

Optimizing Department of Homeland Security

Defense Investments:

Applying Defender-Attacker (-Defender) Optimization

To Terror Risk Assessment and Mitigation¹

Gerald G. Brown, W. Matthew Carlyle, and R. Kevin Wood

Operations Research Department

Naval Postgraduate School, Monterey, CA 93943

The U.S. Department of Homeland Security (DHS) is investing billions of dollars to protect us from terrorist attacks and their expected damage (i.e., risk). We present prescriptive optimization models to guide these investments. Our primary goal is to recommend investments in a set of available defense options; each of these options can reduce our vulnerability to terrorist attack, or enable future mitigation actions for particular types of attack. Our models prescribe investments that minimize the maximum risk (i.e., expected damage) to which we are exposed. Our “Defend-Attack-Mitigate risk-minimization model” assumes that terrorist attackers will observe, and react to, any strategic defense investment on the scale required to protect our entire country. We also develop a more general tri-level “Defender-Attacker-Defender risk-minimization model” in which (a) the defender invests strategically in interdiction and/or mitigation options (for example, by inoculating health-care workers, or stockpiling a mix of emergency vaccines) (b) the attacker observes those investments and attacks as effectively as possible, and (c) the defender then optimally deploys the mitigation options that his investments have enabled. We show with simple numerical examples some of the important insights offered by such analysis. As a byproduct of our analysis we elicit the optimal attacker behavior that would follow our chosen defensive investment, and therefore we can focus intelligence collection on telltales of the most-likely and most-lethal attacks.

¹ Appears as Appendix E of: National Research Council, 2008, “*Department of Homeland Security Bioterrorist Risk Assessment: A Call for Change*,” National Academies Press, Washington, DC.

INTRODUCTION

Since 9/11, the U.S. Department of Homeland Security (DHS) has marshaled significant resources to assess the risk to our populace from terrorist attacks of all kinds. The work we report here is directly motivated by just one such risk assessment: pursuant to Homeland Security Presidential Directive 10 (HSPD 10) [The White House, 2004], DHS has conducted an extensive bio-terrorism risk-assessment exercise, referred to here as the Bio-Terror Risk Assessment (BTRA) [DHS 2006]. BTRA estimates risks of many bio-terror attack possibilities, and classifies a list of particular bio-terror agents as *most-, intermediate-, and least-threatening*.

The BTRA risk assessment depends upon subject-matter experts (SMEs) advising, with perfect knowledge, the probability that the “attacker” (terrorist or terrorist group), or “defender” (the federal government), will choose some particular option at each stage of an 18-stage probability risk assessment tree.

We contend that representing intelligent adversarial decisions with static probabilities elicited from SMEs is an untenable paradigm: Not only can experts make mistakes, but static probabilities make no sense when the attacker can observe and react, dynamically, to any earlier decisions made by the defender.

We also hold that the business of DHS lies not just in assessing risks, but also in wisely guiding investments of our nation’s wealth to reduce these risks. These are strategic *decisions* that must be made now, in a deliberative fashion.

Here, we try to adopt the same problem context as BTRA to recoup its estimable investment in risk modeling. But, we distinguish between (a) strategic investment decisions that DHS makes that are visible to terrorists, (b) the decision a terrorist makes to attempt an attack and, finally, (c) the after-attack mitigation efforts that prudent DHS investments will have enabled.

Our work applies equally well to any category of threat that concerns DHS enough to warrant investments so significant they cannot be hidden from our taxpayers, and thus not from terrorists, either.

Such threats cover biological, nuclear, chemical, and conventional attacks on our infrastructure and citizens, as well as sealing our borders against illegal immigration, and a host of military topics.

The modeling presented here has been motivated and validated by more than one hundred worldwide infrastructure vulnerability analyses conducted since 9/11 by the military-officer students and the faculty of the Naval Postgraduate School [Brown et al. 2005a, Brown et al. 2006a]. Some of these studies have been developed into complete decision-support systems:

- Salmerón, et al. [2004] have received DHS and Department of Energy support to create the Vulnerability of Electric Grids Analyzer (VEGA), a highly detailed, optimization-based decision-support system. VEGA can evaluate, on a laptop computer, the vulnerability and optimal defense of electrical generation and distribution systems in the U.S., where risk is measured as expected unserved demand for energy during any repair-and-recovery period.
- We have developed a decision-support system to advise policy regarding the interdiction of a proliferator's industrial project to produce a first batch of nuclear weapons [Brown et al. 2006b, 2007].
- The U.S. Navy has developed a decision-support system to optimally pre-position sensor and defensive interceptor platforms to protect against a theater ballistic missile attack [Brown et al. 2005b].

The message here is that, with experience, we have gained confidence that these new mathematical methods produce results that exhibit the right level of detail, solve the right decision problems, and convey useful advice and insight to policy makers. Such capabilities have not been available before.

THE MODEL, "MXM"

BTRA uses a descriptive model. Our focus is prescriptive, rather than descriptive: our models suggest prudent investment and mitigation plans for biodefense, and we strive to provide a realistic representation of the attack decisions made by an intelligent adversary.

As the *defender*, we seek to allocate a limited budget among biodefense investment options to form a defense strategy that minimizes the maximum *risk* to a terrorist *attacker*'s actions. We might define risk as the expected number of fatalities, or as the expected 95-th percentile of fatalities, or as any other gauge that appeals. Risk is a somewhat ambiguous term when used to discuss our bilateral view of conflict between intelligent adversaries, so we hereafter substitute “expected damage to the defender.” We assume that an intelligent adversary will attempt to inflict maximum expected damage. The following, simplified model minimizes a reasonable upper bound on expected damage; we discuss generalizations later.

Indices

$d \in D$	defense strategy, e.g., stockpile vaccines A and B, but not C
$a \in A$	attack alternative, e.g., release infectious agent V
$m \in M$	after-attack, mitigation activity, e.g., distribute vaccine A
$m \in M_d$	mitigation activities enabled by defense option d , e.g., distribute vaccine A, distribute vaccine B
$d \in D_m$	defense strategies that enable mitigation activity m
$k \in K$	resource types used by mitigation activities, e.g., aircraft for distributing vaccine, personnel for administering vaccine

Data

$damage_{d,a}$	expected damage if defense strategy d and attack alternative a are chosen, given no mitigation
$mitigate_{d,a,m}$	expected damage reduction of after-attack mitigation effort m , given investment strategy d and attack a , (assumes additive reduction and $\sum_m mitigation_{d,a,m} \leq damage_{d,a}$)
$r_{k,d}$	total mitigation resource of type k available if defense strategy d is chosen

$q_{k,d,m}$ consumption of mitigation resource k provided by defense option d for mitigation activity m

Decision Variables

w_d 1 if defense strategy d chosen, else 0

x_a probability attacker chooses attack alternative a ($0 \leq x_a \leq 1$)

