The Use of Stochastic Games for a Homeland Security Decision Problem

by

Erim Kardeş

Acknowledgments

This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number 2007-ST-061-000001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security

Abstract

This paper presents a stochastic game model for a homeland security decision problem. We illustrate this method on the man portable surface-to-air missiles (MANPADS) problem. This approach allows us to model the decision problem from the perspective of both opponents and to perform sensitivity analyses for both decision makers.

Introduction

Since September 11, 2001, the threat of catastrophic terrorism has motivated multi-billion dollar investments in the United States and elsewhere, with the goal of improving safety and security. Investments of this magnitude demand careful consideration of the costs of implementation, operation and maintenance, as well as the benefits derived from a reduction in exposure to future losses.

Unlike naturally occurring or accidental events – such as floods, earthquakes or system failures – terrorism is fundamentally adversarial and adaptive. Thus investments designed to protect against one type of terrorism threat (e.g., against blasts, biological agents, or radiological devices) have potential to elevate the risk of other types of terrorism. Furthermore, interventions to protect against a certain type of terrorism threat category may influence the terrorists' selection of alternative methods within the same threat category. On the other hand, investments targeted at reducing the general effectiveness of terrorist organizations, or targeted at the willingness of individuals to engage in terrorism, may protect against multiple types of terrorism.

The purpose of this paper is to present a modeling approach to a homeland security decision problem that could also be used for other type problems. This modeling approach is illustrated on the MANPADs problem. MANPADS are man portable surface-to-air missiles. There have been publicized MANPADS attacks on large civilian aircrafts in Kenya and Iraq, which increased the fear of such attacks on the US soil or outside the US ([von Winterfeldt and Sullivan(2006)]). A countermeasure that could be installed on aircraft to deflect such missiles incoming to an aircraft is called direct infrared countermeasures (DIRCMs). DIRCMs jam the heat seeking device of a MANPAD and deflect its course away from the airplane. The heat seeking MANPADs are called infrared missiles (IR).

In this paper, we use stochastic games to model the decision processes of two opponent decision makers in the MANPADS problem. The two decision makers are the attacker and the defender. Game theory has been used in the terrorism context by numerous researchers (see [Sandler and M.(2003)]), as it captures strategic interaction among the decision makers, where actions are interdependent and neither of the sides can be considered passive. It is stated in [Sandler and M.(2003)] that the model the authors used in that paper would benefit from a multi-stage analysis, a key feature of the model we develop in this paper.

[Bier et al.(2005)Bier, Nagaraj, and Abhichandani] applies game theory and reliability analysis to identify optimal defenses against intentional threats to system reliability. Authors also mention in this paper that a modeling approach that allows attackers to come back and re-attack would be worthwhile to investigate. Other research efforts that include a one-shot game model between an attacker and a defender include [Major(2002)], [Zhuang and Bier(in press)].

Stochastic Games

In our attacker-defender stochastic game model, we consider the game process from the perspective of the defender. Therefore, we consider that any parameter of the stochastic game associated with the attacker is the defender's subjective judgment about an attacker, who plays the game in an adversarial manner.

In this game, the play proceeds from one state to the other according to transition probabilities controlled jointly by the attacker and the defender. It consists of states and actions associated with each player. Once the game starts in a state, each player chooses their respective actions. The play then moves into the next state with some probability and continues from thereon. The probability that the game moves into the next state is determined by the current state and the actions chosen in the current state.

Let a state be denoted by s. We denote the set of players by $I = \{1, 2\}$, first player being the attacker and the second being the defender. Suppose that in state s, the attacker chooses an action b, and the defender chooses an action c. We therefore have an action tuple $a = (b, c) \in A$ in state s. Here, A denotes the set of all possible action pairs of the players. At each state in this game, players choose their alternatives secretly. Then the game moves into a state k with probability $P_{smk} \geq 0$. If there are M states in the game, the probabilities of transitions from state s to each of the other states must add up to 1, that is, we have, $\sum_{k=1}^{M} P_{smk} = 1$.

Players use strategies at each stage in this game. A strategy is a vector of probabilities that a player assigns to his alternatives in a given state. If a player assigns the probability of 1 to an alternative and 0 to others, then this strategy is called a pure strategy.

