Game theory in infrastructure security

V.M. Bier & S. Tas
Department of Industrial and Systems Engineering, University of Wisconsin–Madison, Madison, WI 53706, USA

Abstract

Game-theoretic security models have gained popularity in infrastructure security in recent years, due to the fact that game theory is suitable for dealing with intelligent threats. In this chapter, we briefly discuss some of the key concepts in game theory, categorize game-theoretic models in infrastructure security and give some examples, and finally discuss some of the limitations of game-theoretical models.

Keywords: Game Theory, Attacker–Defender Games, Sequential Games, Infrastructure Security.

1 Introduction

There has been increasing use of game-theoretic models in infrastructure security, especially after September 11, 2001. This is appropriate because game theory considers the intelligent and adaptive nature of an adversarial threat. Therefore, in this chapter we review the application of game-theoretic models to infrastructure security. We first explain some of the important concepts of game theory. Then, we discuss several categories of game-theoretic models in infrastructure security and present some examples. Finally, we discuss some of the limitations of game-theoretic models.

A *game* is a formal description of the strategic interactions of multiple agents (in infrastructure security, typically an attacker and a defender). These interactions can be between those defending a system and those attacking it or between multiple defenders. Game theory assumes that each agent or "player" in a game wishes to find its best strategy given the strategies adopted by the other player(s). This assumption makes it possible for an analyst to make predictions about which strategies players would be likely to choose (under the assumption of rationality); for example, in game theory, an agent would never choose a strategy if it is strictly dominated by another strategy, in the sense that

WIT Transactions on State of the Art in Science and Engineering, Vol 54, © 2012 WIT Press
www.witpress.com, ISSN 1755-8336 (on-line)
doi:10.2495/978-1-84564-562-5/06

the other strategy performs better for all possible strategies of other player(s) in the game.

In game theory, an *equilibrium* is any set of strategies where no player has an incentive to change its strategy, if all other players continue to play their equilibrium strategies. An equilibrium can be either *pure* (where each agent has a unique and deterministic equilibrium strategy) or *mixed* (where at least one agent is assumed to choose probabilistically from among multiple equilibrium strategies).

Games can be classified along multiple dimensions. With regard to payoffs, in a *constant-sum game*, the sum of the agents' payoffs is the same for all possible outcomes of the game. A special case of constant-sum games is *zero-sum games*, in which one agent's loss is equal to the other agent's gain, so the sum of the agents' payoffs is zero. For zero-sum games, there is ensured to be at least one equilibrium, although it may be a mixed rather than a pure equilibrium [1]. These games are relatively straightforward to analyze but will not always be realistic in practice. Therefore, *non-zero-sum games* may often be needed; however, some non-zero-sum games have no equilibrium.

With respect to the timing of play, in *simultaneous games*, agents choose their actions without knowing the actions of the other players. (Note that the moves of the agents do not necessarily have to be simultaneous, as long as no player can observe the actions of any other player.) By contrast, in *sequential games*, the agents act in a certain order, instead of simultaneously. In these games (also known as leader–follower games in economics and attacker–defender (AD) or defender–attacker (DA) games in security), the leader moves first and the follower moves second, generally after observing the action(s) of the leader. The leader generally has a first-mover advantage, since the choices made by the leader can limit the options available to the follower(s). In infrastructure security, decisions about observable capital investments are typically modeled as sequential games (since an attacker can often observe such defensive investments before choosing an attack strategy), while decisions that can be changed easily and rapidly (like allocation of police patrols) may be modeled as simultaneous games.

Games can also be classified with respect to how much players know. A game is one of *complete information* if the payoffs for each combination of actions chosen by the various players are *common knowledge* to all players. By contrast, when some information is not common knowledge (in other words, when the players do not share all information about one another's preferences or behavior), then the game is one of *incomplete information* (also known as a *Bayesian game*). In a Bayesian game, some players may have only probabilistic information about the preferences of other players (e.g., a probabilistic distribution over what "type" another player is). Likewise, in a game of *perfect information*, all players know all past moves in the game at any given point in time. By contrast, in a game of *imperfect information*, at least one player does not know all past moves of the other player(s).

2 Game-theoretic models

We can categorize as follows the game-theoretic security models that have been discussed in the literature:

1. Simultaneous AD games
2. Sequential DA games
3. Sequential AD games
4. Sequential defender–attacker–defender (DAD) games
5. Simultaneous defender–defender (DD) games

Each of these games is discussed in detail below. (One should note that players may be decentralized; for example, players may attack or defend only specific parts of an infrastructure system, instead of centralized attack and defense of the entire system.)

