the decomposed problems is proportional to $2^{N/k}$. Although these smaller problems have to be solved multiple times to construct the piecewise linear functions, the reduction in size can help reduce the overall computing time significantly as we show in our experimental Section 5.

3.1 Convexity

Furthermore, with the additional assumption of convexity of the $C_j^l(B_j)$, model (1) can be restructured as a linear programming problem and solved very easily. Convexity can not be guaranteed due to the integrality of some of the y_i variables (representing the Y/N countermeasures) and possible contingency constraints, which would create local concave corners along the curve. However, a general convex profile of the relationship between C_j and B_j is intuitive, because the cost-effectiveness of the allocation is expected to progressively decrease with increasing B_j, and could be reasonable assumed for $C_j^l(B_j)$ if the intervals between consecutive values of B_j are sufficiently larger than the c_i, so that the small concave corners would 'disappear' (see example in Figure 3).

Figure 3 Example of convex Cjl(Bj) (see online version for colours)

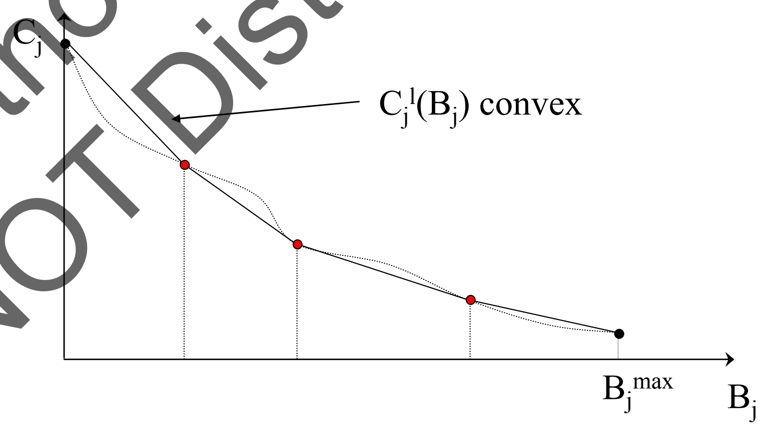

So while smaller intervals between consecutive values of B_j would guarantee a better adherence of the piecewise linear functions to the actual $C_j(B_j)$, larger intervals would instead lead to convexity which would allow the model (2) to be solved even faster. It would be best to have the smallest possible intervals as long as convexity holds and a suggested practical approach on how to construct the piecewise linear functions would be:

- solve the model for $B_j = 0$ and $B_j = B_j^{max}$ to obtain the extreme points of $C_j^l(B_j)$
- solve for intermediate points in the middle of existing points: first for $B_j = B_j^{max}/2$, then for $B_j = B_j^{max}/4$ and $B_j = 3B_j^{max}/4$, etc., while convexity of $C_j^l(B_j)$ holds or a maximum predetermined number of interval is reached.

3.2 Multiple layers

To obtain each corner of the piecewise linear functions $C_j^l(B_j)$ would require solving the sub-model for each selected value of B_j and can still be very hard or too long to obtain, because of the complexity and nonlinearity of the sub-model itself. However, we could obtain an approximate solution using the same decomposition procedure within each sub-model to create several 'sub-sub-models', which could be easier or faster to solve. This decomposition procedure may be used indefinitely, as long as the needed Assumptions 1, 2 and 3 of independency are acceptable. See Figure 4 as an example: the entire set of scenarios S is divided in n subsets; the subset S_j is then divided in several sub-subsets and the sub-subset S_{jh} is also further divided.

