The NISPAcee Journal of Public Administration and Policy, Vol. XIII, No. 2, Winter 2020/2021

10.2478/nispa-2020-0024

Using Game Theory in Public Domains: The Potential and Limitations of Security Games

Zuzana Špačková[1], David Špaček[2]

Abstract

Since its origins, when it was mainly connected to the field of economics, game theory has brought important theoretic insights into many domains. Besides biology, philosophy or computer science, its findings have been applied to various fields of public policy. One specific area of public policy is that of security. Within the last two decades we have been witnesses to a significant increase in efforts to model security issues using tools of game theory and to derive political implications. The paper deals with the model of a Stackelberg security game and its real-world applications in security domains. The main aim and purpose of the paper is to provide a survey of selected cases of real-world deployed applications of the game-theoretic Stackelberg model in the area of public security and, based on the literature analysis, to discuss the potential and limitations of the model for policy- and decision-makers that are dealing with security measures on various governmental levels. Existing cases clearly indicate that the model can contribute to a better design and implementation of the security policy and help implement a better allocation of resources and thus potentially improve the effectiveness of security measures. On the other hand, the paper also discusses some limitations and potential future adjustments of the model together with points for further research.

Keywords:
game theory, sequential games, Stackelberg game, security systems, national security

1 University of Defence in Brno, Czech Republic.

2 Faculty of Economics and Administration, Masaryk University, Czech Republic.

1. Introduction

Security is a concern of major importance to governments (Tsai et al. 2009) and therefore an important topic for public policy on the national as well as the local level. Protecting inhabitants and national infrastructure are crucial tasks for each government. Airports, public-transportation networks or large public events are a few examples of potential targets vulnerable to terrorist attacks. When deciding which target to protect, security forces face the problem of limited security resources, which prevent non-stop coverage of all strategic targets. Taking, for example, an airport with multiple terminals, it has to be decided by security decision-makers which terminal to patrol in which moment, because usually it is not possible to cover all terminals, or areas, on a permanent basis. Let us suppose that police decide to patrol one terminal (let us consider the most important one) constantly and leave the others uncovered. An adversary who seeks to harm the airport observes that only one terminal is patrolled and decides to strike at an uncovered one. If police change the strategy in a way that they switch to protect another terminal, the adversary will discover the new pattern as well and will attack a currently uncovered terminal (probably the most important one). Therefore, a way to improve the security of the airport is via a randomization of police actions, i.e. creating randomized patrolling schedules which would cast the adversary into uncertainty about the coverage of terminals. The idea behind this is that in the case of randomized police actions, it is impossible for the adversary to detect any pattern of patrolling, thus making it difficult to choose a target to attack. And this can be used in various fields of national security, which is a rather complex public policy.

Other examples of situations in which patrolling schedules need to be designed in order to deter (or apprehend) potential adversaries are problems of green security (e.g. protection against poachers), deterrence of fare evaders in urban transportation or protection of vulnerable urban areas. Available literature and the cases selected for this paper indicate that these situations can be modelled using means of game theory, and as such, game theory may contribute to designing and implementing security policies in various public domains. Game theory provides useful analytic tools in many domains: economics, biology, law, public policy and others. Global threats of terrorism, drug-smuggling, and other crimes have led to a significant increase in research on game theory for security (Tambe 2011). This is important, because, for instance, realizing homeland security depends on a certain vision of security threats (Danila 2013). Game theory may help visualize and describe scenarios that can be used in the formulation of policies and instruments for achieving security, public safety etc., and for their actual use.

Security and crisis/emergency management policies and measures are usually an interdisciplinary theoretical and practical field (as can clearly be seen in explanations of homeland/national security policies – e.g. Noftsinger et al. 2007; Morag 2011; Cross, 2007). Such policies are also intergovernmental (May et al. 2011) – in

order for their measures to be effective and more efficient they require cooperation and coordination of multiple administrative levels and other actors. Therefore not only the role of national government in homeland/national security policy has been discussed in the security-policy literature, but also the role of local governments and local actors. For instance, with regards to the United States, it was pointed out that the greater the national security threats, the more important the role of local policy (Clarke and Chenoweth 2006). Also the role of the supra-national level has been emphasized – e.g. because the EU has increased its efforts to build common crisis-management capacity across the continent and has been seeking to both coordinate national crisis and disaster authorities and build its own supranational capacities (Rhinard and Boin 2009). Literature has also discussed whether security measures of individual governmental levels are mutually compatible and what determines the potential effectiveness of security policy across federal and decentralized systems (e.g. Friedmann and Cannon 2007; Gerber et al. 2007; Chappell and Gibson 2009; Birkland 2009). Thus, researchers may discuss if using game-theory tools can help increase the compatibility and the effectiveness of the implementation of security measures, e.g. through better allocation of shared resources in dealing with security issues. As Reddick (2008) concluded, security management can be improved when there is organizational collaboration. Cooperation and coordination between the European states and various governmental levels has been on the rise, and these intergovernmental aspects of the emergency and security policies can also be seen in the current COVID-19 measures (Špaček 2020, forthcoming).

Where security is concerned, the value of the Stackelberg model has been frequently discussed in available literature. This model represents a natural approximation of real-world situations in which adversaries first engage in observing security measures taken by security forces and then commit to an attack. Algorithms based on a Stackelberg security model can help policy- and decision-makers as well as public managers in the area of security in deciding which protective strategy to adopt in situations where the protection of multiple targets (combined with limited security resources) is concerned. They can also allow for more policy learning (Kamra et al. 2018).

The methodology of this paper is exploratory and based on secondary data analysis, not on our own empirical research. We did a literature search on Web of Science, Scopus and Google Scholar in order to find and obtain academic papers and book chapters dedicated to explanations and testing of the Stackelberg security games. But we did not mean the paper to be an exhaustive study of existing literature. Our literature search clearly indicated that growing attention has been given to the Stackelberg model in available literature. We prepared this paper with the main aim to provide a survey of selected cases of real-world deployed applications of the game-theoretic Stackelberg model and, also, to discuss some further potential applications based on it in a specific public domain, security. Another motivation to prepare the paper was the fact that the Stackelberg model has been applied es-

pecially in Western democracies (and especially in the USA – Kar et al. 2017), and such cases may be inspiring for creators of public policies in the security areas in the Central and Eastern European (CEE) region.

The paper is divided into a theoretical and an applied part. In its theoretical part, the model of Stackelberg security games is presented in detail, together with the necessary basis of game theory in order to be more readable for a larger audience. Selected real-world applications of the model are discussed in the applied part of the paper. All these applications are based on a general model of a (Bayesian) Stackelberg security game, incorporating its more or less important modifications and adjustments. Possibilities of evaluating these real-world deployed security systems are discussed at the end of the paper together with implications for future research.