Let (f_s, g_s) be the strategies for the attacker and the defender, respectively. For example, $f_s(b)$ denotes the probability that the attacker assigns to his alternative b. At each state in this game, players choose their strategies secretly. Hence, if players play f_s and g_s in state s, then the joint probability that an alternative combination $a = (b, c)$ being chosen is given by

$$\pi_s(a) = f_s(b) g_s(c).$$

If an alternative pair $a = (b, c)$ is chosen in state s, then the attacker and the defender are incurred immediate payoffs denoted by C^1_{sa} and C^2_{sa}, respectively.

The value of the stochastic game to a player $i \in 1, 2$ in state s is defined as the summation of the expected immediate payoff obtained at that state and the expected future values to be obtained by moving on to the subsequent states. In other words, if players play (f_s, g_s) in state s, the value of the game to a player i starting the game in state s, v_s^i, is given by

$$v_s^i = \sum_{a \in A} f_s(b) g_s(c) \{C_{sa}^i + \beta \sum_{k=1}^{M} P_{sak} v_k^i\}. \tag{1}$$

In general, each player in the stochastic game wishes to minimize his own value. Therefore, for the defender, we obtain

$$\min v_s^2 = \min_{g(s)} \sum_{a \in A} f_s(b) g_s(c) \{C_{sa}^2 + \beta \sum_{k=1}^{M} P_{sak} v_k^2\}. \tag{2}$$

The interpretation of equation (2) is that if a player knew how to play optimally from the next stage on, then, at the current stage, he would select the strategy that minimizes the expected immediate cost at the current stage plus the total expected future costs. Hence, player i is not only concerned with the immediate outcome of his actions but also with the future consequences of his strategies in the current stage.

Definition. A strategy pair (f^*, g^*) is a Nash equilibrium in a stochastic game if and only if, $\exists \; v = (v^1, v^2)$, such that, ,

$$f_s^* \in \text{argmin}_{f_s} \sum_{a \in A} f_s(b) g_s^*(c) \{C_{sa}^1 + \beta \sum_{k=1}^{M} P_{sak} v_k^1\}. \tag{3}$$

and

$$g_s^* \in \text{argmin}_{g_s} \sum_{a \in A} f_s^*(b) g_s(c) \{C_{sa}^2 + \beta \sum_{k=1}^{M} P_{sak} v_k^2\}. \tag{4}$$

The interpretation of the conditions in the above definition is that players reach an equilibrium if no player can do better by unilaterally changing his/her strategy. For further details on stochastic games, we refer the reader to [Filar and Vrieze(1997)].

Application to MANPADs

In this section, we describe how the stochastic game formulation described in the previous section can be applied to the MANPADS problem. Our model has the following features: First, alternatives for the attacker in the first state are to attack or not to attack using MANPADS. Second, the data used in our model remain in the ranges suggested in [von Winterfeldt and Sullivan(2006)]. Third, if the attacker chooses not to attack, and the defender chooses not to install countermeasures, then we consider the possibility that MANPADS threat could still exist with some probability. Finally, this model is a non-zero sum stochastic game where the defender's loss is not necessarily the gain of the attacker. Therefore, we allow the possibility of having different payoffs for different players.

Our game model is depicted in Figure 1. Although our model has a tree structure, it is fundamentally different from a decision tree approach.

Since it is a game theoretical model, in the first state, players ***plan*** to choose their alternatives secretly, without any order of their possible actions, whereas a decision tree model would require an explicit order of events and/or actions.

Unlike the roll-back procedure used in decision trees to obtain the solution, our stochastic game model is mapped into an optimization problem, the solution of which yields an equilibrium point. Another main difference between our model and a decision tree approach is that the game model allows the possibility of using mixed strategies. For instance, a positive probability that the defender assigns to his action of installing countermeasures could be interpreted as the percentage of commercial airplanes that can be equipped with countermeasures. A major feature of our model is that if the attacker chooses not to attack, and the defender chooses not to install countermeasures, then our game model takes into account the possibility that MANPADS threat could still exist. Finally, the possibility of the players getting different payoffs at a given state is what differs our approach from a decision tree approach, where a single payoff is associated with a consequence.

