Advanced Techniques for Modeling Terrorism Risk

JOHN A. MAJOR

JOHN A. MAJOR
is a senior vice president at Guy Carpenter & Company, Inc., in New York, NY.
john.major@guycarp.com

The attack of September 11, 2001, showed that terrorism is capable of inflicting damages ($40bn+) and loss of life (3,000+) that are multiples of the worst U.S. natural disasters (Hurricane Andrew at $20bn+ and 40-60 lives; Northridge earthquake at $12.5bn and approximately 25 lives) (Bassett and Schroeder [1998]; Mooney [2001]; Pawlowski [2001]). In the wake of this unprecedented disaster, insurers and reinsurers have been excluding terrorism risk from their offerings, with grave consequences for commercial property owners and lenders (Mooney [2001]). Most insurance analysts and actuaries would agree with Munich Re's Christian Kluge: "There is no mathematical model for terrorism" (Fromme [2001]). But the need for one is clear.

Terrorism risk shares features with other forms of catastrophe risk, including a time series of historical events, yet goes beyond them with an extra layer of impenetrability. Defensive studies of terrorism risk resemble risk analyses of complex engineering systems (nuclear power plants, satellite launches, etc.). A particular scenario can be analyzed in terms of the probability of failure of critical subsystems. However, unlike natural disasters, it features human intelligence, and unlike industrial disasters it features human intent. To quantify the risk, much like solving the celebrated "three doors" puzzle (*see Appendix A*), *probability is not enough*. Methods from operations research, including game theory and search theory, as well as certain specialized areas of statistics, may well be needed to construct an adequate modeling framework.

AN ILLUSTRATIVE MODEL

To explore the possible application of game theory to the modeling of terrorism risk, let us consider the following simplified model:

- There is a set of targets indexed by the letter *i*, numbered from 1 to N. Each target *i* has a value V_i.
- An attacker, with total resources A_T, must choose a target and how much resource A_i to assign to it.
- A defender, with total resources D_T, must decide how to allocate resources D_i among the targets.
- The total destruction of target *i* occurs with probability given by a function $p(V_i, A_i, D_i)$.
- The attacker wants to maximize, and the defender wants to minimize, the expected loss EL which is given by the formula:

$$EL = \Sigma_i V_i \cdot p(V_i, A_i, D_i) \qquad (1)$$

We can be justified in using expected value as the criterion if we consider the values V_i to represent *utilities*—and as long as both sides have the same utilities.

In the parlance of game theory (*see Appendix B*), we would say this is a *zero-sum game* with *payoff* EL to the attacker. The attacker's strategy options consist of a choice of target and the assignment of a resource to it.[1] The defender's strategy options consist of the simultaneous assignment of resources to all N targets.[2]

Before going on to consider specific functional forms for the probability function, is there anything meaningful to be said? Indeed, there is.

First, if the probability of a successful attack goes down with an increase of applied defensive resources (and this ought to be true under any plausible formula for *p*), then the defender should use all the defense resources. There is no reason to hold back, because a quantity of unused resource could be applied to reduce the success probability and hence the EL of at least one target.

Second, the *minimax criterion* reveals a solution for the defender. Not knowing how the attacker will choose targets, the defender should choose a strategy that results in the lowest possible worst-case EL, regardless of which target the attacker selects. That implies the resulting EL among (defended) targets will be equal (and the EL among undefended targets will be less). Why? Imagine one target defended in such a way that its EL (if attacked) is greater than another. Then it pays to shift defense resources from the lower-EL target to the higher one, until they have equal ELs (lower than the original high EL). The more valuable targets should thus be equalized in terms of their expected losses. Less valuable targets may be left undefended, because an attack there, even if successful with 100% certainty, will result in a loss that is less than the EL of the defended targets. Call this equilibrium expected loss EL^{o}.

Third, this leaves the attacker with a set of, say, M defended targets, where the best he can do is achieve that same $EL = EL^{o}$ among any one of them. For the undefended targets, he can only achieve their value V, which is less than EL^{o}. His best strategy (a *mixed strategy*) is to choose one of the M targets at random—but with what relative probabilities?

If the attacker can observe that the defender has allocated resources in the optimal fashion, then he is truly indifferent as to the choice of target, and the assignment of probabilities is indeterminate. The usual game-theoretic formulation, on the other hand, assumes that both sides must make their strategic choices with no knowledge of the other side's choice. If the defender might not be allocating resources optimally, then the attacker should use probabilities that protect him from doing any worse than EL^{o} *on average*. This implies that each target's selection probability should be inversely proportional to the marginal effectiveness of defense (the rate of change of EL with respect to changes in D) at that target. (This is taken up in more detail in Appendix D.) If *p* varied (approximately) linearly with D,[3] for example, then selection probability would be (approximately) inversely proportional to target value: $q_i = k/V_i$.