2. Game theory and the Stackelberg Model

The origins of game theory date to the first half of the 20th century and are connected to the personality of John von Neumann who, in 1928, published his paper "Zur Theorie der Gesellschaftsspiele" (Von Neumann 1928), and in 1944, together with Oskar Morgenstern, the famous *Theory of Games and Economic Behavior* (Von Neumann and Morgenstern 1944). Another very significant and well-known personality of game theory formation was John Nash, who, besides other things, provided a solution to non-cooperative games which bears his name and has widely become known as the Nash equilibrium.

At the time of its formation, game theory was mainly connected to the field of economics. Since then, it has brought important theoretic insights to other fields as well, e.g. biology, political science, philosophy, computer science etc. Within the last two decades we have borne witness to a significant increase in efforts to model security issues using tools of game theory in order to derive practical (political) implications based on them.

The games are most commonly modelled using so-called the normal form of a game (see, e.g., Tadelis 2013). A normal-form game consists of three sets: the set of players ("participants" of the game), the set of all potential strategies of each player and the set of players' individual payoff functions. Each game player chooses a set of his possible strategies (actions that he may take), also referred to as pure strategies. The payoff function of a player defines his payoff depending on his strategy and the strategies of all other players of the game. We suppose that all players seek to maximize their (expected) individual payoffs and that they possess full information about their co-players' strategy sets and payoff functions.

Normal-form games are most commonly represented by an n-dimensional array. In the case of two players, the game is represented by a matrix referred to as a payoff matrix. A payoff matrix includes information about the number of play-

ers (i.e. two), their sets of pure strategies and the payoffs for each player relative to each possible outcome (i.e. the combination of player one's strategy and player two's strategy). The goal of each player in the game is to maximize their expected payoff by choosing the optimal strategy. A dominant (pure) strategy of a player is the one that provides him the best outcome compared to all his other strategies (the dominated ones) regardless of the strategy chosen by a co-player. If there is such a strategy for a player, they will play it with certainty, as it is their optimal (pure) strategy. In case all players have optimal pure strategies the result of the game and the payoffs are deterministic. If there is no optimal pure strategy for a player, they commit to a mixed strategy which is a probability distribution on the set of their pure strategies (Owen 2013).

The Stackelberg model was first introduced for studying duopoly competition (Von Stackelberg 1934) as a dynamic extension of the Cournot model and was represented by a non-symmetric sequential game in which two firms made decisions, in a given order, about their quantities (Webster 2014). In a Stackelberg game, two players are considered: a leader and a follower. These players are not necessarily individuals; they may as well be groups cooperating on the basis of a common strategy (e.g. police on one side and a criminal organization on the other side). The leader commits to a strategy at first. In the second step, the follower chooses his strategy in order to optimize his expected reward. Since he can observe the leader's choice and respond to it (in a way to optimize his expected payoff), one may consider that the follower is somehow advantaged. However, the advantage is on the side of the leader, as demonstrated by the payoff matrix given in Table 1.

Table 1
Payoff matrix of a leader-follower game

		Follower	
		c	d
Leader	a	3, 1	5, 0
	b	2, 0	4, 2

Source: Author

The leader (a row player) has two possible pure strategies: a and b, and the follower's (a column player's) strategy set contains c and d. The players' individual payoffs corresponding to all possible outcomes of the game (2x2 possibilities) are given by the payoff matrix. (For example, if the leader decides to play strategy a and the follower chooses d, the leader's payoff is 5 and the follower's 0.)

One can observe that if the game is played as a simultaneous one (both players deciding simultaneously without knowing the co-player's decision), it has only

one pure strategy, the Nash equilibrium, with the leader playing a and the follower playing c. In the equilibrium, the leader's payoff is 3 and the follower obtains 1. For the leader strategy b is strictly dominated by strategy a, which means that whatever strategy the follower chooses, strategy a always guarantees a better outcome for the leader. (The follower does not have a dominant or dominated strategy. However, as they know that the leader is going to play a, their best response is c.)

In the case of playing a sequential game, the leader may commit to playing strategy b, and the follower, after having observed the action chosen by the leader, responds by playing d (which ensures them a higher payoff). This results in payoff 4 for the leader and 2 for the follower. If the leader commits to a mixed strategy of playing a and b with the same probability of 0.5, the follower will maximize their expected payoff by playing d (ensuring them the expected reward of 1 against 0.5 if choosing c) and the leader's expected payoff will increase to 4.5. The leader's strategy randomization thus leads to an increase in their expected reward while increasing the follower's uncertainty at the same time.

The solution of the Stackelberg game is called Stackelberg equilibrium. A strong Stackelberg equilibrium assumes that the follower has perfect knowledge of the leader's mixed strategy and that they react fully rationally to it in order to maximize their expected reward. In a strong Stackelberg equilibrium, none of the players have incentives to change the strategy, as for the leader it is given by their optimal mixed strategy and for the follower it is their optimal response to the defender's mixed strategy.

3. Modelling real-world security scenarios as a Stackelberg Game: security games

The Stackelberg Security Game model has been immensely influential in security research since it was introduced roughly a decade ago (Sinha et al. 2018). Stackelberg games are well suited to studying security issues as they provide a natural approximation of the real-world security scenarios (An et al. 2012).

So-called security games are non-zero-sum games motivated by real-world security domains, built upon the model of the Stackelberg game. The motivation behind this type of games is that we have limited security resources, which must protect multiple potential infrastructure targets of varying importance (Tambe 2011). Our resources do not allow us to protect (cover) all potential targets at all times so we must deploy them selectively. However, our adversaries may observe our actions (i.e. the schemes of protection we have adopted) and adapt their conduct (e.g. plan of attacks) in order to improve its effectiveness. The way here how to weaken the position of potential adversaries is via a randomization of (protective) actions. Unpredictability over defensive actions increases the adversaries' uncertainty and, consequently, improves the protection of the targets.

Modelling such defender-attacker paradigms through game theory may help policy- and decision-makers as well as public managers in the area of security. They may help policy-makers as well as implementors (including security forces) think about the allocation of limited resources to protect infrastructure while considering different weights of various targets and an adversary's responses to different protective strategies. As a solution to the problem, they may obtain a weighted randomization strategy that aims at increasing efficiency and effectiveness of resources allocated to public security.