Figure 1 depicts our game model with original baseline data values. The first state in the model includes two alternative options for each player: The alternatives of "Attack" and "No Attack" for the attacker and the alternatives of "CM" and "No CM" that stand for the decisions to install DIRCM countermeasures on civilian airliners. The pairs of numbers in the cells of this 2x2 matrix indicate the baseline payoffs associated with each alternative combination (in billions of dollars). For instance, if the defender chooses to install countermeasures and the attacker chooses to attack, the attacker is incurred a cost of attacking, which is \$1 billion, whereas the defender is incurred the installation and maintenance cost of countermeasures of \$10 billion over a horizon of ten years. Based on the possibly chosen alternatives in the first state, the process then moves into subsequent states where additional payoffs are incurred to the two players. For instance, if the defender chooses to install countermeasures and the attacker chooses to attack, with the possibility of a hit, the process moves on to state 2 with 0.25 probability, or to state 3 if there is no hit with a probability of 0.75. If there is a hit, the two subsequent states represent the possibilities of safe landing and fatal crash. If there is a miss, the game ends, but there is a possibility that it will be repeated represented by the arc from state 1 to itself. The payoff and probability figures in the model are precisely the baseline values suggested in [von Winterfeldt and Sullivan(2006)]. Furthermore, every payoff and probability value in this model are parametrized and could be changed for sensitivity analysis purposes. The attacking cost for the attacker is taken to be 1% of the fatal crash cost to the defender. Except for the first state, the cost of the defender is a gain for the attacker. Hence, this model is non-zero sum and captures the cost of attacking from the attacker's perspective. If the players

