ROLE OF ECONOMIC POLICIES IN THE SECURITY OF

CRITICAL INFRASTRUCTURES

by

Carlos Alfredo Barreto Suarez

APPROVED BY SUPERVISORY COMMITTEE:

Alvaro A. Cardenas, Chair

Daniel G. Arce M.

R. Chandrasekaran

Murat Kantarcioglu

To my family and friends
for their unconditional support.

ROLE OF ECONOMIC POLICIES IN THE SECURITY OF CRITICAL INFRASTRUCTURES

by

CARLOS ALFREDO BARRETO SUAREZ, BS, MS

DISSERTATION

Presented to the Faculty of

The University of Texas at Dallas

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY IN

COMPUTER SCIENCE

THE UNIVERSITY OF TEXAS AT DALLAS

May 2018

ACKNOWLEDGMENTS

I'm deeply grateful to my advisor, Alvaro A. Cárdenas, for his support, advice, and patience during the last five years. I'm in debt to him for all his help, his active engagement in my research, and the freedom he gave me to explore many research problems (as this document attests).

I'm also grateful to my mentors and friends, Nicanor Quijano and Eduardo Mojica-Nava, for their constant support and camaraderie. They instilled in me the importance of mathematical formalism and introduced me to game theory; our early work on mechanism design became the foundation of Chapter 4.

I have benefited greatly from my collaboration with Professor Alain Bensoussan. It has always been a pleasure to talk to him, both for his keen observations and his sense of humor. Taking his class on stochastic dynamic programming was crucial for developing Chapters 5 and 6. I also collaborated with Galina Schwartz to investigate how to manage cyber-risks in cyber physical systems. Chapters 2 and 6 contain some of the ideas that we discussed and the main results from this collaboration will be published soon. I'd like to thank the members of my dissertation committee for their valuable suggestions and comments.

This work was possible thanks to the resources offered at The University of Texas at Dallas. I'm really thankful to the Office of Graduate Studies because their workshops helped me to improve my writing skills. I'm also grateful to the staff members of the Computer Science department, in particular, to Carina Legorreta and Douglas Hyde. Mr. Hyde helped me with a last-minute mishap the day of my defense. It is easy to become careless when you are surrounded by competent and kind people.

I also benefited from stimulating discussions with many colleagues from UT Dallas, in particular with Jairo Giraldo, Francisco Combita, Mustafa Faisal, Junia Valente, and Esmail Babakrpour. My time in Dallas has been particularly hilarious thanks to my friend Jorge Medellin.

Last but not least, I thank my parents and sisters for their unconditional support.

April 2018

ROLE OF ECONOMIC POLICIES IN THE SECURITY OF

CRITICAL INFRASTRUCTURES

Carlos Alfredo Barreto Suarez, PhD
The University of Texas at Dallas, 2018

Supervising Professor: Alvaro A. Cardenas, Chair

In the last few years we have witnessed the development of sophisticated attacks that target critical infrastructures. Such attacks can cause catastrophic damage; for instance, attacks on the electricity system can impact a variety of industrial, commercial, and residential customers. Protecting critical infrastructures remains a challenge, because the cyber threats evolve in time and these systems have both correlated risks and information asymmetries. Moreover, many security problems arise due to improper economic incentives, rather than technical difficulties. In this research we investigate how economic policies affect the security of critical infrastructures.

First, we illustrate the importance of economic incentives showing how policies designed to protect systems have the opposite effect. In particular, we analyze how a company exploited flaws in contractual policies (asymmetric information) to profit by sponsoring attacks. We also show how to redesign the policies to prevent these situations.

Second, we analyze attacks that leverage the market's infrastructure to manipulate the demand of users. We find that an attacker with enough influence can either increase his profit (protecting his anonymity) or cause blackouts. The attacker can succeed in markets with both centralized and distributed structures; however, attacks on distributed systems produce less profit, but also make it more difficult to detect and penalize attacks.

Third, we investigate the optimal allocation of resources to protect systems against cyber threats that evolve in time. We model the evolution of threats with a Markov process and contemplate three protection schemes: prevention (e.g., secure code development), detection (intrusion detection systems), and risk transfer (e.g., cyber insurance). We find that uncertainties in the system's state make insurance more attractive as a risk management tool, but still,

the defenders need incentives to purchase cyber insurance. Moreover, insurance can improve the investment in either prevention or detection; however, policies with indemnity subsidies and unlimited coverage can introduce perverse incentives that degrade the investments in security.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

During the past decade our power grid, vehicles, medical devices, buildings, and many other systems we interact with have been modernized with embedded computing systems. Such systems that connect the physical world to cyber space are usually referred to as Cyber Physical Systems (CPSs). While these systems provide new societal benefits, they may also allow cyber attackers to affect our physical world causing kinetic effects. For example, cyber attacks against the power grid can cause blackouts, attacks against modern vehicles can cause accidents, and attacks against medical devices can harm their users (Koppel, 2016; Greenberg, 2015; Newman, 2015; Leverett et al., 2017).

Most industries face different challenges investing in security protections, because they cope with different threats. In particular, there is a difference between industries that use conventional Information Technology (IT) systems and industries that work with CPSs. On one hand, companies that manage traditional IT assets (e.g., have a web-presence, or handle any financial transaction) have experience dealing with increasingly sophisticated and well-organized criminal groups that try to compromise their systems for financial gains. These companies cannot underestimate their cyber risks; therefore, they constantly upgrade and test their systems to minimize losses.

Although attacks targeting physical processes exist,[1] they are still rare and not openly reported. Therefore, *most* industries in the CPS domain have never seen attacks sabotaging their physical processes and they do not see a clear business case for investing in information security protections (Langner and Pederson, 2013).

Firms often use risk management practices to reach an acceptable level of risk balancing the cost mitigation and the benefits. This attitude makes sense from a business perspective, but can be insufficient in the protection of CPS (Langner and Pederson, 2013). For example, the risk assessments of critical infrastructures can ignore intangible damage or negative externalities suffered by the society. Therefore, corporative decisions can fail to (completely) handle risks that affect the society.

As the U.S. Department of Energy (DoE) stated in their Energy Delivery Systems Cyber Security Roadmap (Batz et al., 2011) "Making a strong business case for cyber security

[1]Examples of cybernetic attacks with physical consequences are Stuxnet (Zetter, 2014), the attacks against the power grid in Ukraine (Zetter, 2016), and the Triton malware attacking safety systems in the middle east (Finkle, 2017).

investments is complicated by the difficulty of quantifying risk in an environment of rapidly changing, unpredictable threats with consequences that are hard to demonstrate." This has left our physical critical infrastructures fairly vulnerable to computer attacks and with technology that is decades behind the current security best practices used in other IT domains.

Unlike other national security issues, the government cannot deal with some cyber threats through diplomatic or military actions, due to difficulties exercising retaliation against cyber adversaries.[2] Moreover, market incentives alone have not created momentum on industries in CPS to improve their security posture. The failure to properly secure CPS (especially critical infrastructures) has resulted in several calls for government regulation of cyber security (Schneier, 2017; Fu, 2016; Vardi, 2017; Stark, 2017; Pagliery, 2016). However, the industry in general has pushed back against regulation,[3] arguing that they can stifle innovation.[4]

Alternatively, instead of mandating regulations, the government can create economic policies that incentivize investments in security. A common economic incentive used in public policies consists in making companies liable for the negative effects of their activities (e.g., environmental damage), so that they adopt precautions to avoid fines. However, this scheme can fail if the authorities cannot determine the culpability of each company. In particular, interdependencies and correlations among companies can difficult the assessment of liabilities of cyber events (Böhme and Schwartz, 2010; Laszka et al., 2014). For example, companies can suffer attacks crafted using vulnerabilities of a third party software, rather than vulnerabilities in their own systems (Cherepanov, 2017; Krebs, 2014). Such interdependencies can worsen security problems, because many individuals who do not suffer the consequences of attacks remain unprotected, affecting other parties. This occurs to the owners of *internet of things* (IoT) device, who aren't concerned for their security because other parties bear the cost of *denial of service* (DoS) attacks (Anderson, 2001). Similarly, software companies usually adopt contractual policies that shield them from liabilities.

[2]Retaliation is difficult and dangerous, because 1) adversaries usually protect their identities covering their traces and/or impersonating third parties; and 2) cyber attacks against a particular target can affect other parties, since systems are highly interconnected. A controversy among the U.S. and Germany offers a clear example of the risks of retaliation (Nakashima, 2017). The dispute started because the U.S. Cyber Command carried out an operation of sabotage against the Islamic State, which involved hacking servers located in Germany.

[3]E.g., mandating the compliance to specific security standards of security

[4]Designing and enforcing regulations, having into account the variety of industries, can be expensive and can have unexpected effects, such as discourage research (Harrington and Morgenstern, 2007).

The security problems described before relate more with bad incentives, rather than traditional information security (Anderson, 2001). In this research we investigate how to improve the security of critical infrastructures, specifically, by analyzing the impact of economic policies in the security of these systems. We find that economic policies can affect the security of systems in many ways. First, we learnt that parties involved in the protection of systems can leverage the anonymity of attacks to profit by sponsoring attacks. Nevertheless, an appropriate design of the protection policies can prevent such situations. Second, we find that the structure of power markets can affect both attackers and defenders. In particular, although we can increase the difficulty to implement attacks, it becomes more difficult to detect and penalize them. Third, we find that uncertainties about the security of the system change the investments in protection. For example, uncertainties make insurance more attractive; however, the defender accepts insurance with cost lower than the fair premium. Also, insurance with subsidies and full coverage can create perverse incentives that reduce investments in security.

The document is organized as follows: we give some background on cyber risks in CPS in Chapter 2. Chapter 3 illustrates the importance of economic incentives showing how policies designed to protect systems can introduce perverse incentives, which have the opposite effect. In Chapter 4 we analyze attacks on the power grid that use the market's infrastructure to manipulate the demand of users. In Chapter 5 we investigate the optimal allocation of resources to protect systems against cyber threats that evolve in time. We model the evolution of threats with a Markov process that describes the dynamic interaction between an attacker and a defender. In Chapter 6 we extend the model introduced in Chapter 5 to contemplate cyber insurance as an investment to manage risks. We conclude the document in Chapter 7.

CHAPTER 2

BACKGROUND

2.1 Cyber Risks: Heritage from Modernity

In general, a cyber risk includes anything that causes harm exploiting vulnerabilities of information systems. Cyber incidents usually involve damage to information; however, in recent years some sophisticated attacks have surpassed cybernetic barriers affecting the physical world. In this section we introduce cyber risks on information systems and critical infrastructures and their relations with other risks.

2.1.1 Cyber Risks

Cyber threats emerged as byproducts of technological developments, since pressure to commercialize products quickly (and some public policies[1]) led to insecure systems that expose the confidentiality, integrity, or availability of information[2] (Anderson, 2001). Most of the attackers seek financial benefits through criminal activities, such as stealing information (Karpesky Lab, 2014), or extorting individuals either by encrypting their data (ransomware) (Smith, 2017) or by disabling their web services through DoS attacks (Krebs, 2016) (see Fig. 2.1). Other attackers pursue political interests, such as hacktivism, espionage, sabotage, terrorism, or war[3] (Koppel, 2016). Although other threats proceed from unintentional mistakes, (Ponemon Institute, 2016) confirms the intuition that malicious attacks have larger cost than system glitches and human errors.

Cyber crime has thrived because it is a profitable activity with relatively low risk (attackers usually remain anonymous or reside in countries where they cannot be prosecuted) (Greene, 2006). Also, it is easy to attack multiple targets simultaneously and the attackers do not need advanced knowledge, in part because they can find information and tools in hacker forums

[1] Jim Gettys explains in (Gettys, 2018) that despite the concerns on security, many systems remained unprotected on purpose to avoid export regulations imposed around cryptography. These regulations restricted the export of cryptographic technologies and devices due to their military value. Hence, incorporating strong cryptographic authentication in systems would impede their distribution.

[2] A survey on different industries (Ponemon Institute, 2016) found that most data breaches occur due to malware, criminal insiders, social engineering, and SQL injection.

[3] War and terrorism are acts of violence that can use the same methods; however, they have some differences. A war is a conflict between nation-states, while terrorism is a political act that targets civilians to inflict terror and compel governments to meet some demands.

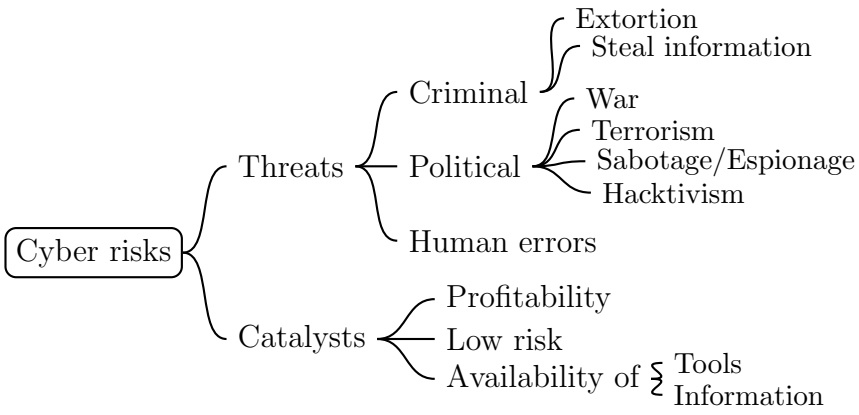

Figure 2.1: Characteristics of cyber risks.