The two players in a Stackelberg security game are a defender (in reality, the given type of security forces) on one side and an attacker (adversary) on the other side. As pointed out in Tambe (2011), "security games have the characteristic that what is good for the attacker is bad for the defender and vice versa". However, we do not require the loss to the attacker to be symmetric with the gain to the defender (and vice versa), i.e., in general, we do not consider this game to be zero-sum.[3] There has been a discussion in the literature related to the question why these games are non-zero-sum (see, e.g., Powell 2007); an elementary explanation being, for example, that the adversary may view some targets as especially important for his audience, while these may not be of such importance to the defender (police). Or, as pointed out in Tambe (2011), even a failed attack may not be seen by the adversary as the worst outcome because of publicity related to this attempt or the fear it generates. The defender has several targets of varying importance to protect (cover) and limited resources restricting them to protect only some of them at a time (against the attacker). Each possible allocation of the resources represents one pure strategy on the defender's part. The attacker chooses a target from a target set and each target chosen corresponds to one pure strategy. The basic logic behind the defender's and attacker's payoffs is that if the attacker attacks a target which is protected by the defender, their payoff is worse compared to a situation in which the target was not covered. And, reversely for the defender, if a non-protected target is attacked, the defender's payoff is worse compared to a situation in which the target was covered.

To model such a situation using game theory, Tambe (2011) introduces a simple example of a small airport having two terminals (1 and 2) and only one police unit and one adversary. The police unit can protect only one terminal at a time while the adversary can attack one target. Terminal 1 represents a more important target than Terminal 2. The situation is illustrated by Table 2.

3 A game is called zero-sum if the sum of the payoff for all players equals zero in any outcome. In the case of a two-player game this means that the gain of one player is equal to the loss of the other player (see, e.g., Prisner 2014 or Shor 2005).

Table 2
Payoff matrix of a security game

		Adversary	
		Terminal 1	Terminal 2
Defender	Terminal 1	5, –3	–1, 1
	Terminal 2	–5, 5	2, –1

Source: Tambe (2011)

If the police protect Terminal 1 constantly (as it is the more important one), the adversary detects this deterministic strategy and commits to attacking Terminal 2, which yields him a payoff of 1 (and to the defender a loss of –1). If the police switch to protect Terminal 2 all the time, the adversary will observe this pattern again and will attack Terminal 1 (resulting in payoffs amounting to –5 for the defender and 5 for the attacker). The solution here for the police is to randomize their actions and commit to a mixed strategy. If the police protect Terminal 1 within 60 % of the days (and Terminal 2 during 40 %) it will bring them a better result. The adversary will have observed that the police protects Terminal 1 with a probability of 0.6 and Terminal 2 with a probability 0.4, however, he does not know which terminal will be protected at the very moment of his planned attack. The adversary's uncertainty rises and so does the expected reward for the police. (The adversary's best response in this case would be to attack Terminal 1, which would lead to an expected payoff of 1 for the police.)

It needs to be stressed here that this model assumes the adversary's perfect knowledge of the defender's mixed strategy, and that they will react rationally to this strategy, i.e., the adversary will act in order to maximize their expected payoff. The key question to the defender is which allocation of limited security resources (i.e. which vector of mixed strategy) will be the optimal one. In the simple example above (which represents a 2x2 game), the problem may be quite easily solved by hand. However, the task becomes more complex in real situations in which we may consider hundreds of targets and multiple police units.

In real-world settings, a security force (a defender) usually must face multiple potential attackers who have different evaluations over targets' importance. In "Bayesian Stackelberg games" the uncertainty over different adversary types adds to the model. In a two-player game, there is one leader (defender) type and several follower (attacker) types. The leader is not familiar with the follower's type; however, he knows the probabilities of each follower type's appearance in the game. Consequently, instead of one, we obtain multiple payoff matrices, each corresponding to a different adversary type (and different valuation of targets).

Tambe (2011) demonstrates this situation by a scheme given in Table 3.

Table 3
Payoff matrices of a Bayesian Stackelberg Game

		Adversary TYPE 1	
		Terminal 1	Terminal 2
Defender	Terminal 1	5, −3	−1, 1
	Terminal 2	−5, 5	2, −1

		Adversary TYPE 2	
		Terminal 1	Terminal 2
Defender	Terminal 1	1, −2	−2, 3
	Terminal 2	−3, 5	3, −1

		Adversary TYPE 3	
		Terminal 1	Terminal 2
Defender	Terminal 1	4, −2	−3, 3
	Terminal 2	−5, 5	2, −2

Source: Tambe (2011)

Knowing the probabilities for the appearance of each follower type in the game, a Bayesian game can be transformed into a normal-form game using Harsanyi transformation (see, e.g., Wilczyński et al. 2016). This way we obtain a transformed game comprising one payoff matrix with the leader having two strategies and the single follower, eight.

The key challenge here is, again, to find the optimal mixed strategy (i.e. optimal randomized allocation of limited resources) to the defender which maximizes his expected utility. As has been stated, security forces (police, army, security agencies etc.) have limited resources, which prevents them from protecting all strategic targets at all times. The targets to be covered must thus be selected, which leaves the unprotected targets vulnerable to adversaries. Moreover, adversaries may carry out surveillance over the strategies implied by the defender and aim their attack towards the targets which show up to be unprotected. The way to improve the security

of targets is via increasing the adversary's uncertainty over the defender's actions. The defender may achieve this goal by randomizing his actions. In order for this randomization to be effective, the weights of importance of the targets should be taken into consideration. When simple situations (such as the example of a small airport) are concerned, this randomization may be made by hand. However, the task becomes more difficult with an increasing number of possible targets (to cover) and/or adversary types. Besides, human scheduling relates to several shortcomings. In addition to the burden to officials executing such manual scheduling there is a risk for them to tend to adhere to patterns and routines. Computational techniques and fitting software prove to be more precise and efficient. Suitable software performs with considerably higher speed, is adapted to properly weigh the costs and benefits of different actions and is ready to process various input variables.

The following sub-section provides a survey of selected real-world cases in which such software based on Stackelberg games contributes to optimally allocating limited security resources. All these real-world deployed applications are based on the general model of the Bayesian Stackelberg security game using more or less important modifications of or adjustments to it.

4. Real-world applications of a Stackelberg Game

a) The Los Angeles Airport and ARMOR

The Bayesian Stackelberg model lies at the heart of the Assistant for Randomized Monitoring Over Routes (ARMOR) implemented by the Los Angeles International Airport since 2007. ARMOR represented the very first real-world deployed application of the model. This software assistant helps the police by providing a randomized pattern for setting up checkpoints and canine unit patrols.

Los Angeles International Airport is the world's fourth busiest airport and the second busiest airport in the USA, serving more than 80 million passengers per year (for the last three years, see https://www.lawa.org/en/lawa-investor-relations/statistics-for-lax/volume-of-air-traffic) which makes of it an attractive potential target to attackers.