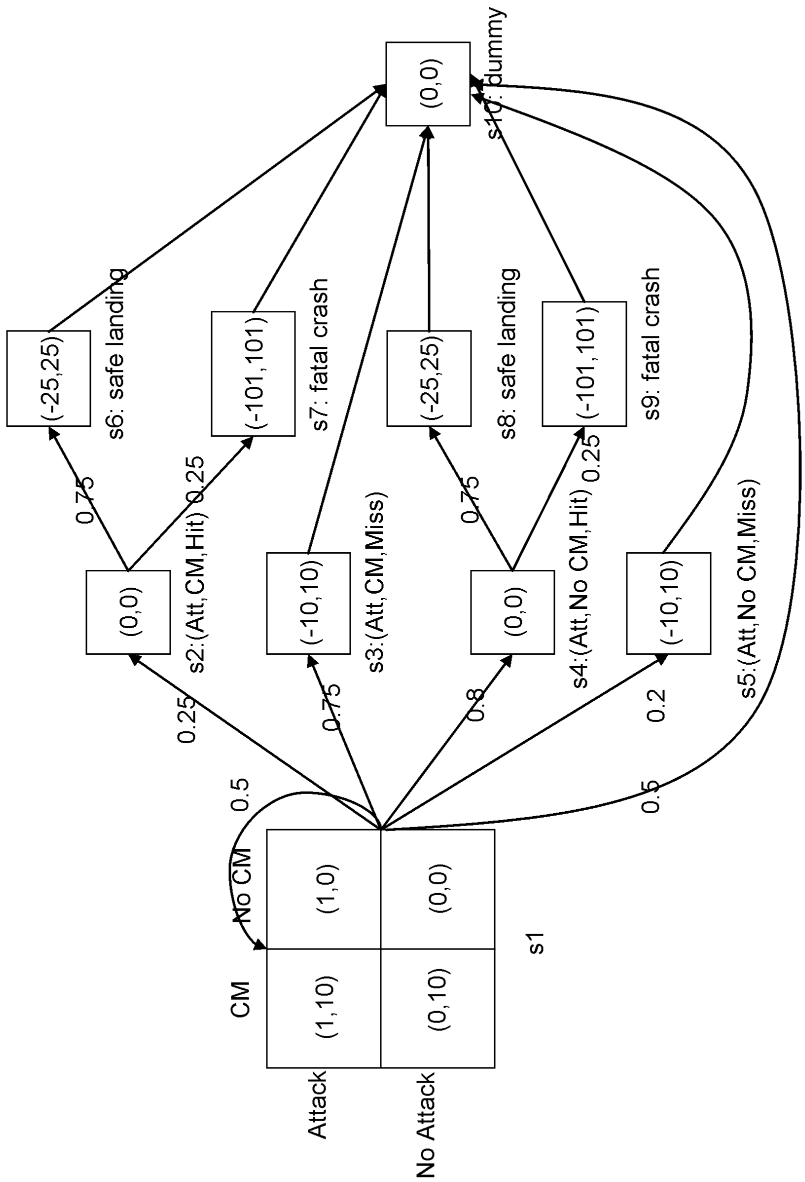

Figure 1: The MANPADs Model

choose not to attack and not to install countermeasures, as a baseline probability value, we assume a 50% chance that the same game will be played in the next ten years. This implies that there will be a 50% chance that the MANPADS threat will exist in the next ten years, if nothing happens. This value is parametrized as well, and we investigate its effect in the sensitivity analysis. Table 1 depicts the equilibrium solution of this stochastic game with baseline data.

Players	CM	No CM	Attack	No Attack	Value (Billions)
Defender	1	0	-	-	28.2
Attacker	-	-	1	0	-17.2

Table 1: Equilibrium Solution for the main MANPADS model

Since this is a nonzero sum game, the values to the players starting the game in state 1 are $28.2 billion and $-17.2 billion, where the minus sign indicates a gain to the attackers. At equilibrium, the attacker and the defender choose to attack and to install countermeasures respectively, with certainty. This is so since the cost of attacking to the attackers is quite small (1$ billion), whereas the expected gain to the attacker from attacking is large. The baseline data in our model makes attacking a very attractive option for the attacker; therefore, the equilibrium solution prescribes the attacker to attack. Therefore, it would be beneficial not only to analyze the equilibrium strategies of the players (or at least of the attackers') but also to consider the best response strategies of the defender against non-equilibrium strategies of the attackers'. We perform such analyses in the next section.

Sensitivity Analyses

In this section, we first run two-way sensitivity analysis by changing the fatal crash cost and the probability of attack (chosen by the attacker). To this end, we fix the attempt probability of attackers at various points between 0 and 1 and for each fixed attempt probability and a given fatal crash payoff, we solve an optimization problem for the defender. We use the following optimization problem given below in a generic form. The use of the following formulation is not restricted to our MANPADs model, and could be used for different problems as well for sensitivity analyses purposes.

$$\min_{w_s^2, f_s} \sum_{s \in S} w_s^2$$

s.t.

$$\sum_{a \in A} f_s(b) g_s(c) \{C_{sa}^i + \beta \sum_{k=1}^{M} P_{sak} v_k^i\} = v_s^i, \text{ for all } s \in S$$

$$\sum_{b \in B} f_s(b) = 1, \text{ for all } s \in S$$

$$f_s(b) \geq 0, \text{ for all } b \in B, \text{ for all } s \in S$$

In this formulation, B represents the set of actions of the defender in a state. We allow this set to be different for each state. Recall that $g_s(c)$ denotes the probability that the attacker assigns to his alternative c. The above formulation treats $g_s(c)$ as a given parameter, and w_s^2 and f_s as variables. It therefore defines the defender's best response problem. To see this, note that the last two constraints in the above formulation ensures that the mixed strategy of the defender is a probability distribution over his set of alternatives. The first constraint satisfies the definition of a value given in section 2. Finally objective value captures the fact that the defender wishes to minimize his value of the game by choosing his best strategies.

AMPL is used for modeling the mathematical program and the nonlinear solver KNITRO is used for calculating best responses for the defender.

We depict the values of fatal crash costs and attempt probabilities that that do not favor installing countermeasures in the area below the function in Figure 2. Countermeasures are preferred in the area above this function. For instance, if the attack probability is 0.2, and the total economic costs of fatal crash is less than about \$150 billion, then countermeasures are not preferred. Figure 3 depicts the values to the defender given by the best response strategies against the attacker's attack probabilities and various economic costs for a fatal crash. For instance if the attack probability is 1 and if the fatal crash costs are ignored (or assumed to be zero), the best response value to the defender is about \$19 billion, which comprises the economic costs associated with "safe landing" and/or "miss" states. Next, we run the same analysis with a fixed countermeasure cost of \$30 billion instead of \$10 billion. Figure 4 plots the areas where the fatal cost and attack probability