$y_{d,m}$ fraction of defense strategy d effort devoted to mitigation activity type m

Formulation: MIN-MAX-MIN (MXM)

(Defender-Attacker-Mitigator)

$$z^* = \min_{w_d} \max_{x_a} \min_{y_{d,m}} \sum_{d,a} \text{damage}_{d,a} w_d x_a - \sum_{d,a,m} \text{mitigate}_{d,a,m} x_a y_{d,m} \quad (\text{D0})$$

$$\sum_d w_d = 1 \quad (\text{D1})$$

$$\sum_a x_a = 1 \quad (\text{A1})$$

$$\sum_{d,m} q_{k,d,m} y_{d,m} \leq \sum_d r_{k,d} w_d \quad \forall k \in K \quad (\text{M1})$$

$$y_{d,m} \leq w_d \quad \forall d \in D, m \in M_d \quad (\text{M2})$$

$$w_d \in \{0,1\}, x_a \geq 0, y_{d,m} \geq 0 \quad \forall d \in D, a \in A, m \in M_d$$

Description

The order of appearance of the operators, min, followed by max, followed by min, in the objective function (D0) represents the sequential nature of the decisions we are modeling, from the outside to the inside. The coefficient $\text{damage}_{d,a}$ in the objective accounts for any interdiction effects that strategy d has on attack a , effects that are independent of any mitigation activities. (For example, vaccinating emergency and health-care providers falls under the category of “interdiction”: after an attack, no follow-up mitigation efforts apply to this vaccination.) The right-most minimization term, over $y_{d,m}$, subtracts from expected damage if a mitigating effort has been enabled by the defense plan, and if some amount of

that mitigation is applied. For simplicity of exposition, we assume that mitigation results are additive and restricted to sum to some value not exceeding total expected damage. (See the definition of $mitigate_{d,a,m}$.) Constraint (D1) simply limits the defender to choosing one defense strategy. Constraint (A1) limits the attacker to choosing a mixed attack strategy, which of course admits a pure attack as well. Constraints (M1) are joint resource constraints on mitigation efforts; constraints (M2) stipulate that mitigation efforts are permitted only if the enabling defense strategy has been chosen. Constraints (M1) subsume those of type (M2), but we keep these separate for later clarity. The attack variables, x_a , and the mitigation variables, $y_{d,m}$, are continuous. If the attacker variables are restricted to be integer (for instance, they might be binary variables indicating whether or not the terrorists decide to fully develop and deploy a particular pathogen in an attack), then the resulting analysis becomes significantly more complicated than that which we present here. Although dealing with bioterrorist attacks might be most naturally modeled using integer attacker variables, our model with continuous attack (y_a) variables will at least provide a conservative estimate of the defender's objective; i.e., the attacker's abilities to inflict damage are over-estimated by our model.

Discussion of MXM

Figure E 1 depicts a tree showing the sequential actions of the defender (selecting a defense strategy), the attacker (choosing attack alternatives), and the defender (mitigating damage with resources put in place by the defense strategy). (We use the generic term “tree” to represent the sequence of defender and attacker decisions we model. The “decision tree” of Raiffa [1968] pits a single decision maker against Mother Nature, while here we have two opponents trying to shape an outcome governed by Mother Nature. The term “game tree” [Kuhn, 1953] is a more appropriate term for our bioterror situation.) Each defense strategy has an immediate effect on the maximum damage of any attack, reflected in $damage_{d,a}$; it can also enable the capability to reduce after-attack damage by as much as $mitigate_{d,a,m}$, if the chosen defense strategy permits a full allocation of mitigation resources to mitigation action m . Given a fixed defense strategy, we assume the attacker will first observe this

strategy and then respond with a mixed strategy over the set of possible attacks. As we have said, this might be a relaxation of the original optimization problem faced by the attacker, and therefore grants him more attack capability than he really has in this sequential decision-making. In general we cannot tell how weak this relaxation is, but for specific cases (especially those with a moderate number of feasible attacker decisions) we can use enumeration to bound the effect of this relaxation on the optimal objective function value.




Figure E 1. This tree depicts, left-to-right, a leading defense strategy choice w_d , consisting of component defense investment options, and visible to an attacker, followed by attack alternative choice(s) x_a that (each) inflict expected damage $damage_{d,a}$. Square nodes indicate defender decisions, and circle nodes indicate attacker decisions. We only illustrate a mitigation subtree ($y_{d,m}$ decisions) for one (w_d, x_a) pair. For a given defense strategy $w_d = 1$, the optimization recommends a mixed attack strategy for the attacker and a mixed mitigation response $y_{d,m}$ from the defender. The defense strategy establishes all mitigation resources that can be used after an attack. That strategy is seen by the attacker when he develops his attack plan. Enabled mitigation resources can reduce expected damage through $-mitigate_{d,a,m} x_a y_{d,m}$. (Our conservative model does not allow the defender to observe the precise type of attack, however, so the mitigation response may not be optimal.)

A “mixed attack strategy” means that the optimal attacker decision includes multiple attacks and then we choose mitigation responses, and this results in some damage that can only be estimated, and some part of that estimation can involve an expectation. (For example, the damage could involve an expectation taken over a probability distribution for the time between when an attack is launched to when

it is discovered.) Thus, integrating damage over one or more probability distributions yields an objective function that measures “expected damage.”

Solving MXM

Temporarily fixing $w = \hat{w}$ in **MXM**, we take the linear-programming dual (hereafter referred to simply as “the dual”) of the innermost minimizing linear program, using dual variables α_k for constraints (M1), and $\beta_{d,m}$ for constraints (M2). This converts the inner “max-min problem” into a “max-max problem,” which is a simple maximization:

Formulation: MAX-ATTACKER-LP(\hat{w})

$$\begin{aligned}
z_{\max} &= \max_{\alpha, \beta} \sum_{d,a} \text{damage}_{d,a} \hat{w}_d x_a - \sum_k r_{k,d} \alpha_k - \sum_{d,m \in M_d} \hat{w}_d \beta_{d,m} \\
\text{s.t.} \quad & \sum_a x_a = 1 & \text{(A1)} \\
& \sum_k q_{k,d,m} \alpha_k + \beta_{d,m} \geq \sum_a \text{mitigate}_{d,a,m} x_a & \forall d \in D, m \in M_d & \text{(DM1)} \\
& \alpha_k \geq 0 & \forall k \in K \\
& \beta_{d,m} \geq 0 & \forall d \in D, m \in M_d
\end{aligned}$$

Now, leaving $w = \hat{w}$ as shown in MAX-ATTACKER-LP, we take the dual of this linear program, using dual variables \mathfrak{R} for constraint (A1) and $y_{d,m}$ for constraints (DM1), and then release w to vary as before, to achieve the following integer linear program which is essentially equivalent to MXM:

Formulation: MIN-ILP

(Defender-Attacker-Mitigator)

$$\begin{aligned}
z_{\min} &= \min_{\mathfrak{R}, w_d, y_{d,m}} \mathfrak{R} & \text{(DILP0)} \\
\text{s.t.} \quad & \mathfrak{R} \geq \sum_d \text{damage}_{d,a} w_d - \sum_{d,m} \text{mitigate}_{d,a,m} y_{d,m} \quad \forall a \in A & \text{(DILP1)} \\
& \sum_d w_d = 1 & \text{(D1)} \\
& \sum_{d,m} q_{k,d,m} y_{d,m} \leq \sum_d r_{k,d} w_d & \forall k \in K & \text{(M1)} \\
& y_{d,m} \leq w_d & \forall d \in D, m \in M_d & \text{(M2)} \\
& w_d \in \{0, 1\}, y_{d,m} \geq 0 & \forall d \in D, m \in M_d
\end{aligned}$$

The optimal solution to MIN-ILP prescribes among other things a choice for the defense strategy, w^* , to be implemented immediately by the defender, before an attack occurs. Given optimal incumbent solution w^* , we recover the attacker’s optimal strategy x^* by solving **MAX-ATTACKER-LP**(w^*).

A Numerical Example of MXM

We provide a small numerical example to illustrate the features of **MXM**.

We introduce a number of *defensive investment options*, programs that can be composed in groups into *defense strategies*. Table E 1 displays defensive investment options and costs.

i	cost _i
i01	2
i02	3
i03	5

Table E 1. Defensive investment options and costs. For example, option “i03” costs 5. Total budget, logical, and perhaps political considerations will limit the combinations of these options that can comprise admissible defense strategies.

In our example, the defensive investment options are denoted “i01,” “i02,” and “i03.” From this set, policy makers have determined 6 combinations that comprise the subset of admissible defense strategies whose implementation will depend on the available budget; see Table E 2. Table E 3 displays expected damage resulting from each defense strategy and each attack alternative, i.e., the terms $damage_{d,a}$.

		Investment options		
		i01	i02	i03
Defensive strategies	d00			
	d01	x		
	d02		x	
	d03			x
	d04	x	x	
	d05		x	x

Table E 2. Defensive investment options in each potential defense strategy. Strategy “d00” makes no investment at all. Defense strategy “d05” includes investment options “i02” and “i03.” Logical, political, or other considerations preclude some of the strategies, for example, {“i01,” “i03”}. The total available budget, not yet specified, can also preclude certain strategies. For instance, {“i02” and “i03”} cannot be selected if the total budget is less than 8.

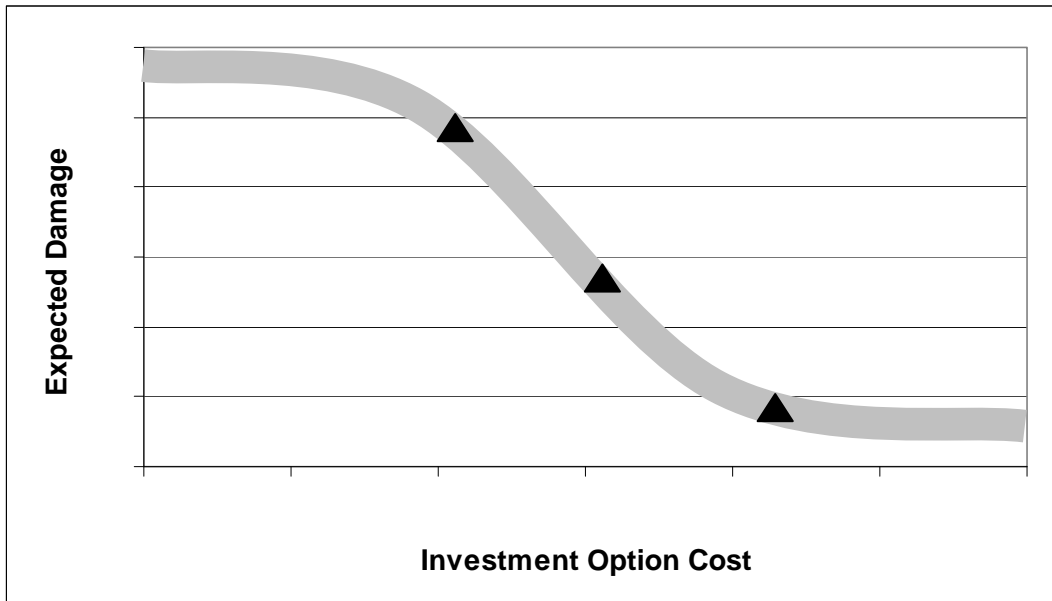


Figure E 2. The purpose of DHS defensive investment options is to reduce expected damage before an attack occurs, and/or allow mitigation of expected damage after one occurs. The generic relationship illustrated here conjectures little to no effect at low investment levels, followed by increased effectiveness, and eventually leveling off with diminishing returns. The triangles represent points we might use as alternate investment options to adequately represent the entire function.

Figure E 2 illustrates the generic relationship relating investment options to the ability to reduce expected damage from any terrorist attack before it is carried out, and/or mitigate damage after an attack occurs. This is a complicated function, neither convex nor concave, but our sampling of representative points can be used to represent this in characterizing component investment options in defense strategies.

Damage estimates in Table E 3 include any synergies among or interference between component investment options in each defense strategy preparing for each attack. This is key. BTRA makes a point of such dependencies, and we represent these in complete, realistic detail here.

	a01	a02	a03
d00	10	10	10
d01	10	5	7
d02	6	8	7
d03	6	6	6
d04	4	3	5
d05	5	5	4

Table E 3. Expected damage resulting from each defense strategy (row) and each attack alternative (column), accounting for interdiction but not mitigation. (This table gives the values for $damage_{d,a}$ for MXM. We use integral data to permit reproduction of our results.)

Table E 4 represents estimated mitigation capabilities. These mitigation estimates correspond to a single, “full-strength” mitigation effort being applied to a single attack alternative. If the attacker chooses a mixed attack strategy, we may need to spread mitigation effort across multiple activities, reducing the expected effectiveness of each activity accordingly.

m=m1	a01	a02	a03	m=m2	a01	a02	a03
d00	0	0	0	d00	0	0	0
d01	1	0	0	d01	1	0	0
d02	0	1	1	d02	0	2	0
d03	0	0	1	d03	0	0	1
d04	1	1	1	d04	0	1	2
d05	0	1	1	d05	0	0	2

Table E 4. Maximum expected damage reduction from a mitigation activity enabled (prior to an attack) by a defense strategy (and applied after an attack). These tables specify $mitigate_{d,a,m}$ for MXM, for each of two mitigation options (“m01” and “m02”), for each combination of defense and attack. For example, with defense option “d04” and attack “a03,” if we choose mitigation “m01” we reduce the damage by one unit, but if we choose mitigation “m02” we reduce the expected damage by two units (circled values).