There are six inbound roads into Los Angeles International Airport, and limited security resources prevent police from running permanent checkpoints at each of them. Similarly, there are nine terminals and the police do not have enough canine units (searching for potential explosives) to operate at all terminals all day. Access roads and terminals are potential targets to adversaries, while there are multiple types of the latter.

Modelling the tasks in question (i.e. to optimally set checkpoints or canine patrol routes at the Airport) as a Bayesian Stackelberg game makes it possible to randomize police actions, taking into consideration their weights (determined by

their complex costs and benefits) as well as uncertainty over adversary types. In terms of game theory, the goal is to find an optimal mixed strategy for the leader to commit to, given that the follower may know this mixed strategy when choosing his strategy (Pita et al. 2008). In order to do so, ARMOR relies on an algorithm called DOBSS (Decomposed Optimal Bayesian Stackelberg Solver), which operates directly on the Bayesian representation without a need for Harsanyi transformation (see, e.g., Wilczyński et al. 2016) and solves the problem using linear programming.

The architecture of DOBSS consists of four key components: a front-end interface through which an officer can insert the information, a method for creating Bayesian Stackelberg game matrices, an implementation of DOBSS and a method for producing suggested schedules (Pita et al. 2008).

b) Federal Air Marshals Service and IRIS

While ARMOR generates randomized schedules for controlling vehicles arriving at the Los Angeles International Airport and canine patrols at the airport, other crucial security measures in transportation networks (airplanes, buses, trains) are represented by onboard patrols. Basically, from a game-theoretic point of view, we face a similar problem of limited security resources (i.e. security staff available) over multiple potential targets. However, whilst in the case presented above the number of targets amounted to less than ten (six inbound roads or eight terminals), there are hundreds or thousands of vehicles to protect in transportation networks.

An example of a deployed system in transportation networks based on game theory is the one used by Federal Air Marshal Service (FAMS) in the USA since 2009. Federal Air Marshals are armed federal law enforcement officers deployed on passenger flights worldwide to protect airline passengers and crew against the risk of criminal and terrorist violence (see https://www.tsa.gov/about/jobs-at-tsa/federal-air-marshal-service-and-law-enforcement). The software scheduling assistant used by the FAMS is called Intelligent Randomization In Scheduling (IRIS) system. In a similar way to ARMOR, it models the problem in question (that is, how to optimally allocate available officers on flights) based on a Stackelberg game. However, there are additional challenges relative to transportation networks which need to be taken into consideration. These challenges relate to (1) the big number of potential targets that need to be processed by the system (tens of thousands of commercial flights per day), (2) lots of values to be entered (tens of thousands of payoff values) and (3) hard scheduling constraints (the targets are moving vehicles). The DOBSS algorithm lying at the heart of ARMOR would not be able to handle such extensions in an efficient way. The IRIS system incorporates all three challenges mentioned above.

Initially, IRIS used the ERASER-C (Efficient Randomized Allocation of SEcurity Resources – Constrained) algorithm, a mixed-integer linear programme (for

details see Tsai et al. 2009) before it switched to ASPEN (Accelerated SPars ENgine), based on the branch and price framework (see Jain et al. 2010).

The IRIS system consists of four modules: (1) an input module in which four classes of data need to be inserted (so that a representative Stackelberg game is generated): resource data (coverage ability and number of FAMs), target data (flight data), data required for risk assessment (number of passengers, flight path etc.) and supplementary data; (2) a back-end module comprising six components: forming of the target definition for the game, payoff generation process, translation into a Stackelberg game, model solving, generation of a randomized schedule of probability weights for each target and creation of actual sample schedules fort the FAMS. This information is displayed via (3) a display/output module. (4) A project manager is a project-based system in which the input files can be stored (Tsai et al. 2009).

c) United States transportation security administration and GUARDS

A case of national-scale security deployment of a game-theoretic approach is that of the software system GUARDS (Game-theoretic Unpredictable and Randomly Deployed Security) used by the United States Transportation Security Administration (TSA). TSA was created in 2011, in reaction to terrorist attacks in New York on 11 September, with the mission to protect the nation's transportation systems to ensure freedom of movement for people and commerce (see https://www.tsa.gov/). It focuses primarily on airport security and the prevention of aircraft hijacking. (The above-mentioned Federal Air Marshals Service comes under the TSA's supervision.) Thus the TSA mission covers more than 450 U.S. airports.

The software scheduling assistant, GUARDS, designs the task of allocating the TSA's limited resources over hundreds of security activities as a Stackelberg game in which the TSA acts as a defender (leader) having a set of targets to protect, a number of security activities for protecting each target and limited resources to be assigned to these security activities (Pita et al. 2011). An attacker observes the TSA's resource allocation and then chooses a target to attack. What is different compared to the above-mentioned approaches, is that the model employed allows for both heterogeneous security activities and threats. There are multiple areas at the airport (ticketing areas, waiting areas, cargo holding areas) for which the defender chooses among several security activities (perimeter patrols, screening cargo and employees etc.) while he may allocate more than one resource per area. Similarly, an adversary may execute heterogeneous attacks on an area (chemical weapons, active shooters, bombs). An improved model taking account of these extensions is referred to as Security Circumvention Games (SCGs). In SCGs, the TSA must choose some combination of security activities to execute within each area, and the attacker must reason over both which area to attack and which method of attack to execute based on the defender's strategy (Pita et al. 2011). When computing the solutions, a compact representation of SCGs is used (in order to reduce the number of the de-

fender's strategies to be enumerated). Another challenge that the GUARDS system copes with is that of acquiring the appropriate knowledge needed to be considered. A two-phase knowledge acquisition process is employed, combining a centralized approach (gathering information common among all airports) and a decentralized one (inputting customized information by individual airports).

The system consists of three modules: 1) an input module in which quantifiable inputs (area data, security activities data and resource data) need to be inserted; 2) a back-end module which generates the game, solves it and returns a sample schedule; and 3) a display/output module.

d) United States coast guard and PROTECT

Another deployed game-theoretic system is that used by the United States Coast Guard, in which a real maritime patrolling problem is modelled as a Stackelberg game.

The United States Coast Guard (USCG) look after the security of the coasts, ports and inland waterways in the U.S.A., which comprise many potential targets to attackers. Obviously, the USCG do not possess enough resources to cover all these targets permanently, and adversaries may observe any patterns in patrolling actions.

