# Game Theoretic Analysis of Security and Sustainability

**Bo An**
School of Computer Science and Engineering
Nanyang Technological University, Singapore
boan@ntu.edu.sg

## Abstract

Computational game theory has become a powerful tool to address critical issues in security and sustainability. Casting the security resource allocation problem as a Stackelberg game, novel algorithms have been developed to provide randomized security resource allocations. These algorithms have led to deployed security-game based decision aids for many real-world security domains including infrastructure security and wildlife protection. We contribute to this community by addressing several major research challenges in complex security resource allocation, including dynamic payoffs, uncertainty, protection externality, games on networks, and strategic secrecy. We also analyze optimal security resource allocation in many potential application domains including cyber security. Furthermore, we apply game theory to reasoning optimal policy in deciding taxi pricing scheme and EV charging placement and pricing.

## 1 Introduction

Security is a critical concern around the world that arises in protecting our ports, airports, transportation or other critical national infrastructures from adversaries, in protecting our wildlife and forests from poachers and smugglers, and in curtailing the illegal flow of weapons, drugs and money; and it arises in problems ranging from physical to cyber-physical systems. In all of these problems, we have limited security resources which prevent full security coverage at all times; instead, limited security resources must be deployed intelligently taking into account differences in priorities of targets requiring security coverage, the responses of the adversaries to the security posture and potential uncertainty over the types, capabilities, knowledge and priorities of adversaries [An *et al.*, 2017].

Game theory is well-suited to adversarial reasoning for security resource allocation and scheduling problems. Casting the problem as a Stackelberg game, novel algorithms have been developed for efficiently solving such games to provide randomized patrolling or inspection strategies [Tambe, 2011]. These algorithms have led to some initial successes in this challenging problem arena, and are now deployed in multiple applications: ARMOR has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals [Pita *et al.*, 2008]; IRIS has been used for randomized deployment of the US Federal Air Marshals (FAMS) since 2009 [Tsai *et al.*, 2009]; PROTECT is deployed for generating randomized patrol schedules for the US Coast Guard in Boston, New York, Los Angeles and other ports around the US [An *et al.*, 2011; Shieh *et al.*, 2012]. Moreover, recent work on "green security games" has led to testing our decision aids for protection of fisheries with the US Coast Guard and protection of wildlife at sites in multiple countries, and "opportunistic crime security games" have focused on suppressing urban crimes [Fang *et al.*, 2016; Zhang *et al.*, 2014].

These initial successes point the way to major future applications in a wide range of security arenas, including network security, cyber security, adversarial machine learning and so on; with major research challenges in various aspects of complex security resource allocation, such as dynamic payoffs, observation uncertainty, protection externality, as well as the strategic secrecy. The goal of this paper is to briefly describe our recent progress in applying game theoretic models to address these challenges.

## 2 Security Game Model

We first provide an overview of the basic security game framework, which includes two players, a defender who first decides how to use $m$ identical resources to protect a set of targets $T$ ($m < |T|$), and an attacker who observes the defender's strategy before choosing a target to attack. A defender's pure strategy is a subset of targets from $T$ such that at most $m$ targets from $T$ are protected. An attacker's pure strategy is a target from $T$ which will be attacked. A mixed strategy allows a player to play a probability distribution over pure strategies. Formally, the defender's mixed strategy can be compactly represented as a coverage vector $\mathbf{c} = \langle c_t \rangle$ where $c_t$ is the probability that target $t$ is covered. The attacker's mixed strategy $\mathbf{a} = \langle a_t \rangle$ is a vector where $a_t$ is the probability of attacking target $t$.

The defender's payoff for an attack on an uncovered target $t$ is denoted as $P_t^d$, and $R_t^d$ for an attack on a covered target. Similarly, $R_t^a$ and $P_t^a$ are the attacker's payoffs respectively. A commonly used assumption is that $R_t^d > P_t^d$ and $R_t^a > P_t^a$ in order to model the fact that the defender would always prefer the attack to fail, while the attacker would prefer it to

succeed. For a strategy profile $\langle \mathbf{c}, \mathbf{a} \rangle$, the expected utilities for both players are given by:

$$U_d(\mathbf{c}, \mathbf{a}) = \sum_{t \in T} a_t U_d(\mathbf{c}, t) \text{ where } U_d(\mathbf{c}, t) = c_t R_t^d + (1 - c_t) P_t^d$$

$$U_a(\mathbf{c}, \mathbf{a}) = \sum_{t \in T} a_t U_a(\mathbf{c}, t) \text{ where } U_a(\mathbf{c}, t) = c_t P_t^a + (1 - c_t) R_t^a$$