2.1 Simultaneous AD games

Simultaneous AD games are a special case of AD games where the attacker and the defender select their strategies independently, without knowing the strategy chosen by the other player. In other words, the attacker does not know the defender's decision when the attacker makes his own decision, and the same is true for the defender.

Bier *et al.* [2] used simultaneous AD games (where the attacker has no information about the defensive investments made by the defender) to identify optimal strategies for protecting the components of a simple series system. In this model, the attacker wants to maximize the expected loss experienced by the defender. The defender is assumed to minimize the expected loss (plus the cost of any defensive investments, in the unconstrained version of this model). Bier *et al.* found that the defender has greater flexibility in allocating defensive investments cost effectively when the attacker cannot observe those investments, compared to games in which the attacker can observe the defensive investments. This result shows the potential benefits of secrecy for the defender. Zhuang and Bier [3] also modeled secrecy as a simultaneous game.

Similarly, in Hausken *et al.* [4], the defender makes tradeoffs between protecting against terrorism only, natural hazards only, or both (all hazards). The authors considered a simultaneous version of this game (in which the adversary and the defender do not know each other's actions) and determined under what conditions the defender would prefer to play a simultaneous rather than a sequential AD game (and, likewise, when the attacker would prefer to play a simultaneous rather than a sequential DA game).

Simultaneous games have also been used to analyze intrusion detection for decentralized (i.e., peer-to-peer or *ad hoc*) networks. For example, Patcha and Park [5] considered a simultaneous game between a sender and a receiver. This

is a Bayesian game because the sender can be of two types, either malicious or not. The objective of the attacker (the malicious type of sender) is to successfully send a malicious message without being detected by the defender's intrusion-detection system. The defender wants to intercept these attacks, while minimizing the rate of false alarms for regular messages.

Unfortunately, the simultaneous-move assumption (while simple) is often not realistic in infrastructure security since, for example, the defender may engage in costly infrastructure improvements over time and at least some of those defensive investments may be observable by the attacker. Therefore, in Section 2.2, we consider sequential DA games, which are able to overcome some of the aforementioned limitations of simultaneous AD games.

2.2 Sequential DA games

Sequential DA games determine how a defender should optimally protect a system if the attacker can observe its defensive actions. The sequence of play in such games is as follows: The defender moves first by implementing an optimal defense; the attacker than observes the defense and identifies the best possible attack strategy given that defense. See Figure 1 for a typical DA game.

Bier *et al.* [2] used a sequential DA game to determine how to protect both series and parallel systems (in addition to the simultaneous model for series systems discussed earlier). They found that for series systems, it is optimal in a sequential game to equalize the expected damage from attacks on all defended components, as opposed to the simultaneous game (in which the marginal reduction in expected damage from incremental investment in any component is the same). Thus, in contrast to the simultaneous case, the defender has less ability to choose the most cost-effective defensive investments in the sequential game.

Azaiez and Bier [6] extended this work to systems with more general structures (combined series/parallel systems). They also revised the defender's objective function from minimizing expected loss to maximizing the cost of an attack to the attacker; in this model, defensive investments are assumed to increase the cost of an attack but not to decrease the success probability of the attack. As in Bier *et al.* [2], the results suggest that defensive investment should equalize the attractiveness of all defended components in any series subsystem of the overall system; however, the definition of component attractiveness is somewhat different, due to the more complex structure of the systems being analyzed and the different nature of the objective function.

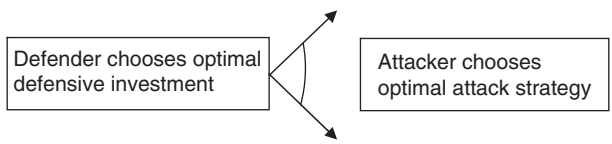

Figure 1: A sequential DA game.

Hausken and Levitin [7] similarly used game theory in their analysis of combined series/parallel systems where the components have different performance and reliability characteristics, and applied their model to an illustrative power-substation system. The objective of the defender in their model is to find optimal separation strategies for the elements within the system, in order to ensure that any single attack can damage only a subset of components.