The choice of defense strategy is limited by a total budget, which we vary over the integers from 0 to 11. We allow full employment of either mitigation effort, or any convex combination of them.

Because the defender is minimizing the optimal objective function value of a *maximization* problem, the optimal solution invests to reduce the expected damage, given future mitigation capability, of the most-threatening mixed attack. This requires that the defender invest in a defense strategy that enables him to mitigate several very-damaging attacks, and not just the worst one.

Figure E 3 shows minimized maximum expected damage as a function of total defense budget, and Table E 5 summarizes the solutions for each budget break-point. For instance, with a budget of 3, the optimal defense plan in MXM is to choose defense option “d02.” The terrorists’ optimal attack is a mixed strategy, with a probability of 0.50 of choosing “a02” and probability 0.50 of choosing “a03”. The resulting expected damage, after mitigation, is 6.5. Analysis of this simple case reveals that we have optimally allocated our mitigation effort among the two worst attacks, reducing the expected damage in each attack to the same value, 6.5. We can do no better than this, given our conservative approximation.

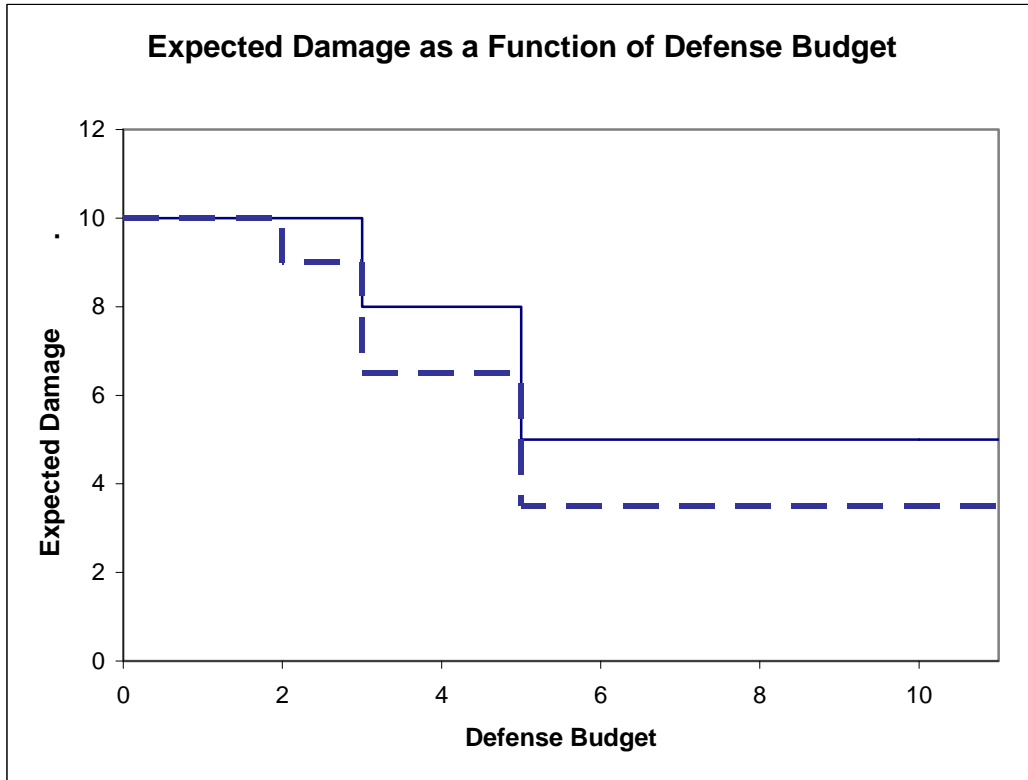


Figure E 3. Expected damage as a function of defense budget. This display is for policy makers: as we devote more and more defense budget, we achieve less and less expected damage. Because the defender’s investment options here are discrete, each improvement appears as a staircase drop as soon as sufficient budget permits some new, improved cohort of investment defense options, i.e., a new defense strategy. The law of diminishing returns is evident: expected damage reduced by each budget dollar decreases as budget increases. Policy-makers can usually put their finger on the spot that appeals in an illustration such as this, perhaps based on criteria not part of the underlying modeling. The uppermost, solid line displays the expected damage when all mitigation $_{d,a,m}$ values are set to zero (i.e., we have no mitigation capability) and only consider the expected damage from adopting a defense strategy, and then suffer the worst-case attack per expected damage in Table E 3. The dashed line illustrates the expected damage from MXM, the tri-level optimization.

Budget	MXM with $y = \mathbf{0}$			MXM			
	w	x	z^*	w	x	y	z^*
0	d00	a01	10	d00	a01	-	10
2	d01	a01	10	d01	a01	m01	9
3	d02	a02	8	d02	a02(.50) a03(.50)	m01(.50) m02(.50)	6.5
5	d04	a03	5	d04	a01(.50) a03(.50)	m01(.50) m02(.50)	3.5

Table E 5. For each budget just sufficient to afford a new defense strategy, we show the Defender-Attacker solution and expected damage (i.e., for MXM with $y = \mathbf{0}$), the Defender-Attacker-Defender solution (for MXM) and expected damage. For example, with a budget of 3, the optimal defense strategy in MXM is “d02”. The terrorists’ optimal attack is a mixed strategy, choosing alternative “a02” with probability 0.50, and “a03” with probability 0.50. We anticipate responding accordingly with “m01,” the optimal response to “a01,” with the same probability (0.50), and similarly with “m02” with probability 0.5. The resulting expected damage, after optimal mitigation in each case, is 6.5.

Generalizing Beyond Tri-level Decision Problems

The DHS biological threat risk assessment (BTRA) consists of an 18-stage probability risk assessment tree, where each decision has been replaced by an *a priori* probability; for example, see NRC [2007]. In the case of the each opponent, these probabilities are determined by subject-matter experts assessing how terrorists might make each decision, and how well DHS will do thwarting a bioagent attack at some intermediate stage of its development.

We could instead model BTRA as a 19-stage defender-attacker-defender model, with a new stage zero describing how DHS can invest in strategic biological defense strategies, and each of the intermediate stages represented by a set of decision variables that prescribe attacker or defender behavior, and solve a multi-stage defender-attacker-defender(-attacker-...) model to determine optimal stage-zero investment decisions to minimize expected damage assuming each opponent makes the optimal decision at each node of the corresponding tree. To fully represent the sequential nature of these decisions, we would require all decisions (except maybe those in the final stage) to be modeled with integer variables. However, solving such a model for just two stages of integer decisions is difficult.