Fourth, given the above, and if we knew what that equilibrium EL^{o} value was and could work out the pattern of attacker selection probabilities q_i, we could then derive the overall probability distribution of loss from an attack. Each of the M defended targets has a q_i probability of being attacked, and each such target *i*, if attacked, has a probability of losing its value V_i equal to $p_i = EL^{o}/V_i$.

Note, nowhere did we need to refer to the central limit theorem, independent increments, extreme value theory, or any other such probabilistic assumptions and tools that are often applied to the study of nature or complex processes. The conclusions flowed from our assumptions about human intentions and rational[4] behavior.

A NUMERICAL EXAMPLE

In this section, we use a particular function $p(V_i, A_i, D_i)$ which allows us to engage in specific computations and carry out a numerical example. The function is:

$$p(V_i, A_i, D_i) = \exp\left(-\frac{A_i \cdot D_i}{\sqrt{V_i}}\right) \cdot \left(\frac{A_i^{\,2}}{A_i^{\,2} + V_i}\right) \qquad (2)$$

This consists of two terms multiplied together. The first term represents the probability of a planned attack escaping detection. The second represents the probability of the attack succeeding in its technical execution, given that it goes undetected. (See Appendix C for details.)

Note that in the extreme of no defense (D = 0), the probability of escaping detection (represented by the first term) is 100%. The probability of success is then a matter of having assigned attack resources sufficiently greater than the square root of V to bring the second term near to 100% as well. An assignment of the full budget A_T of attack resources is optimal. However, if there is some defense assigned, this may not be the case. Past a certain point, attack resources become counterproductive because they raise the probability of detection faster than they raise the probability of technical success. This is illustrated in Exhibit 1 for a target with value of 1.

EXHIBIT 1

Increasing Attack Resources Improves Success Probability When There Are No Defenses But Can Be Counterproductive in the Face of Defenses

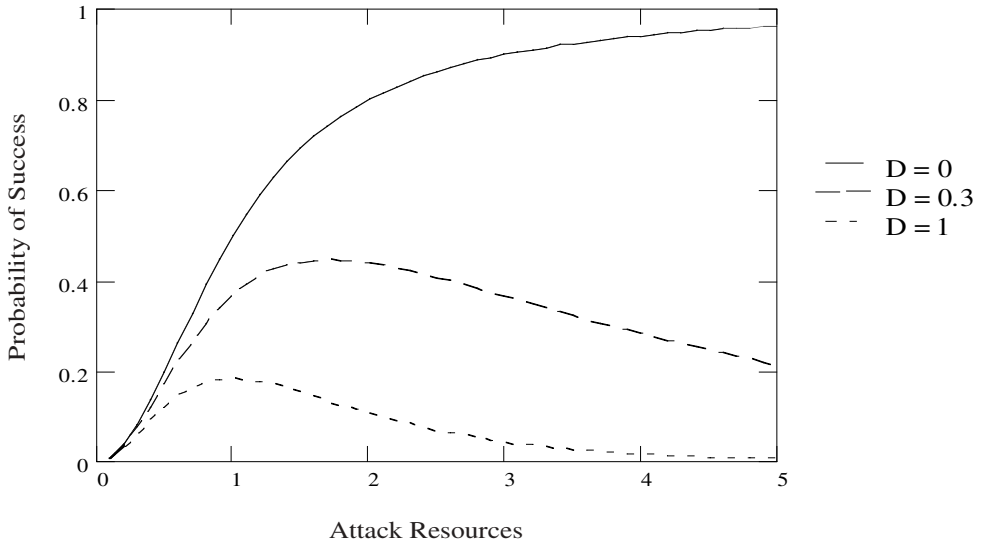

For a given value V and level of defense D, there is an optimal attack resource, call it A^o, which may or may not be greater than the attack budget A_T. In assigning defense resources, the defender needs to know what that attack budget is in order to compute the equilibrium EL^o. However, if the attacker has very large resources, or has the potential for increasing them, it may suffice to assume unlimited attack resources, $A_T = \infty$. In that case, only the optimal attack resource levels A^o matter. We may consider this a "conservative" approach to defense resource assignment.

Consider the following example: There are 20 targets, each of which has value 1.5 times the one preceding it, with the biggest target having a value of 1. The attacker has essentially unlimited resources and the defender has 20 units to be allocated among the targets.

What are the optimal defense and attack strategies, the resulting equilibrium expected loss, and the resulting probability distribution of losses?

Exhibit 2 depicts the target values as the diagonal line of "+" symbols with dots in between them, with the biggest target being #20 on the right-hand side, having value 1 (topmost value in the vertical scale).

At this point, we need to justify why a pure strategy, and not a mixed strategy, is optimal for defense. In this numerical example, it follows from the convexity of *p* in *D*. Any mixed strategy for defense is dominated by its "average" pure strategy; therefore no mixed strategy can be optimal. The optimal defense strategy is depicted in the square boxes connected by solid lines. To fit on this scale, the defense numbers were divided by 10. Thus, the defense allocation for target #20 is really 5.2, not 0.52. Notice that only the most valuable 10 targets (#11-20) are defended. The smaller targets are left with zero defense allocation.