(Ablon et al., 2014; Karpesky Lab, 2014). As a consequence, in recent years the number of information breaches has increased, as well as the number of records compromised in each breach (Latham & Watkins, 2014).

Some consequences of security breaches (for individuals) are identity theft and credit/debit card frauds. Besides, firms can suffer both well-known *direct costs* and long term and uncertain *intangible costs*. Among direct costs we find technical investigations and and identity recover. On the other hand, the biggest concern for firms comes from intangible costs, such as revenue losses, long-term cost of loosing customers, and devaluation of the firm's brand name (Ponemon Institute, 2016; Verizon, 2017). According to (Mossburg et al., 2016), more than 90% of the total costs come from intangible factors.

2.1.2 Cyber Risks of Critical Infrastructures

In the last few years we have witnessed the development of sophisticated attacks that target critical infrastructures. Stuxnet is the first known case of a computer worm designed to harm a physical process. Stuxnet sabotaged centrifuges at the Natanz uranium enrichment plant in Iran in 2010. It manipulated valves to increase the pressure of centrifuges, damaging these devices and spoiling the enrichment process. The attack surprises by its subtlety, because it deceived the operators making them believe that the process was operating normally (Zetter, 2014). Four years later, a second cyber attack disrupted the control systems of a steel mill in Germany (Zetter, 2015). This attack prevented the proper shut down of a blast furnace, causing massive damage to the system.

In 2015, the first confirmed attack on a power grid was launched on the electricity system of Ukraine (Zetter, 2016). The aggressors launched a sophisticated attack using the operator's

workstations to open breakers and interrupt the power flow of around sixty substations. They also overwrote the firmware of control devices, leaving them unresponsive to remote commands. On top of that, the attack was designed to delay the report of incidents and to impede remote actions of operators. In 2017, a group of hackers targeted nuclear facilities from the Wolf Creek Nuclear Corporation in the U.S. (Perlroth, 2017). Although the attackers didn't cause harm in the system (they seemed interested in collecting intelligence about nuclear facilities), damage in the system could cause explosions, fires, and spills of dangerous material.

2.1.3 Characteristics of Attacks on CPS

The threats to critical infrastructures affect both information and hardware components. Hence, hazardous events can lead not only to information leaks, but also to large scale disasters. Attacks on the power grid can have devastating consequences due to the universal dependence on the electricity system. Besides, recovering from an attack on the power grid can take months, because some parts of the infrastructure are not easily replaced. For example, voltage overloads can cause damage to large transformers, which are neither interchangeable nor kept in backup inventories.[4] According to (Koppel, 2016), produce, deliver, and replace one of these transformers would take around two years.

The Cambridge centre for risk studies has studied the potential effects of terrorist attacks on the power systems from the U.S. (Cambridge Centre for Risk Studies, 2015) and the U.K. (Kelly et al., 2016), considering damage generators and substations, respectively. The studies concluded that such attacks would cause total losses from $243 bn to $1 trn in the U.S. and $15 bn to $110 bn in the U.K.[5] The estimated cost of these attacks is comparable with some of the most devastating natural disasters to date, such as Hurricane Katrina, which caused losses of $172bn (Swiss Re, 2017)

2.1.4 Relation with Other Risks

Cyber risks share the following properties with the risk of terrorism (see Fig. 2.2 and Table 2.1):

- The events are caused by an intelligent adversary, rather than by change or by acts of the nature. The main challenge is that intentional actions can damage important components, while natural events are random.

[4]Each transformer has particular specifications and can cost between $3 and $10 million.

[5]The losses include the potential impact on other industries affected by blackouts.

Risk on CPS

- Adversaries are intelligent and pursue financial and political goals
- Threats evolve in time
- Events can affect more people and industries than other threats (aggregate risk)
- Damage is not immediate but surpasses geographical boundaries

Figure 2.2: Characteristics of cyber risk on CPS.

- The attack techniques evolve to circumvent countermeasures of defenders.
- The risk is difficult to predict and control, it does not have geographical boundaries, and can affect multiple industries. Consequently, a vulnerability can allow attacks at a global scale, such as the Petya ransomware attack (Smith, 2017).
- The responsible for the attacks can remain anonymous or out of reach of the authorities.[6] Hence, the government cannot use retaliation as a mechanism to prevent attacks (Petersen, 2008).
- The events can have aggregate risk (correlated risk).[7] For instance, an event such as the 9/11 can raise claims of business interruption, property damage, workers compensation, and life and health insurance simultaneously.

We can find differences between cyber risks and other risks. On one hand, unlike terrorist attacks or natural disasters, cyber attacks don't cause large immediate physical damage, although a cyber attack can comprise a larger area affecting much more people (Koppel, 2016). The greatest concern is that an cyber attack can disable critical infrastructures, for

[6] According to (Risk Management Solutions, Inc. and Centre for Risk Studies, University of Cambridge, 2016): "Conviction rates for cyber criminals are much lower than for many other criminals."

[7] Despite the consensus on the importance of correlation in cyber risks, some studies have found low correlation in cyber events (Biener et al., 2015). Furthermore, (Eling and Wirfs, 2016) reports that costs are higher when only one firm is involved.

long periods of time. On the other hand, (Eling and Wirfs, 2016) finds that cyber risks have both frequent and infrequent events, unlike other threats, such as terrorism. Frequent events with low cost (e.g., DoS attacks and data breaches) are referred as *short tail risk*. Events with low frequency and high cost (e.g., a blackout in the power grid) are referred as *long tail risks* or extreme events.

Some studies have found that events affecting physical and information assets have comparable costs. For instance, the survey in (Ponemon Institute, 2015) found that the maximum losses on physical and information assets are almost the same. However, damage on information assets causes roughly twice the cost on business than damage on physical assets. On the other hand, (Eling and Wirfs, 2016) finds that cyber risks have a lower tail and cause lower losses than other operational risks. In particular, the calculated VaR for non-cyber risks is more than twice the VaR for cyber risks.

2.1.5 Why CPSs Remain Unprotected?

Despite of their importance, critical infrastructures suffer a disturbing lack of security. In particular, a survey on the security of critical infrastructures around the world (Ponemon Institute, 2014) found that most companies suffered a data breach in the last 12 months.[8] Moreover, few companies consider protection (e.g., risk reduction) as one of the main objectives. For instance, in (Ponemon Institute, 2014) most of the companies interviewed (55%) have only one person assigned to security and in few cases (16%) they are aware of the vulnerabilities of the ICS/SCADA. Furthermore, most of the respondents consider that their company is unlikely to update their legacy devices to the next security state. Below we expose some reasons for the low protection of critical infrastructures.

1. Cyber risks are difficult to estimate and evolve in time. Hence, companies that use risk management approaches either disregard the risk (believing that they won't raise the interest of attackers) or accept it, when the cost protecting the system seems larger than the potential consequences (Langner and Pederson, 2013). In both cases the firms do not invest in security.

2. CPS, unlike other IT systems, are not flexible enough to allow the incorporation protection schemes in short periods of time (e.g., days or weeks). Hence, dealing

[8]The survey lists among the targets of the attacks databases, personal devices (e.g., personal computers and smart phones), servers, industrial control systems, and SCADA, among others.

with evolving threats constitute a major difficulty, therefore, anticipating potential vulnerabilities before they raise becomes crucial to protect CPS (Langner and Pederson, 2013).

3. Lack of (or bad) incentives also affect the protection of systems. In particular, Langner and Pederson mention that "... risk management is not a technical approach but a business..." (Langner and Pederson, 2013). Hence, firms and/or technology providers can reduce investments in cyber security to increase their profits.[9] For example, managers with pressure to meet earning benchmarks can reduce the expenditures associated with safety (Caskey and Ozel, 2017).

4. Some firms believe that the traditional insurance covers cyber risks; however, the traditional insurance specializes in risks of *tangible property*, which excludes the risk from cyber threats. Disagreements in the interpretation of insurance policies concerning cyber incidents have led to legal disputes settled in court rooms. In particular, (Latham & Watkins, 2014) shows some examples of lawsuits where the court support insurers who refuse to cover losses caused by information breaches (Romanosky et al., 2017).

[9]Many companies can regard investments in security as expenses, since the society bears most of the losses arising from attacks.

Table 2.1: Characteristics of risks.

	Natural Disasters	**Capital Markets**	**Terrorism**	**Cyber risks** IT	CPS
Threat					
Source	Nature (random)	Market forces (random)	Intelligent adversary		
Motive	Accidental		Political	Criminal Political Accidental	Political Accidental
Scope					
Spatial	Geographical area	Market	Geographical location	Information systems	Physical systems
Insurance					
Type	Standard	Unavailable	Backed by government	Limited coverage	Limited coverage

CHAPTER 3

PERVERSE INCENTIVES IN SECURITY CONTRACTS: A CASE STUDY IN THE COLOMBIAN POWER GRID

3.1 Introduction

In the last four decades, Colombia suffered one of the longest periods of sustained internal conflict. During this period, guerrilla groups targeted most of the country's critical infrastructures as part of their political and economic agenda. In particular, the electricity infrastructure has been one of the main targets. According to data compiled by the National Memorial Institute for the Prevention of Terrorism, between 1994 and 2004 67% of the global attacks to electricity infrastructures occurred in Colombia, while other countries accounted for less than 7% each (Zimmerman et al., 2005).

Fig. 3.1 shows the total number of attacks in Colombia during the last years.[1] The number of attacks decreased presumably thanks to the strengthening of the military forces and anti-terrorism policies. Furthermore, electric companies also developed strategies to cope with attacks, that is, reduce the duration of service interruptions. In particular, electric companies gained expertise repairing transmission towers, which allowed them to reduce the number of pending repairs (see Fig. 3.1). For instance, reparations took around 13 days in 2004, while in 2009 they took 6 days in average. Moreover, the companies can reestablish the service installing provisional towers until they finish repairs.

The experience of public and private sectors operating such critical infrastructures under constant attacks can provide insights to protect other systems. In this case, we analyze a real scenario that illustrates how economic incentives can deteriorate the operation of the power grid. In particular, authorities discovered that Electroservicios (a contractor in charge of repairing transmission towers damaged by guerrilla attacks) colluded with guerrilla groups to destroy electricity towers. Electroservicios paid guerrilla members \$8 million pesos (approximately $\$4K$ USD[2]) to bring down transmission towers, and ISA (a transmission company in Colombia) would pay Electroservicios \$150 million pesos (approximately $\$75K$ USD) to repair each tower. As a result, the guerrilla militants attacked approximately 215

[1]This work uses data extracted from annual reports of an electricity company (Interconexión Eléctrica S.A. E.S.P. (ISA), 2018) and news reports (Semana, 2008; Caracol radio, 2009).

[2]The exchange rates are from 2008.

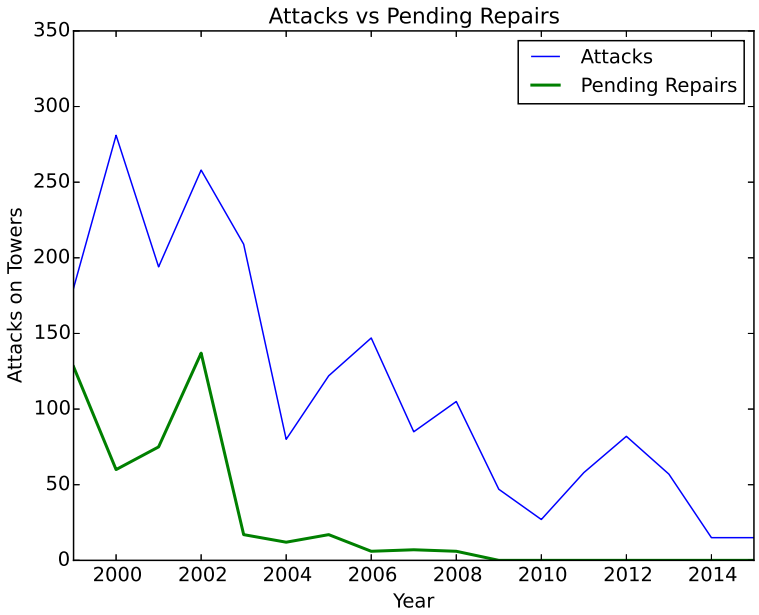

Figure 3.1: Number of attacks and pending repairs of transmission towers in Colombia during 1999-2015.

electric towers between 2005 and 2008 in the region where Electroservicios operated (Semana, 2008; Caracol radio, 2009).