An application called Port Resilience Operational/Tactical Enforcement to Combat Terrorism (PROTECT) has assisted the USCG in allocating its patrolling resources at the Port of Boston since 2011. Building on a Stackelberg game framework, PROTECT provides the solution as randomized patrol patterns which take into account the importance of different targets, the adversary's observation of actions taken by the USCG and their anticipated response to them.

In comparison with the software types presented above, PROTECT introduces several advances. The most significant one is that it departs from the assumption of adversaries' perfect rationality. Adopting findings from behavioural game theory, it replaces this assumption with a model of a boundedly rational adversary using a quantal response model. As stated in An et al. (2012), this model suggests that instead of strictly maximizing utility, individuals respond stochastically in games: the chance of selecting a nonoptimal strategy increases as the cost of such an error decreases.

An algorithm called PASAQ (see Yang et al. 2012) is used and improved by a compact representation of defender strategies based on an analysis of their equivalence and dominance.

PROTECT generates randomized schedules of patrols including the time when they should begin, the targets to be visited by each patrol and activities to be performed at each of them. Besides experimental analyses, a real-world data provided by the USCG prove this system to be successful (An et al. 2012).

e) Los Angeles metro rail system and TRUSTS

The deployment of controls in a proof-of-payment transit system where passengers may enter and travel without actually having purchased a ticket is another problem which can be modelled using a Stackelberg game. Of course, the character of "danger" is completely different to the above-mentioned situations. In this case, human lives are not in danger when fare evasion occurs, but public finance is. Otherwise, the assumptions to model the problem using a leader-follower game are met: the government do not have enough resources to cover all transit lines with controls, and so patrols are set selectively. Any patterns in such a setting may be observed and exploited by potential fare-evaders. The only possibility to discourage passengers from fare-evading, given the limited resources, is to efficiently randomize the controls and minimize their predictability this way. Randomized schedules are often human-made; however, this constitutes a considerable burden to the officers involved and, considering all the complexities in question, the outputs hardly represent an optimal allocation. Modelling the problem using game theory helps generate more efficient patrol schedules.

An example of a proof-payment transit system is the Los Angeles Metro Rail system. In 2012 an application called Tactical Randomization for Urban Security in Transit System (TRUSTS) was deployed there. TRUSTS models the situation as a Stackelberg game with one leader (the Los Angeles Sheriff's Department) and many followers, while each of them takes a fixed route at fixed time. The leader commits to a mixed patrol strategy, and passengers, after observing it, decide whether to buy a ticket or not so that they minimize their expected total cost. The leader's objective is to maximize total revenue (tickets sold and penalties). What is new to the model is a great number of possible pure patrol strategies, each being subject to spatial and temporal constraints. In order to cope with this difficulty, TRUSTS uses the transition graph capturing spatial and temporal structure of the domain and calculates the optimal flow through this graph using linear programming. What is more, an extension is applied of a history-duplicate transition graph, which forbids too long patrols and penalizes patrols having too many switches (Yin et al. 2012).

f) Green security games

Another domain where game theory (more specifically Stackelberg games) may be applied is that of so-called green security. This domain covers, for example, protecting fisheries form over-fishing or protecting wildlife from poaching. As Fang (2015) points out, there are several key features characteristic to this domain, deviating from assumptions of the basic model. There are multiple adversaries whose actions (attacks) are repeated and frequent, which exceeds the one-shot model of Stackelberg security games. Furthermore, the attackers generally do not invest themselves in extensive surveillance and spend less time and effort in each attack. It is thus important to take account of their bounded rationality and bounded surveillance

when modelling the situation. Another key feature is that there are more attack data available in green security domains compared to infrastructure domains, which enables us to study the attackers' decision-making from the data (Fang 2015).

It is stressed that the Stackelberg assumption (i.e. adversaries fully observe the defender's mixed strategy conducting extensive surveillance before each attack) can be unrealistic in green security domains, and as there may be a lag in the attacker's observing, it may be valuable to the defender to plan ahead regarding their actions. They introduce a novel model called Green Security Games (GSGs) in which it is assumed that the attacker's understanding of the defender's strategy may not be up-to-date, and thus it can be approximated as a convex combination of the defender's actions in recent rounds. As such, CSGs provide more generality and allow the defender to plan their strategies (Fang 2015).

In Fang (2015), the author further proposes two algorithms which plan ahead; one planning a fixed number of steps ahead and the other providing a short sequence of strategies for repeated execution.

g) Network security games

Moving assets through a hostile transportation network is another type of situation which can be modelled using game theory. According to Okamoto et al. (2012) such hostile network environments are characterized by six key features:

(1) the topology of the network creates exponentially sized strategy spaces (it is no longer possible to find optimal strategies using normal form techniques),

(2) security is not the only criterion which needs to be considered (performance objectives need to be taken into consideration),

(3) patterns of behaviour may be detected by the adversary,

(4) adversaries are rational agents (balancing the harm they can inflict with attack costs),

(5) information on the adversaries (their capabilities, payoffs and costs) is rarely available,

(6) there may be multiple adversaries with various abilities.

The problem is modelled as a game between a sender and an adversary in which the sender chooses a flow through the network, via which he would move the assets, and the adversary chooses one or more attacks that do harm to one or more link in the network. The sender seeks to minimize the harm while the adversary strives to maximize the harm reduced by the attack costs.

The authors follow up on previous work in network security games and extend it with several new approaches and contributions. They assume non-zero payoffs because of attack costs borne by the attacker. Furthermore, they introduce a re-

finement of the Stackelberg equilibrium (the locally optimal inducible Stackelberg equilibrium), which is more suitable for network security games. The authors also consider games of incomplete information (about the adversary's capabilities, payoffs and costs) and formulate them as Bayesian games (Okamoto et al. 2012).

Similarities can be found in Ge et al. (2019). In their paper they found that the unequal information between the attack and defence sides of the network is not taken into account in the research of network-security defence technology in the game model. They argue that by establishing the network model, the strong equilibrium strategy algorithm of a Stackelberg Security Game with incomplete information is used to generate the optimal defence strategy of the network. The defender can actively release the incentive strategy without reducing the accuracy of defence strategy generation. According to them, the experimental results provide a method for the selection of the optimal defence strategy and can strengthen the security of the system.

h) Security at Large Public Events

Stackelberg games may as well serve as a basis when modelling protection at large public events. Such events, mainly in major cities, where masses of people take part, represent attractive targets for terrorist attacks. (Yin et al. 2014 cite an example of the Boston Marathon bombings in April 2013.) It is thus greatly important to deploy limited security resources in a way to maximize the protection of these events.