We do not have the technology to handle three, much less 18, stages of alternating integer decisions. Allowing continuous decision variables in each of the stages except stage zero (our defense decision variables) would again be a relaxation of the restrictions on the attacker, and could, in some cases, yield extremely weak bounds on our defensive capability.

We now show in the case of a two-stage model how this relaxation from integer to continuous variables reduces the sequential decision problem to a simultaneous two-person zero-sum game.

Consider the bi-level attacker-defender, max-min optimization formulation: $(\mathbf{A}_L\mathbf{D}_L)$, where the subscript “L” denotes a linear program (i.e., continuous decision variables, and objective and constraints that are linear in those decision variables):

$$\begin{array}{llll}
 \max_x \min_y & g^T x + x^T Q y + c^T y & & \text{[dual variables]} \\
 \text{s.t.} & Ax & \leq b & [\pi] \quad (\text{B1}) \\
 & & Dy & \geq d \quad [\mu] \quad (\text{B2}) \quad (\mathbf{A}_L\mathbf{D}_L) \\
 & x \geq 0 & & \\
 & & y \geq 0 & .
 \end{array}$$

$(\mathbf{A}_L\mathbf{D}_L)$ is a more general version of the model used by Fulkerson and Harding [1977] and Golden [1978] for their work on continuous network interdiction models.

Take the dual of the inner (defender, “y”) problem in $(\mathbf{A}_L\mathbf{D}_L)$:

$$\begin{array}{llll}
 \max_{x,\mu} & g^T x + d^T \mu & & \\
 \text{s.t.} & Ax & \leq b & [\pi] \quad (\text{B1}) \\
 & -Q^T x + D^T \mu & \leq c & [y] \quad (\text{D2}) \quad (\mathbf{A}_L\bar{\mathbf{D}}_L) \\
 & x \geq 0 & & \\
 & & \mu \geq 0 & .
 \end{array}$$

This is our standard way to convert a “max-min” problem, for which there is no conventional optimization method, into an equivalent “max-max” problem that is nothing more than a conventional linear program.

Now, reverse the order of play in $(\mathbf{A}_L\mathbf{D}_L)$ to $(\mathbf{D}_L\mathbf{A}_L)$:

$$\begin{array}{llll}
\min_y \max_x & g^T x + x^T Q y + c^T y & & \text{[dual variables]} \\
\text{s.t.} & Ax & \leq b & [\pi] \quad (\text{B1}) \\
& & Dy & \geq d \quad [\mu] \quad (\text{B2}) \quad (\mathbf{D_L A_L}) \\
& x & \geq 0 & \\
& & y & \geq 0 \quad .
\end{array}$$

This variation on $(\mathbf{A_L D_L})$ is formulated as if the defender makes his decision first.

Take the dual of the inner, attacker, (“x”) problem in $(\mathbf{D_L A_L})$:

$$\begin{array}{llll}
\min_{y, \pi} & b^T \pi + c^T y & & \\
\text{s.t.} & A^T \pi - Q y & \geq g & [x] \quad (\text{D1}) \\
& & Dy & \geq d \quad [\mu] \quad (\text{B2}) \quad (\mathbf{D_L \bar{A}_L}) \\
& \pi & \geq 0 & \\
& & y & \geq 0 \quad .
\end{array}$$

This formulation is equivalent to $(\mathbf{D_L A_L})$, and is also a linear program.

We observe that $(\mathbf{A_L \bar{D}_L})$ and $(\mathbf{D_L \bar{A}_L})$ are linear programming duals of each other, and thus (assuming both are feasible) have the same optimal objective-function values, which is the same as the optimal objective value of $(\mathbf{A_L D_L})$. Therefore, the sequence in which the decisions are made (either attacker first, followed by defender, or defender first, followed by attacker) has no impact on the optimal objective-function value.

We have therefore proved the following:

Theorem 1: for any attacker-defender model in the form $(\mathbf{A_L D_L})$, we can exchange the order of decisions without affecting the optimal objective function value.

Theorem 1 is a simple extension of von Neumann's [1928] minimax theorem for polyhedral feasible regions using a proof technique similar to Ville [1938], but using the more modern technology of LP duals directly. This exchange argument, along with the observation that any two consecutive decision stages controlled by the same decision maker are equivalent to a single stage, (since both stages are either a maximization or both are a minimization over a set of decision variables, this is equivalent to a single maximization, or minimization, over all of those variables simultaneously), can be repeated for any

number of consecutive stages with continuous decision variables. The final model obtained in this manner is a simple maximization or minimization problem.

Specifically, if we were to apply this to the 18-stage BTRA model, (i.e., the model we would solve for any fixed, known defense decision in stage zero) we would aggregate adjacent attacker stages (and adjacent defender stages, if there are any) and reduce the 18-stage BTRA tree to eight stages. We would then require that all decision variables be continuous, and then swap adjacent defender-attacker pairs of stages until we obtain a model having all of the attacker decisions in stage one and all of the defender decisions in stage two. This resulting model is equivalent to model $(A_L D_L)$, above, and hence is equivalent to a simultaneous game.

The optimal solution would prescribe mixed strategies for the attacker and defender, eliminating the sequential nature of the real decisions that must be made. In general the results from such an analysis might not be very accurate, as every relaxation of a block of integer variables to continuous and the subsequent interchange and aggregation of adjacent stages can result in a significant relaxation of attacker restrictions; in some models these approximations could get *significantly* less informative with each additional stage exchanged in this manner.

However, if the sequencing of two adjacent attacker-defender stages is not a critical component of the formulation then the optimal solution of the relaxation might not be far off from that of the original model. As a simple example, if the attacker chooses which pathogen to load into a truck, and the defender then chooses whether or not to emplace transportation blockades, relaxing the decision variables and exchanging these two stages might not be as significant a relaxation as in a situation where the attacker decides whether or not to release a pathogen, and the defender then chooses whether or not to employ his stockpile of a certain vaccine that can treat the attacker's pathogen. In the former case the blockades will work against the truck regardless of the pathogen chosen, while in the second example committing to use the vaccine before a pathogen is released is clearly a bad idea, and allows the attacker to cause significantly more damage.

How to Generalize BTRA to a Decision Model Prescribing Defense Investments

If we are to leverage the considerable effort that went into the development of BTRA, we must use the data obtained, and elicit subject-matter-expert input, to develop a two- or three-stage sequential decision model of defensive investments, attacks, and mitigation responses such that the relaxation obtained by allowing continuous attacker variables, as in MXM, is at least a reasonable approximation.

If we are successful in our new modeling effort, then the decisions at each stage except our new stage-zero will be continuous (and, more specifically, interpreted as mixed strategies), but now the values of these mixed-strategy probabilities will be *prescribed* by the optimization model: for the stage under control of the terrorists, these will represent the *worst-case mix of attack decisions the terrorists can devise*; in the mitigation stage, under DHS control, these will represent the best response to each of the attacker's possible decisions in the previous stages.