Bier *et al.* [8] studied DA games in which a defender must allocate defensive resources to a collection of individual assets or "locations" and then the attacker chooses a location to attack. They concluded that the defender's problem gives rise to negative externalities between locations because increasing the resources allocated to one location increases the likelihood of attacks at other locations. In fact, the defender exploits these externalities to manipulate the attacker's behavior, by protecting the more valuable and vulnerable locations, in order to deflect attacks to locations that will be less damaging to the defender. It is important to note that due to the first-mover advantage, the defender in this model prefers his or her defensive allocation to be public rather than secret. This is because defense of high-valued targets not only reduces the success probability of attacks on those targets but also deflects attacks to less-valued targets.

By contrast, Dighe *et al.* [9] found conditions under which partial secrecy (i.e., disclosure of the total defensive investment but secrecy regarding the allocation of that investment to the various assets or locations) can be advantageous to the defender in a sequential DA game. This result is due to the fact that the success probability of an attack is a non-convex function of the level of defensive investment in their model. Pita *et al.* [10] took advantage of similar results to achieve more cost-effective protection in an application to allocation of guards and checkpoints at the Los Angeles International Airport. Moreover, Zhuang and Bier [3] found conditions under which secrecy and deception may be advantageous to the defender even in the absence of such non-convex success probabilities; their results suggest that secrecy and deception are more likely to be desirable when defense is only marginally justified (i.e., not so costly as to be clearly not worthwhile but not so cheap as to be obviously worth implementing).

Bier *et al.* [11] applied the model developed by Bier *et al.* [8] to the defensive budgets of the top 10 urban areas in the United States using various measures of target attractiveness to the attacker. In addition to expected fatalities and expected economic losses, they also used data on infrastructure (in particular, average daily bridge traffic and number of airport departures) as illustrative measures of target attractiveness.

Sequential DA models have also been applied to telecommunication networks. For example, Cox [12] focused specifically on resilient network design, in which an attempt to disrupt traffic leads to rerouting of the traffic. The defender first chooses a set of defensive measures, and the attacker then interdicts some number of links or nodes of the telecommunication network. The goal of the defender is to ensure that the network has sufficient capacity and path diversity so that service can continue with no disruption even after an attack of a given size.

2.3 Sequential AD games

Sequential AD games determine what an attacker should do if he gets to move first and how the defender should optimally respond to an observed attack. The sequence of play in such games is as follows: Given the attacker's constraints, the attacker launches an optimal attack; the defender then observes the attacker's strategy and the resulting damage, and identifies the best possible response. See Figure 2 for the sequence of decisions in a typical AD game.

In many such models, the attacker has the advantage of surprise, that is, deciding when, where, and how to attack. Due to this first-mover advantage, such AD games can be considered a worst-case situation from the viewpoint of the defender.

Many examples of sequential AD games in the literature deal with networked infrastructure. In particular, for example, network interdiction problems are a special case of sequential AD games, where the attacker maximizes the defender's minimum operating cost (see, for example, Kunturska *et al.* [13]). Such games can be used to identify and assess network vulnerabilities. Here, the attacker launches an optimal attack and disables some components of the network; the defender then determines how to optimally operate the network given its remaining capability.

For example, Israeli and Wood [14] developed an interdiction algorithm in which the attacker maximizes the shortest path on a directed power network. Similarly, Salmeron *et al.* [15] modeled the interdiction of critical system components (transmission lines, generators, and transformers) in an electric transmission system, using a heuristic algorithm to solve the attacker's optimal interdiction strategy (in other words, the most critical components in the network). The objective of the defender is to minimize the sum of generating and load-shedding costs. For the same problem, Bier *et al.* [16] used a greedy algorithm that interdicts the components with the maximum flow and obtained similar results to those of Salmeron *et al.* [15]. Salmeron *et al.* [17] improved on the heuristic algorithm in their earlier paper, with the result that they can generate faster and better solutions for considerably larger electric power grids. Sequential AD models have also been applied to oil supply chains [18], road-network vulnerability [19], and information systems [20].

In some interdiction problems, the defender can also interdict the attacker's network. For example, see Wood [21] for a defender interdiction problem where the defender minimizes drug trafficking in a capacitated network.

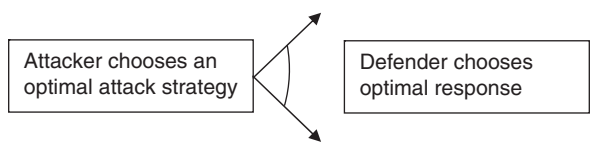

Figure 2: A sequential AD game.