Figure 4 Decomposition in several layers

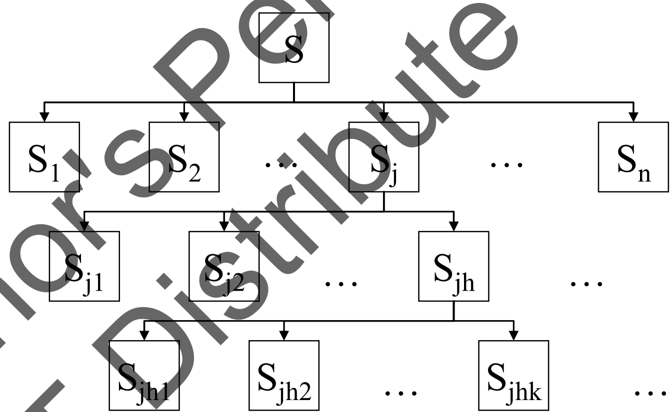

The original comprehensive model can, therefore, be ultimately solved in a bottom-up fashion, constructing the piecewise linear functions of the sub-problems at the bottom leaves of the diagram (S_{jh1}, S_{jh2},…, S_{jhk}, etc.), which would allow to solve their 'parent' set S_{jh} for different values of the budget B_{jh} and thus construct the piecewise linear function $C_{jh}^l(B_{jh})$ along with $C_{j1}^l(B_{j1})$, $C_{j2}^l(B_{j2})$, etc. Similarly, we can find all the $C_j^l(B_j)$ and ultimately solve for the entire set S.

4 Subsets selection

In this section, we illustrate how the above rigorous decomposition methodology could be carried out in practice and what are the pros and cons in selecting the subsets in one way or another.

4.1 Geographical decomposition

Until 2005, the federal government allocated the funds following a multiple-layer distribution model reflecting the organisational hierarchy of the nation, which naturally mirrors a geographical partition of the US: funds were first distributed across states, which then distributed their share across counties and then across cities and local jurisdiction. The share of funds allocated at each step was roughly proportional to the population of the recipient. In 2006, states have been replaced by UAs, which do not necessarily coincide with them, but are of comparable size. However, the distribution is performed in a similar fashion.

Thus, the practical application of the decomposition procedure explained above could naturally be applied following the geographical separation of the US, where the subsets in the first layers are the UAs, the second layer are the counties, then the cities and so on, as the following Figure 5 describes.

Figure 5 Geographical models decomposition

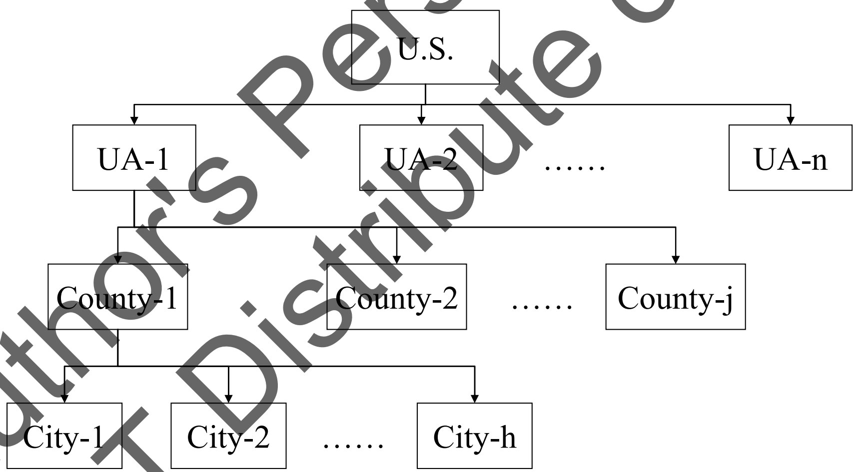

Practically, a county, for example, would ask its cities to provide a set of proposals for different budgets. Thus, each city would provide the ‘parent’ county with the piecewise linear relationship indicating the effectiveness in reducing the overall cost function for different possible values of the budget. This would allow each county to build their own piecewise linear relationship to feed its ‘parent’ UA, etc. This information would finally allow the Department of Homeland Security to solve the entire approximate model (2) and allocate the funds, now in a top-down fashion. Of course, cities could need further decompositions if their allocation problem results are too complex. Conversely, a county (or even an UA) could decide and be able to solve its entire allocation problem without further decompositions.

Thus, the underlying structure and the fund distribution 'channels' would not differ very much from what is currently in place, but the mechanism and the rationale behind it is much more rigorous, so that the amount of funds distributed from the top to the bottom at each step of the above diagram follows a risk-based logic and can be quite different from the current ones and theoretically more correct towards the goal of reducing the overall risk. The accuracy of the risk-based distribution will be proportional to the level of confidence in the data and risk assessments performed to build the model (2).