In a Stackelberg game model, the defender chooses its strategy first, and the attacker chooses a strategy after observing the defender's choice. The attacker's response function is $g(\mathbf{c}) : \mathbf{c} \to \mathbf{a}$. We assume that $g(\mathbf{c})$ is unique to every $\mathbf{c}$. The standard solution concept is Strong Stackelberg Equilibrium (SSE) [Leitmann, 1978]. A pair of strategies $\langle \mathbf{c}, g(\mathbf{c}) \rangle$ form an SSE if they satisfy the following: 1) The defender plays a best-response: $U_d(\mathbf{c}, g(\mathbf{c})) \geq U_d(\mathbf{c}', g(\mathbf{c}'))$ for any $\mathbf{c}'$. 2) The attacker plays a best-response: $g(\mathbf{c}) \in F_a(\mathbf{c})$ where $F_a(\mathbf{c}) = \arg\max_{\mathbf{a}} U_a(\mathbf{c}, \mathbf{a})$ is the set of follower best responses. 3) The attacker breaks ties optimally for the defender: $U_d(\mathbf{c}, g(\mathbf{c})) \geq U_d(\mathbf{c}, \mathbf{a}')$ for any $\mathbf{a}' \in F_a(\mathbf{c})$.

The above basic security game model is extended for modelling more complex security scenarios, e.g., a pure defender strategy could be a path/flow on a graph and a pure attacker strategy could be a connected subgraph.

## 3 Analyzing Complex Adversarial Interaction

One major focus of our research is to study the practical and challenging issues arising in realistic complex security resource allocation problems.

**Dynamic Payoffs** The impact of an attack on attractive targets, reflected by payoffs, might change over time. Motivated by the protection of public events (e.g., Boston Marathon), Yin *et al.* [2014] propose a game model for such domains in which the defender can dynamically reallocate security resources. Accordingly, each agent has a continuous and infinite strategy space and the payoff of an attack is time-dependent. Novel algorithms are proposed to efficiently compute the optimal deterministic patrolling strategy. Recent work also considers the adversary's surveillance capability due to more frequent interactions, in which case an optimal mixed defender strategy is preferred [Yin *et al.*, 2015]. When the defender can move security resources among targets without time delay, our algorithm represents mixed strategies in a compact way and avoids traversing the whole strategy space by exploiting the properties of the optimal defender strategies. Then from the compact mixed strategy, we sample 'implementable' pure strategies as the defender can only make a finite number of transfers of resources in a time period in the real world, despite that she can transfer resources at any time point in this period. We also consider a more general case in which it takes nonnegligible time to move resources among targets. We develop an efficient approximation scheme to compute the near-optimal defender mixed strategy.

**Uncertainty** The classic security game model presented in Section 2 assumes that both agents are fully rational and the attacker can exactly learn the defender's mixed strategy, which means that the observation of the attacker is unlimited. These assumptions are not realistic and various uncertainties exist in real-world security domains, such as uncertainty of the attacker's observation, bounded rationality of the attacker, and the defender's limited prior knowledge about the attacker and the environment. An *et al.* [2013] consider the cost of surveillance and model the attacker's decision making as a Markov decision process (MDP) where each state denotes the observation history and the attacker can either attack the best target according to the belief on the defender's strategy, or make another observation. Novel algorithms including backward induction with forward search and an approximation approach are proposed to solve the MDP. Guo *et al.* [2017b] study the repeated network interdiction where the defender has no prior knowledge about the attacker's behavior and the attacker can be irrational to any extent. Moreover, the defender only has partial information about the attacker's action in past interactions. The online learning approach is applied to minimize the defender's regret of not sticking to the best fixed resource allocation and an efficient algorithm with sublinear regret is proposed.

**Protection Externality** It is widely assumed in security games that one security resource assigned to a target only protects that target. However, in many important real-world security scenarios, a resource exhibits protection externality such that it also protects the neighbouring targets. Gan *et al.* [2015b] investigate such Security Games with Protection Externalities (SPEs) and prove the NP-hardness of computing SSE for SPE. On the positive side, they propose a novel column generation based approach CLASPE that is able to scale up to realistic-sized SPE problem instances. A more general model called Security Game on a Plane (SGP) is proposed for practical security settings where defense resources can be located on a continuous plane [Gan *et al.*, 2017]. They find that computing an SSE of an SGP is NP-hard even for zero-sum games, and these are inapproximable in general. On the positive side, they find an exact solution technique for general SGPs based on an existing approach, and develop a PTAS for zero-sum SGP to more fundamentally overcome the computational obstacle.