2.4 Sequential DAD games

In sequential DAD games, the defender designs an optimal network before the attack, assuming that the attacker will launch an optimal attack, and the defender will then choose the best possible response to that attack. In other words, we can describe the interaction between the attacker and the defender in three phases:

1. Preattack (defender's network-design problem): The defender protects the infrastructure network (e.g., through hardening, redundancy, surveillance, pre-emption, and deterrence).
2. Attack: The attacker chooses and executes an optimal attack strategy.
3. Postattack (defender's best-response problem): The defender responds to the attack (e.g., by rerouting the flow through the network).

See Figure 3 for the sequence of decisions in a typical DAD game.

Yao *et al.* [22] defined sequential DAD models as including active defense (since the defender foresees an optimal attack and designs her network accordingly), rather than only passive defense (where the defender merely responds to an attack after it happens, as in sequential AD models). Yao *et al.* [22] also proposed a solution procedure for DAD games and applied this procedure to a problem involving defense of power networks.

Smith *et al.* [23] considered a sequential DAD game for a generic transportation or telecommunication network, in which an attacker attempts to minimize the maximum possible post-interdiction profit achievable by the defender. They considered two heuristic attack strategies (interdicting the arcs with the highest capacities or with the highest initial flows) but found that these did significantly less well than the optimal attack strategy. Note that this is different than the results in Bier *et al.* [16], where the maximum-flow heuristic attack strategy worked well. The reason for this difference is unclear, but it may be because the electricity networks in Bier *et al.* [16] are heavily capacity constrained, while the randomly generated networks in Smith *et al.* [23] generally have large excess capacity.

In Church and Scaparra [24], the goal of the defender is to first identify and then fortify (protect) critical facilities in a network of service facilities. The attacker wants to maximize the weighted average service distance for the entire system, where if service is lost at one facility, it will be provided by other facilities.

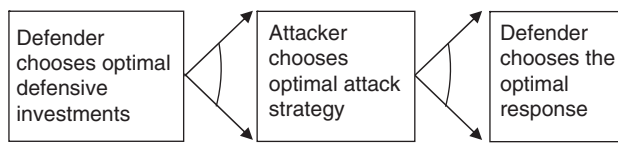

Figure 3: A sequential DAD game.

The authors tested their model for two moderate-sized cases and found that the solutions of their fortification problems contained at least one facility that was also part of the optimal interdiction strategy of the attacker.

Brown *et al.* [18] developed a decision-support system to identify the critical components in an electric power network. The system begins with a sequential AD model, as in Salmeron *et al.* [15] and Brown *et al.* [25], to identify the critical components in the network. However, it then identifies near-optimal defender hardening strategies, thus extending the model to a DAD game. Similarly, Bier *et al.* [16] also extended the AD game discussed in the Section 2.3 to include a hardening algorithm, in which the defender hardens a subset of the possible targets identified for interdiction. See Brown *et al.* [26] for a discussion of sequential DAD games and their potential application to bioterrorism.

2.5 Simultaneous DD games

Simultaneous DD games involve the defensive investments of multiple agents in a system, where the threat may sometimes be modeled as "exogenous" (i.e., unrelated to the defender decisions). Kunreuther and Heal [27] described these games as interdependent security games. In principle, defensive investments by any one agent can create either positive externalities (e.g., if one agent uses antivirus software, reducing the risk to other agents) or negative externalities (e.g., if one target is hardened, deflecting attacks to other targets) for other agents.

Kunreuther and Heal [27] began by focusing on applications with positive externalities (e.g., vaccination, computer security, and baggage screening at airports). In such problems, if defense is low cost for any given agent, that agent will choose to invest in security (S), as shown in Figure 4; if defense is costly, agents will generally not invest (N). Interestingly, Heal and Kunreuther [28] found that there is a region (for intermediate investment costs) where there are multiple equilibria (either N,N or S,S); in other words, each player prefers to

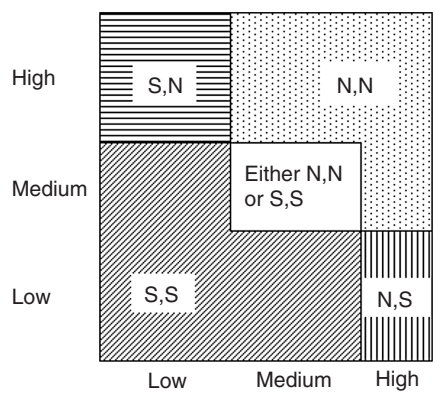

Figure 4: Security investment decision of two firms (modified from Ref. [28]).