We believe that the lessons learned from this experience can help to protect other systems in similar scenarios. For example, (Krebs, 2017) reported that a DoS protection company carried out attacks on Minecraft servers (using the Mirai botnet) to push sales of its services. Such attacks can succeed because cyber attacks are even harder to attribute than the physical attacks discussed here.

This chapter is organized as follows. Section 3.2 describes the contract scheme used to award repair contracts and how contractors could exploit them to profit. In Section 3.3 we show changes in contracts that can reduce the incentives of contractors to commit fraud. We illustrate the efficacy of the new contracts with a numerical example in Section 3.4. In Section 3.5 we analyze the incentives that other parties, e.g., local workers, have to attack the system and how to mitigate them. We summarize and comment future directions in Section 3.6.

3.2 Background

In this section we construct a model of typical contractual policies; later we introduce the fraud carried out by Electroservicios and show how collusion with guerrillas affect the contractual policies.

3.2.1 Model of Lawful Contractors

Due the large volume of attacks, transmission companies in Colombia had to hire third parties to repair transmission towers (see Fig. 3.1).[3] According to the Colombian contracting code of public administrations (Congreso de Colombia, 1993, 2013), public biddings must follow reverse auctions.[4] In a reverse auction the transmission company wants to buy a service (tower repair) and multiple sellers (who must satisfy the contract specifications) offer bids on the contract (e.g., repair costs) and the seller with the lower bid wins the contract. The reverse auction can have many stages in which bidders make offers using closed envelopes. The bids at each stage start at the lowest bid offered in the previous stage; thus, the sellers compete reducing their bids.

Here we assume that N contractors bid for the contract, where the i^{th} contractor offers repair services with a cost $c_i \geq 0$ and obtains a profit $U_i \geq 0$, with $i = 1, \ldots, N$. According to the reverse auction mechanism the contractor with the lowest bid wins the contract; therefore, the transmission company has to pay $p = c_{min}$ (per tower repaired), where $c_{min} = \min_{i \in \{1,\ldots,N\}} c_i$. Without loss of generality we can assume that $c_1 \leq c_2 \leq \cdots \leq c_N$, hence, $p = c_{min} = c_1$.[5]

With each attack the transmission company deals with repairs and additional operational costs caused by the interruption of the electricity flow to regions that the transmission company must serve. Although the Colombian regulations do not penalize failures to deliver electricity caused by terrorist attacks, transmission companies still must purchase more

[3] In the original setting, the transmission company designated a single contractor to repair the towers in a given region.

[4] Auctions are mechanisms that allow a seller to elicit the private information from buyers and assign a good to the buyer willing to pay the largest quantity (Nisan et al., 2007). A reverse auction follows the same principle, but inverting the roles of the parties.

[5] The repair cost isn't always the same, because it depends on the damage to the towers. For simplicity we assume that terrorist attacks cause the maximum damage to the towers.

expensive sources of electricity (if available), such as carbon-based fuels.[6] We generalize the additional operational costs arising from attacks with the parameter $o \geq 0$.

In summary, with honest contractors the transmission company pays $p = c_1$ to repair each tower. Therefore, the cost of θ attacks for the transmission company is

$$\theta(p + o) \tag{3.1}$$

and the benefit for the contractor is

$$\theta U_1.$$

3.2.2 Fraud in Repair Contracts

In 2007 public authorities started investigations because 93% of all attacks in the country took place in the same region, called Cauca (see Fig. 3.2). The inquiries revealed that the attacks had the following characteristics:

- All towers belonged to the same transmission company (ISA).
- The attackers' modus operandi was the same (e.g., they deployed the explosives in the same place).
- The same contractor repaired all the towers in the region called Cauca.

Repair costs per tower ranged from \$50 to \$150 million pesos ($25K - 75K$ USD). The estimated loses for ISA (the electricity transmission company) were approximately $\$16K$ million pesos (around \$8 million dollars at the time).

In 2008 the authorities infiltrated the contractor and obtained a confession from one of the executives. They found out that the contractor's business boomed since 2005 thanks to sponsored attacks. The contractor did not attack the electric towers directly, instead, they hired four guerrilla militants and paid each one of them \$2 million pesos (around $1K$ USD). The contractor used the following criteria to implement the attacks:

- They chose towers with easy access to facilitate the escape of militants and the arrival of contractors, so they could arrive fast to the site of the repair,

[6]More than 70% of the electricity in Colombia is generated by hydropower, and when attacks limit the transmission of this type of energy, the transmission company needs to satisfy the demand with more expensive carbon-based power.

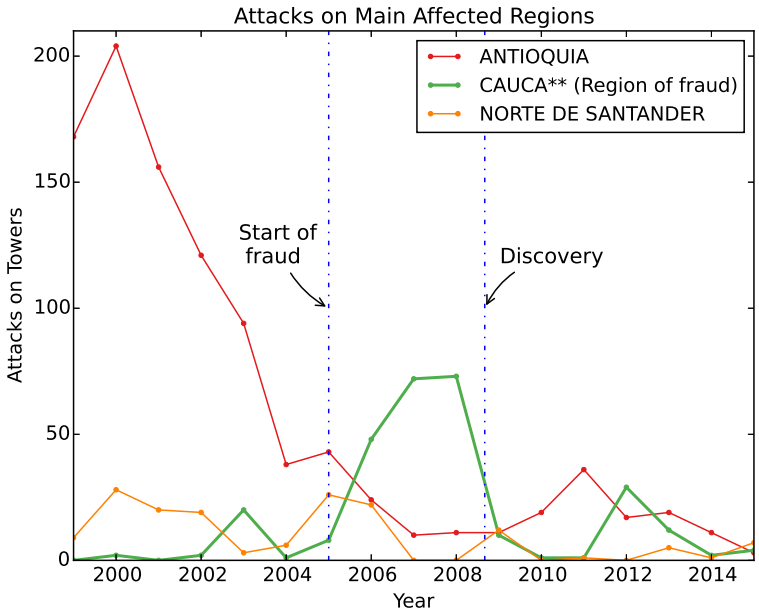

Figure 3.2: Number of attacks from 1999 to 2015 in the three regions with more attacks in Colombia.

- The guerrilla militants had instructions to partially damaged the towers to allow both cheap and fast repairs.
- The militants carried out the attacks only on weekdays, so that the contractor avoided paying overtime its employees.

Thanks to the careful planning of attacks the contractor increased its profits with each repair.

3.2.3 Modeling the Electroservicios Case

Here a contractor sees the opportunity to hire militants to commit attacks on specific towers and increase the frequency of their repair services. Let $\tilde{\theta}_i \in \mathbb{Z}^*$ be the number of attacks sponsored by the i^{th} contractor and $b(\tilde{\theta}_i)$ be the bribe or cost to launch $\tilde{\theta}_i$ attacks. We assume that the bribe $b(\cdot)$ increases with the number the attacks because 1) the capture risk increases with the frequency of the attacks (e.g., the military forces would increase the frequency of patrols) and 2) the militant's opportunity cost (e.g., guerrillas carry out other activities and they can ask higher compensations to spend more time and resources attacking electric towers). Hence, we define the bribe as convex function $b : \mathbb{Z}^* \to \mathbb{R}_+$.

Let use assume that, if a contractor i charges c_i to repair a tower damaged with a sponsored attack, its profit is $\tilde{U}_i$. Since carefully planned (sponsored) attacks reduce the

repair expenses, the contractor will earn a larger profit sponsoring attacks, that is, $\tilde{U}_i \geq U_i$.[7] Let us denote the excess benefit with sponsored attacks as $L_i = \tilde{U}_i - U_i$.

We now parameterize the way in which a contractor can change its bids if he can sponsor attacks. In particular, a contractor can use the additional benefit from sponsored attacks, L_i, to lower its bid and become more competitive, that is, increase the chances to win the contract. For example, if a contractor decides to accept a lower benefit (per tower) $\tilde{U}_i - \gamma L_i$ instead of $\tilde{U}_i$, with $\gamma \in [0, 1]$, then it can charge $c_i - \gamma L_i$ to repair towers damaged through sponsored attacks. However, the bid offered to the transmission company must account for both legitimate and sponsored attacks. Hence, if the contractor sponsors $\tilde{\theta}_i$ attacks, its modified bid, denoted $\tilde{c}_i$, must satisfy

$$\theta c_i + \tilde{\theta}_i(c_i - \gamma L_i) = (\theta + \tilde{\theta}_i)\tilde{c}_i \tag{3.2}$$

The left hand side of Eq. 3.2 shows the payments required to repair θ legitimate attacks and $\tilde{\theta}_i$ sponsored attacks. From the previous expression we have

$$\tilde{c}_i = c_i - \frac{\tilde{\theta}_i}{\theta + \tilde{\theta}_i}\gamma L_i \geq c_i - \gamma L_i.$$

Observe that the new bid $\tilde{c}_i$ decreases with γ and the number of sponsored attacks, i.e., $c_i \geq \tilde{c}_i$. Moreover, the contractor reaches the minimum bid $\tilde{c}_i = c_i - \gamma L_i$ when $\tilde{\theta}_i$ is much larger than θ.

In addition, the profit of the contractor that sponsors $\tilde{\theta}_i$ attacks becomes

$$\theta U_i + \tilde{\theta}_i(\tilde{U}_i - \gamma L_i) - b(\tilde{\theta}_i). \tag{3.3}$$

Thus, if $\gamma = 0$ the contractor does not offer reduced bids, and its benefit per sponsored attack is $\tilde{U}_i$. On the other hand, if $\gamma = 1$, then the contractor accepts the typical benefit U_i (instead of $\tilde{U}_i$) and reduces its bid to the lowest value. The optimal number of attacks, denoted by $\tilde{\theta}_i^v$, solves the following maximization problem

$$\begin{aligned} \underset{\tilde{\theta}_i}{\text{maximize}} \quad & \theta U_i + \tilde{\theta}_i(\tilde{U}_i - \gamma L_i) - b(\tilde{\theta}_i) \\ \text{subject to} \quad & \tilde{\theta}_i \in \mathbb{Z}^*. \end{aligned} \tag{3.4}$$

[7] In practice, the transmission company can inspect the damage on each tower to adjust the payment p. However, here we assume the worst case in which the transmission company makes the same payment independently of the type of attack (legitimate or sponsored).

Observe that a contractor would sponsor attacks when the profit with at least one attack ($\tilde{\theta}_i = 1$) exceeds the cost, that is

$$\tilde{U}_i - \gamma L_i > b(1).$$

From the case of Electroservicios we know that the profit from attacks is much larger than the cost. For this reason, it is necessary to redesign the contracts to avoid incentives of contractors to increase the number of attacks.

Example 1 (Optimal Number of Attacks)**.** *Let us define the cost of one attack as* $b(1) = b_0 + \lambda$, *where* $b_0, \lambda \geq 0$. *We assume that the bribe* $b(\cdot)$ *increases with the number of attacks, in particular, we define the cost of a second attack as* $b(2) - b(1) = b_0 + \lambda(1+\alpha)$, *with* $\alpha > 0$. *In this case, the second attack costs* $\lambda\alpha$ *more than a single attack. Using the previous considerations we define the cost of the* k^{th} *attack as*

$$b(k) - b(k-1) = b_0 + \lambda_k,$$

where b_0 *is a fixed cost and* λ_k *is a variable cost defined with the recursion* $\lambda_k = \lambda_{k-1}(1+\alpha)$, *where* $\lambda_1 = \lambda$. *Now we can define the total bribe for* $\tilde{\theta}_i$ *attacks with the following function:*

$$b(\tilde{\theta}_i) = \sum_{j=1}^{\tilde{\theta}_i} (b_0 + \lambda_j) = \tilde{\theta}_i b_0 + \lambda + \lambda(1+\alpha) + \cdots + \lambda(1+\alpha)^{\tilde{\theta}_i - 1} = \tilde{\theta}_i b_0 + \lambda \frac{(1+\alpha)^{\tilde{\theta}_i} - 1}{\alpha} \quad (3.5)$$

The right hand side equality follows because $\sum_j \lambda_j$ *is a geometric series. Since* $\alpha > 0$, *then* $b(\tilde{\theta}_i)$ *is a strictly increasing convex function, as required by our assumptions.*

We can use Eq. (3.5) *to find the optimal number of sponsored attacks* $\tilde{\theta}_i^v$ *that solves the optimization problem in Eq. (3.4). Since* $b(\tilde{\theta}_i)$ *is convex, then the objective function in Eq. (3.4) is a concave function, and the solution satisfies the following First Order Condition (FOC):*

$$\frac{\partial}{\partial \tilde{\theta}_i} \left(\tilde{\theta}_i (\tilde{U}_i - \gamma L_i) - b(\tilde{\theta}_i) \right) \Big|_{\tilde{\theta}_i = \tilde{\theta}_i^v} = 0.$$

Solving the previous equation we have

$$\tilde{\theta}_i^v = \ln \left(\frac{\alpha(\tilde{U}_i - \gamma L_i - b_0)}{\lambda \ln(1+\alpha)} \right) \Big/ \ln(1+\alpha) \quad (3.6)$$

3.3 Designing Contracts to Disincentivize Attacks by Contractors

The transmission company would prevent attacks reducing asymmetries in information that allow frauds; for example, Yardstick competition helps identifying anomalous behaviors by comparing the costs of similar firms (Shleifer, 1985). This mechanism can help to identify contractors who bid much lower; however, a malicious contractor aware of the regulation can offer bids consistent with the bids in other markets. Here we discuss a contract scheme designed to reduce incentives to sponsor attacks in the power system infrastructure. The model is based on how the transmission company of Colombia changed the contractual policies after the case of Electroservicios came to light.