**Network Security Games** Many security domains involve network structures such as transportation networks and social networks. Unlike the target-based security games where the attacker's strategy space is the set of targets, the attacker's pure strategies on a network are much more complex, such as paths, flows, and subgraphs. The network structure causes an exponential-sized or even infinite attacker's strategy space, and the scalability becomes a critical concern. Guo *et al.* [2016a] study the optimal strategy to prevent a set of cooperative terrorist groups through cutting off their connections. The cooperation is modeled by a coalitional skill game on a network and each subgraph denotes one possible coalition. The defender's optimization is formulated as a mixed-integer linear program with an exponential number of variables, each corresponding to one coalition. In order to scale up to realistic-sized problems, a novel branch and price algorithm is proposed. Wang *et al.* [2016] discuss the optimal strategy for the police to monitor potential terrorists to prevent planned terrorist plots, such as Paris shootings on January 7, 2015. The terrorist planner can strategically choose to arouse several members within a terrorist network and plan an attack. A novel double oracle approach is provided to address the exponentially large strategy spaces of both the defender and the attacker by iteratively generating candidate pure strategies of both agents. Guo *et al.* [2016b] study the optimal network flow interdiction problem, which captures

many smuggling scenarios on transportation networks. The defender plays a randomized checkpoint allocation and the smuggler chooses a network flow to maximize his own profit. A column and constraint generation method is proposed by iteratively generating candidate checkpoint allocation for the defender and source-target path for the smuggler.

**Strategic Secrecy** There is a longstanding dilemma in security games: given the theoretical advantages of commitment, why is it that real-world security forces often use secrecy (e.g., plainclothes)? We address this dilemma by introducing the private information (number of resources) of the defender and adopting perfect Bayesian equilibrium as the solution concept [Guo *et al.*, 2017a]. We model these aspects as Disguised-resource Security Game (DSG) and analyze the strategic secrecy where the defender strategically deceives the attacker by disguising some of her resources. The number of the revealed resources is modeled as a signal which can only be sent by the defender with enough resources. We compare the expected utility of strategic secrecy with the value of public commitment and formally show that they have different advantages depending on the payoff structure. We provide a support set enumeration approach for computing PBE solutions and show that there is a fundamental tradeoff between secrecy and commitment by theoretical analysis and experimental evaluation. We conclude that the boundary of such trade-offs is close to zero-sum games which confirms the practical use of plainclothes police due to the approximate zero-sum nature of many security scenarios

## 4 Addressing Issues in Potential Applications

Another important focus of our research lies in addressing research challenges in new application domains. Our approaches can be adapted or extended for those potential application scenarios.

**Cyber Security** Modern cyber crimes are becoming more and more complex and players in many cyber security domains are strategic. Game theory is a powerful tool for modeling cyber attack scenarios, analysing attacker behaviors and providing defense actions. Zhao *et al.* [2016] consider spear phishing attackers who make sequential attack plans based on the outcome of previous attacks. We formulate a bilevel optimization problem for the defender and show that the attacker's problem (i.e., lower level problem) can be solved by a linear program. Solving the linear program is computationally consuming since the number of variables and constraints grows exponentially with the number of users. We find a simplified representation of the defender's utility and thus reduce the defender's bilevel program into a single level binary combinatorial optimization program by exploiting the structure of the attacker's MDP. We also extend the single-credential case to a more general case where there could be multiple sensitive credentials. Li *et al.* [2017a; 2017b] extend the Stackelberg security game model to analyze the Man-In-The-Middle (MITM) attack.

**Transportation Networks** Preventing crimes or terrorist attacks in urban areas has drawn extensive attentions in recent years. Law enforcement officers need to respond quickly to catch the attacker on his escape route, which is subject to time-dependent traffic conditions on transportation networks. The attacker can strategically choose his escape path and driving speed to avoid being captured. Existing work on security resource allocation has not considered such scenarios with time-dependent strategies for both players. Zhang *et al.* [2017] study the problem of efficiently scheduling security resources for interdicting the escaping attacker. The problem is shown to be NP-hard. An efficient double oracle algorithm to compute the optimal defender strategy is proposed, which combines mixed-integer linear programming formulations for best response problems and effective approximation algorithms for improving the scalability of the algorithms.