Nevertheless, we note that the geographical decomposition might not be the best choice to avoid interdependencies among subsets. While Assumptions 1 and 2 could be easier to accept with careful modelling, Assumption 3 could instead be unacceptable or at least very weak, jeopardising the accuracy of the obtained final allocation. In fact, as mentioned, the main sources of dependency among the subsets are the game-theoretic strategies involving the threat variables and defending a target by implementing a countermeasure would lead terrorists to divert their attention elsewhere, which can very well be in another city or county or UA and this would contradict with Assumption 3.

On the other hand, the geographical decomposition can be very suitable for the allocation problem to counter natural disasters, which can be seen as 'stupid terrorists' and would not be affected by game-theoretic issues at all. Therefore, the independency of the subset could be guaranteed.

4.2 Decomposition based on similarity

A possibly better alternative in choosing the subset decomposition strategy for terrorist events can be developed by the following rationale. Suppose that a target suddenly becomes less attractive to the terrorists for a specific planned type of attack, because several countermeasures have been adopted to protect it. From the terrorist point of view, preparing any attack takes time and effort and it is not easy to quickly change strategy in response to the new adopted defences. There is little time for them to completely modify strategy. Therefore, it is more likely that a similar target suitable for a similar attack could instead be sought. Basically, the strongest game-theoretic dependencies will exist among similar scenarios, where similarity is referred to the type of target (a stadium, a chemical plant, an airport, etc.) and/or the mean of attack (chemical weapon, conventional bomb, etc.) rather than the location.

Ignoring game-theoretic dependencies (that can happen with a geographically based subset selection) can significantly weaken the assumption of the decomposition methodology, especially within the terrorism context. Therefore, a wiser decomposition would suggest grouping similar scenarios with the strongest game-theoretic ties among each other. For example, grouping all the stadiums in a subset, all the chemical plants in another subset, all the airports in a third subset, etc., would allow modelling the (possibly strong) game-theoretic dependencies between similar targets. Thus, Assumption 3 would be more reasonable to accept. See Bier (2005) and Bier et al. (2005) for risk-based resource allocation with strong game-theoretic components.

5 Experiment

In this section, we provide an example to illustrate the decomposition procedure. The figures and assessments are fictitious (we would not be able to collect real data, since most of them are classified), but kept within realistic ranges. We consider 32 potential terrorist attack scenarios and a set of 70 possible countermeasures. A total budget of B = \$2.5M needs to be allocated. (Scenarios can be: 'missile attack on a chemical plant', 'bio-chemical attack on the football stadium', 'suicide bomb attack in a restaurant', 'conventional bomb attack on an aircraft', etc. Countermeasures can be: 'build a more effective fence around the plant', 'double the security in the stadium for all events', 'improve the bio-chemical and biological emergency response squads', 'improve intelligence to detect incoming threats', etc.). Table 1 summarises the threat (p_s), vulnerability (v_s), consequence (K_s) and resulting risk ($r_s = p_s v_s K_s$) assumed for each scenario s before implementing any countermeasure. For simplicity, we assume that all the terms of K_s have been converted by appropriate weights to economic consequences (M\$).

Table 1 Threat (p_s), vulnerability (v_s), consequence (K_s) and risk (r_s) for each scenario s

s	p_s	v_s	K_s (M\$)	r_s (M\$)	s	p_s	v_s	K_s (M\$)	r_s (M\$)
1	0.2	0.1	100	2	17	0.2	0.1	200	4
2	0.2	0.1	500	10	18	0.2	0.1	700	14
3	0.3	0.2	1,000	60	19	0.3	0.2	800	48
4	0.3	0.1	300	9	20	0.3	0.1	350	10.5
5	0.2	0.1	150	3	21	0.2	0.1	750	15
6	0.2	0.1	450	9	22	0.2	0.1	450	9
7	0.3	0.2	800	48	23	0.3	0.2	200	12
8	0.3	0.1	500	15	24	0.3	0.1	550	16.5
9	0.2	0.1	150	3	25	0.2	0.1	180	3.6
10	0.2	0.1	600	12	26	0.2	0.1	650	13
11	0.3	0.3	800	72	27	0.3	0.3	1,000	90
12	0.3	0.1	700	21	28	0.3	0.1	260	7.8
13	0.2	0.1	150	3	29	0.2	0.1	150	3
14	0.2	0.1	600	12	30	0.2	0.1	600	12
15	0.3	0.3	800	72	31	0.3	0.3	800	72
16	0.3	0.1	700	21	32	0.3	0.1	700	21

The following Table 2 lists the countermeasures (i), their costs (c_i) and their effect, if implemented, in reducing p_s, v_s and/or K_s for some scenario s. We assume that the countermeasures are all of the Y/N type associated with a binary variable y_i and that they are not linked by any logic constraints. We also ignore possible synergistic or antagonistic interdependencies among countermeasures. These simplifications do not significantly affect the effectiveness of the decomposition approach; they only help to greatly ease the calculations involved in this illustrative example.