3.3.1 Contractor Side

The basic idea of the new contract structure consists in selecting n contractors and assign them specific repairs randomly.[8] In this way, contractors that sponsor attacks cannot know if they will repair the attacked towers, which give them additional profit.[9] In this case the transmission company uses an auction to select n contractors with the lowest bids. With the new contract the expected profit of contractors that sponsor $\tilde{\theta}_i$ attacks becomes

$$\frac{\theta U_i + \tilde{\theta}_i(\tilde{U}_i - \gamma L_i)}{n} - b(\tilde{\theta}_i). \tag{3.7}$$

Thus, the optimal number of sponsored attacks with the new contract, denoted by $\tilde{\theta}_i^m(n)$, depends on the number of contractors n and solves the following maximization problem:

$$\begin{aligned} \underset{\tilde{\theta}_i}{\text{maximize}} \quad & \frac{\theta U_i + \tilde{\theta}_i(\tilde{U}_i - \gamma L_i)}{n} - b(\tilde{\theta}_i) \\ \text{subject to} \quad & \tilde{\theta}_i \in \mathbb{Z}^*. \end{aligned} \tag{3.8}$$

Observe that the new contract scheme with $n > 1$ reduces the number of sponsored attacks ($\tilde{\theta}_i^v \geq \tilde{\theta}_i^m(n)$), because the expected benefits of contractors decrease with respect to n. For instance, with $n = 2$ the contractor's benefits decrease more than a half. The reduction in benefits will be greater for small values of n, and less significant as n increases.

[8] The contractors can incur in additional costs if they have to repair towers in many regions, rather than a single one.

[9] The new contracts would fail if a large set of contractors collude in attacks, but as far as we are aware, that level of corruption hasn't been encountered in Colombia.

Since the objective function in Eq. (3.8) is concave, the optimal number of attacks $\tilde{\theta}_i^m$ satisfies the following FOC:

$$\frac{\partial}{\partial \tilde{\theta}_i}\left(\frac{\tilde{\theta}_i}{n}(\tilde{U}_i - \gamma L_i) - b(\tilde{\theta}_i)\right)\Bigg|_{\tilde{\theta}_i=\tilde{\theta}_i^m} = 0. \tag{3.9}$$

The electricity transmission company can prevent sponsored attacks selecting n large enough to guarantee that the cost of (at least) one attack exceeds the expected profit of the contractors, that is,

$$\frac{\tilde{U}_i - \gamma L_i}{n} < b(1). \tag{3.10}$$

3.3.2 Utility Side

Selecting n contractors (instead of just one with the lowest bid) increases the costs for the electric transmission company. In particular, the payment for individual repairs is larger in this second contract because it must cover the repair expenses of all the selected contractors.[10] Therefore, we define the payment as $\hat{p} = c_n$ (we assume that the contractors report truthfully their bids). Furthermore, the transmission company must pay also the expenses of moving personnel and equipment to the tower. Such expenses will depend on the location of both the company and the attack. For simplicity, let us assume that the transmission expenses c_t are constant for all the contractors. Thus, the repair payment in the new contract is equal to

$$\hat{p} = c_n + c_t.$$

Hence, the additional cost with respect to the original contract is

$$p_r(n) = \hat{p} - p = c_n - c_1 + c_t.$$

Thus, the expected cost for the transmission company becomes

$$(\theta + \tilde{\theta}_i^m(n))(p + p_r(n) + o). \tag{3.11}$$

The transmission company would choose n companies to make attacks unprofitable with minimum expenses. We express the problem of selecting the number contractors as

$$\begin{aligned} \underset{n}{\text{minimize}} \quad & (\theta + \tilde{\theta}_i^m(n))(p + p_r(n) + o) \\ \text{subject to} \quad & n \geq 1, \\ & \text{Eq. (3.10)}. \end{aligned}$$

[10] The new contract can also increase the repair times, but we do not contemplate that cost here because the transmission company does not pay penalties for service interruptions caused by terrorist attacks. However, users suffer higher losses if the repair times increase.

We include Eq. (3.10) as a restriction to guarantee that n is large enough to avoid attacks; this occurs when public policies decree zero tolerance with terrorist actions. However, without such restriction the transmission company can decide whether to implement the new contract evaluating its expected cost. In particular, in some cases it can be cheaper to allow attacks, that is, when the cost of with the original contract is lower than the cost with the new mechanism

$$(\theta + \tilde{\theta}_i^v)(p + o) < (\theta + \tilde{\theta}_i^m(n))(p + p_r(n) + o).$$

If n satisfies the Eq. (3.10) then we have

$$\tilde{\theta}_i^v(p + o) < \theta(c_n - c_1 + c_t)$$

If the bids of contractors are close (i.e., $c_n - c_1 \approx 0$), then the transmission expenses c_t and the number of legitimate attacks θ will determine the convenience of the mechanism.

3.4 Numerical Example

3.4.1 Estimation of Parameters

Since we do not have enough information to estimate all the parameters of the model, we extract some parameters from news reports and make further assumptions to give values to other parameters of the model.

The news report by Caracol (Caracol radio, 2009) mentions that approximately 215 attacks on energy towers were sponsored in 3 years. The report mentions that the transmission company paid around \$150 million pesos (\$83,333 USD) to repair each tower (labor costs are less expensive in Colombia than in the US or Europe). Hence we assume that the cost to repair a tower that suffers a legitimate attack is

$$c_1 = p = \$83,333,$$

where c_1 includes both net repair expenses E_1 and the expected benefit U_1, such that

$$c_1 = E_1 + U_1.$$

If we assume a rate of return of 10%, then the contractor expects a benefit $U_1 = 0.1E_1$ from an investment of capital E_1. Therefore, the repair cost is $c_1 = 11U_1$. Hence, we estimate that the benefit of the contractor with legitimate attacks is

$$U_1 = c_1/11 \approx \$7,576.$$

On the other hand, we assume that careful attacks can lower the damage of the towers. The report in (Semana, 2008) mentions that the minimum repair payment was \$50 million pesos (\$27,778 USD). Hence, we assume that sponsored attacks reduce the repair costs to the minimum. If we denote the minimum repair cost as $\underline{c_1} = 27,778$, then the benefit is $\underline{U_1} = \underline{c_1}/11 \approx \$2,525$ and the minimum expenses are $\underline{E} = \$25,253$. Moreover, we consider that the transmission company does not know the exact damage of the tower, then it will make the usual payment p, leaving the contractor with a benefit per tower of

$$\tilde{U}_1 = p - \underline{E} = \$58,081.$$

Here the benefit with sponsored attacks $\tilde{U}_1$ is more than seven times the benefit received by contractors when they repair electric towers with "regular" (i.e. not sponsored) attacks.

Now, let us define the bribe required to attack one tower as $b(1) = \$4,444$ (recall that the attacks were made by 4 militants, whose fee was \$1,111 USD). Let us assume that $b(1) = b_0 + \lambda$, with a variable cost equal to the 20% of the constant cost, that is, $\lambda = 0.2b_0$. Consequently, $b(1) = 1.2b_0$ and $b_0 = \$3,704$. If we assume that the number of sponsored attacks was optimal, then we can estimate the average number of attacks during one year as $\tilde{\theta}_i^v = 215/3 \approx 72$. Besides, in a competitive auction the contractor would have to reduce its bid the most it can to increase its chances to get the contract. Hence, we can assume that $\gamma = 1$. Moreover, the only parameter that remains is α, which can be estimated from Eq. (3.6) as $\alpha = 0.0234$. Finally, we assume that the transmission cost c_t and operational losses o are equal to zero.

3.4.2 Number of Sponsored Attacks

We are interested in observing the optimal number of attacks with each contract. Fig. 3.3 shows that (in the initial contract) the number of attacks $\tilde{\theta}_i^v$ (see Eq. (3.6)) increases as γ decreases. This happens because with small γ the contractor has more benefits per tower, and thus, more incentives to attack. However, getting more benefits per tower can prevent them from offering a competing bid in the first place, so this is something the attacker needs to balance when submitting the bids.

We now investigate changes in the number of attacks with the modified contracts, designed to reduce the perverse incentives of contractors. Solving Eq. (3.9) we have that optimal number of attacks $\tilde{\theta}_i^m(n)$ in the new contract is

$$\tilde{\theta}_i^m(n) = \frac{1}{k_3}\left(\ln k_1 + \ln\left(\ln(k_2 - b_0 n) - \ln n\right)\right), \tag{3.12}$$

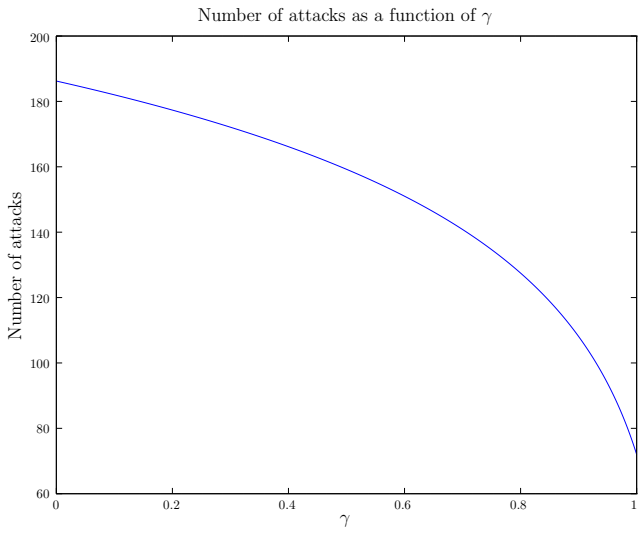

Figure 3.3: Number of attacks as a function of the bid reduction determined by γ.

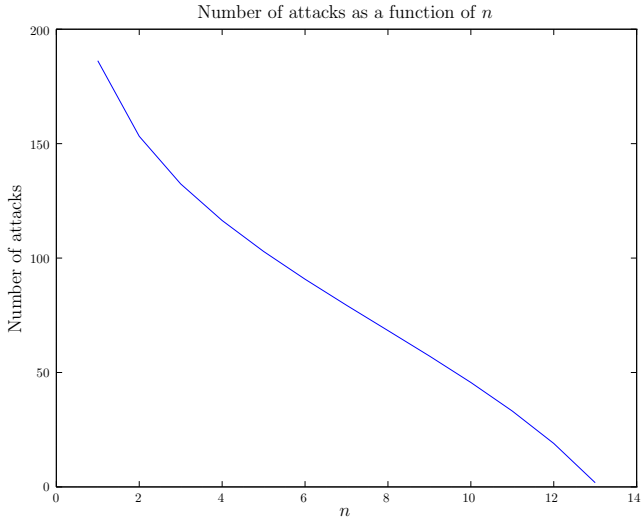

Figure 3.4: Number of attacks as a function of the number of companies n.

were $k_1 = {}^{\alpha}/_{\lambda \ln(1+\alpha)}$, $k_2 = \tilde{U}_i - \gamma L_i$, and $k_3 = \ln(1+\alpha)$.

Fig. 3.4 shows the optimal number of attacks $\tilde{\theta}_i^m(n)$ in a contract with the proposed mechanism (see Eq. (3.12)). In this case we assume that $\gamma = 0$, which results in the best scenario for the contractor. In this experiment the number of attacks is greater than one if $n \leq 13$.

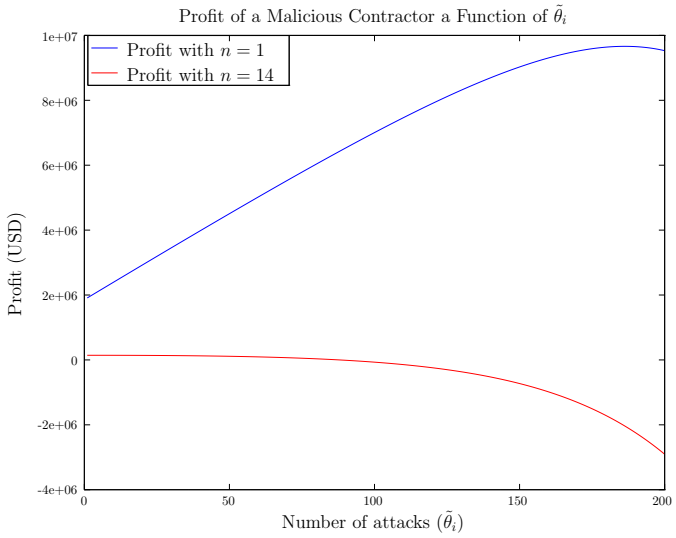

Figure 3.5: Profit of a contractor that sponsors attacks in contracts with either 1 or 14 contractors. The inclusion of more contractors decreases the optimal number of attacks.