Optimal attack resources are traced by the *x*'s in the upper right. They are mostly between 0.3 and 0.35, being constrained not by total attack resources but by the need to escape detection by defense.

The resulting equilibrium EL is 0.018, which is maintained for all defended targets. This is shown in the trace of diamond symbols in the lower half of the exhibit. Undefended targets have EL less than the equilibrium EL because their values, which they can lose with 100% probability, are lower. That is why the EL curve coincides with the values for targets #1-10.

Optimal attack probabilities are approximately proportional to the square root of the value of the target. (See Appendix D for derivation.) These are shown in Exhibit 2 as dashed lines.

As outlined previously, for defended targets, the probability of an attack being successful is EL/V. Combining these results, if an attack is attempted, the probability dis-

EXHIBIT 2

Target Attack-Defense Example (explained in text)

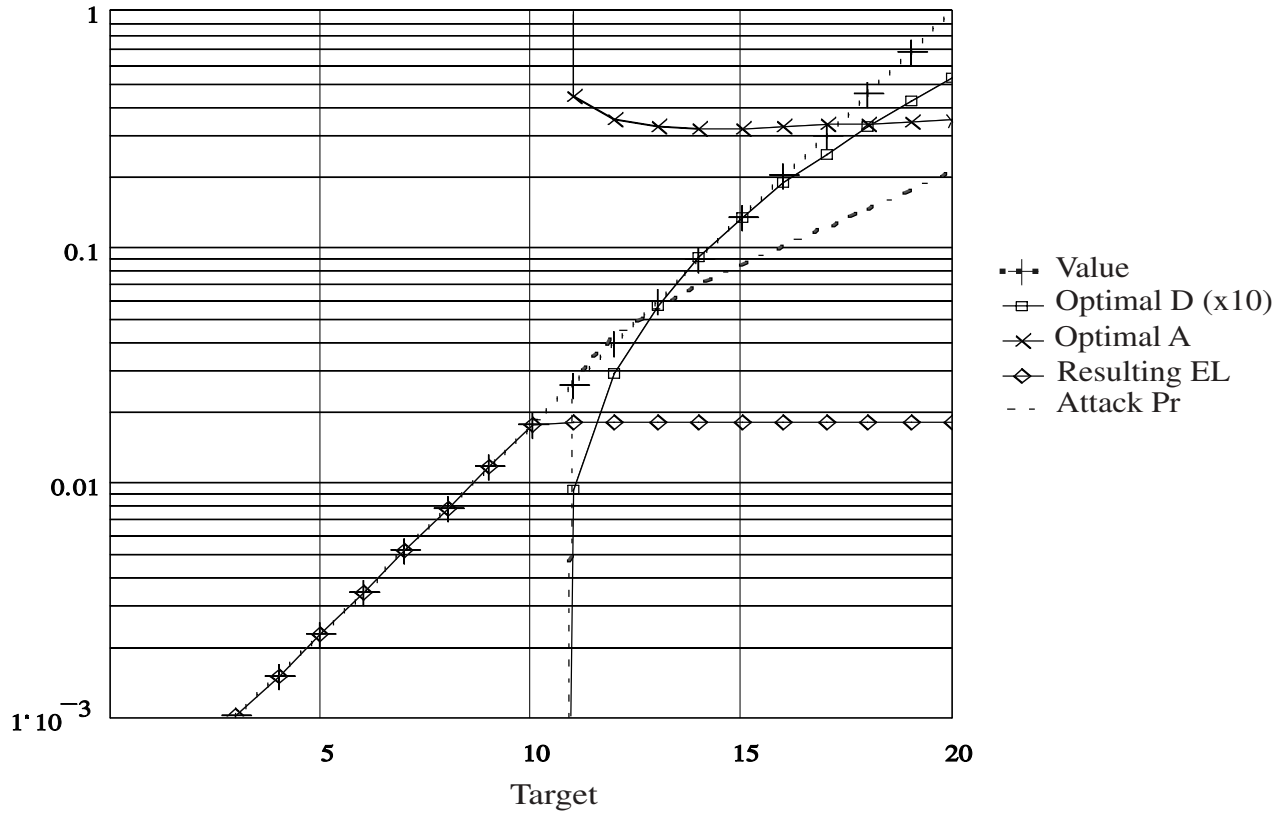

EXHIBIT 3

Probability Distribution of Losses

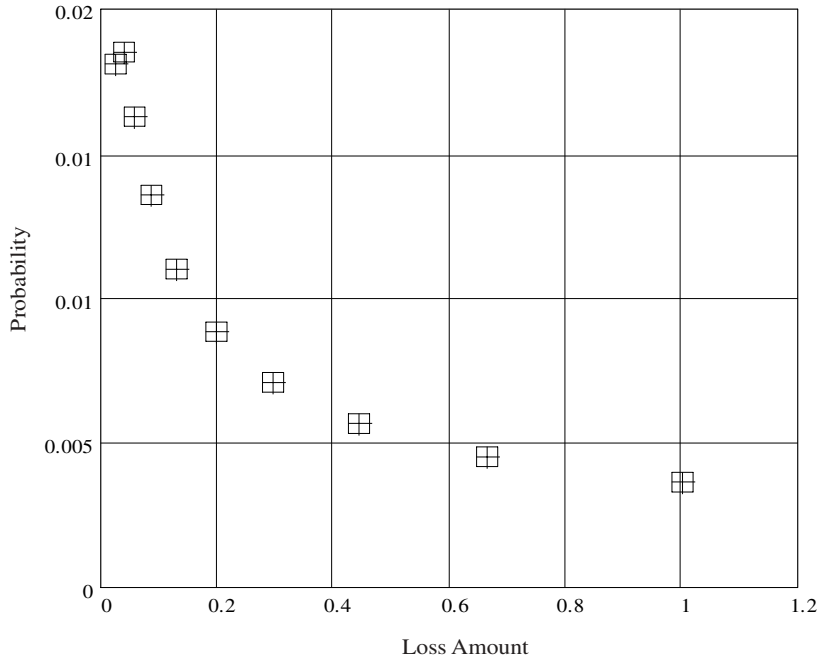

tribution of losses is as presented in Exhibit 3. There is an overall 89.3% probability that the attack is not successful. The remaining 10.7% is not spread evenly over the defended target values, however; it is more likely to be a successful attack on a smaller target. The probability of the largest loss (a value of 1 from the largest target) is only 0.36%.

FROM ILLUSTRATION TO USEFULNESS

The previous exposition was intended to give a sense of what kind of reasoning and analysis would be required for a risk model appropriate to the terrorism hazard. In order to make such a model useful for actuarial purposes, an enormous amount of work still needs to be done.

First, the model as presented is really only a severity model. It says, *if* an attack is attempted, *then* losses will occur with such and such probability. It says nothing about how often an attack will be attempted, which is the role of a frequency model. Gordon Woo [2002] discusses in depth how the characteristics of a terrorist organization, in particular its organizational structure, influence the frequency of planned attacks.

How well does the model represent reality? This is the central concern for any model. While remarkably simple models are often very useful (think of lattice models in finance), the goal is to capture some essential truth about the situation being modeled. Here are some issues that need to be addressed in this regard:

- Does it make sense to consider defense resources as being assigned to targets? In reality, considerable counterterrorism resources are devoted to intelligence gathering and other non-target-specific activities. This model completely ignores them. Is that acceptable?
- To be applied, a model needs to have its components operationally defined. Specifically, how are we to measure the value of a target? Dollars are the usual measure of value, but perhaps lives or even media air time are more important to the attacker. Assigning "utilities" is even more problematic. How to measure attack and defense resources? Again, a monetary unit is a candidate, but number of people devoted to the task ("FTE" in human resources jargon) might be as good or better.