Table 2 Costs (c_i) in \$10K for the countermeasures (i) and reduction of p_s, v_s, K_s

i	c_i	*Mitigation effect (reduction of)*	i	c_i	*Mitigation effect (reduction of)*
1	8	p_1 by 0.05, p_2 by 0.05	28	15	p_{15} by 0.1
2	20	p_3 by 0.1, v_3 by 0.05, v_4 by 0.05	29	20	v_{14} by 0.02
3	30	K_1 by \$50M	30	30	p_{14} by 0.05, K_{14} by \$200M
4	30	p_1 by 0.1, p_2 by 0.1, p_3 by 0.1, p_4 by 0.1	31	23	K_7 by \$300M
5	15	v_1 by 0.05	32	11	K_{12} by \$250M
6	20	v_2 by 0.05	33	24	p_{11} by 0.1
7	15	p_5 by 0.05, v_7 by 0.05, K_5 by \$50M	36	10	p_{17} by 0.05, p_{18} by 0.05
8	20	p_5 by 0.1, p_6 by 0.1, p_7 by 0.1, p_8 by 0.1	37	15	p_{19} by 0.1, v_{19} by 0.05, v_{20} by 0.05
9	4	p_6 by 0.05	38	6	K_{17} by \$150M
10	7	v_{11} by 0.05	39	25	p_{17} by 0.1, p_{18} by 0.1, p_{19} by 0.1, p_{20} by 0.1
11	15	p_7 by 0.1, p_8 by 0.1	40	12	v_{17} by 0.05
12	15	v_5 by 0.05, v_6 by 0.05	41	8	v_{18} by 0.05
13	3	p_9 by 0.05	42	15	p_{21} by 0.05, v_{23} by 0.05, K_{21} by \$100M
14	16	p_9 by 0.1, v_{15} by 0.05, K_{11} by \$300M	43	21	p_{21} by 0.1, p_{22} by 0.1, p_{23} by 0.1, p_{24} by 0.1
15	21	p_{10} by 0.05, K_9 by \$100M	44	17	p_{22} by 0.05
16	10	p_{10} by 0.1	45	14	v_{27} by 0.05
17	22	p_{11} by 0.1, K_{16} by \$250M	46	24	p_{23} by 0.1, p_{24} by 0.1
18	17	p_{13} by 0.1, v_9 by 0.05, K_{15} by \$300M	47	35	v_{21} by 0.05, v_{22} by 0.05
19	9	p_{12} by 0.1, K_{13} by \$100M	48	7	p_{25} by 0.05
20	15	p_{12} by 0.1, p_{14} by 0.1, v_{10} by 0.02	49	14	p_{25} by 0.1, v_{31} by 0.05, K_{27} by \$250M
21	12	p_{15} by 0.1, v_{16} by 0.05, K_2 by \$250M	50	3	p_{26} by 0.05, K_{25} by \$10M
22	15	v_{13} by 0.05, K_4 by \$100M	51	7	p_{26} by 0.1
23	7	p_{16} by 0.1, K_6 by \$100M	52	33	p_{27} by 0.1 K_{32} by \$250M
24	12	p_{16} by 0.1, K_8 by \$100M	53	21	p_{29} by 0.1, v_{25} by 0.05, K_{31} by \$300M
25	11	v_8 by 0.05, K_{10} by \$200M	54	34	p_{28} by 0.1, K_{29} by \$100M
26	16	v_{12} by 0.05	55	22	p_{28} by 0.1, p_{30} by 0.1, v_{26} by 0.02
27	15	p_{13} by 0.05	56	20	p_{31} by 0.1, v_{32} by 0.05, K_{18} by \$150M