3.4.3 Profit of the Parties

Fig. 3.5 shows the maximum profit of a contractor (i.e., profit when $\gamma = 0$) as a function of the number of attacks $\tilde{\theta}_i$ in both the original contract and the new contract designed to prevent perverse incentives (see Eqs. (3.3) and (3.7), respectively). The number of attacks that maximizes the contractor's profit in the original contract (or a contract with $n = 1$) is $\tilde{\theta}_i^v = 186$. However, if the transmission company implements a contract with $n = 14$ contractors, then the optimal number of attacks becomes $\tilde{\theta}_i^m = 0$. Thus, random selection of contractors reduces the incentives for sponsored attacks.

Fig. 3.6 shows the cost for the transmission company (see Eqs. (3.1) and (3.11)) as a function of the number of contractors n. We show an example in which the new contract is not convenient because raises the costs of repairs. In particular, if $p_r(n) = 30p$, then the cost for the transmission company is higher with the new contract.

3.5 Worker's Incentives to Sponsor Attacks

The contractors often hire non-specialized workforce from the region of the accident to reduce costs. A concern is that these workers would demolish the towers to be hired. In particular, individuals who are unemployed or have a low salary profit by sponsoring attacks. In this section we analyze the conditions in which workers can sponsor attacks and how to make these attacks unprofitable.

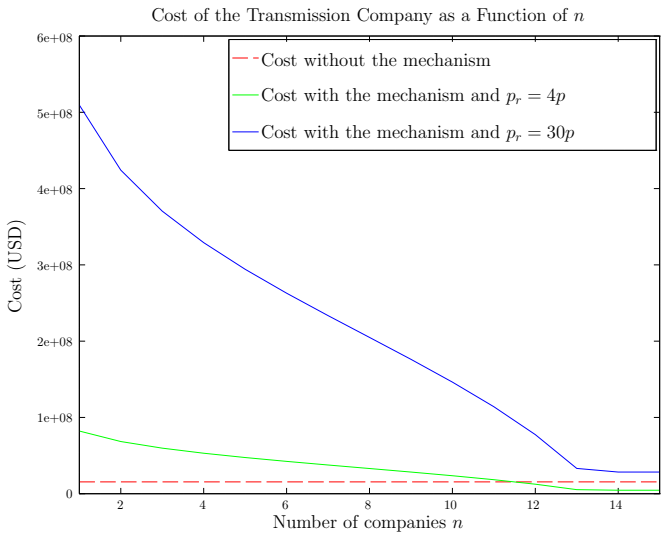

Figure 3.6: Cost for the transmission company as a function of the number of contractors n.

3.5.1 Incentives of Workers

Let us consider the conditions that make attacks profitable for a single worker. Let S be the salary paid by the repair company, S_{min} the minimum salary (either the salary of its current job or the minimum salary accepted), and τ the time that takes to repair a tower. We assume that repairing towers give a higher compensation, i.e., $S \geq S_{min}$, because repairing towers is a risky activity. Specifically, workers are exposed to bombs hidden close to the towers and they can be kidnapped by the guerrilla groups (El Tiempo, 2000).

Thus, sponsoring attacks on θ_i towers is profitable if the profit with an attack is higher than its cost. We can express this condition as

$$(S - S_{min})\tau\theta_i > b(\theta_i) \tag{3.13}$$

This can be rewritten as

$$S - S_{min} \geq \frac{b(\theta_i)}{\tau\theta_i}.$$

From this expression we can see that longer repair periods increase the interest in sponsoring attacks. We know that in the worst case repairs can take $\tau = 13$ days. Consequently, at least one attack is profitable if

$$S - S_{min} \geq \frac{b(1)}{13} = \$341.85.$$

The minimum daily wage in Colombia during 2005 was $S_{min} = \$7.89$. Hence, an individual worker can sponsor an attack if his payment is around 40 times the minimal wage, which

seems unlikely. On the other hand, if the company hires m workers, then the workers can form coalitions to share the cost of sponsoring an attack. We can use Eq. (3.13) to find that a profitable attacks satisfies

$$S - S_{min} \geq \frac{b(\theta_i)}{m\tau\theta_i}.$$

We know that, in the worst case, a tower requires 4 quadrilles (teams) to repair it. Each quadrille is composed by 25 persons, of which 14 are specialists. Hence, we assume that 11 persons per quadrille can be hired from the local region. Consequently, we assume that the company hires $m = 44$ local workers. With these parameters we calculate the minimum salary that incentives at least one attack

$$S - S_{min} \geq \frac{4444}{44 \cdot 13} = \$7.77.$$

Note that in the worst case, unemployed workers have $S_{min} = 0$. Since contracts cannot offer salaries lower than the minimal wage, then a coalition of workers in a region can get some profit with one attack, because the minimal wage ($7.89) exceeds the minimum salary $7.77 required to sponsor an attack.

The transmission company would avoid attacks selecting the salary as $S = S_{min}$. However, this approach fails because the workers can request higher compensation than the minimum wage to repair the towers. Also, the transmission company can ignore S_{min}.

3.5.2 Incentives with Random Selection of Workers

Alternatively, we can think in a raffle to choose workers. Let us assume that the total number of possible workers is M. If m workers plot an attack expecting to be hired by the company, k workers of the coalition have the following probability of being hired:

$$B(M, m, k) = \frac{\binom{m}{k}\binom{M-m}{m-k}}{\binom{M}{m}}.$$

The expected number of selected workers from the coalition is less than m if $M > m$. That is,

$$\bar{m} = \sum_{k=0}^{m} B(M, m, k)k < m.$$

Furthermore, $\bar{m}$ decreases as M increases. Thus, the condition for a profitable attack becomes

$$S - S_{min} \geq \frac{b(\theta_i)}{\tau\theta_i\bar{m}} > \frac{b(\theta_i)}{\tau\theta_i m}$$

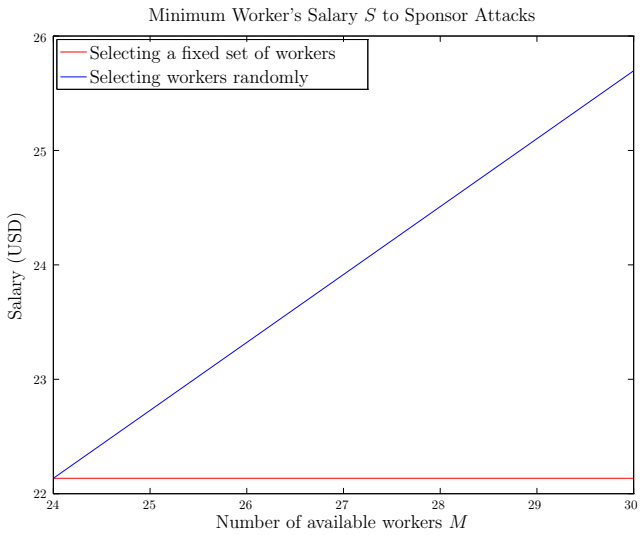

Figure 3.7: Worker's salary that allow them to sponsor attacks (as a function of the number of available candidates M).

We can see that with a raffle is harder for the attackers to coordinate attacks, because the repair salary must be higher than the case without raffles. Fig. 3.7 shows the salary S that allows a coalition of $m = 44$ workers to get profit from attacks. Note that the minimum salary to sponsor attacks increases by selecting workers randomly.

3.6 Conclusions

In this chapter we model a series of attacks that happened in the Colombian power system, and the actions the electric transmission company took to minimize future contractors and workers from launching similar attacks. This research has some limitations. On one hand, we assume that only one contractor engages in fraud, that is, the bidders do not form coalitions. However, in practice multiple contractors can associate to defraud the electric company. On the other hand, we didn't model trade off between the number of companies and the cost for the electricity company. It can occur that it is convenient to allow attacks (at least some), when having a large pool of contractors increases excessively the repair costs.

We believe that the strategic nature of attackers, defenders, and the ecosystem of industries and other agents involved in the protection of large critical infrastructures in Colombia can serve to find analogies for the protection of critical infrastructures against cyber attacks. For example, an anti-DDoS service provider launched DDoS attacks on Minecraft servers (using the Mirai botnet) force its owner to use its service (Krebs, 2017). Similar to the case studied

in this chapter, to prevent these types of attacks companies can hire the services of multiple anti-DDoS companies, who wouldn't know in advance whether they will be hired to deal with a particular incident.

CHAPTER 4

IMPACT OF THE MARKET STRUCTURE IN THE SECURITY OF SMART GRIDS

Authors – Carlos Barreto, Alvaro A. Cardenas, Nicanor Quijano, and Eduardo Mojica-Nava

The Computer Science Department, EC 31

The University of Texas at Dallas

800 West Campbell Road

Richardson, Texas 75080-3021

Corresponding author: Carlos Barreto.
The content of this chapter is reprinted with permission from: 1) C. Barreto, A. A. Cárdenas, N. Quijano, and E. Mojica-Nava, "CPS: Market Analysis of Attacks Against Demand Response in the Smart Grid," In Proceedings of the 30th Annual Computer Security Applications Conference, ©2014 ACM, Inc. http://doi.acm.org/10.1145/2664243.2664284; and 2) C. Barreto and A. A. Cárdenas, "Detecting fraud in demand response programs," In 2015 54th IEEE Conference on Decision and Control (CDC), ©2015 IEEE.

4.1 Introduction

The electricity market is an interesting target because it introduces economic incentives, which extend the possible objectives of adversaries; for example, adversaries can attempt to change the market's equilibrium, rather than breaking the whole system. Furthermore, successful attacks on the market can have as much impact as some physical attacks, because the market affect the decisions of users, which ultimately affect the physical components of the power grid.

Although we don't have records of cybernetic attacks targeting markets, modern power grids (called *smart grids* (SG)) introduce technologies that can extend the reach of attackers. In particular, *demand response* (DR) programs use new technologies envisioned to coordinate the demand of customers and improve the efficiency of the market. Unfortunately, such technologies can become tools that allow the attackers to pursue more sophisticated enterprises.

In this work we analyze how DR programs can give attackers a new way to defraud the electricity system without the risks of being identified. Specifically, by attacking the DR signals sent by the DR system (as shown in Fig. 4.1), the attacker adds a layer of indirection hindering his identification. For example, the attacker can instruct a subset of consumers $\mathcal{V}$ (set of victims) to reduce their electricity consumption, which in turn will reduce the cost of electricity. In this way, other users (from a set of attackers $\mathcal{A}$) can benefit consuming larger amounts of electricity at reduced prices. If the set $\mathcal{A}$ is large enough, then a forensic analysis will reveal many suspects, who can claim *plausible deniability*, hindering the precise identification of the culprit.

In this work we formulate the aforementioned problem, analyze ways to detect attacks (i.e., detect that an attack is occurring, but not who is responsible for it), and propose ideas to reduce economic incentives to defraud the system. We also show how a malicious attacker can cause peaks in demand to cause blackouts.

We model a general scenario in which a central planner implements a DR scheme designed to achieve efficient outcomes, that is, maximize the *customer surplus*. We model two attackers that select their attack strategy in a principled way to achieve specific goals: 1) maximize its profits and protect its identity, and 2) damage the system. We focus on two previously proposed DR schemes with two general control approaches, namely centralized and decentralized control. Our objective is to evaluate how centralized and decentralized systems affect the implementation and detection of attacks.

We find that attackers can carry out successful attacks on both centralized and decentralized systems. However, attacks on decentralized systems lead to less profit; furthermore,

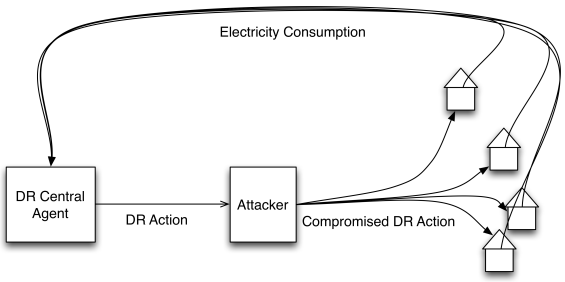

Figure 4.1: Adversary Model: by compromising DR signals, the attacker can affect the behavior of a large sector of the population, and instruct them to behave in a manner beneficial for the attacker (e.g., force them to reduce electricity consumption so the attacker can get electricity at reduced rates).

information asymmetries of decentralized systems prejudice the detection of attacks (i.e., differentiate attacks from faults) and the estimation of penalties.

4.1.1 Literature Review

Here we investigate how false data injection attacks can affect electricity markets (Liu et al., 2009). Previous works focus on the bulk electricity market (Liyan et al., 2012; Negrete-Pincetic et al., 2009); however, it is more likely that attacks will happen in the retail market, since it has more participants with highly varying levels of trustworthiness, which increases the difficulty to attribute attacks. Thus, attackers will have higher incentives for attacking retail markets than bulk-electricity markets.