It is not lost on us that some of the BTRA probabilistic risk assessment tree's probabilities exhibit dependence on the outcomes of some prior stages in the tree. A reformulation to a two- or three-stage sequential decision model would necessarily require some reworking of this data. For brute-force permutation of (potentially aggregated) stages, we could unwind the conditional probabilities with Bayes Theorem (just as DHS already does when they split the single BTRA tree into 28 independent trees, one for each bioagent, where selection of bioagent is the third terrorist stage in the original tree).

However, we hope to move away from subject-matter-expert elicitations of highly dependent probabilities as follows. These dependencies are presumably due to the influence of prior stages on the state of the terrorist (or DHS) in terms of exhaustion of limited resources. MXM would explicitly guide strategic defensive investment in stage zero, and subsequently offer all the explicit resource-limiting features of a linear program for all the attacker decisions, and in parallel all the defender's mitigation decisions that consume the mitigation resources provided by stage zero. Linear programming has long been widely applied to planning industrial and military operations that precisely mimic a bioterror-agent production program, or a defense plan.

We recommend eliciting from SMEs an explicit assessment of the resources and capabilities of each opponent, and the way and rate at which various alternate activities would consume these. This is, in fact, the way BTRA reports that the SMEs explained their reasoning to support probability assessments. We advise using these technological estimates as explicit inputs, and letting MXM determine attacker mixed-strategy probabilities and expected consequences as outputs. This would be much more transparent modeling, provide better documentation, and be less likely to be influenced by poor SME guesses about high-dimensional decisions governed by complicated resource limitations. This also avoids the current step where SMEs convert capabilities assessments into just a few discrete, qualitative probability classes (e.g., “not likely”=0.2, “likely”=0.5, “very likely”=0.8).

The initial linear integer program, and subsequent pair of linear programs afford us a great deal of flexibility and fidelity in describing the actions of each opponent, and we can solve these at very large scale with off-the-shelf optimization software. Also, solutions to such optimization models can be analyzed to discover the “why” as well as the “what” of each plan. Powerful, effective sensitivity and parametric analysis techniques are well-known for these optimization models.

We represent defensive investment strategy selection simply, as we think realistic and politically palatable during this early phase of homeland security capital planning. We anticipate that this will eventually mature to more closely resemble classic military capital planning [e.g., Brown, et al. 2004].

We present a deterministic model that minimizes the maximum expected risk. If stochastic evaluation proves essential, our model can be used within a simulation. Banks and Anderson [2006] demonstrate such exogenous simulation with a two person, zero-sum game. Tintner [1960] shows this for a linear program. Our integer linear program is amenable to such simulation.

Secrecy in Planning

If, as the defender, we strongly believe that we are able to conceal some of our defensive capability from the attacker, then the transparency of model MXM is likely to be inappropriate for determining optimal defense decisions. Instead, we find ourselves in an *asymmetric* conflict: the attacker and the

defender *do not agree on the objective function*. This more general case falls in the domain of bilevel and multi-level programming (see, for example, Candler and Townsley [1982], Bard and Moore [1992], and Migdalas et al. [1998]), and the associated mathematical models are more difficult to solve than those we have presented here.

In an extreme case, for example, we might believe that even though the attacker can observe our strategic defensive investments, he is completely unaware of our mitigation capabilities. We could then assume that he will make his decisions based only upon the $damage_{d,a}$ values, whereas, given that we are perfectly aware of our mitigation capabilities, we will make our investment decisions based on $damage_{d,a}-mitigate_{d,a,m}$ values. This would be formulated as a tri-level integer programming model, the most general versions of which are difficult to solve. However, a straightforward heuristic for solving our problem would solve an attacker-defender version of the problem with no mitigation options (i.e., by fixing $y_{d,m}=0$), and then choose the optimal mitigation decision for whatever defense and attack decisions are made.

Clearly this can lead to a suboptimal defense investment, especially when there are defense options that do not directly reduce expected damage (i.e., $damage_{d,a}$ might be high for those defenses) but that enable mitigation efforts that are significantly more effective than those available under other defensive investments. We can use the stockpiling of a vaccine as an example; creating the stockpile will not reduce the damage of any attack, but the mitigation activity of distributing the vaccine and inoculating the susceptible population can be extremely effective. In this case, the optimal defense and the resulting worst-case attack damage can differ significantly from the myopic defense. There are other, more effective heuristics for multilevel optimization in the literature, the breadth of which is beyond the scope of this paper.

In the case where the “secret” objective values maintain the same relative ranking between each pair of feasible defense and attack combinations as disclosed by the “public” objective function, then the optimal defense and resulting worst-case attack do not change. For example, if the mitigation effects $mitigate_{d,a,m}$ are always a fixed percentage of $damage_{d,a}$, then the optimal defensive investments, and the

corresponding worst-case attack, will be the same, and the overall expected damage will be reduced by that fixed percentage. In this case, (and similar cases, in which the mitigation efforts do not produce drastically different results from each other relative to the defense-and-attack combination they are applied to), it makes no sense to take extreme measures to conceal our mitigation capability. In fact, we should broadcast it widely, in hopes that it will deter attacker efforts.

However, in the case where our mitigation capabilities are much more (or less) effective for one (or a small number of) attacks than for the rest, and this fact fundamentally changes the worst-case attack decision for each of our defense options, then we conjecture that we should conceal this capability to maintain our advantage (or conceal our weakness) for that attack, and hopefully “shape” the attacker’s decisions towards the attacks that we are more capable of handling. However, every situation is different, and it is extremely hard to predict what the effect any given “secrecy policy” will have on the optimal outcome, much less on the actual attacker behavior. More research in this area is required.

Solving MXM At Very Large Scale with Decomposition

Although we have solved large attacker-defender models of the same *form* as MXM, [Brown et al., 2005b], if instances of MXM become too large to solve using commercial off-the-shelf integer linear programming software, we can use (and have used) a version of Benders decomposition [e.g., Bazaraa, Jarvis, and Sherali, 1990, pp.366-367] to solve MIN-ILP, with integer stage-zero investment decisions and continuous mitigation decisions in the master problem, and the resulting attacker LP subproblems. Israeli and Wood [2002] explicitly develop such a decomposition for the case of shortest-path network interdiction problems.