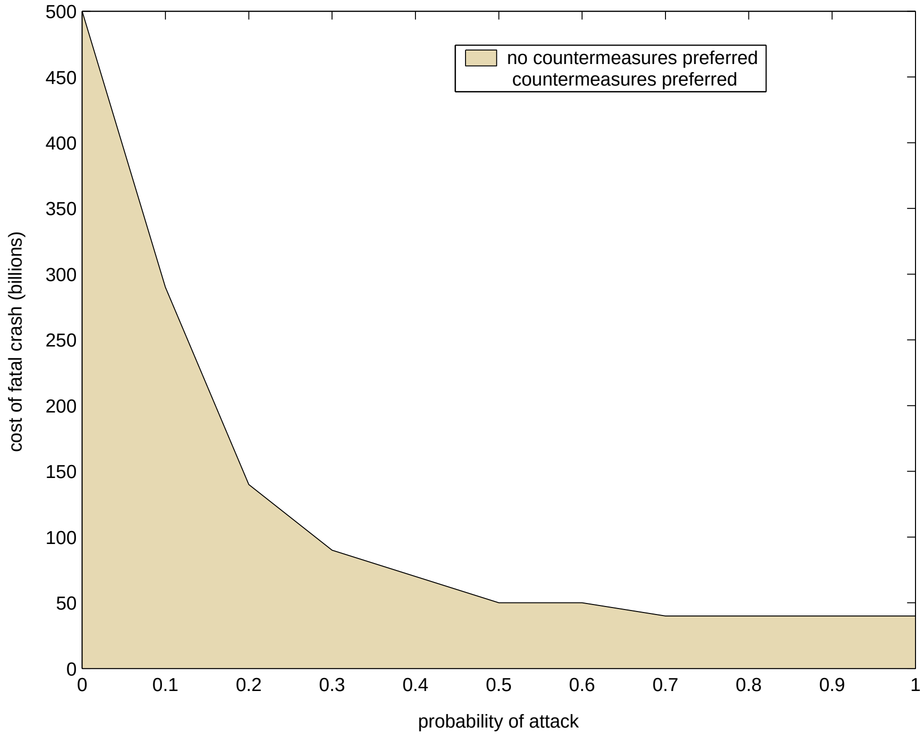

Figure 2: 2-way sensitivity analysis on the cost of fatal crash and the probability of attempt

data favor to install countermeasures or not. As depicted in the figure, the higher countermeasure cost of $30 billion results in higher fatal crash costs that favor countermeasures.

We next run two-way sensitivity analyses by changing the strategies of the defender and the cost of fatal crash, and by defining the problem from the attacker's point of view. In this analysis, we keep the payoffs associated with safe landing and attacking at their original values. Analyses indicate that the optimization of the best response function for the attacker for any combination of the defense strategies and positive fatal crash costs ($0-$400 billion dollars) prescribes the attacker to attack. This is so, since the gain that the attacker achieves from attacking is substantial, compared to the cost of attack and to the benefits of not attacking. We then analyze the problem from the attacker's point of view by assuming that the fatal crash costs range from -$400 to $400 billion, which indicates a gain to the defender, possibly due to the benefit to the society caused by thwarting an attack. In this case, if the defender chooses to install countermeasures with some probability in the range of 0 to 0.4, then the attacker chooses to attack if the cost of a fatal crash to the attacker is less than $50 billion. If the defender considers installing countermeasures with some probability between 0.5 and 0.8, the threshold for the cost of fatal crash to the attacker, under which the attacker is willing to attack, increases to $100 billion.