choose whatever the other player chooses. This creates a need for mechanisms to coordinate the investment decisions of the various agents in order to achieve the social optimum. Kunreuther and Heal [27] suggested internalizing these externalities by coordination mechanisms (e.g., insurance, liability, fines or subsidies, regulations, third-party inspections, and recommendations by trade associations). Zhuang *et al.* [29] extended this basic model from a "snapshot" model (in which the threat is manifested immediately) to a time-dependent model (in which the threat is realized only over time). They found that in this case, the presence of a single myopic player in the game can make it undesirable for non-myopic players to invest in security, even when it would otherwise be in their interests to do so.

However, there can be also cases with negative externalities. For example, investing in the security of the aviation system may increase the risk to other modes of transportation. Similarly, investing in the security of the U.S. Postal Service may increase the threat to private mail carriers such as Federal Express [30].

Keohane and Zeckhauser [31] considered various types of externalities among potential victims as a result of the strategic nature of the terrorist threat. For example, if a tall building is hardened, this may either decrease the probability of an attack or decrease the consequence of an attack. This may endanger other buildings if the hardening deters attacks on the protected building, creating negative externalities for the owners of the other buildings. At another level, as the number of people exposed to a threat decreases (e.g., through protection of individual buildings), the attractiveness of the entire region to potential attackers may decrease, creating positive externalities for all residents of the region. Keohane and Zeckhauser [31] identified four relevant effects: discouragement of attacks (if protection against threat decreases the probability of attack, creating a positive local externality), diversion of attacks (if precautions shift the attack to other targets in the same location, creating a negative local externality), displacement of attacks (if local precautions shift the attack to other locations, creating a negative global externality), and finally containment (if protection against a threat benefits others, creating a positive global externality), as in the case of screening or vaccination [27].

Hausken [32] incorporated game theory into probabilistic risk analysis. He used two hypothetical infrastructure examples to illustrate series and parallel systems, respectively. For a series system, he borrowed the example of dams from Hirshleifer [33]. In this example, there is a circular island, where each citizen owns a "pie-shaped" slice or wedge of land. Storms may flood the island, in which case all citizens would be affected. The only protection is to build dikes around the entire island. Each citizen makes his own decision about how high of a dike to build on his slice of land. This creates positive externalities, since investment by any one citizen also benefits other citizens; in this case, each defender wishes to invest in defense only if other players also invest. In order to avoid the equilibrium where nobody invests, Hausken proposed that the defenders should coordinate their efforts so that everybody invests and everybody benefits.

For a parallel system, Hausken used the example of antimissile defense from Hirshleifer [33]. In this example, each citizen of a city may erect an antimissile defense battery on the outskirts of the city, again creating positive externalities (since any one citizen's defense also benefits other citizens). However, due to the parallel (or redundant) nature of the defences in this case, only a single successful defense is necessary to destroy an incoming missile. Therefore, the social optimum may be for only one citizen (or a small number of citizens) to erect defences. However, one equilibrium is that no citizen erects any defences, since they all want to shift the burden of defense to their neighbors.

3 Limitations of game-theoretic models

Although game-theoretic models provide a rigorous and mathematically sophisticated way of incorporating the actions of an intelligent and adaptive adversary into security decision making, they also have some limitations. In particular, real-world problems may not necessarily fit the assumptions that a standard game-theoretic model requires, as discussed in the following.

Rationality of the agents It may be unrealistic to assume that the attacker is perfectly rational and has unlimited computational ability. In particular, game-theoretic models disregard the fact that strategic reasoning (and the supporting calculations) requires minimum levels of skill (education and problem-solving ability) and resources (time, tools, computers, data, etc.), which might be either unavailable to attackers or too costly to be worthwhile. Thus, Ezell *et al.* [34] criticized game theory as being primarily a normative technique for determining how players should play in an idealized world, rather than a descriptive technique that represents how players might actually play. One possible way of addressing this limitation is by considering heuristic rather than optimal decision making by the attacker; see, for example, Bier *et al.* [16].