On the other hand, (Tan et al., 2013) analyzed the impact of integrity attacks on the retail DR market and showed attacks on price signals (scaling attacks and delay attacks) that cause severe oscillations of the electricity demand. Our work departs from (Tan et al., 2013) in several ways. On one hand, (Tan et al., 2013) modeled the market with a single-input single-output linear system. Instead, we incorporate market interactions of a multi-agent system where each agent has a nonlinear valuation of electricity, similar to (Roozbehani et al., 2012, 2010; Huang et al., 2012; Samadi et al., 2011; Chen et al., 2010; Ibars et al., 2010; Gellings, 2009; Fahrioglu and Alvarado, 1998; Li et al., 2011). Moreover, one contribution of this work is to model a more powerful attacker that can select an arbitrary attack signal.

Our work is closely related to (Liu et al., 2016), which analyzes attacks on the retail market that attempt to either reduce prices (benefiting the attacker) and create peaks in demand. In addition, (Liu et al., 2016) proposes a detection scheme using partially observable Markov decision processes. Our work departs from (Liu et al., 2016) in the following aspects.

First, we consider attackers with specific objectives, which allows us to define precise attack signals, rather than heuristic attacks. Second, the detection scheme in (Liu et al., 2016) relies on historical data to estimate the security of the system. In contrast, we adopt a pessimistic posture assuming that the utility does not have reliable data about previous attacks. Hence, our detection mechanism only assumes knowledge of the normal demand of users.

In this work we extend our previous works (Barreto and Cárdenas, 2015a; Barreto et al., 2014) formalizing the properties of the attack, the detection mechanism, and the penalties.

4.1.2 Outline

In Section 4.2 we introduce the market model of the electricity system and two DR schemes. Section 4.3 introduces the goal of the attacker and the attack strategies for each DR scheme. We introduce detection schemes and penalties for each DR system in Sections 4.4 and 4.5, respectively. We summarize the main conclusions and comment future directions in Section 4.6.

4.2 Background: Market Models

We consider a market with three participants, namely users, generators, and a DR operator. The users, $\mathcal{P} = \{1, \dots, N\}$, consume $q_i \geq 0$ units of energy and obtain a benefit defined by the function $v_i(q_i)$, for $i \in \mathcal{P}$. The system can have multiple generators who supply the total demand of energy. Although the cost of producing energy depends on both the type of the generator (e.g., thermal, hydroelectric, or nuclear) and the distribution losses, here we assume that the cost depends only on the total demand. In particular, we consider that the function $C(g)$ determines the cost of producing g units of energy. Moreover, the DR operator manages the market and coordinates the distribution of power to guarantee an equilibrium between generation and demand. For simplicity, we ignore the losses of distribution, therefore, the total generation equals the total demand, i.e., $g = \sum_{i \in \mathcal{P}} q_i$. The DR operator also chooses the tariff of energy p charged to users.

In summary, we can express the surplus of the i^{th} user as

$$U_i(q_i, p) = v_i(q_i) - q_i p,$$

where the vector $\boldsymbol{q} = [q_1, \dots, q_N]$ represents the demand of the population and $\|\cdot\|$ is the 1-norm, i.e., $\|\boldsymbol{q}\| = \sum_{i \in \mathcal{P}} q_i$. We make the following assumptions about the market's parameters

Assumption 1. *The valuation function $v_i : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is twice differentiable, strictly concave, non-decreasing, and satisfies $v_i(0) = 0$.*

The cost function $C : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is differentiable, strictly convex, and non-decreasing.

The electricity tariff p and the behavior of users determine the equilibrium of the system. In the traditional power grid, users cannot observe changes in prices, therefore, we assume that they are *non-strategic*, that is, they do not consider the effect of their actions in the prices. However, new technologies of the smart grids improve the information and decision capabilities of users, so they can become *strategic*, in other words, they make decisions anticipating their impact in the market.

The existence of natural monopolies in the electricity system compel price regulations to guarantee efficiency in the system. Here we assume that the DR operator plays the role of a regulator who chooses the price tariff p to protect users from monopolies. In particular, the DR operator can use *average cost prices*, defined as

$$p(g) = \frac{C(g)}{g},$$

which guarantee that payments by users cover the production costs, allowing fair return to the generators (included in the cost function) (Laffont and Tirole, 1993).

4.2.1 Market with Strategic Users

Strategic users make decisions anticipating changes in the prices, which in turn depend on the total demand. Therefore, the profit of strategic users becomes[1]

$$U_i\left(q_i, \|\boldsymbol{q}\|\right) = v_i(q_i) - q_i p\left(\|\boldsymbol{q}\|\right), \quad i \in \mathcal{P}. \tag{4.1}$$

Here we assume that users know the price function $p(\cdot)$ beforehand, hence, they only need to observe the total demand in the system $\|\boldsymbol{q}\|$ to make decisions. In particular, rational agents will choose the demand that maximizes their profit, therefore, they choose the demand q_i to solve the following optimization problem

$$\begin{aligned} \underset{q_i}{\text{maximize}} \quad & U_i\left(q_i, \|\boldsymbol{q}\|\right) \\ \text{subject to} \quad & q_i \geq \underline{Q_i}, \end{aligned} \tag{4.2}$$

Observe that the optimal demand of users changes when they become strategic and this can produce inefficient outcomes, such as *the tragedy of the commons* (Barreto et al., 2013; Hardin, 1968) or the *price of anarchy* (Papadimitriou, 2001). Hence, it is necessary to coordinate the actions of users to improve the efficiency of the system. Below we introduce mechanisms to prevent inefficient outcomes.

[1]Users only need the total demand, rather than the precise demand of all users. Therefore, users reveal their consumption, which is private information, only to the utility company.

4.2.2 Demand Response Schemes

With new technology we can aspire to improve the efficiency of the system, because users become active participants in the market. DR schemes seek to coordinate the actions of strategic users to reach the social optimal outcome, defined as the demand $\boldsymbol{\mu}$ that maximizes the customer surplus, therefore, $\boldsymbol{\mu}$ solves

$$\begin{aligned} \underset{\boldsymbol{q}}{\text{maximize}} \quad & f(\boldsymbol{q}) = \sum\nolimits_{i\in\mathcal{P}} U_i(q_i, \|\boldsymbol{q}\|) \\ \text{subject to} \quad & q_i \geq \underline{Q_i}, \quad i \in \mathcal{P}. \end{aligned} \tag{4.3a}$$

Assumption 2. *We assume that in the optimal outcome all users have a positive demand that satisfies $\mu_i > \underline{Q_i}$ for all $i \in \mathcal{P}$.*

The customer surplus is continuous, differentiable, and concave (see Assumption 1). Furthermore, from Assumption 2, the constraint in Eq. (4.3a) is not binding, therefore, the optimal solution $\boldsymbol{\mu}$ satisfies the following (Boyd and Vandenberghe, 2004; Kuhn and Tucker, 1951; Karush, 1939)

$$\frac{\partial}{\partial q_i} f(\boldsymbol{q})\big|_{\boldsymbol{q}=\boldsymbol{\mu}} = \dot{v}_i(\mu_i) - p\left(\|\boldsymbol{\mu}\|\right) - \|\boldsymbol{\mu}\|\, \dot{p}\left(\|\boldsymbol{\mu}\|\right) = 0, \tag{4.4}$$

for all $i \in \mathcal{P}$.

Below we consider two demand response models that lead to the optimal outcome (or demand) $\boldsymbol{\mu}$ in the equilibrium using centralized and decentralized approached, namely *direct load control* and *dynamic prices* (Barreto et al., 2013; Vardakas et al., 2015).

Direct Load Control (DLC)

This approach consists in implementing a centralized control of the loads. In particular, users reveal their private information (valuation function) and give control of their thermostats and other devices to the utility, which in turn computes and applies the optimal demand (see Fig. 4.2a). DLC can be implemented using the Vickrey-Clarke-Groves (VCG) mechanism, which allocates resources maximizing the social interest, guaranteeing that users report truthfully their private information (Vickrey, 1961; Clarke, 1971; Groves, 1973). The exact implementation of DLC is out of scope in this work, because we are interested mainly in the general characteristics of this approach. We refer the interested reader to the implementation of DLC in (Vardakas et al., 2015).

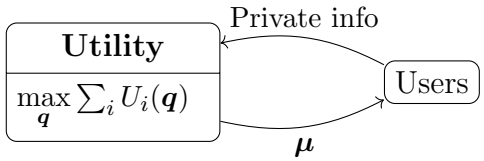

(a) Centralized DR system in which users reveal their preferences to the utility, who computes the optimal demand and send it to users.

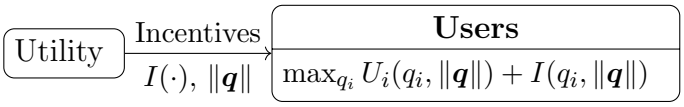

(b) Decentralized DR system that uses an incentive function $I_i(\cdot)$ to align the objectives of users.

Figure 4.2: Demand response schemes.

Dynamic Prices (DP)

Dynamic pricing is a decentralized control approach (see Fig. 4.2b), which unlike DLC, allows users to manage their consumption using the information provided by the utility (e.g., prices or total demand). In this case we consider the DR scheme in (Barreto et al., 2013; Barreto and Cárdenas, 2015b), which modifies the user's profit function in Eq. (4.1) adding the incentive function

$$I_i(\boldsymbol{q}) = \left\|\boldsymbol{q}_{-i}\right\| \left(p\left(\frac{N}{N-1} \left\|\boldsymbol{q}_{-i}\right\| \right) - p\left(\|\boldsymbol{q}\|\right) \right),$$

where $\left\|\boldsymbol{q}_{-i}\right\| = \|\boldsymbol{q}\| - q_i$ is the total demand without including the i^{th} user. The incentive function charges fees to users with high demand and gives incentives to the users with low demand. In this way the incentives internalize the burden caused by users to the system.

With the incentive $I(\cdot)$ the user's optimization problem in Eq. (4.2) becomes

$$\begin{aligned} \underset{q_i}{\text{maximize}} \quad & W_i(q_i, ||\boldsymbol{q}||) = U_i(q_i, ||\boldsymbol{q}||) + I(q_i, \|\boldsymbol{q}\|) \\ \text{subject to} \quad & q_i \geq \underline{Q}_i, \quad i \in \mathcal{P}, \end{aligned}$$

which has the same *first order conditions* (FOC) that the optimal solution Eq. (4.4), hence, the dynamic prices scheme leads to the optimal outcome $\boldsymbol{\mu}$. As a drawback, this mechanism requires external subsidies, because the total amount of incentives is positive, i.e., $\sum_i I(q_i, \|\boldsymbol{q}\|) > 0$.

4.3 Adversary Model

In this section we analyze attacks that exploit the DR infrastructure to change the state (demand) of the system. We consider two types of adversaries, a *fraudster* and a *malicious* attacker. The fraudster is a user who attempts to profit from the attack, while the malicious attacker is an external individual who tries to damage the system. We assume that both attackers can control the signals sent by the DR system to the users (see Fig. 4.1). Hence, with DLC the attacker modifies directly the demand of users, while with DP the attacker modifies the incentives sent to users. Below we describe the objective of each attacker and show that, although the attackers have different objectives, they can use the same techniques to achieve their goals.

4.3.1 Fraudster Attacker

The objective of the fraudster consists in maximizing his profit while avoiding being identified. We consider the following restrictions: 1) The attacker cannot modify his own electricity bill, for instance, by hacking his smart meter to report less consumption.[2] 2) The attacker has as much information as the DR operator. This means that, depending on the DR scheme, the attacker can access either only the consumption (in DP) or the consumption and the valuation function of users (in DLC).

Since the attacker cannot modify his bill directly, he can try to compromise the DR system to force a favorable state. For example, the attacker can use the DR system, which originally maximizes the customer surplus (see Eq. (4.3)), to maximize his own profit,[3] that is,

$$\begin{aligned} \underset{q_1,\dots,q_N}{\text{maximize}} \quad & U_i\left(q_i, \|\boldsymbol{q}\|\right) \\ \text{subject to} \quad & q_i \geq \underline{Q_i}, \quad i \in \mathcal{P}. \end{aligned} \tag{4.5}$$

A disadvantage of this attack lies in its risk, because once the attack is detected, the utility can identify the perpetrator precisely, since only the fraudster profits from the attack. The fraudster can reduce the risk of the attack implementing another attack that both protects

[2]The utility company knows the total energy distributed to the users and their reported consumption. Therefore, imbalances between the generated energy and the reported consumed energy could lead to investigations to trace the cause, revealing the identity of the attacker.