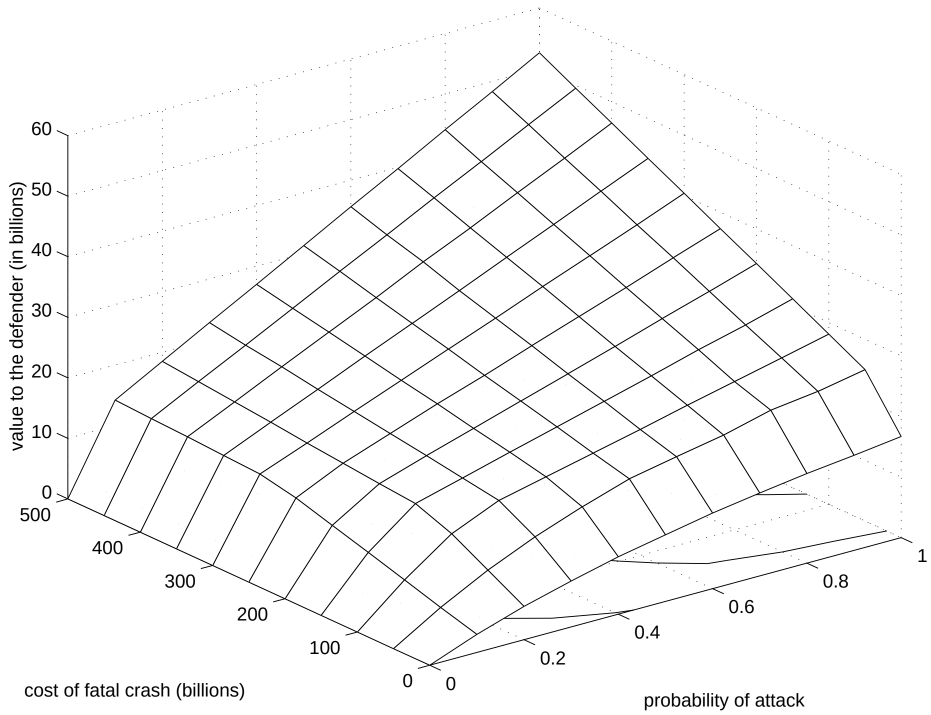

Figure 3: Best response values to the defender

If the defender considers installing countermeasures with some probability between 0.8 and 1, this threshold increases to \$150 billion. This means that as the defender puts more weight to his countermeasure installation option, the attacker is willing to increase his threshold for his cost of fatal crash that will prescribe him to attack. However, we see a limit to this threshold, which is \$150 billion. The attacker is willing to increase his threshold since the probability of safe landing is high when there are countermeasures installed, and the gain to the attacker from safe landing is taken to be \$25 billion dollars. Another reason for the attacker to increase his threshold is that the cost of attacking is only \$1 billion. Therefore, we next run sensitivity analyses by changing the cost of fatal crash in the range of -\$400 to \$400 billion, while keeping the cost of attack as much as the countermeasure cost (\$10 billion), and while safe landing causes a \$25 billion cost to the attacker, rather than the defender. Under this data scenario, the attacker chooses to attack if the gain he achieves from attacking exceeds \$150 billion, no matter the defender chooses to install countermeasures or not. We see from these analyses that when the payoffs to the attacker are at their original values, the attacker is willing to increase his threshold for his cost of fatal crash that will prescribe him to attack as the defender puts more weight to his countermeasure installation option. On the other hand, we see that the attacker could require a minimum gain from attacking, if the cost of attacking is high, and safe landing incurs costs on the attacker.