A significant feature of large public events is their dynamic nature. Firstly, the importance of targets changes over time. In the case of a marathon, the number of participants and spectators at any specific area changes in the course of the race (as the runners advance continuously through it). Secondly, the attacker may attack at any moment, and the defender is able to relocate the resources among targets at any time. This means that the strategy space of both the defender and the attacker is continuous and infinite. And finally, as these events take place quite rarely, the attacker may not be able to observe the defender's (mixed) strategy and respond to it. Therefore, Yin et al. (2014) propose algorithms which compute the optimal pure (dynamic) strategy for the defender as a robust solution to the problem. They design a game model minimizing the worst-case loss of the defender in which both agents have continuous and infinite strategy spaces and the payoff of an attack varies over time.

The authors propose an algorithm computing the optimal dynamic allocation of the defender's resources (which can be relocated among targets without any delay in time), i.e. dynamic patrolling, called SCOUT-A (Scheduling seCurity resOurces in pUblic events with no relocating delAy) and an algorithm SCOUT-C to deal with general cases (Yin et al. 2014).

i) Urban Crime

Zhang et al. (2015) use a game-theoretic model for computing optimal patrol strategies to deter crimes in urban areas. As they state, these crimes are mostly opportunistic in nature, i.e. the criminals are less strategic in planning attacks (compared to, e.g., terrorists) but are more flexible in executing them. The presence of police patrols aims to deter such crimes. Just as in the examples (cases) cited so far, the criminals may observe potential patterns in patrol distribution and then strike at uncovered areas (since police have limited resources, which prevents permanently covering all areas). Human (manual) planning of patrol schedules shows not to be the most efficient approach since it is time-consuming, and humans may tend towards patterns. Automated planners may bring a remedy to this problem.

The authors deal with the problem of generating efficient patrol strategies against opportunistic criminals. The main novelty of their approach is that they learn the criminal behaviour from real data. Instead of modelling the adversary's (criminal's) decision-making explicitly, they learn their interaction with the defender (patrol officers) using real world data. (The models discussed so far assume a perfectly rational attacker or attackers with bounded rationality.) This interaction is modelled as a Dynamic Bayesian Network (DBN, see, e.g., https://www.bayesserver.com/docs/introduction/dynamic-bayesian-networks). Another contribution of Zhang et al. (2015) is a sequence of modifications of this DBN model allowing for its compact representation, which leads to better learning accuracy and increased speed of learning of the expectation maximization algorithm. (The maximization expectation algorithm serves to learn unknown parameters in the DBN from given learning data.) The authors further introduce two planning algorithms for computing the optimal patrol strategy. They also propose to frequently update the adversary model because criminals may adapt to deployed defender strategy.

5. Discussion and concluding remarks

Stackelberg security games are considered to be a natural approximation of real-world security scenarios, and the cases outlined in the text above show that it may provide a useful analytic tool for finding the optimal strategy in the area of public security. Typically, in these scenarios, a security force (e.g. the police) seeks to optimally allocate its limited security resources (officers) over multiple targets to protect against a rational adversary. As, according to the model, the security force (defender) commits to a strategy first and the adversary (attacker) responds to this strategy (so to maximize their expected payoff) after having thoroughly observed it, an efficient security system requires a randomization of the defender's action. Unpredictability of the defender's actions increases adversaries' uncertainty and thus improves the level of security. Algorithms based on the Stackelberg model calculate this optimal (randomized) mixed strategy to the defender.

The cases presented in this paper prove that real-world deployed software assistants built upon Stackelberg security games are helpful in situations where (a limited number of) police patrols need to be allocated in order to protect multiple strategic targets, such as airport terminals, flights or coasts. In these situations, these software assistants provide an optimal solution how to efficiently use these limited resources. The model is useful as well in cases of green security, urban transport, security of large public events etc. The cases collected and outlined in this paper clearly show that security games can be used by public policy- and decision-makers in projecting and realizing measures aimed at dealing with goals of security policies (especially with deterrence and prevention, protection and response – Donley and Pollard 2002).

On the other hand, an essential prerequisite of the efficient functioning of these mechanisms is that the initial model approximates the real-world setting as accurately as possible. For instance, the assumptions of the very basic model of a Stackelberg game are very strong and restrictive (e.g. the attacker is fully rational and has perfect knowledge of the defender's strategy). In every situation studied, its specific features and characteristics need to be taken into consideration by the model. Many variations of the basic Stackelberg security game model and solution concepts have been studied to handle different types of adversaries (Sinha et al. 2018). As Wang et al. (2019), who suggested a revised approach for the area of wildlife protection, recently summarized, researchers, inspired by successful deployments of Stackelberg Security Game in real life, were working hard to optimize the game models to make them more practical. Researchers have also tried to integrate the original models with old and new theories (e.g. Trejo et al. 2015; Kamra et al. 2018). Some suggested that understanding how individuals perceive risk is vital to understanding the behaviour of attackers (Ridinger et al. 2016). Not long ago, for instance, it has been pointed out that in the cyber-security domain the interaction between the defender and the adversary is quite complicated with a high degree of incomplete information and uncertainty (Kar et al. 2017).

On the other hand, with the increasing complexity of models the algorithms searching for optimal solution become more complex and sophisticated as well, and it is not easy to provide a fast algorithms that can scale to very large and complex problems (Pita et al. 2009). So the question is: How can one evaluate the functioning of the algorithms employed and the effectiveness of various scheduling assistants? Tambe (2011) points out that there are two primary difficulties related to attempts to evaluate deployed security systems. Firstly, the real adversaries will hardly cooperate in evaluation, and, secondly, there is often very little data available (e.g. about terrorist attacks).

However, there are several means of evaluation, theoretical or empirical ones. One possibility is testing such generated schedules against simpler randomization approaches (such as uniform randomization); this type of evaluation was used, for

example, in Pita et al. (2008), Tsai et al. (2009) and further. They can also be compared with previous scheduling practices as executed in Pita et al. (2008) or An et al. (2012). Where possible, real-world data should serve as a basis for evaluation. As example, Yin et al. (2012, see above), when designing a randomized system of patrols in order to deter passengers from fare evasion, performed simulations based on actual ridership data from Los Angeles Metro train lines. The results of these simulations suggested a possibility of significant fare-evasion deterrence in case their TRUSTS assistant had been deployed.

However, when considering an anti-terrorist security system, it is difficult, or rather impossible, to measure directly the deterrence effect of deployed mechanisms (which is a key goal of many security systems). Although we are capable of quantifying the number of attacks realized under the use of the mechanism, it is impossible to say how many attacks would have occurred without it. One possibility is to compare data on arrest records; however, as there are multiple factors influencing these numbers, such analysis would not tell a lot.