Common knowledge Many game-theoretic models rely on the assumption of common knowledge, in which all agents share the same understanding of the system, and also know the objective functions and payoffs of other agents, or at least prior distributions for the types of the other agents in a Bayesian game. See, for example, the discussion in Bier *et al.* [35]. In practice, however, it is a daunting task to determine agents' utility functions and payoffs; in fact, one reason for the existence of intelligence services is precisely because the defender may not know the attacker's preferences (and vice versa). However, Bayesian games can still be used when some players have incomplete information about the characteristics of other players.

Modeling challenges Most models in the literature assume that both the attacker and the defender have at least roughly the same objective function. It is also common to use linearized objective functions to simplify nonlinear models (such as linear DC load flow models to simulate nonlinear power flows); for example, see Brown *et al.* [25]. Another challenge is that the attacker's optimization problem

is often discrete (e.g., whether to attack or which target to attack) rather than continuous, which can make the attacker's optimization problem difficult to solve; some techniques for solving such problems are discussed in Brown *et al.* [18].

Excessive conservatism Ezell *et al.* [34] noted that in many game-theoretic models, the attacker's objective is assumed to be to maximize the consequences of an attack. This means that the optimal defender strategy in such games will typically be to defend against only the most severe possible attack(s). However, this may leave the defender vulnerable to lesser attacks that may be attractive to attackers with slightly different objective functions, due to insufficient defensive "hedging."

As a result of problems such as these, it is generally accepted that game theory does not always provide accurate predictions of how players behave empirically. Therefore, as an alternative to game theory, Ezell *et al.* [34] proposed the use of probabilistic risk analysis to represent attacker choices as uncertain events. However, Parnell *et al.* [36] noted that probabilistic risk assessment often involves dauntingly large numbers of hypothetical event sequences, with subjective probabilities that must be assessed by elicitation of expert judgment, even though terrorist attacks are so rare that subjective assessment of their likelihoods may be of limited accuracy. The National Research Council [37] also noted that probabilistic risk analysis may not be adequate for capturing the behavior of intelligent adversaries.

Parnell *et al.* [36] also argued that probabilistic risk assessment may understate the likelihood of severe events (like the attack on September 11) by treating the various steps or choices leading up to such an event as being probabilistically independent. By contrast, game-theoretic models avoid this pitfall by focusing on the purposive or intentional nature of the intelligent adversary's actions, rather than just assigning probabilities to those choices.

Cox [38] emphasized that it is not necessary to pick one of these methods over another and that game theory can in fact be integrated with risk analysis. For example, risk analysis can support game-theoretic models by providing probability distributions (rather than point estimates) of consequences for every pair of AD actions [35]; thus, Banks and Anderson [39] combined game theory and risk analysis in a bioterrorism example, in which they used risk analysis to generate reasonable probability distributions for the payoff matrices of various AD combinations for a smallpox threat. Similarly, Guikema [40] discussed cases in which game theory is incorporated into reliability analysis of complex systems. For additional discussions of the pros and cons of game theory for security, see Parnell *et al.* [41], Ezell and von Winterfeldt [42], and Guikema and Aven [43].

4 Conclusion

Game theory provides a rigorous mathematical way to account for the actions of intelligent adversaries, and as a result, game-theoretic models are being increasingly used in infrastructure security. In this chapter, we categorized different

game-theoretic models in the literature on infrastructure security and provided illustrative examples.

Of course, as noted by Box and Draper [44], no model represents the real world perfectly. However, game-theoretic models can still yield useful insights into what both attackers and defenders may do if they take into account each other's actions and preferences.

Applying game theory to infrastructure networks of realistic size and complexity can still be challenging. However, with the advancement of computing technologies and algorithms, it will become increasingly feasible to solve complex infrastructure security problems using game theory.

References

[1] Von Neumann, J., Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, **100**, pp. 295–320, 1928. Translated by S. Bargmann as On the theory of games of strategy in contributions. *Contributions to the Theory of Games IV*, eds Tucker, A. and Luce, R.D., *Annals of Mathematics Study*, **40**, Princeton University Press: New Jersey, pp. 13–42, 1957.

[2] Bier, V.M., Nagaraj, A., & Abhichandani, V., Protection of simple series and parallel systems with components of different values. *Reliability Engineering and System Safety*, **87(3)**, pp. 315–323, 2005.

[3] Zhuang, J. & Bier, V.M., Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation, *Defence and Peace Economics*, **22(1)**, pp. 43–61, 2010.