- Where does one acquire a realistic list of targets and values? The major catastrophe modeling firms have inventories of commercial and residential building stock in the U.S., as well as infrastructure information (bridges, tunnels, port facilities, etc.). These huge databases might be able to supply a realistic distribution of target values. An alternative is to estimate the distribution from aggregate statistics, say by a multifractal allocation technique (Lantsman et. al. [1999]; Major and Lantsman [2001]). This could obviate a large amount of computational effort.
- The model's defense resource constraint must be interpreted as total societal resources devoted to guarding and protecting valuable properties. This goes well beyond what can be read in the newspaper as the latest appropriation from Congress for military homeland defense, because it includes state and local civilian police forces as well as private-sector security resources.
- An attack budget[5] is the most speculative item, requiring, for most accuracy, information that might be classified, and therefore unavailable to the private sector analyst.
- What if the attacker and defender do not have the same utilities? Then we are out of the realm of *zero-sum* games and need to look to the *Nash equilibrium* for a solution. (*See Appendix B.*) This could complicate the analysis considerably.
- What if the optimal defense cannot be executed perfectly, or even approximately? What is the attacker's optimal strategy if defense weaknesses are known? If they can only be known after resources are expended searching for them? Analytical approaches to these questions could become quite complex. Paul Kleindorfer[6] suggested an intelligent-agent or cellular-automaton model where attackers move through a grid seeking out attack opportunities. Such an approach could take us out of the realm of analytical models altogether and into full-blown simulation. Weaver et al. [2001] discuss a highly detailed simulation approach that incorporates realistic terrorist scenarios and the use of non-zero-sum game theory.
- The model considers one attack at a time, applying the attacker resource constraint to each attempt. In reality, the attacker's resources are dynamic, being acquired and expended over time. Even in this simple model, it often turns out to be to the attacker's advantage to launch two less-well-funded attacks simultaneously than one optimally-funded attack. This expands the attacker's options and complicates the analysis.[7]

- Given a believable structure to the model, how is one to fit parameters? Clearly, analysis of historical data will play a role. Even more so than in hurricane and earthquake studies, however, the need for keen insight and understanding (read: expertise), guiding the selection and adjustment of data, is acute.