A way to obtain "human data" is through controlled laboratory experiments. These enable capturing behavioural features not included in theoretical models; however, experimental findings suffer from certain shortcomings, as well. One traditional criticism to experimental findings relates to the fact that the most frequent subjects of laboratory experiments are university students who may not correctly represent the real actors involved in tested situation. (E.g., the defender and the attacker in real-world security situations are hardly both students.)

There are various means of security-system evaluation, considering the costs and benefits. According to Tambe (2011), we distinguish between direct benefits which can be measured and indirect benefits. While direct benefits are, for example, reduced security costs or an increased number of attackers, the indirect benefits include deterred attacks or increased planning time and requirements for a successful attack. The costs, on the other hand, may be measured quite easily. These comprise, for instance, costs of system implementation and its maintenance.

In addition to the quantitative methods mentioned so far, an important source of information are qualitative expert evaluations. Domain experts may point to key benefits of deployed mechanisms and their effectiveness in boosting security. This can be a valuable input into discussions of researchers who have put the original model under discussion also.

One may also suggest that the presented models do not consider the transaction costs of switching from one strategy (to defend/attack airport terminal 1) to another one (defend/attack terminal 2) and that the transaction costs may be highly decisive for setting the appropriate security (public) policy. We have not found any literature that would deal with this. On the other hand, the model of a game does not assume for players to switch from one strategy to another in the course of a game. This may be the reason why switching costs are not considered. Once the

defender chooses their strategy, they stick with it, and the same applies for the attacker. The defender may, of course change their strategy for a following game (that is, to switch to another strategy between games), however, potential costs of such change are not considered by their utility function in a given game.

For an attacker we do not consider the possibility of switching the strategy, either. The model assumes that the adversary surveys the defender's action (over a period) and then launches an attack (in game-theoretic terms, they choose their pure strategy). So, they does not switch between their strategies.

This simplification does not deviate importantly from a real setting where a concrete allocation of resources (that is, for example, a decision to protect a concrete terminal) is set for a period of time (e.g., a day) and does not change in the course of it. Consider a concrete example in which we need to decide which terminal is going to be patrolled the next day (supposing, e.g., two terminals but only one disposable patrol). The "transaction costs" between sending the patrol to the first terminal, or the second one, may be negligible. But further research may focus on this definitely.

References

An, B., E. Shieh, R. Yang, M. Tambe, C. Baldwin, J. DiRenzo, B. Maule and G. Meyer. 2012. "PROTECT: A Deployed Game Theoretic System for Strategic Security Allocation for the United States Coast Guard." *AI MAGAZINE* 33(4), 96–110, doi: 10.1609/aimag.v33i4.2401.

Birkland, T. A. 2009. "Disasters, Catastrophes, and Policy Failure in the Homeland Security Era." *Review of Policy Research* 26(4), 423–438.

Chappell, A. T. and S. A. Gibson. 2009. "Community Policing and Homeland Security Policing Friend or Foe?" *Criminal Justice Policy Review* 20(3), 326–343.

Clarke, S. E. and E. Chenoweth. 2006. "The Politics of Vulnerability: Constructing Local Performance Regimes for Homeland Security." *Review of Policy Research* 23(1), 95–114.

Cross, M. K. D. 2007. "An EU Homeland Security? Sovereignty vs. Supranational Order." *European Security* 16(1), 79–97.

Danila, V. B. 2013. "Homeland Security: Designing Security in South-East European Countries." In Editor (ed.). *Applied Social Sciences: Economics and Politics*. Newcastle upon Tyne: Cambridge Scholars Publishing, 111–118.

Donley, M. B. and N. A. Pollard. 2002. "Homeland Security: The Difference between a Vision and a Wish." *Public Administration Review* 62 (Special Issue), 138–144.

Fang, F. S., Stone, P. and M. Tambe. 2015. "When Security Games Go Green: Designing Defender Strategies to Prevent Poaching and Illegal Fishing." In Yand, Q. and M. Wooldridge (eds.). *Twenty-Fourth International Joint Conference on Artificial Intelligence*. Available at https://www.aaai.org/ocs/index.php/IJCAI/IJCAI15/paper/viewFile/10763/11025 (last accessed 7 October 2020).

Friedmann, R. R. and W. J. Cannon. 2007. "Homeland Security and Community Policing: Competing or Complementing Public Safety Policies." *Journal of Homeland Security and Emergency Management* 4(4), doi: 10.2202/1547-7355.1371.

Ge, X., T. Zhou and Y. Zang. 2019. "Defense Strategy Selection Method for Stackelberg Security Game Based on Incomplete Information." In Tao, Q., Zhou, Y., and Jie Huang (eds.). *AICS 2019: Proceedings of the 2019 International Conference on Artificial Intelligence and Computer Science*. New York: Association for Computing Machinery, 555–561, https://doi.org/10.1145/3349341.3349467 (last accessed 30 April 2020).

Gerber, B. J., D. B. Cohen and K. B. Stewart. 2007. "U.S. Cities and Homeland Security: Examining the Role of Financial Conditions and Administrative Capacity in Municipal Preparedness Efforts." *Public Finance and Management* 7(2), 152–188.

Jain, M., E. Kardes, C. Kiekintveld, F. Ordóñez and M. Tambe. 2010. "Security Games with Arbitrary Schedules: A Branch and Price Approach." In Fox, M. and D. Poole (eds.). *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*. Palo Alto: Association for the Advancement of Artificial Intelligence, 792–797, doi: 10.1017/CBO9780511973031.009.

Kamra, N., U. Gupta, F. Fang, Y. Liu and M. Tambe. 2018. "Policy Learning for Continuous Space Security Games Using Neural Networks." In Zhang, S. (ed.). *Proceedings of the 2018 conference of the Association for the Advancement of Artificial Intelligence*. Available at https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16525 (last accessed 30 April 2020).

Kar, D., T. H. Nguyen, F. Fang, M. Brown, A. Sinha, M. Tambe and A. X. Jiang. 2017. "Trends and Applications in Stackelberg Security Games." In Başar, T. and G. Zaccour (ed.). *Handbook of Dynamic Game Theory*. Cham: Springer, 1–47.

May, P. J., A. E. Jochim and J. Sapotichne. 2011. "Constructing Homeland Security: An Anemic Policy Regime." *The Policy Studies Journal* 39(2), 285–307.

Morag, N. 2011. "Does Homeland Security Exist Outside the United States?" *Homeland Security Affairs* 9(11), 1–5.

Noftsinger, J. B., K. F. Newbold and J. K. Wheeler. 2007. *Understanding Homeland Security*. Basingstoke: Palgrave.