[3] The individual optimization problem in Eq. (4.2) differs with the problem in Eq. (4.5) in its decision variables. Specifically, in Eq. (4.2) the attacker controls only the his own demand q_i, while in Eq. (4.5) the attacker controls the demand of all users.

his identity and allows him to benefit from the attack. Specifically, we consider attacks in which the fraudster tries to conceal his identity sharing the benefits of the attack with other users, preventing the precise identification of the culprit. Let us represent the fraudster's objective with the following optimization problem (which is solvable with the DR schemes mentioned before (Barreto et al., 2014)):

$$\begin{aligned} \underset{\boldsymbol{q}}{\text{maximize}} \quad & f_a(\boldsymbol{q}) = \lambda \sum\nolimits_{i\in\mathcal{A}} U_i(\boldsymbol{q}) + \sum\nolimits_{j\in\mathcal{V}} U_j(\boldsymbol{q}) \\ \text{subject to} \quad & q_i \geq \underline{Q_i}, i \in \mathcal{P}. \end{aligned} \tag{4.6}$$

In this case the attacker's objective function $f_a(\cdot)$ partitions the population $\mathcal{P}$ in two sets, $\mathcal{A}$ and $\mathcal{V}$, and it provides benefits to the users in set $\mathcal{A}$. For λ large enough, the DR system will (for practical purposes) maximize the profit of users in the set $\mathcal{A}$. In this way, we consider multiple users who benefit from the attack— either because of attacker-coalitions or because the fraudster shares benefits of the attack with other users in an attempt to remain untraceable.[4]

For simplicity we make the following assumption that simplifies the characterization of attacks.

Assumption 3. *The number of users in the sets $\mathcal{A}$ and $\mathcal{V}$ determine the outcome of the attack, rather than the particular users in each set. Therefore, an attacker obtains the same benefit using the partitions $\{\mathcal{A}_1, \mathcal{V}_1\}$ and $\{\mathcal{A}_2, \mathcal{V}_2\}$ if $|\mathcal{A}_1| = |\mathcal{A}_2|$ and $|\mathcal{V}_1| = |\mathcal{V}_2|$.*

Remark 1. *Assumption 3 implies that all users have similar characteristics, therefore, their individual characteristics do not affect the impact of attacks.*

From Assumption 3 we can characterize an attack with the tuple (λ, γ), where $\lambda \geq 1$ is the severity of the attack and $\gamma \in [0, 1]$ is the proportion of users who profit from the attack. Let us denote the number of attackers as $N_A = \lceil \gamma N \rceil = |\mathcal{A}|$ and the number of victims as $N_V = N - N_A = |\mathcal{V}|$.[5]

Example 2. *Let us consider a homogeneous population with $N = 100$ users, who have valuations of the form*

$$v_i(q_i) = v(q_i) = \alpha \log(1 + x_i),$$

[4]We call attackers to all the users that profit form an attack, although some of them might be unaware of the attack.

[5]Attacks with $\gamma \in \{0, 1\}$ do not have any effect, because either $\mathcal{A}$ or $\mathcal{V}$ is empty, making the solution of Eq. (4.6) equal to $\boldsymbol{\mu}$.

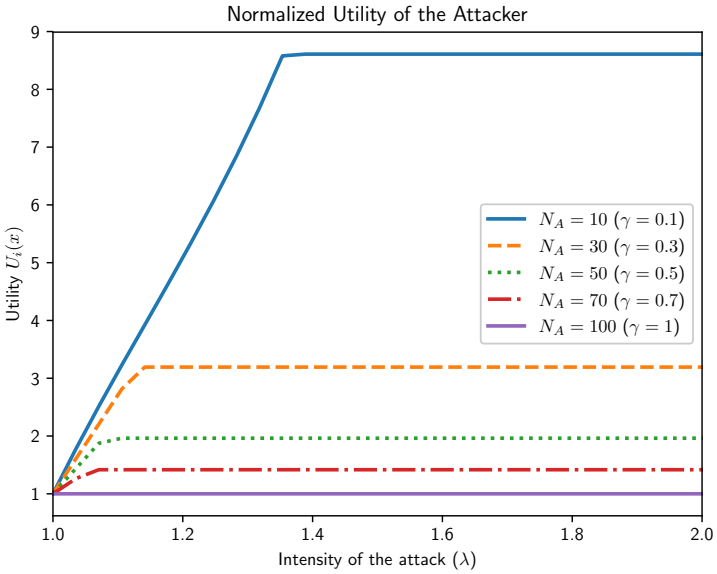

Figure 4.3: Normalized utility of a fraudster as function of the intensity of the attack λ, for different number of attackers. The profit increases with λ, but the number of victims limits the maximum profit achieved increasing λ.

where $\alpha = 4$. Moreover, we define the minimum demand $\underline{Q_i} = 0$ for all users and the parameters of the price function are $\beta = b = 1$. Fig. 4.3 shows the normalized profit of the attacker with different attack parameters (we normalize the profit with respect to the profit without attacks, i.e., $\gamma = 1$). Observe that the attacker's profit increases with the degree of the attack λ; however, it reaches an upper bound for large values of λ. This happens because the victims reduce their demand as the intensity of the attack increases, but eventually they reach their lower demand $\underline{Q_i}$. Moreover, by sharing the benefits with less users (i.e., with lower γ), the attacker can increase his profit, because the attack can produce larger demand reductions.

Attacks on DLC

Attacks on centralized systems have almost no restrictions. For instance, in the recent attacks to the Ukraine's power system (Greenberg, 2017; Zetter, 2016), the perpetrators gained access over the grid's control centers and were able to take off multiple substations. Similarly, in our case the attacker can compromise the DR system that solves the optimization problem in Eq. (4.3) and modify it to solve an alternative function that captures his objectives. Thus, an attack on a DLC system can allow the attackers to implement the solution of the optimization problem in Eq. (4.6).

Attack on DP

In a decentralized system the attacker has more restrictions because users make decisions independently. However, the fraudster can implement his optimal attack modifying the messages sent by the utility to the users, in particular, the incentive function $I(\cdot)$ (see Fig. 4.1). Leveraging the theory of mechanism design we can show that an attacker can incentivize all agents to adopt $\boldsymbol{x}$ by sending the following false incentives:

$$\tilde{I}_j(\boldsymbol{q}) = (\|\boldsymbol{q}_{\mathcal{V}}\| - q_j + \lambda \|\boldsymbol{q}_{\mathcal{A}}\|) \left(\frac{N}{N-1} p\left(\left\|\boldsymbol{q}_{-j}\right\|\right) - p(\|\boldsymbol{q}\|) \right), \tag{4.7}$$

for all $j \in \mathcal{V}$ and

$$\tilde{I}_i(\boldsymbol{q}) = \left(\frac{1}{\lambda} \|\boldsymbol{q}_{\mathcal{V}}\| + \|\boldsymbol{q}_{\mathcal{A}}\| - q_i \right) \left(\frac{N}{N-1} p\left(\left\|\boldsymbol{q}_{-i}\right\|\right) - p(\|\boldsymbol{q}\|) \right), \tag{4.8}$$

for $i \in \mathcal{A}$. The fake incentive functions in Eq. (4.7) and Eq. (4.8) deceive users showing either higher prices (in the case of victims) or lower costs (in the case of attackers).

Properties of the Attack

Let us define the benefit of attackers as

$$\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) = \sum\nolimits_{i \in \mathcal{A}} \{U_i(\boldsymbol{x}) - U_i(\boldsymbol{\mu})\}$$

and the losses of victims as

$$\Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) = \sum\nolimits_{j \in \mathcal{V}} \{U_i(\boldsymbol{\mu}) - U_j(\boldsymbol{x})\}.$$

The following result shows that both attackers and victims have non-negative benefits and losses, respectively; however, the losses exceed the benefits.

Proposition 1. *Let $\boldsymbol{\mu}$ be the optimal social outcome (the solution to Eq. (4.3)) and $\boldsymbol{x}$ the be outcome with an attack (λ, γ) (solution to Eq. (4.6)). If $\lambda > 1$, then $\Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq \Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq 0$.*

Proof. From the optimality of $\boldsymbol{\mu}$ and $\boldsymbol{x}$ we have

$$\lambda \sum\nolimits_{i \in \mathcal{A}} U_i(\boldsymbol{x}) + \sum\nolimits_{j \in \mathcal{V}} U_j(\boldsymbol{x}) \geq \lambda \sum\nolimits_{i \in \mathcal{A}} U_i(\boldsymbol{\mu}) + \sum\nolimits_{j \in \mathcal{V}} U_j(\boldsymbol{\mu}) \tag{4.9}$$

and

$$\sum\nolimits_{i \in \mathcal{A}} U_i(\boldsymbol{x}) + \sum\nolimits_{j \in \mathcal{V}} U_j(\boldsymbol{x}) \leq \sum\nolimits_{i \in \mathcal{A}} U_i(\boldsymbol{\mu}) + \sum\nolimits_{j \in \mathcal{V}} U_j(\boldsymbol{\mu}). \tag{4.10}$$

From Eq. (4.9) we obtain

$$\lambda\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq \Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma)$$

and from Eq. (4.10) we obtain

$$\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \leq \Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma).$$

From the previous expressions we have

$$\lambda\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq \Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma),$$

which requires that $\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq 0$, otherwise the inequality is not satisfied. □

Let us denote by $\boldsymbol{x}$ the demand resulting from an attack (λ, γ), and by the vectors $\boldsymbol{x}_{\mathcal{V}}$ and $\boldsymbol{x}_{\mathcal{A}}$ the consumption of victims and attackers, respectively. Thus, $\|\boldsymbol{x}_{\mathcal{A}}\| = \sum_{i\in\mathcal{A}} x_i$ and $\|\boldsymbol{x}_{\mathcal{V}}\| = \sum_{j\in\mathcal{V}} x_j$.

The Lagrangian associated with the problem in Eq. (4.6) is

$$L(\boldsymbol{q}, \boldsymbol{\nu}) = \lambda \sum\nolimits_{i\in\mathcal{A}} U_i(\boldsymbol{q}) + \sum\nolimits_{j\in\mathcal{V}} U_j(\boldsymbol{q}) + \sum\nolimits_{h\in\mathcal{P}} \nu_h \cdot q_h,$$

for slack variables $\nu_i \in \mathbb{R}$. Thus, the demand profile under an attack, denoted by $\boldsymbol{x}$, must satisfy the following optimality conditions:

$$\dot{v}_i(x_i) - \beta\left(||\boldsymbol{x}|| + \|\boldsymbol{x}_{\mathcal{A}}\| + \frac{1}{\lambda}\|\boldsymbol{x}_{\mathcal{V}}\|\right) - b + \frac{1}{\lambda}\nu_i = 0, \tag{4.11}$$

$$\dot{v}_j(x_j) - \beta\left(||\boldsymbol{x}|| + \|\boldsymbol{x}_{\mathcal{V}}\| + \lambda\|\boldsymbol{x}_{\mathcal{A}}\|\right) - b + \nu_j = 0, \tag{4.12}$$

$$x_h \geq 0, \qquad \nu_h \geq 0, \qquad (x_h - \underline{Q_h})\nu_h = 0, \tag{4.13}$$

for all $i \in \mathcal{A}$, $j \in \mathcal{V}$, and $h \in \mathcal{P}$.

The consequences of the attack are: i) reduction of demand by victims; and ii) increased demand by attackers; These properties are formally proved in the following proposition:

Proposition 2 (Proposition 1 in (Barreto and Cárdenas, 2015a))**.** *Let $\boldsymbol{\mu}$ be the ideal equilibrium (Eq. (4.3)) and $\boldsymbol{x}$ be the equilibrium with an attack (Eq. (4.6)). If there is an attack with $\lambda > 1$, then the consumption of attackers increases and the consumption of victims decreases with respect to the ideal case. That is, $x_i > \mu_i$ and $x_j < \mu_j$, and $\|\boldsymbol{x}\| < \|\boldsymbol{\mu}\|$ for every attacker $i \in \mathcal{A}$ and victim $j \in \mathcal{V}$.*

Proof. We can evaluate the derivative of the attacker's objective function (Eq. (4.6)) in the ideal outcome $\boldsymbol{\mu}$ to obtain

$$\lambda \left(\dot{v}_i(\mu_i) - p(||\boldsymbol{\mu}||) - \beta \left\| \boldsymbol{\mu} \right\| \right) + (\lambda - 1)\beta \left\| \boldsymbol{\mu}_{\mathcal{V}} \right\|,$$

$$\dot{v}_j(\mu_j) - p(||\boldsymbol{\mu}||) - \beta \left\| \boldsymbol{\mu} \right\| + (1 - \lambda)\beta \left\| \boldsymbol{\mu}_{\mathcal{A}} \right\|.$$

Note that the left hand side of the previous equations is precisely the FOC of the original optimization problem (Eq. (4.4)). Thus, the derivative with respect to q_i is

$$\frac{\partial}{\partial q_i} f_a(\boldsymbol{q}) \big|_{\boldsymbol{q}=\boldsymbol{\mu}} = (\lambda - 1)\beta \left\| \boldsymbol{\mu}_{\mathcal{V}} \right\| > 0,$$

and the derivative with respect to q_j is

$$\frac{\partial}{\partial q_j} f_a(\boldsymbol{q}) \big|_{\boldsymbol{q}=\boldsymbol{\mu}} = -(\lambda - 1)\beta \left\| \boldsymbol{\mu}_{\mathcal{A}} \right\| < 0.$$

Hence, we know that $x_i > \mu_i$ and $x_j < \mu_j$. □

From Assumption 2 we have $\mu_i > 0$, which together with Proposition 2 results in $x_i > \mu_i > 0$ for $i \in \mathcal{A}$. Therefore, the slack variables for the attackers (see Eq. (4.13)) are equal to zero, that is, $\nu_i = 0$ for $i \in \mathcal{A}$.