We modify **MIN-ILP**, replacing equations (DILP1) with a set of constraints (DILP-CUTS), and calling the resulting model **MIN-ILP-DECOMP**($\{\hat{x}^N\}$), where $\{\hat{x}^N\}$ represents the set of all attacker plans from completed decomposition iterations: $\{\hat{x}^N\} \equiv \{\hat{x}^n, n = 1, \dots, N\}$.

$$\mathfrak{R} \geq \sum_{d,a} damage_{d,a} w_d \hat{x}_a^n - \sum_{d,a,m} mitigate_{d,a,m} \hat{x}_a^n y_{d,m} \quad n = 1, \dots, N \quad (\text{DILP-CUTS}).$$

The complete decomposition algorithm is as follows:

Algorithm DHS-MXM-DECOMP

Input: Data for bio-terror defense problem, optimality tolerance $\varepsilon \geq 0$;

Output: ε -optimal (**MXM**) defender plan (w^*, y^*) ;

- 1) Initialize best upper bound $z_{UB} \leftarrow \infty$, best lower bound $z_{LB} \leftarrow 0$, define the incumbent, null (**MXM**) defender plan $(w^* \leftarrow \hat{w}^1 \equiv "d00", y^* \leftarrow y^1 \leftarrow 0)$ as the best found so far, and set iteration counter $N \leftarrow 1$;
- 2) **Subproblem:** Using $\hat{w} = \hat{w}^N$, solve the linear program subproblem **MAX-ATTACKER-LP**(\hat{w}) to determine the optimal attack plan \hat{x}^N ; the bound on the associated total expected target damage is $z_{\max}(\hat{x}^N)$;
- 3) If $(z_{UB} > z_{\max}(\hat{x}^N))$ set $z_{UB} \leftarrow z_{\max}(\hat{x}^N)$ and record improved incumbent **MXM** defender plan $(w^*, y^*) \leftarrow (\hat{w}^N, \hat{y}^N)$;
- 4) If $(z_{UB} - z_{LB} \leq \varepsilon)$ go to **End**;
- 5) Given attack plans $\{\hat{x}^N\}$, attempt to solve master problem **MIN-ILP-DECOMP**($\{\hat{x}^N\}$) to determine an optimal defender plan $(\hat{w}^{N+1}, \hat{y}^{N+1})$.

The bound on the total expected target damage is $z_{\min}(\hat{w}, \hat{y})$;

- 6) If $z_{LB} < z_{\min}(\hat{w}, \hat{y})$ set $z_{LB} \leftarrow z_{\min}(\hat{w}, \hat{y})$;
- 7) If $(z_{UB} - z_{LB} \leq \varepsilon)$ go to **End**;
- 8) Set $N \leftarrow N+1$ and go to step (2) (**Subproblem**);
- 9) **End:** Print “ (w^*, y^*) is an ε -optimal (**MXM**) defender solution,” and halt.

The optimal attacker plan x^* can be recovered by solving **MAX-ATTACKER-LP**(w^*).

Each instance of **MAX-ATTACKER-LP**(\hat{w}) is a linear program of a form we expect to be easy to solve even at large scale.

MIN-ILP-DECOMP($\{\hat{x}^N\}$) is easy to solve, but might get more challenging if embellished with too many more linear constraints. For a difficult instance, or at very large scale, we can solve **MIN-ILP-DECOMP**($\{\hat{x}^N\}$) with an approximate, but very fast heuristic, and our decomposition is still valid.

The iterative behavior of the decomposition is instructive. Set a defense plan, and observe the attack response. Set another defense plan that is robust with respect to the attack response observed, and then observe another attack response. As such iterations continue, the defender learns more about the attacker, and refines his defense plan accordingly. Ultimately, the defender learns enough to declare that his best defense plan is (ε -) optimal against the best possible attacker plan, and attains a mathematical certificate of the quality of his defense preparations. (See Table E 6.)

Iteration	Defense Strategy	Lower Bound	Attack Alternative	Upper Bound	
n	MIN-ILP-DECOMP	z_{LB}	MAX-ATTACKER-LP	z_{UB}	Mitigation
1	“d00”	0	“a03”	10	“m01”
2	“d05”	2	“a01”	5	“m02”
3	“d04”	3.5	“a01”(0.5) “a03”(0.5)	3.5	“m01”(0.5) “m02”(0.5)

Table E 6: Decomposition iterations reveal learning by opponents. Here, the defender starts with defense strategy “d00” (do nothing), the attacker responds with his most-damaging alternative “a03” inflicting damage 10. Subsequent iterations adjust defense strategy based on elicited attacker behavior, until neither opponent can take another turn for any further improvement. Our subject-matter experts (SMEs) are now optimization models. The last iteration yields the same optimal solution as shown in Table E 5. Instead of using a “do-nothing” solution to initialize the algorithm, we can just as easily take any feasible incumbent proposed by any decision maker as our first attempt: the algorithm will evaluate this solution, and then either obtain a certificate of its optimality, or find a better incumbent. This is the distinguishing advantage of viewing these decomposition algorithms as “learning” methods that iteratively improve upon an incumbent, possibly suboptimal, solution.

The decomposition mathematically represents two opposed sets of subject-matter experts: a Blue Team (defender), and Red Team (attacker). The decomposition iterations mathematically mimic a wargame between these opponents, where the defender suffers the disadvantage of not being able to hide his defense strategy, but the players play the game again and again, honing their respective strategies, until neither opponent can improve.

At ultra-large scale, we can nest decompositions. We do not anticipate this will be necessary for this application.

We have implemented MXM and our decomposition algorithm for solving it in GAMS [2007]. All model instances have been solved optimally. The complete implementation is available from the authors.

How do we get here from a descriptive risk assessment (e.g., DHS BTRA)?

First, we must recognize and accept that each event-tree path in BTRA consists almost exclusively of a set of *decisions*—these are *not* random events. There are 18 successive “events” in the NRC rendition of BTRA [NRC 2007, Table E 3-4]. From start to finish, we show each event number, using parentheses to distinguish defender actions, and brackets for Mother Nature at the end: the BTRA event sequence is 1-5, (6), 7-11, (12), 13, (14-15), 16, (17), [18]. The first attacker event sequence addresses selection of agent, target method of dissemination, and acquisition, the next attacker sequence involves details of agent production and processing, the following attacker sequence describes transport and storage, and the last estimates repeated attacks. These attacker sequences are interrupted by opportunities for the defender to interdict. The last stage [18] represents Mother Nature influencing consequences. For our purposes, there are merely four alternations from attacker to defender, followed by one truly random event governed by Mother Nature at the end.

Second, we decide how to reckon $damage_{d,a}$ as a function of defense strategy d and attack alternative a . This is not a glib statement, but rather a meta-design guide to return to the foundations of BTRA and critically review the assumptions of sequence-dependence and level of detail.

In theory, this could be achieved by setting a defender option d , and estimating the consequences of this action on BTRA for each pure attacker response. This is no harder than for BTRA, and if we concentrate on estimating $damage_{d,a}$ as a function of defense option d and a more palatable (i.e., unlike BTRA, a less minutely-detailed and less overwhelmingly numerous) set of attack alternatives a , we would create a risk-calculation engine that is at once credible, and efficient.