Even more so than in the case of hurricane and earthquake modeling, it cannot be overemphasized that building a model such as the one outlined above is an exercise in futility at best (and self-delusion at worst) without adequate input from terrorism experts. As Woo [2002] puts it, "Any probabilistic framework for quantifying terrorism risk, however logically designed, will ultimately have to involve a measure of expert judgement."

CONCLUSION

We saw how the terrorism risk differs in kind from natural and man-made (accidental) catastrophes because of the elements of intelligence and intent. As a consequence of that, in modeling the terrorism risk, probability is not enough. Analysis techniques borrowed from wartime operations research, especially game theory, are at least as valuable as the trusted standbys of convolution and Poisson distributions.

A highly simplified model was presented, revealing some counterintuitive maxims about defending targets, and showing a direct, if perhaps surprising, route to the probability distribution of losses.

Numerous issues standing between the illustrative model and a truly usable terrorism risk model were outlined. Access to terrorism expertise is a crucial ingredient. Despite the drastic simplification involved, however, a model very much in the spirit of the one presented here has the potential to offer useful insights to the insurance profession.

APPENDIX A
The Three Doors Puzzle

> Suppose you're on a game show, and you're given the choice of three doors. Behind one door is a car, behind the others, goats. You pick a door, say number 1, and the host, who knows what's behind the doors, opens another door, say number 3, which has a goat. He says to you, "Do you want to pick door number 2?" Is it to your advantage to switch your choice of doors?

This puzzle appeared in *Parade* magazine's "Ask Marilyn®" column a few years back and generated enormous controversy among people with extensive knowledge of probability theory (Weiner [1995]; vos Savant [1996]). Most agreed that the probability of having the right door was 1/3 before the host opened the second door. After that, two conflicting answers dominated.

One camp argued that simple common sense (or, if one insisted, the application of conditional probability calculations) showed that the probability was split 50-50 between the remaining two unopened doors, so there was no advantage to switching. The other camp argued that simple common sense (or, if one insisted, an enumeration of all possible outcomes) showed that since the probability of holding the right door was 1/3, the host's action did not change that probability, so the remaining door had a 2/3 probability of revealing the prize and it was advantageous to switch.

To understand the root of the controversy, one must first realize that (under the conventional "frequentist" interpretation) probability only makes sense as the proportion of outcomes in a large number of similar situations. One must imagine playing the game over and over in order to attach a concrete meaning to the concept of probability. In playing the game over and over, one must ask, *How is the game show host playing this game*?

Those in the first camp, applying the rules of conditional probability, are assuming that the host chooses one of the two doors at random. His selection *in this particular instance* happens to reveal a goat. If this is how the game is played, then it is true that the probability is one-half. Those in the second camp assume that the host always chooses a door with a goat behind it. If that is the case, then the last door does indeed have a 2/3 chance of revealing the prize. Since this is the "conventional" interpretation of the game,[8] it was the answer favored by Ms. vos Savant.

Notice however, that the problem *as stated* did not specify how it was that the host came to open the door. Indeed, we could imagine a range of possible strategies on his part. He could first flip a coin or roll dice to decide which of the two strategies imagined by camps 1 and 2 he would use, for example. The probabilities would change accordingly.

Indeed, it isn't necessarily the case that the host would always open another door. What if he only gave you the opportunity to switch if you hadn't picked the right door? Then switching would give you a 100% chance of winning. Conversely, if he only offered you the switch when you had picked the right door, then switching would guarantee a loss. If we don't know how the host is playing the game, then *probability is not enough*. We need to analyze the problem from the perspective of *game theory*. (*See Appendix B.*)

A game theory analysis might approach this as a *two-person, zero-sum game* with the *payoff* being the probability of obtaining the prize. You want to maximize the probability and assume, defensively, that the host wants to minimize it.[9] In this

interpretation, the solution of the game consists of a *saddle point*. Your strategy, maximizing your probability of winning regardless of which strategy the host is following, it is to hold on to your original door. This guarantees you have (at the very least) a 1/3 probability of winning the prize.

APPENDIX B
Game Theory

The development of game theory is credited to the mathematician John von Neumann,[10] with his first paper published in 1928 and first extensive book treatment (von Neumann and Morgenstern [1944]) sixteen years later. A quite accessible introduction to the theory is presented in Williams [1954].