**Adversarial Machine Learning** Adversarial machine learning can be viewed as a game between the learner and the attacker. The attacker's strategy includes poisoning the training data and manipulating the testing data. The learner's strategy includes selecting appropriate learning algorithms and parameters, and additional defense actions such as giving penalty on detected attack behavior. Zhao *et al.* [2017] study the label contamination attacks, where the attacker can manipulate part of the labels of the training data to make the learned model beneficial to him. Existing work assumes that the attacker has full knowledge of the victim learning model, whereas the victim model is usually a black-box to the attacker. They develop a projected gradient ascent algorithm to compute label contamination attacks on a family of empirical risk minimizations and show that an attack on one victim model can also be effective on other models.

**Elections** Democratic institutions rely on the integrity of the voting process. A major threat to this integrity is the possibility that the process can be subverted by malicious parties to their own goals [Bhattacharjya, 2010]. This problem is both a fundamental theoretical problem in social choice, and a major practical concern for democratic institutions. Yin *et al.* [2016] model the control of election as a denial-of-service (deletion) attack on a subset of voting groups. They show that for plurality voting, election control through group deletion to prevent a candidate from winning is in P, while it is NP-hard to prevent such control. They then present a double-oracle framework for computing an optimal prevention strategy, developing exact mixed-integer linear programming formulations for both the defender and attacker oracles (both of these subproblems are shown to be NP-hard), as well as heuristic oracles.

**Nuclear Smuggling** Maritime container shipping has been a critical measure for terrorists and smugglers to transport illegal goods including weapons of mass destruction and even nuclear materials. Given the huge number of containers transported, it remains a significant challenge to efficiently deploy the inspection facilities and security officers [Leonard *et al.*, 2015]. While existing work neglects the sophisticated behavior of the smuggler, Wang *et al.* [2017] propose a novel container inspection model which formulates the smuggler's sequential decision behavior as an MDP. The special structure of the problem results in a non-convex optimization problem, which cannot be addressed by existing approaches. They first use a linear relaxation approximation with guarantee of solution quality which reformulates the model as a bilinear optimization problem. Algorithms inspired by the multipleparametric disaggregation technique are provided to solve the reformulated bilinear optimization problem and compute efficient and robust solutions.

**Coral Reef Ecosystems** Coral reefs are precious natural resources, which form some of the world's most productive ecosystems. However, some human activities, like coral mining, can severely damage the coral reef ecosystems. Therefore, many countries have built Marine Protected Areas (MPAs) to restrict potentially damaging activities by patrolling in the MPAs. It is a great challenge to efficiently protect the MPAs through patrolling since protection agencies usually have to protect a large open water area using very limited resources. Yin and An [2016] consider the time-dependent strategies of both agents and the time duration of attacks, and propose a Stackelberg game model to formulate the problem of protecting MPAs. A compact-strategy double-oracle algorithm on graphs is proposed.

## 5 Computational Sustainability

We also put effort into computational sustainability research, which is a vital field aiming at applying computational techniques to address environmental, economic, and social problems arising from the needs of sustainable development. Human beings play a central role in such domains through strategically deciding how to use public resources such as roads. Our main focus is to provide optimal policies for the government in consideration of people's strategic behavior.

Taxi service is an indispensable part of public transport in modern cities. However, its efficiency is greatly hindered by the decentralized operation mode. There are more than 60,000 licensed taxis in Beijing serving nearly 20 million citizens. However, despite rising customer demand during peak hours, most taxi drivers act counter-intuitively, intentionally avoiding working during those periods. It turns out that the improper distance-based pricing scheme is the main cause of this situation, also called the peak-time dilemma. Low travel speed during peak time as a result of heavy traffic causes low or even negative revenue generation for taxi drivers, leaving them to pursue the only option that makes them money: not working during peak time. Gan *et al.* [2013; 2015a] formulate how taxi drivers respond to government's pricing scheme in consideration of many inter-dependent factors. The objective is to compute the optimal dynamic time-dependent fare structure to incentivize taxi drivers to work so that more people can be served. With convex polytope representation techniques and strategy expanding methods, the proposed algorithms can efficiently deal with realistic scenarios and handle arbitrary scheduling constraints.