[4] Hausken, K., Bier, V.M., & Zhuang, J., Defending against terrorism, natural disaster, and all hazards (Chapter 4). *Game Theoretic Risk Analysis of Security Threats*, eds Bier, V.M. and Azaiez, M.N., Springer: New York, pp. 65–97, 2009.

[5] Patcha, A. & Park, J.M., A game theoretic formulation for intrusion detection in mobile ad hoc networks. *International Journal of Network Security*, **2(2)**, pp. 131–137, 2006.

[6] Azaiez, N. & Bier, V.M., Optimal resource allocation for security in reliability systems. *European Journal of Operational Research*, **181(2)**, pp. 773–786, 2007.

[7] Hausken, K. & Levitin, G., Minmax defense strategy for complex multi-state systems. *Reliability Engineering and System Safety*, **94**, pp. 577–587, 2009.

[8] Bier, V.M., Oliveros, S., & Samuelson, L., Choosing what to protect: Strategic defense allocation against an unknown attacker. *Journal of Public Economic Theory*, **9(4)**, pp. 563–587, 2007.

[9] Dighe, N., Zhuang, J., & Bier, V.M., Secrecy in defensive allocations as a strategy for achieving more cost-effective attacker deterrence. *International Journal of Performability Engineering*, special issue on System Survivability and Defense against External Impacts, **5(1)**, pp. 31–43, 2009.

[10] Pita, J., Jain, M., Western, C., Portway, C., Tambe, M., Ordóñez, F., Kraus, S., & Paruchuri, P., Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *AAMAS-08* (*Industry Track*), Proceedings of the International Foundation for Autonomous Agents and Multiagent Systems, pp. 125–132, 2008.

[11] Bier, V.M., Haphuriwat, N., Menoyo, J., Zimmerman, R., & Culpen, A.M., Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Analysis*, **28(3)**, pp. 763–770, 2008.
[12] Cox, L.A. Jr., Making telecommunications networks resilient against terrorist attacks (Chapter 8). *Game Theoretic Risk Analysis of Security Threats*, eds Bier, V.M. and Azaiez, M.N., Springer: New York, pp. 175–198, 2009.
[13] Kunturska, U., Schmöcker, J., Fonzone, A., & Bell, M.G.H., Improving reliability through multi-path routing and link defense (Chapter 9), *Game Theoretic Risk Analysis of Security Threats*, eds Bier, V.M. and Azaiez, M.N., Springer: New York, pp. 13–32, 2009.
[14] Israeli, E. & Wood, K. Shortest-path network interdiction. *Networks*, **40**, pp. 97–111, 2002.
[15] Salmeron, J., Wood, K., & Baldick, R., Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems*, **19(2)**, pp. 905–912, 2004.
[16] Bier, V.M., Gratz, E.R., Haphuriwat, N.J., Magua, W., & Wierzbicki, K.R., Methodology for identifying near-optimal interdiction strategies for a power network transmission system. *Reliability Engineering and System Safety*, **92(9)**, pp. 315–323, 2007.
[17] Salmeron, J., Wood, K., & Baldick, R., Worst-case interdiction analysis of large-scale electric power grids. *IEEE Transactions on Power Systems*, **24(1)**, pp. 96–104, 2009.
[18] Brown, G.G., Carlyle, W.M., Salmerón, J., & Wood, K., Defending critical infrastructure. *Interfaces*, **36(6)**, pp. 530–544, 2006.
[19] Bell, M.G.H., Kanturska, U., Schmöcker, J.D., & Fonzone, A., Attacker–defender models and road network vulnerability. *Philosophical Transactions A*, **366**, pp. 1872–1893, 2008.
[20] Liu, D., Wang, X.F., & Camp, J., Game-theoretic modeling and analysis of insider threats. *International Journal of Critical Infrastructure Protection*, **1(1)**, pp. 75–80, 2008.
[21] Wood, R.K., Deterministic network interdiction. *Mathematical and Computer Modelling*, **17(2)**, pp. 1–18, 1993.
[22] Yao, Y., Edmunds, T., Papageorgiou, D., & Alvarez, R., Trilevel optimization in power network defense. *IEEE Transactions on Systems, Man and Cybernetics-Part C: Applications and Reviews*, **37(4)**, pp. 712–718, 2007.
[23] Smith, J.C., Sudargho, F., & Lim, C., Survivable network design under various interdiction scenarios. *Journal of Global Optimization*, **38(2)**, pp.181–199, 2007.
[24] Church, R.L. & Scaparra, M.P., Protecting critical assets: The r-interdiction median problem with fortification. *Geographical Analysis*, **39(2)**, pp. 129–146, 2007.
[25] Brown, G.G., Carlyle, W.M., Salmerón, J., & Wood, K., Analyzing the vulnerability of critical infrastructure to attack and planning defenses. *Tutorials in Operations Research*, pp. 102–123, 2005.
[26] Brown, G., Carlyle, W.M., & Wood, R., Optimizing Department of Homeland Security defense investments: Applying defender–attacker (–defender) optimization to terror risk assessment and mitigation (Appendix E). *Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change*, The National Academies Press: Washington, DC, pp. 90–102, 2008.
[27] Kunreuther, H. & Heal, G., Interdependent security. *Journal of Risk and Uncertainty*, **26(2)**, pp. 231–249, 2003.
[28] Heal, G. & Kunreuther, H., Modeling interdependent risks. *Risk Analysis*, **27(3)**, pp. 621–634, 2007.