Quade [1966] writes, "The theory of games is a mathematical treatment of planning under conflict.... [It] does not cover all the diverse factors which enter into behavior in the face of a conflict of interest. There are certain important limitations.... But ... its contributions to policy analysis are possibly far greater [than linear programming] for it tells us how to think about situations of conflict...."

A *two-person, zero-sum game* is conceptualized as a matrix where each row represents the choice of action (strategy) that player 1 can make, each column the strategy for player 2, and the entries in the matrix representing a payoff from player 1 to player 2. (That every point lost by one player is gained by the other is the reason for the term "zero-sum.")

A foundational concept in game theory, the *minimax criterion* holds that each opponent should act in such a way as to *mini*mize his *max*imum losses (as long as this cannot be exploited by the opponent). This is quite a conservative principle, pessimistically assuming that the opponent will play as best he can, and avoiding unnecessary risks, no matter how attractive.

When the minimax strategy, applied to both sides, leads to a stable choice of strategies, the single resulting outcome is known as a *saddle point*. When a game does not have a saddle point, then the players can do better by using more than one strategy. In those cases, there will be an optimal probability distribution they should follow in randomly choosing which strategy to use. The "minimax theorem" of game theory asserts that when such *mixed strategies* are allowed, there is always a stable solution.

When there is more than one player, or the payments are not constant sum (e.g., there are outcomes where both lose or both gain), the theory becomes much more complex (McCain [1999]). Players may cooperate (collude) by forming coalitions. There may be side-payments between players. Multiperson game theory is often used to analyze foreign policy or economic behavior.

In non-constant-sum games, the concept of minimax often does not apply, and the previous analysis techniques do not work. Mathematician John F. Nash, Jr.[11] (Nash [1994]; O'Connor and Robertson [2001]) invented the concept of *Nash equilibrium*. If there is a set of strategies with the property that no player can benefit by changing strategy while the other players do not, then that set of strategies (with corresponding payoffs) is known as the Nash equilibrium. This is the idea of Pareto optimality applied to game theory. Some games have multiple Nash equilibria, making them quite "interesting."

C. J. Hitch (in Quade [1966]) comments on the counterintuitive conclusions that game theory brings to the question of defending targets:

> Suppose you have your defenses deployed as well as you can. Now you get more defenses. How do you deploy them? Well, my intuition told me (and so did most people's) that you deploy them mainly to protect additional targets ... that you did not previously have enough stuff to defend. Game theory says no. You use additional defenses mainly to increase the defense of targets already defended. In fact, over a wide range, the more you have, the more you concentrate it.

We do see something quite like this in our simplified model, where an increase in defense resources is allocated mostly to already heavily-defended targets.

APPENDIX C
A Success Probability Function

What influences the success of a terrorist attack? Woo [2002] cites two principal obstacles: detection (whether through prior intelligence or detection at the time of the attack) and technical or logistical shortcomings. We model these separately and combine them in the end.

Search Theory

To model detection, we will take our cue from search theory (Frost [1999]).

Search theory can be traced back to the work of B.O. Koopman [1946] in World War II. Then, the object of concern was detecting enemy submarines. Subsequently, much of the development was aimed at search and rescue (SAR) operations. Here, we take a simplified approach to the question of terrorist detection.

For our model, we consider D defenders (guards, say) patrolling a target (a building) and A attackers (terrorist infiltrators) entering the area. Abstract this to points placed on a grid. Say there are G grid locations. If the defenders and attackers are randomly placed on the grid, what is the probability that a type A point and a type D point end up at the same grid location?

Start with D = A = 1. The probability is clearly equal to 1/G, because there is only one of the G locations where the defender is, and that is the one out of G chance that the attacker has of coinciding with the defender.

If D > 1 and A = 1, it is tempting to think the answer is D/G, but it is a bit more complicated because the defenders might not all be located on distinct grid locations. The answer is really $1 - (1 - 1/G)^D$, reflecting the fact that each of the D defenders has an independent 1/G chance of coinciding with the attacker. This is the complement of the probability that *all* D defenders independently miss the attacker, $(1 - 1/G)^D$.

Similarly, in the general case that D > 1 and A > 1, the probability that the attack goes undetected is equal to $(1 - 1/G)^{A*D}$. This represents the conjunction of the A independent events of all defenders missing a particular attacker.

We go on to assume that the "size of the search space" G is equal to the square root of the value of the target—while more valuable (read: bigger) targets need more defenders, it should not go up linearly. Further, we use the well-known exponential approximation and set:

$$\Pr\{\text{Escape Detection}\} = \exp(-A \cdot D/V^{1/2}) \qquad \text{(C-1)}$$

Dose-Response Models

For technical/logistical shortcomings, we take inspiration from dose-response modeling (Derr [2000]).