Next, we run two-way sensitivity analysis on the probability of re-play and cost

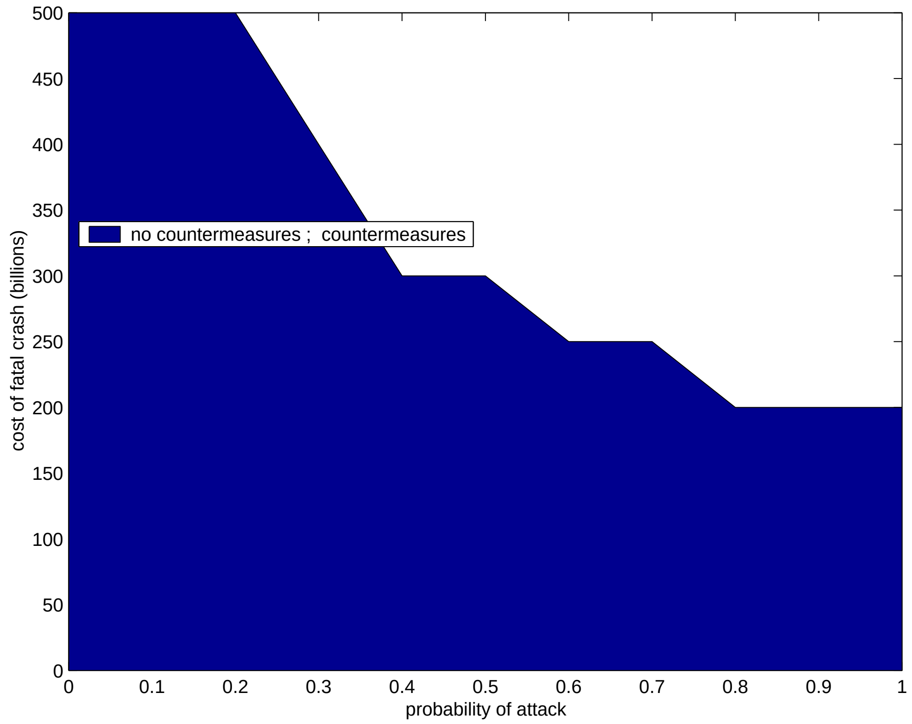

Figure 4: 2 way sensitivity analysis with countermeasures costs set to $30 billion

of fatal crash, when the chances for the attack probability are 50%. The area where the countermeasures are preferred is plotted in Figure 5.

Figure 5 indicates that if the re-play probability is between 0.1 and 0.4, then countermeasures should be preferred if the economic costs of a fatal crash exceeds $100 billion, when the cost of countermeasures is at $10 billion. Furthermore, we see that as the re-play probability increases, the costs of a fatal crash that favors countermeasures decreases as a step function of the re-play probabilities. Figure 5 outlines where these steps occur. Figure 5 also shows, for a $10 billion countermeasure cost, that if the re-play probability is between 0.5 and 0.9, then countermeasures should be preferred if the cost of a fatal crash exceeds $50 billion.

Next, we run the same analyses for an attack probability of 0.20 and show the plot in Figure 6. In this case, for a $10 billion countermeasure cost, if the re-play probability is 0.2, countermeasures are not preferred unless the cost of fatal crash exceeds $250 billion, which is more than twice the fatal crash cost when the attack probability is 0.50.

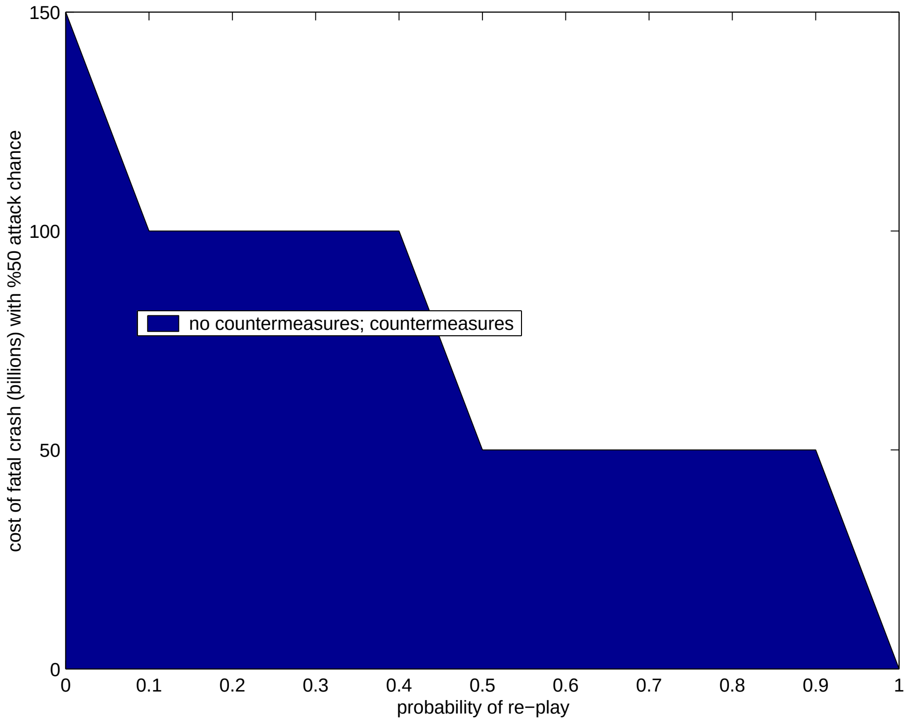

Figure 5: 2 way sensitivity analysis on the probability of re-play and fatal crash cost (with a fixed attack probability of 0.50)