[29] Zhuang, J., Bier, V.M., & Gupta, A., Subsidies in interdependent security with heterogeneous discount rates. *The Engineering Economist*, **52(1)**, pp. 1–19, 2007.
[30] Bier, V.M., Choosing what to protect. *Risk Analysis*, **27(3)**, pp. 607–620, 2007.
[31] Keohane, N.O. & Zeckhauser, R.J., The ecology of terror defense. *Journal of Risk and Uncertainty*, **26(2)**, pp. 201–229, 2003.
[32] Hausken, K., Probabilistic risk analysis and game theory. *Risk Analysis*, **22(1)**, pp. 17–27, 2002.
[33] Hirshleifer, J., From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, **41(3)**, pp. 371–386, 1983.
[34] Ezell, B.C., Bennett, S.P., von Winderfeldt, D., Sokolowski, J., and Collins, A.J., Probabilistic risk analysis and terrorism risk. *Risk Analysis*, **30(4)**, pp. 575–589, 2010.
[35] Bier, V.M., Cox, L.A., & Azaiez, M.N., Why both game theory and reliability theory are important in defending infrastructure against intelligent attacks (Chapter 1). *Game Theoretic Risk Analysis of Security Threats*, eds Bier, V.M. and Azaiez, M.N., Springer: New York, pp. 1–11, 2009.
[36] Parnell, G.S., Borio, L.L., Brown, G.G., Banks, D., & Wilson, A.G., Scientists urge DHS to improve bioterrorism risk assessment. *Biosecurity and Bioterrorism*, **6(4)**, pp. 353–356, 2008.
[37] National Research Council, *Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change*, The National Academies Press: Washington, DC, 2008.
[38] Cox, L.A. Jr., Game theory and risk analysis. *Risk Analysis*, **29(8)**, pp. 1062–1068, 2009.
[39] Banks, D.L. & Anderson, S., Combining game theory and risk analysis in counterterrorism: A smallpox example. *Statistical Methods in Counterterrorism: Game Theory, Modeling, Syndromic Surveillance, and Biometric Authentication*, eds Wilson, G., Wilson, G.D., and Olwell, D.H., Springer: New York, pp. 9–22, 2006.
[40] Guikema, S.D., Game theory models of intelligent actors in reliability analysis (Chapter 2). *Game Theoretic Risk Analysis of Security Threats*, eds Bier, V.M. and Azaiez, M.N., Springer: New York, pp. 13–32, 2009.
[41] Parnell, G.S., Borio, L.L., Cox, L.A., Brown, G.G., Pollock, S., & Wilson, A.G., Response to Ezell and von Winterfeldt. *Biosecurity and Bioterrorism*, **7(1)**, pp. 111–112, 2009.
[42] Ezell, B.C. & Winterfeldt, D., Probabilistic risk analysis and bioterrorism risk. *Biosecurity and Bioterrorism*, **7(1)**, pp. 108–110, 2009.
[43] Guikema, S.D. & Aven, T., Assessing risk from intelligent attacks: A perspective on approaches. *Reliability Engineering & System Safety*, **95**, pp. 478–483, 2010.
[44] Box, G.E.P. & Draper, N.R., *Empirical Model-Building and Response Surfaces*, Wiley: New York, 1987.