An important application of biostatistics is the analysis of life-or-death responses to drug toxicity or other treatment protocols. Say outcomes are represented by Y = 0 for survival and Y = 1 for death. Let π_i represent the probability that $Y_i = 1$ where *i* indexes the experimental subjects. Let x_i represent a (possibly vector-valued) treatment variable (e.g., amount of drug administered). A model frequently used to relate treatments to outcomes is the *linear logistic model*:

$$\ln(\pi_i/(1 - \pi_i)) = \alpha + x_i \cdot \beta \qquad \text{(C-2)}$$

where ln( ) is the natural logarithm and α and β are parameters to be fit to observed data. The expression on the left-hand side is known as the *log-odds* or *logit*. Notice how it goes to infinity as π_i goes to 1 and to negative infinity as π_i goes to zero.

For our application, we take Y = 1 to represent the success of an undetected attack and use ln(A) and ln(V) as components of the treatment variable. Since we are not fitting data, but rather building an illustrative model, we are free to use whatever values of α and β we choose. Specifically, we choose $\alpha = 0$ and $\beta = (2, -1)$ for the model

$$\ln(\pi_i/(1 - \pi_i)) = 2 \cdot \ln(A_i) - \ln(V_i) \qquad \text{(C-3)}$$

Algebraically rearranging this, we get

$$\Pr\{\text{Success} \mid \text{Escape Detection}\} = \pi = (A^2/(A^2 + V)) \qquad \text{(C-4)}$$

Note that zero attack resources result in zero success probability, and that success probability approaches 100% asymptotically as attack resources increase without bound. In dose-response parlance, the *LD50* refers to the "lethal dose" at which 50% of the subjects die. Here, the LD50 for attack success is the square root of V. As in the case of detection, we have chosen a relationship that puts A in the same scale as the square root of V. While a more valuable target (bigger building) may need more resources for a successful attack (bigger bomb), that relationship should not go up linearly.

The Success Probability Function

Putting the pieces together, we get an expression for the probability of a successful attack:

$$\Pr\{\text{Success}\} = \Pr\{\text{Escape Detection}\} \cdot \Pr\{\text{Success} \mid \text{Escape Detection}\} \qquad \text{(C-5)}$$

$$p(V_i, A_i, D_i) = \exp\left(-\frac{A_i \cdot D_i}{\sqrt{V_i}}\right) \cdot \left(\frac{A_i^2}{A_i^2 + V_i}\right) \qquad \text{(C-6)}$$

APPENDIX D
Optimal Attack Probabilities

The usual game-theoretic formulation assumes that both sides must make their strategic choices with no knowledge of the other side's choice. This is realistic in our model, because even if the attacker has some knowledge of defense allocations, it is likely to be imperfect. The principle of minimax applies here in the form of an imperative to guarantee the best of "worst-case" *average* results, regardless of what the defender does.

If q_i denotes the probability that target *i* is attacked, then the average EL is given by:

$$EL(\delta) = \sum_i q_i \cdot V_i \cdot p\left(V_i, A_i^o, D_i^o + \delta_i\right) \qquad \text{(D-1)}$$

where δ is a vector of defense allocation perturbations summing to zero and the q_i are zero for targets #1-10, but otherwise sum to one. A first-order Taylor approximation gives:

$$EL(\delta) \approx EL^o + \sum_i q_i \cdot V_i \cdot \frac{\partial}{\partial D_i} p\left(V_i, A_i^o, D_i^o\right) \cdot \delta_i \qquad \text{(D-2)}$$

Therefore, a mixed strategy (choice of q_i values) that "immunizes" the attacker against perturbations in the defense allocation would set the summation to zero for any such δ. This is accomplished by equating the coefficients of δ_i, setting

$$q_i = \frac{-k}{V_i \cdot \frac{\partial}{\partial D_i} p\left(V_i, A_i, D_i\right)} = \frac{k}{V_i \cdot A_i^o \cdot EL^o \cdot V_i^{-\frac{3}{2}}} = \frac{k \cdot \sqrt{V_i}}{A_i^o \cdot EL^o} \tag{D-3}$$

where k is a normalizing constant. Since the optimal attack assignments are approximately constant in this example, the optimal attack probabilities are approximately proportional to the square root of the value of the target.

This differs qualitatively from the result in Gross [1950] where target selection probability is *inversely* proportional to value. The reason for this is that Gross's model, in effect, sets $p = \max(0, A - D)$ with the marginal effectiveness of defense $V*\partial p/\partial D$ equal to V. In our model, however, high-value targets with larger defense allocations (and low success probabilities) are closer to "saturation"—a large increment to D can only reduce the probability by a small amount, making $V*\partial p/\partial D$ decrease with increasing V.

ENDNOTES

[1]Technically, a point in the space $\{1, \ldots, N\} \times (0, A_T]$.

[2]Technically, a point in the bounded simplex $\{D_i: \Sigma D_i \leq D_T\}$ embedded in $[0, D_T]^N$.

[3]As it does in Gross [1950].