Electric vehicle (EV) is environmentally-friendly and promising to release the concern of fossil fuel shortage, and thus has been popular in recent years. A prerequisite of top-priority for EV diffusion is to deal with the limited battery storage with deployment of supporting facilities, i.e., charging stations. Specialized EV charging station, which provides more than 10 times faster charging speed than domestic charging, is therefore a critical element for successful EV promotion. Considering the EV drivers' charging behavior and its influence on the performance of charging stations, Xiong *et al.* [2015] divide the EV drivers into different types of congestion game players, and then model the charging station placement problem to compute the optimal assignment of charging stations. They first formulate the problem as a bilevel optimization problem, which is subsequently converted to a single-level optimization problem by exploiting structures of the EV charging game. Properties of the problem are analyzed and exploited to compute the optimal allocation of charging stations. Xiong *et al.* [2016] also approach the management of EV charging stations from the pricing perspective as a more flexible and adaptive complement to established charging station placement. They formulate the pricing problem as a mixed integer non-convex optimization problem, and propose a scalable algorithm to solve it.

## 6 Conclusions

Game theoretic analysis has become a powerful measure to provide the optimal decisions in security and sustainability. This paper highlights our recent progress in analyzing complex adversarial interaction, addressing challenges in potential security application domains, and providing optimal policies for the government in consideration of people's strategic behavior in the sustainability domain.

While the deployed applications have advanced the state of the art, significant future research remains to be done. We need to build more realistic model (including the environment, human behavior, interactions between different players, and uncertainty) to make our solutions effective in practice. With more realistic models, we need to develop new algorithms that are able to efficiently compute optimal solutions of real-world security scenarios which often come with an exponential (even infinite) number of pure strategies for all the players. We also need to further improve solutions' robustness due to uncertainty about the knowledge, rationality, and capability of players. For domains with plenty of data, relying too much on data for making decisions might be dangerous as learning could be exploited by adversaries. Thus, how to exploit data by learning and how to combine learning based approach with model based approach need to be carefully investigated. Furthermore, how to integrate real time information into optimal allocation of security resources over a long period remains to be carefully studied. Finally, new application domains often introduce new challenges in modelling and algorithm design.

## Acknowledgments

The author gratefully acknowledges support from NTU, NRF, MOE, and MSRA.

## References

[An *et al.*, 2011] Bo An, James Pita, Eric Anyung Shieh, Milind Tambe, Christopher Kiekintveld, and Janusz Marecki. GUARDS and PROTECT: Next generation applications of security games. *SIGECOM*, 10(1):31–34, 2011.

[An *et al.*, 2013] Bo An, Matthew Brown, Yevgeniy Vorobeychik, and Milind Tambe. Security games with surveillance cost and optimal timing of attack execution. In *AAMAS*, pages 223–230, 2013.

[An *et al.*, 2017] Bo An, Milind Tambe, and Arunesh Sinha. *Improving Homeland Security Decisions*, chapter Stackelberg security games (SSG) basics and application overview. Cambridge University Press, 2017.

[Bhattacharjya, 2010] Satarupa Bhattacharjya. Low turnout and invalid votes mark first post war general polls. http://www.sundaytimes.lk/100411/News/nws16.html, 2010.

[Fang *et al.*, 2016] Fei Fang, Thanh H. Nguyen, Rob Pickles, Wai Y. Lam, Gopalasamy R. Clements, Bo An, Amandeep Singh, Milind Tambe, and Andrew Lemieux. Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *IAAI*, pages 3966–3973, 2016.

[Gan *et al.*, 2013] Jiarui Gan, Bo An, Haizhong Wang, Xiaoming Sun, and Zhongzhi Shi. Optimal pricing for improving efficiency of taxi systems. In *IJCAI*, pages 2811–2818, 2013.

[Gan *et al.*, 2015a] Jiarui Gan, Bo An, and Chunyan Miao. Optimizing efficiency of taxi systems: Scaling-up and handling arbitrary constraints. In *AAMAS*, pages 523–531, 2015.

[Gan *et al.*, 2015b] Jiarui Gan, Bo An, and Yevgeniy Vorobeychik. Security games with protection externalities. In *AAAI*, pages 914–920, 2015.

[Gan *et al.*, 2017] Jiarui Gan, Bo An, Yevgeniy Vorobeychik, and Brian Gauch. Security games on a plane. In *AAAI*, pages 530–536, 2017.

[Guo *et al.*, 2016a] Qingyu Guo, Bo An, Yevgeniy Vorobeychik, Long Tran-Thanh, Jiarui Gan, and Chunyan Miao. Coalitional security games. In *AAMAS*, pages 159–167, 2016.

[Guo *et al.*, 2016b] Qingyu Guo, Bo An, Yair Zick, and Chunyan Miao. Optimal interdiction of illegal network flow. In *IJCAI*, pages 2507–2513, 2016.