By whatever means, we must estimate $damage_{d,a}$ for each defense option d and each attack alternative a . *If we cannot estimate risks at this fidelity, we have no business doing risk analysis.*

We would prefer to be able to choose a number of defense strategies, rather than just one. But, current risk analysis produces a single damage estimate distribution for each attack scenario. We assume these damage estimates are neither additive nor separable between and among attacks, so we must rely on

the simplified risk analysis we have. Accordingly, we endow each defense strategy with the number of defense investment options reflected in each BTRA path.

Our attack alternatives have not specified any particular agent. Our methods can accommodate attacks by classes of agents that include engineered and future agents not yet known.

Solving the tri-level model achieved here isolates an optimal defense strategy, and all its component investment options. Because this optimal strategy dominates every attack by any agent, we have presented an intrinsic risk analysis that highlights the most-critical, achievable defense strategy. We can trivially rule out this best strategy, and solve for the second-best, and so forth. *This renders an explicit, unambiguous prioritization of defense strategies*

Mere probabilistic risk assessment is not enough

What we are proposing here responds directly to the explicit language of HSPD 10:

“the United States requires a continuous, formal process for conducting routine capabilities assessments to guide prioritization of our on-going investments in biodefense-related research, development, planning, and preparedness.”

Further, we could not agree more with this:

“Successful implementation of our program requires optimizing critical cross-cutting functions”

Recently, HSPD 18 [The White House, 2007] has further clarified our direction:

“optimize the investments necessary for medical countermeasures development, and ensure that our activities significantly enhance our domestic and international response and recovery capabilities”

Further:

“Mitigating illness and preventing death are the principal goals of our medical countermeasure efforts.”

Moving beyond mere descriptive risk analysis, we want to address:

- “(a) Target threats that have potential for catastrophic impact on our public health and are subject to medical mitigation;
- (b) Yield a rapidly deployable and flexible capability to address both existing and evolving threats;
- (c) Are part of an integrated weapons of mass destruction consequence management approach informed by current risk assessments of threats, vulnerabilities, and capabilities; and
- (d) Include the development of effective, feasible, and pragmatic concepts of operation for responding to and recovering from an attack.”

We can see from these policy directives that the highest-level DHS problem is *planning investments*—huge investments—to prepare to *mitigate* the consequences of any attack.

The material presented here follows both the letter and the spirit of this direction.

REFERENCES

- Banks, D. and Anderson, S., 2006, “Combining Game Theory and Risk Analysis in Counterterrorism: A Smallpox Example,” Statistical Methods in Counterterrorism (Wilson, A., Wilson, G., and Olwell, D., eds., Springer, New York, pp. 9-22.
- Bard, J., and J. Moore, 1992, “An Algorithm For the Discrete Bilevel Programming Problem.” *Naval Research Logistics*, **39**, pp. 419-435.
- Bazaraa, M.S., J. Jarvis, and H.D. Sherali, 1990. *Linear Programming and Network Flows*. John Wiley & Sons. New York, NY.
- Brown, G., Carlyle, M., Salmerón, J. and Wood, K., 2006a, “Defending Critical Infrastructure,” *Interfaces*, 36, pp. 530-544.
- Brown, G., Carlyle, M., Salmerón, J. and Wood, K., 2005a, “Analyzing the Vulnerability of Critical Infrastructure to Attack, and Planning Defenses,” in Tutorials in Operations Research: Emerging Theory, Methods, and Applications, H. Greenberg and J. Smith, eds., Institute for Operations Research and Management Science, Hanover, MD.

- Brown, G., Carlyle M., Harney R., Skroch E., Wood, K., 2006b, "Anatomy of a Project to Produce a First Nuclear Weapon," *Science and Global Security*, **14**, pp. 163-182.
- Brown, G., Carlyle, M., Diehl, D., Kline, J. and Wood, K., 2005b, "A Two-Sided Optimization for Theater Ballistic Missile Defense," *Operations Research*, **53**, pp. 745-763.
- Brown, G., Carlyle, M., Harney, R., Skroch, E. and Wood, K., 2007, "Interdicting a Nuclear Weapons Project," *Operations Research* (to appear).
- Brown, G., Dell, R. and Newman, A., 2004, "Optimizing Military Capital Planning," *INTERFACES*, 34 (6), pp. 415-425.
- Candler, W., and R. Townsley, 1982, "A Linear Two-level Programming Problem." *Computers & Operations Research*, 9, no. 1, pp.59-76.
- Department of Homeland Security, 2006, Bioterrorism Risk Assessment (U), Biological Threat Characterization Center of the National Biodefense Analysis and Countermeasure Center, Washington D.C.
- Fulkerson, D.R., and G.C. Harding [1977], "Maximizing the Minimum Source-Sink Path Subject to a Budget Constraint." *Mathematical Programming*, **13**(1). pp. 116-118.
- GAMS, 2007, "General Algebraic Modeling Language GAMS," homepage. <http://www.gams.com/>, accessed 12 January 2007.
- Golden, B. [1978], "A Problem in Network Interdiction." *Naval Research Logistics Quarterly*, 25(4), pp. 711-713.
- Israeli, E. and Wood, R.K., [2002], "Shortest-Path Network Interdiction," *Networks*, **40**, pp. 97-111.
- Kuhn, H., 1953, "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games*, H. Kuhn and A. Tucker, eds., Vol. II, Princeton University Press, Princeton, NJ, pp. 193-216.
- Migdalas, A., P. M. Pardalos, and P. Varbrand, 1998. *Multilevel Optimization: Algorithms and Applications*. Kluwer. Dordrecht, Germany.

- National Research Council, 2007, Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis (final report) (in review).
- Raiffa, H., 1968, Decision Analysis, Addison-Wesley, Reading, Ma.
- Salmerón, J., Wood, K. and Baldick, R., 2004, "Analysis of Electric Grid Security Under Terrorist Threat," IEEE Transactions on Power Systems, 19(2), pp. 905-912.
- The White House, 2004, "Homeland Security Presidential Directive 10, Biodefense for the 21st Century," <http://www.whitehouse.gov/homeland/20040430.html>. accessed 14 July 2007.
- The White House, 2007, "Homeland Security Presidential Directive 18, Medical Countermeasures Against Weapons of Mass Destruction," <http://www.whitehouse.gov/news/releases/2007/02/20070207-2.html>, accessed 14 July 2007.
- Tintner, G., 1960, "A Note on Stochastic Linear Programming," *Econometrica*, 28(2), pp. 490-495.
- Ville, J., 1938, "Sur la theorie generale des jeux ou interviennent l'habilité des joueurs," in Borel , E., et al. [eds.], *Traite du calcul des probabilites et de ses applications*, Vol. II, pp. 105-113.
- von Neumann, J., 1928, "Zur Theorie der Gesellschaftspiele," *Mat. Ann.* 100, pp. 295-320.