[4]Here, rationality refers to using the appropriate means to achieve a desired end. It is not intended to imply that the desired ends are reasonable.

[5]Budgets, really, because there are multiple terrorist organizations at work.

[6]Private communication.

[7]The attacker's strategy options change from the one-dimensional $\{1, \ldots, N\} \times (0, A_T]$ to the N-dimensional simplex in $[0, A_T]^N$.

[8]The problem as stated is not how the real *Let's Make A Deal* (1963-present) television game worked. In the real game, Monty Hall started with two players choosing different doors, and he opened one player's door to reveal a non-prize, offering the other player a chance to switch. In this circumstance, the probabilities are the reverse of what they are in the stated problem!

[9]Another possibility, however, is to approach the problem as a non-constant-sum game where the host has some incentive, or at least not symmetrical disincentive, to let the player win. However, that would require the use of unstated assumptions, too.

[10]Von Neumann is also remembered as one of the pioneers of computer science and for contributing to the development of the hydrogen bomb.

[11]In four papers between 1950 and 1953, Nash made seminal contributions to both non-cooperative game theory and to bargaining theory. Four decades later, he was awarded the 1994 Nobel Prize in Economic Science for his work on game theory. He is the central character of the Universal Studios motion picture *A Beautiful Mind*. (Von Neumann appears as a character in the movie, too.)

REFERENCES

Bassett, D., and A. Schroeder. *The Climate Canary Report*. U.S. Department of Energy. 1998. (Available at http://www.supramics.com/climate/statement.html)

Derr, R.E. "Performing Exact Logistic Regression with the SAS® System." Cary, NC: SAS Institute, Inc., 2000. (Available at http://www.sas.com/rnd/app/papers/exactlogistic.pdf)

Fromme, H. "Munich Re in Favour of German Terror Pool." Alexander Forbes Group, 2001. (Available at http://www.alexanderforbes.co.uk/afuknews/PressArchive/10182001MunichReInFavourOfGermanTerrorPool.htm)

Frost, J.R. "Principles of Search Theory, Part I: Detection." Soza & Company, Ltd., 1999. (Available at http://www.sarinfo.bc.ca/Library/Planning/ SrchThy1.doc)

Gross, O. "n Targets of Differing Vulnerability with Attack Stronger Than Defense." U.S. Air Force Project RAND Research Memorandum RM-359. The RAND Corporation, 1950.

Koopman, B.O. "Search and Screening." OEG Report No. 56. The Summary Reports Group of the Columbia University Division of War Research, 1946. (Available from the Center for Naval Analyses.)

Lantsman, Y., J.A. Major, and J.J. Mangano. "On the Multifractal Distribution of Insured Property." Guy Carpenter & Company, Inc., 1999. (To appear in *Fractals*, World Scientific Publishers.)

Major, J.A., and Y. Lantsman. "Actuarial Applications of Multifractal Modeling Part I: Introduction and Spatial Applications." CAS Forum, Casualty Actuarial Society, Winter 2001.

McCain, R. "Game Theory: An Introductory Sketch." Drexel University, 1999. (Available at http://william-king.www.drexel.edu/top/eco/game/game.html)

Mooney, S. "Are Terrorism Risks Really Uninsurable?" National Underwriter, October 22, 2001. (Available at http://www.guycarp.com/publications/nu/2001/1022.html)

Nash, J.F. "Autobiography." The Nobel Foundation, 1994. (Available at http://www.nobel.se/economics/laureates/1994/nash-autobio.html)

O'Connor, J.J., and E.F. Robertson. "John Forbes Nash." University of St. Andrews, 2001. (Available at http://www-groups.dcs.st-nd.ac.uk/~history/Mathematicians/Nash.html)

Pawlowski, D. "Life Insurance Industry Loss Estimates from Sept. 11 Events." Fitch IBCA, 2001. (Available at http://www.knowledgedigest.com/Special_Content/Articles_on_9_11_and_Insurance/articles_on_9_11_and_ insurance.html)

Quade, E.S. *Analysis for Military Decisions*. Chicago: Rand-McNally, 1966.

Von Neumann, J., and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1944.

Vos Savant, M. *The Power of Logical Thinking*. New York: St. Martin's Press, 1996.

Weaver, R., et al. "Modeling and Simulating Terrorist Decision-making: A 'Performance Moderator Function' Approach to Generating Virtual Opponents." University of Pennsylvania, 2001. (Available at http://www.seas.upenn.edu:8080/~barryg/terrist.pdf)

Weiner, H. "Marilyn Is Tricked by a Game Show Host." 1995. (Available at http://www.wiskit.com/marilyn/gameshow.html)

Williams, J.D. *The Compleat Strategyst*. New York: McGraw-Hill, 1954.

Woo, G. "Quantifying Insurance Terrorism Risk." In M. Lane, ed., *Alternative Risk Strategies*. London: Risk Books, 2002, pp. 301-318.