Okamoto, S., N. Hazon and K. Sycara. 2012. "Solving Non-Zero Sum Multiagent Network Flow Security Games with Attack Costs." In Albrecht, S. V. and R. Ramamoorthy (eds.). *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*. Valencia, Spain: International Foundation for Autonomous Agents and Multiagent Systems, 879–888). Available at http://www.aamas-conference.org/Proceedings/aamas2012/papers/3E_5.pdf (last accessed 7 October 2020).

Owen, G. 2013. *Game Theory*. Melbourne: Emerald Group Publishing Limited.

Pita, J., H. Bellamane, M. Jain, Ch. Kiekintveld, J. Tsai, F. Ordónez and M. Tambe. 2009. "Security Applications: Lessons of Real-World Deployment." *ACM SIGecom Exchanges* 8(2). Available at https://www.researchgate.net/publication/220553701_Security_applications_Lessons_of_real-world_deployment (last accessed 29 April 2020).

Pita, J., M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri and S. Kraus. 2008. "Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport." In Padgham, L., Parkes, D. and J. P. Muller (eds.). *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008): Industry and Applications Track*. Estoril, Portugal: International Foundation for Autonomous Agents and Multiagent Systems, 125–132, doi: 10.1145/1402795.1402819.

Pita, J., M. Tambe, C. Kiekintveld, S. Cullen and E. Steigerwald. 2011. "GUARDS – Game Theoretic Security Allocation on a National Scale." In Tumer, K., Yolum, P., Sonenberg, L. and P. Stone (eds.). *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems: Innovative Applications Track (AAMAS 2011)*. Taipei, Taiwan: International Foundation for Autonomous Agents and Multiagent Systems, 37–44, doi: 10.1017/CBO9780511973031.006.

Powell, R. 2007. "Defending against Terrorist Attacks with Limited Resources." *American Political Science Review* 101(3), 527–541.

Prisner, E. (2014). *Game Theory: Through Examples*. Washington, District of Columbia: Mathematical Association of America.

Reddick, Ch. G. 2008. "Collaboration and Homeland Security Preparedness: A Survey of U.S. City Managers." *Journal of Homeland Security and Emergency Management* 5(1), doi: 10.2202/1547-7355.1414.

Rhinard, M. and A. Boin. 2009. "European Homeland Security: Bureaucratic Politics and Policymaking in the EU." *Journal of Homeland Security and Emergency Management* 6(1), doi: 10.2202/1547-7355.1480.

Ridinger, G., R. S. John, M. McBride and N. Scurich. 2016. "Attacker Deterrence and Perceived Risk in a Stackelberg Security Game." *Risk Analysis* 36(8), 1666–1681.

Shor, M. 2015. "Dictionary of Game Theory Terms." Available at http://www.game-theory.net/dictionary/ (last accessed 20 July 2019).

Sinha, A., F. Fang, B. An, Ch. Kiekintveld and M. Tambe. 2018. "Stackelberg Security Games: Looking Beyond a Decade of Success." In Lang, J. and J. S. Rosenschein (eds.). *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, 5494–5501, https://doi.org/10.24963/ijcai.2018/775.

Špaček, D. 2020. *COVID-19 – National government approach – Czechia*. Study prepared for the IIAS Special Report: The COVID-19 Pandemic: Early Lessons for Public Governance (forthcoming).

Tadelis, S. 2013. *Game Theory: An Introduction*. Princeton: Princeton University Press.

Tambe, M. 2011. *Security and Game theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge: Cambridge University Press.

Trejo, K. K., J. B. Clempner and A. Poznyak. 2015. "A Stackelberg Security Game with Random Strategies Based on the Extraproximal Theoretic Approach." *Engineering Applications of Artificial Intelligence* 37, 145–153.

Tsai, J., S. Rathi, C. Kiekintveld, F. Ordóñez and M. Tambe. 2009. "IRIS: A Tool for Strategic Security Allocation in Transportation Networks." In Sierra, C. and C. Castelfranchi (eds.). *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*. Budapest, Hungary: International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org), doi: 10.1017/cbo9780511973031.005.

Von Neumann, J. 1928. "Zur Theorie der Gesellschaftsspiele." *Mathematische Annalen* 100(1), 295–320, https://doi.org/10.1007/BF01448847.

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

Von Stackelberg, H. 1934. *Marktform und Gleichgewicht*. Wien: J. Springer.

Wang, B., Y. Zhang, Z.-H. Zhou and S. Zhong. 2019. "On Repeated Stackelberg Security Game with the Cooperative Human Behavior Model for Wildlife Protection." *Applied Intelligence* 49, 1002–1015.

Webster, T. J. 2014. *Analyzing Strategic Behavior in Business and Economics: A Game Theory Primer*. Lanham: Lexington Books.

Wilczyński, A., A. Jakóbik and J. Kołodziej. 2016. "Stackelberg Security Games: Models, Applications and Computational Aspects." *Journal of Telecommunications and Information Technology* 3(3): 70–79. Available at https://www.itl.waw.pl/czasopisma/JTIT/2016/3/70.pdf (last accessed 8 July 2019).

Yang, R., F. Ordonez and M. Tambe. 2012. "Computing Optimal Strategy against Quantal Response in Security Games." In Albrecht, S. V. and R. Ramamoorthy (eds.). *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*. Valencia, Spain: International Foundation for Autonomous Agents and Multiagent Systems, 847–854. Available at http://teamcore.usc.edu/papers/2012/aamas2012_paper12_cameraReady.pdf (last accessed 8 July 2019).

Yin, Y., B. An and M. Jain. 2014. "Game-Theoretic Resource Allocation for Protecting Large Public Events." In Brodley, C. E. and P. Stone (eds.). *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. Palo Alto: Association for the Advancement of Artificial Intelligence. Available at https://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/download/8182/8507 (last accessed 8 July 2019).

Yin, Z., A. X. Jiang, M. Tambe, C. Kiekintveld, K. Leyton-Brown, T. Sandholm and J. P. Sullivan. 2012. "TRUSTS: Scheduling Randomized Patrols for Fare Inspection in Transit Systems Using Game Theory." *AI MAGAZINE* 33(4): 59–72, doi: 10.1609/aimag.v33i4.2432.

Zhang, C., A. Sinha and M. Tambe. 2015. "Keeping Pace with Criminals: Designing Patrol Allocation against Adaptive Opportunistic Criminals." In Bazzan, A. and M. Huhns (eds). *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*. Istanbul, Turkey: International Foundation for Autonomous Agents and Multiagent Systems, 1351–1359. Available at http://teamcore.usc.edu/papers/2015/keep_pace_with_criminal.pdf (last accessed 11 February 2020).