Table 2 Costs (c_i) in \$10K for the countermeasures (i) and reduction of p_s, v_s, K_s (continued)

i	c_i	*Mitigation effect (reduction of)*	*i*	c_i	*Mitigation effect (reduction of)*
57	17	v_{29} by 0.05, K_{20} by \$50M	64	10	v_{30} by 0.02
58	10	p_{32} by 0.1, K_{22} by \$150M	65	25	p_{30} by 0.05, K_{30} by \$200M
59	12	p_{32} by 0.1, K_{24} by \$230M	66	22	K_{23} by \$20M
60	15	v_{24} by 0.05, K_{26} by \$210M	67	16	K_{28} by \$20M
61	15	v_{28} by 0.05	68	20	p_{27} by 0.1
62	17	p_{29} by 0.05	69	10	p_{20} by 0.1
63	18	p_{31} by 0.1	70	18	K_{19} by \$100M

Solving the above global allocation problem as a whole allows us to reach optimality in about 200 seconds, distributing \$2.5M to countermeasures 8, 9, 10, 11, 14, 19, 20, 21, 28, 35, 37, 41, 43, 49, 51, 56, 63, 68 and minimising the overall expected cost C to a value of \$342.7M.

The above set of scenarios is then decomposed in two subsets (scenarios 1 to 16 and countermeasures 1 to 35 in subset I; scenarios 17 to 32 and countermeasures 36 to 70 in subset II) so that Assumptions 1, 2 and 3 are satisfied. In fact, we can assume that the c_i of each countermeasure will be entirely paid by the budget of its subset (Assumption 1), there are no logic constraints linking countermeasures belonging to different subsets (Assumption 2) and there are no relationships between countermeasures in subset I and scenarios in subset II and vice versa (Assumption 3). This allows us to perform the allocation by the decomposition procedure.

Subsets I and II have been solved for different values of their respective budgets B_I and B_{II}. B_I^{max} and B_{II}^{max} are \$2.5M and the interval between consecutive values of B_I and B_{II} is \$312.5K. The resulting convex piecewise linear functions are shown in the following Figures 6(a) and 6(b). These functions would be the output that a ‘child’ subset would need to provide its ‘parent’ with.

For each value of B_I and B_{II}, the solver took 2 seconds or less to reach optimality. The overall approximate model has been then reformulated as a linear program using the piecewise linear functions just obtained and solved to optimality recommending an exact split of the \$2.5M budget among subset I and subset II. With $B_I = B_{II} = \$1.25M$ the solution recommends the following countermeasures to be implemented: 8, 10, 11, 14, 19, 20, 21, 28, 35 in Subset I and 37, 38, 43, 49, 50, 51, 56, 63, 68 in Subset II. The resulting overall cost function is minimised at \$344.07M. The total CPU solution time for the whole decomposition procedure was about 20 seconds. Table 3 summarises the results of the allocation.

Figure 6 Piecewise linear functions for subsets I and II

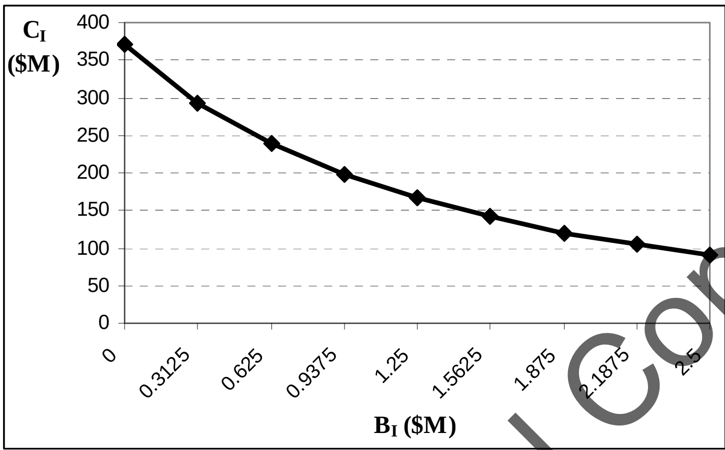

(a)