[Guo *et al.*, 2017a] Qingyu Guo, Bo An, Branislav Bosansky, and Christopher Kiekintveld. Comparing strategic secrecy and Stackelberg commitment in security games. In *IJCAI*, 2017.

[Guo *et al.*, 2017b] Qingyu Guo, Bo An, and Long Tran-Than. Playing repeated network interdiction games with semi-bandit feedback. In *IJCAI*, 2017.

[Leitmann, 1978] George Leitmann. On generalized Stackelberg strategies. *Optimization Theory and Applications*, 26(4):637–643, 1978.

[Leonard *et al.*, 2015] Timothy J Leonard, Philip Gallo, and Simon Véronneau. Security challenges in United States sea ports: An overview. *Transportation Security*, 8(1-2):41–49, 2015.

[Li *et al.*, 2017a] Shuxin Li, Xiaohong Li, Jianye Hao, Bo An, Zhiyong Feng, Kangjie chen, and Chengwei Zhang. Defending against man-in-the-middle attack in repeated games. In *IJCAI*, 2017.

[Li *et al.*, 2017b] Xiaohong Li, Shuxin Li, Jianye Hao, Zhiyong Feng, and Bo An. Optimal personalized defense strategy against man-in-the-middle attack main information. In *AAAI*, pages 593–599, 2017.

[Pita *et al.*, 2008] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, and Sarit Kraus. Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. In *AAMAS*, pages 125–132, 2008.

[Shieh *et al.*, 2012] Eric Shieh, Bo An, Rong Yang, Milind Tambe, Craig Baldwin, Joseph DiRenzo, Ben Maule, and Garrett Meyer. PROTECT: An application of computational game theory for the security of the ports of the United States. In *AAAI*, pages 2173–2179, 2012.

[Tambe, 2011] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, Cambridge, UK, 2011.

[Tsai *et al.*, 2009] Jason Tsai, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe, and Shyamsunder Rathi. IRIS-a tool for strategic security allocation in transportation networks. In *AAMAS*, pages 37–44, 2009.

[Wang *et al.*, 2016] Zhen Wang, Yue Yin, and Bo An. Computing optimal monitoring strategy for detecting terrorist plots. In *AAAI*, pages 637–643, 2016.

[Wang *et al.*, 2017] Xinrun Wang, Qingyu Guo, and Bo An. Stop nuclear smuggling through efficient container inspection. In *AAMAS*, pages 669–677, 2017.

[Xiong *et al.*, 2015] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Ana LC Bazzan. Optimal electric vehicle charging station placement. In *IJCAI*, pages 2662–2668, 2015.

[Xiong *et al.*, 2016] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Yeng Chai Soh. Optimal pricing for efficient electric vehicle charging station management. In *AAMAS*, pages 749–757, 2016.

[Yin and An, 2016] Yue Yin and Bo An. Efficient resource allocation for protecting coral reef ecosystems. In *IJCAI*, pages 531–537, 2016.

[Yin *et al.*, 2014] Yue Yin, Bo An, and Manish Jain. Game-theoretic resource allocation for protecting large public events. In *AAAI*, pages 826–834, 2014.

[Yin *et al.*, 2015] Yue Yin, Haifeng Xu, Jiarui Gan, Bo An, and Albert Xin Jiang. Computing optimal mixed strategies for security games with dynamic payoffs. In *IJCAI*, pages 681–688, 2015.

[Yin *et al.*, 2016] Yue Yin, Yevgeniy Vorobeychik, Bo An, and Noam Hazon. Optimally protecting elections. In *IJCAI*, pages 538–545, 2016.

[Zhang *et al.*, 2014] Chao Zhang, Albert Xin Jiang, Martin B Short, P Jeffrey Brantingham, and Milind Tambe. Defending against opportunistic criminals: New game-theoretic frameworks and algorithms. In *GameSec*, pages 3–22, 2014.

[Zhang *et al.*, 2017] Youzhi Zhang, Bo An, Long Tran-Thanh, Zhen Wang, Jiarui Gan, and Nicholas R. Jennings. Optimal escape interdiction on transportation networks. In *IJCAI*, 2017.

[Zhao *et al.*, 2016] Mengchen Zhao, Bo An, and Christopher Kiekintveld. Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks. In *AAAI*, pages 658–665, 2016.

[Zhao *et al.*, 2017] Mengchen Zhao, Bo An, Wei Gao, and Teng Zhang. Efficient label contamination attacks against black-box learning models. In *IJCAI*, 2017.