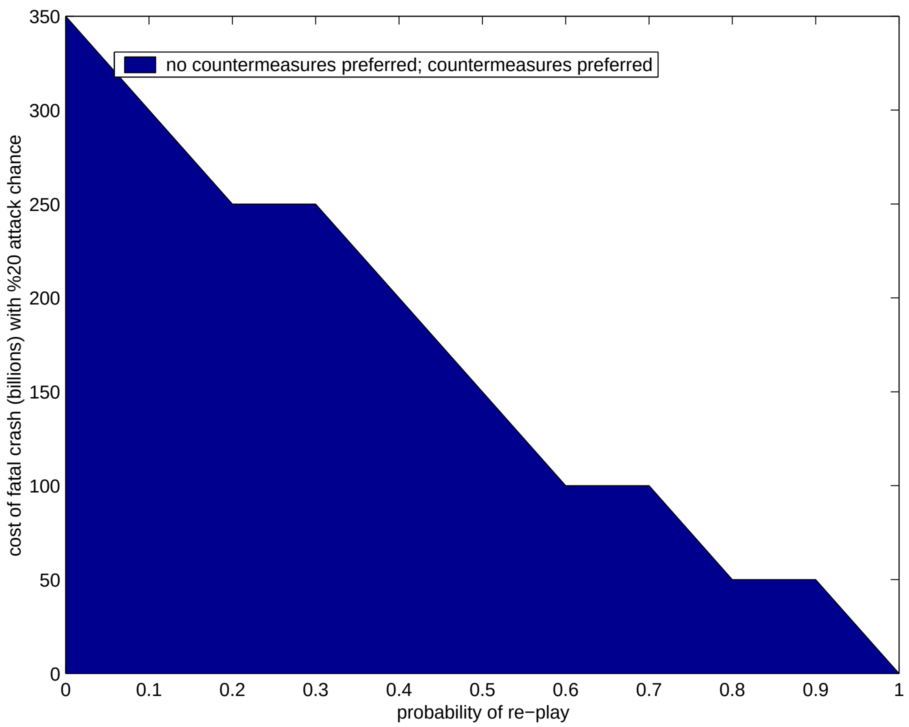

Figure 6: 2 way sensitivity analysis on the probability of re-play and fatal crash cost (with a fixed attack probability of 0.20)

Conclusion

In this paper, we have applied a stochastic game model and performed sensitivity analyses to investigate the strategies of two opponents in the MANPADS problem. Our analyses suggest that the countermeasures are cost-effective if the countermeasures costs over a ten year period is around \$10 billion, the attack probability is high ($>$ \$4 billion), and if the fatal crash cost is more than about \$ 75 billion. This conclusion mirrors the conclusion given in [von Winterfeldt and Sullivan(2006)]. Furthermore, we conclude that if the attack probability is less than 0.4 and the re-play probability is low (around 0.1), then the countermeasures are not cost-effective unless economic costs associated with a fatal crash are very high (above \$250 billion). Finally, our results suggest that assuming the attack probability is around 0.2, countermeasures could be cost-effective if economic costs of a fatal crash is above \$50 billion, and if the MANPADS threat continues to exist with high probability, given that no attacks occur and no countermeasures are installed.

Bibliography

[Bier et al.(2005)Bier, Nagaraj, and Abhichandani] Bier, V., A. Nagaraj, V. Abhichandani. 2005. Protection of simple series and parallel systems with components of different values. *Reliability Engineering and Systems Safety* **87** 315–323.

[Filar and Vrieze(1997)] Filar, J., K. Vrieze. 1997. *Competitive Markov Decision Processes*. Springer-Verlag.

[Major(2002)] Major, J. 2002. Advanced techniques for modeling terrorism risk. *Journal of Risk Finance* **4**.

[Sandler and M.(2003)] Sandler, T., D. G. Arce M. 2003. Terrorism and game theory. *Simulation and Gaming* **34**(3).

[von Winterfeldt and Sullivan(2006)] von Winterfeldt, D., Terry O. Sullivan. 2006. Should we protect commercial airplanes against surface-to-air missile attacks by terrorists? *Decision Analysis* **3** 63–75.

[Zhuang and Bier(in press)] Zhuang, J., V. Bier. in press. Balancing terrorism and natural disasters—defensive strategy with endogenous attacker effort. *Operations Research* .