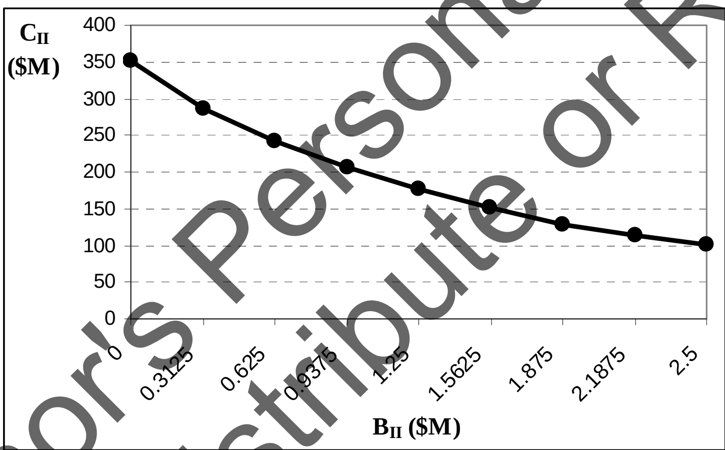

(b)

Table 3 Comparison between global and decomposition solution approach

	Global approach	Decomposition approach	Difference
Overall cost (C)	$342.7M	$344.07M	+0.4%
Countermeasures recommended for implementation	8, 9, 10, 11, 14, 19, 20, 21, 28, 35, 37, *41*, 43, 49, 51, 56, 63, 68	8, 10, 11, 14, 19, 20, 21, 28, 35, 37, *38*, 43, 49, *50*, 51, 56, 63, 68	2 out of 18 countermeasures
Solving time	200 sec	20 sec	–90%

The decomposition approach obtained almost the same results of the global optimisation approach: the overall cost is only 0.4% higher; the recommended countermeasures are the same except two; the CPU solution time is reduced by 90%.

The above example is of course extremely simplified, but shows that if a smart decomposition (either geographical, by scenario similarity or others) is performed, so that Assumptions 1, 2, 3 can reasonably be satisfied, the procedure reaches similar results with respect to the global approach. Also, we notice a significant saving in the CPU solution time, which can be a major issue for much larger problems.

The major challenge in successfully applying the decomposition method resides in selecting proper subsets and its major advantage would likely be in the management of the allocation process as a whole, as it could be practically implemented within the current geographical partitioning of the country, as explained in Section 4.1. In fact, a good decomposition strategy would allow separating the created subsets completely, from data gathering, risk assessments and countermeasures' identification, evaluation, effects and selection. The only 'output' from each subset j to its 'parent' in the 'bottom-up chain' shown in Figure 4 would be the piecewise linear function $C_j^l(B_j)$.

6 Conclusions

The current methodology to allocate federal funds to counter terrorism across the US is based on a '40%/60%' formula, which allocates 40% of the available funds evenly across the States and the remaining 60% following a classified risk-based approach. This mechanism might have been politically more convenient to apply when performing the allocation, but is hard to validate theoretically if the scope of the allocation is to minimise the total risk suffered by the whole nation, as Secretary Chertoff demanded and it has been criticised by most experts.

In this paper, we address the issue by developing a rigorous risk-based funding allocation methodology to counter terrorism and we propose to apply a decomposition methodology for solving it. The underlying basis of the methodology is today's largely accepted definition of risk, which is the product of threat, vulnerability and consequences, which are a weighed sum of four dimensions: economic losses, human health, strategic and psychological impact. The resulting optimisation model is a variation of the Knapsack problem, which is unsolvable to optimality in reasonable time, given the huge dimension of the real nationwide problem. The proposed decomposition solution method provides approximate but faster solutions and it can be implemented to the current funding hierarchical structure, based on geographical partitioning. The corresponding bottom-up procedure asks 'child' subsets (i.e., counties) to provide their 'parents' (i.e., UAs) with piecewise linear functions indicating the cost-effectiveness of allocating different values of the budget in reducing the overall cost. The approximate overall model is then solved to optimality and the allocation is performed now in a top-down fashion among subsets. An example is provided to show the effectiveness of the procedure in terms of adherence to optimality, CPU solution time savings and practical implementation.

While the current geographical administrative partitioning would more easily encourage an implementation of the procedure, the needed independence of the created subsets can be a weak assumption, because of game-theoretic issues among different scenarios. Future research could include estimating the deviation from optimality of the allocation due to the oversight of game-theoretic independence among subsets. We also suggest evaluating a partitioning based on 'similarity' of scenarios, which could reduce this potential problem.

This paper is an attempt to reduce the barrier between theory and applications, by an integration of an effective optimisation methodology into the current societal structure in order to mitigate the risk due to terrorism, a significant challenge faced by our modern society, with the ultimate goal to make the nation safer. We would also emphasise that the methodology and its suggested practical implementation could be applied to reducing risk for natural disasters, for which risk assessments are still quite difficult to perform, but the overall uncertainty is reduced, because of the lack of game-theoretic issues.

Acknowledgements

This research was supported by the USA Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number N00014-05-0630. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the USA Department of Homeland Security.

References

Bier, V.M. (2005) 'Game-theoretic and reliability methods in counter-terrorism and security. mathematical and statistical methods in reliability', in A. Wilson, N. Limnios, S. Keller-McNulty and Y. Armijo, (Eds.): *Series on Quality, Reliability and Engineering Statistics*, World Scientific, Singapore, pp.17–28.

Bier, V.M., Nagaraj, A. and Abhichandani, V. (2005) 'Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries', *Reliability Engineering and System Safety*, Vol. 87, pp.313–323.

Bretthauer, K.M. and Shetty, B. (1995) 'The non-linear resource allocation problem', *Operation Research,* Vol. 43, pp.670–683.

Bretthauer, K.M. and Shetty, B. (2002) 'The non-linear knapsack problem–algorithms and applications', *European Journal of Operational Research,* Vol. 138, pp.459–472.

Brunet, A. (2005) *Protecting Only Part of Our Homeland: Vulnerability Across States and the Allocation of Federal Terrorism Funds,* PhD. Purdue University.

Chankong, A. and Haimes, Y.Y. (1983) *Multiobjective Decision Making,* North Holland, New York, NY.

Congressional Quarterly (2004) *Republican Cox Delivers Remarks at McGraw-Hill Homeland Security Summit.*

Congressional Quarterly (2005) *Chertoff Orders Thorough Review of the DHS structure and Operations.*

De Rugy, V. (2005) 'What does homeland security spending buy?', *American Enterprise Institute for Public Policy Research,* Working paper, No. 107.

Haimes, Y.Y. (2004) *Risk Modeling, Assessment, And Management,* Wiley-Interscience, Hoboken, NJ.

Hall, M. (2003) 'Anti-terror funding questioned', *USA TODAY*, July 14, 2003.

Hammond, J.S., Keeney, R.L. and Raiffa, H. (1999) *Smart Choices: a Practical Guide To Making Better Decisions,* Harvard Business School Press, Boston, MA.

Hwang, L.C. and Yoon, K. (1981) *Multi Attribute Decision-Making: A Methods and Applications,* Springer-Verlag, Berlin.

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

Keeney, R.L. (1992) *Value-focused Thinking: a Path to Creative Decision Making*, Harvard University Press, Cambridge, MA.

Keeney, R.L. and Raiffa, H. (1993) *Decisions with Multiple Objectives: Preferences and Value Tradeoffs,* Cambridge University Press, Cambridge, UK.

Ransdell, T. (2004) *Federal Formula Grant and California – Homeland Security,* Public Policy Institute of California, San Francisco, CA.

Saaty, T.L. (1986) 'Axiomatic foundations of the analytic hierarchy process', *Management Science,* Vol. 32, pp.841–855.

Saaty, T.L. (1994) 'How to make a decision: the analytic hierarchy process', *Interfaces,* Vol. 24, pp.19–43.

Saaty, T.L. and Vargas, L.G. (2001) *Models, Methods, Concepts and Applications of The Analytic Hierarchy Process,* Kluwer Academic Publishers, Boton, MA.

The 9/11 Commission Report (2004) *The National Commission on Terrorist Attacks upon the USA*, pp.396.

Triantaphyllou, E. (2000) *Multi-criteria Decision Making Methods: A Comparative Study*, Kluwer Academic Publishers, Boton, MA.

Weinlein, M. (2004) *Funding for First Responders from a Threat and Prevention Approach*, MS Naval Postgraduate School.

Willis, H.H., Morral, A.R., Kelly, T.K. and Medby, J.J. (2005) *Estimating Terrorism Risk*, RAND Corporation, Santa Monica, CA.

Willis, H.H. (2006) 'Guiding resource allocation based on terrorism risk', Working paper, RAND Corporation, Santa Monica, CA.

Zeleny, M. (1982) *Multiple Criteria Decision Making*, McGraw-Hill, New York, NY.