Optimal Attack

The adversary can design the attack (λ, γ) to guarantee some maximum level of benefit, or equivalently, to limit the impact on the system. In particular, he must pay special attention to the proportion of attackers γ, which limits the maximum benefit form the attack (see our previous example). To find the maximum benefit with some γ we assume that $\lambda \to \infty$ to guarantee that the victims adopt their minimum demand $\underline{Q}_j$, for $j \in \mathcal{V}$.

From Assumption 3 we know that the population is homogeneous (see Remark 1), therefore all attackers will have a similar demand, that is, $x_i = x_a$ for $i \in \mathcal{A}$ and $\|\boldsymbol{x}_{\mathcal{A}}\| = N_a x_a \approx \gamma N$. Moreover, we assume that $\|\boldsymbol{x}_{\mathcal{V}}\| = N_v x_v \approx (1 - \gamma) N x_v$. Therefore, from the equilibrium conditions in Eq. (4.11) we have

$$\dot{v}_i(x_a) = 2\beta\gamma N x_a + \beta(1 - \gamma) N x_v + b.$$

An attacker that knows $\dot{v}(\cdot)$ can estimate the demand of attackers as a function of γ, that is, $x_a(\gamma)$, which can give an estimate of the utility with the attack as a function of γ.

If the attacker does not know the valuation function of the other users, then it won't estimate the impact of an attack.

4.3.2 Malicious Attacker

We assume that the malicious attacker wants to damage the system, but cannot attack the system directly, either because the system has protections or because he resides in another country. Therefore, this attacker would use cyber attacks to target the DR infrastructure and damage the system. The malicious attacker can try to create a sudden (unanticipated) peak of demand, which can cause local blackouts (by tripping a distribution fuse or circuit breaker) or kickstarting blackstart generators. Creating a demand peak with DLC is straightforward, since the attacker can directly send all electricity consumption signals to their maximum value at the same time. On the other hand, the most intuitive attack in systems with DP consists in selecting the DR signals to reduce prices (and increase the demand) at the time with highest demand, denoted as t_{attack}.[6] Specifically, the attacker can compromise the incentive signal and send the following malicious incentives:

$$I_i^m(\boldsymbol{q}) = \begin{cases} I_i(\boldsymbol{q}^t) + \sigma_1 ||\boldsymbol{q}||_1 & \text{if } t = t_{attack}, \\ I_i(\boldsymbol{q}^t) & \text{otherwise,} \end{cases}$$

where $\sigma_1 > 0$.

Alternatively, the attacker can increase the impact of the attack by increasing the prices in hours previous to the attack, and then lower them at t_{attack}. In this way users would accumulate tasks (i.e., demand) until the price becomes favorable. The attacker can use the following incentives

$$I_i^m(\boldsymbol{q}) = \begin{cases} I_i(\boldsymbol{q}^t) + \sigma_1 ||\boldsymbol{q}||_1 & \text{if } t = t_{attack}, \\ I_i(\boldsymbol{q}^t) - \sigma_2 ||\boldsymbol{q}||_1 & \text{if } t \in [t_a, t_b], \\ I_i(\boldsymbol{q}^t) & \text{otherwise,} \end{cases}$$

where σ_1, σ_2 are positive real numbers, $[t_a, t_b]$ is the period during which the attack focuses on reducing the demand.

Example 3. *Fig. 4.4 shows the impact of both the naive and the strategic attacks during the initial transition period.*[7] *Simulations are made with* $\sigma_1 = 50$, $\sigma_2 = 100$, $t_a = 0hrs$, $t_b = 17hrs$,

[6]Here we consider that the market has diverse equilibriums during the day, according to the users' preferences. For example, users consume more energy during the evenings, and less in the early morning. We denote with the vector $\boldsymbol{q}^t$ the demand of the population at time t.

[7]In this example we use population games (Mojica-Nava et al., 2015; Sandholm, 2011; Hofbauer and Sigmund., 1998; Barreto, 2014) to model how users adjust their demand according the incoming information (the system's state of the incentive function). We refer the interested reader to (Barreto et al., 2014) for more details on the implementation.

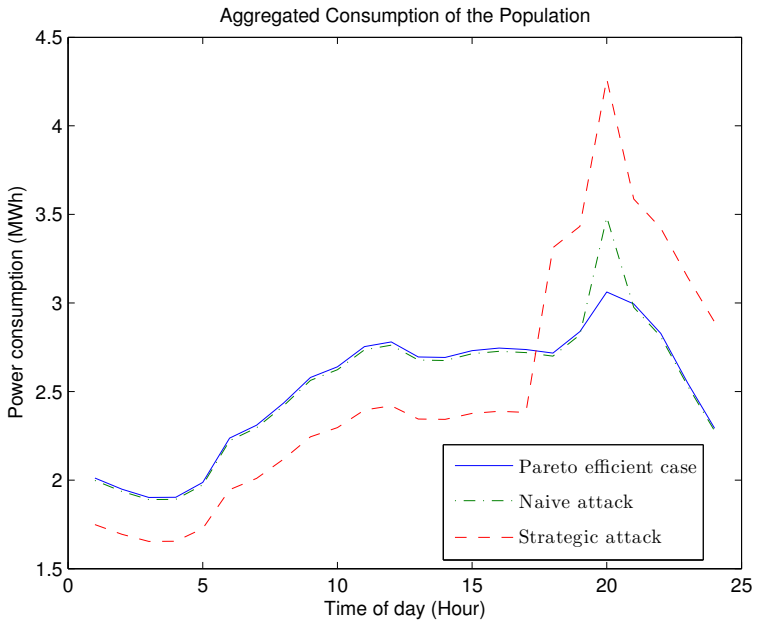

Figure 4.4: Impact of a malicious attack on the population demand for two different attacks 1) attack on a single hour and 2) coordinated attack on various hours of the day.

$t_attack = 20hrs$. *In particular, the attack time coincides with the demand peak in in the Pareto optimal outcome. The naive attack succeeds in causing a increase of the demand at* t_{attack}. *On the other hand, the strategic attack achieves a greater peak by causing demand reduction prior to the attack. Roughly speaking, the strategic attack sets conditions so that the population has more resources to consume at* t_{attack}.

4.3.3 Comparison of Attacks on DLC and DP

A fraudster can implement attacks in both DR systems; however, attacking systems with dynamic prices requires more effort and can generate less profit. On one hand, the incentives scheme charges higher prices to users with higher demand. Hence, the attacker pays higher prices, but he still can get profit from the attack (see Fig. 4.5 for an example of the attacker's utility as a function of λ and γ). Likewise, the impact on the population is less severe with DP, because the victims are partially compensated by the attackers (see Fig. 4.6). On the other hand, the attack in DP requires that each users receives a particular incentive function, therefore, the attacker needs to attack multiple devices.

Moreover, unlike DLC, an attack on DP suffers delays, because the effects take place only once users adjust their demand. Also, with DP users can have more mechanisms to verify if the incoming information is legitimate e.g., by comparing with multiple sources. However,

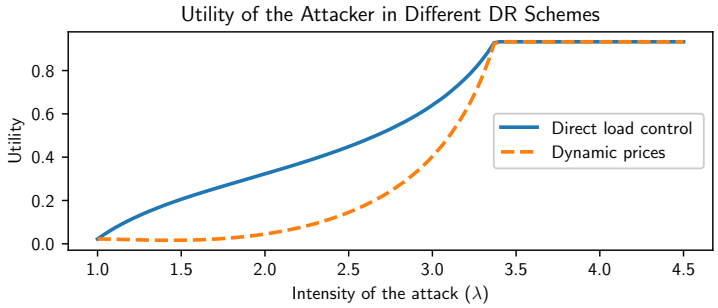

(a) Attacker's profit as a function of the intensity of the attack λ with $N_A = 1$.

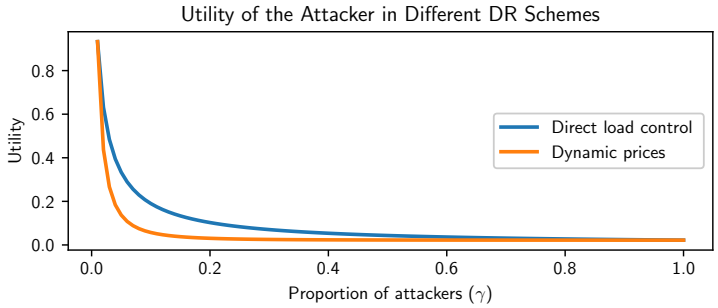

(b) Attacker's profit as a function of the proportion of attackers γ with $\lambda = 3.5$.

Figure 4.5: Utility of the attacker in systems with DLC and DP. Fraudsters obtain more benefits from attacking DLC systems.

with DLC users must carry out the commands sent by the utility, but they can fail to verify whether the commands are legitimate. This occurs because users have limited information to check that the commands maximize the customer surplus.

Attack Requirements

The fraudster can send false information either by compromising the system that broadcasts information or targeting directly the devices that optimize locally the demand. Also, the fraudster can implement the attack if he can discriminate the information sent to the groups of attackers and victims. Although the attacker does not need the valuation functions $v_i(\cdot)$, he needs the total demand of both attackers and victims, $||\boldsymbol{q}_{\mathcal{A}}||$ and $||\boldsymbol{q}_{\mathcal{V}}||$, respectively.[8]

[8]In our assumption the attacker has as much information as the central agent. Hence, although the demand of users is private, the attacker can observe it by compromising the DR management center.

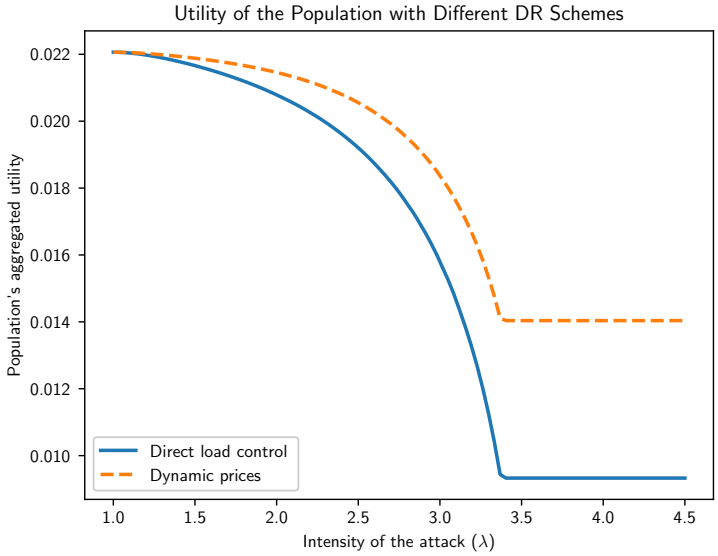

Figure 4.6: Impact of the attack in the customer surplus as a function of the attack severity λ for both the DLC and dynamic pricing schemes with $\gamma = 0.01$.

4.4 Detecting Attacks

In this section we address the problem of determining whether a change in demand occurred due to a fault or an attack (by a fraudster). In particular, we analyze how asymmetric information affects the detection of attacks.

4.4.1 Faults in the System

Here we assume that the utility knows the optimal demand $\boldsymbol{\mu}$, which allows the detection of anomalies. If the utility observes an anomaly $\boldsymbol{q}$, it can use Proposition 2 to classify the users as potential victims (or attackers), depending if they reduce (or increase) their consumption. In other words, deviations from $\boldsymbol{\mu}$ reveal the sets $\mathcal{A}$ and $\mathcal{V}$, which allows to determine the parameter γ of the attack. Accordingly, $\|\boldsymbol{q}_{\mathcal{A}}\| \geq \|\boldsymbol{\mu}_{\mathcal{A}}\|$ and $\|\boldsymbol{q}_{\mathcal{V}}\| \leq \|\boldsymbol{\mu}_{\mathcal{V}}\|$.

Now, let us define the expected demand of users from the set $\mathcal{A}$ in case of attacks or faults. On one hand, consider an attack (λ, γ) with demand $\boldsymbol{\xi}$ that produces the observed demand of users from the set $\mathcal{V}$, that is, $\boldsymbol{\xi}$ satisfies $\|\boldsymbol{\xi}_{\mathcal{V}}\| = \|\boldsymbol{q}_{\mathcal{V}}\|$. In this case, the demand of users from the set $\mathcal{A}$ satisfies Eq. (4.6). On the other hand, we denote with the vector $\boldsymbol{\zeta}$ the demand that honest users would have when a fault reduces the demand of users from the set $\mathcal{V}$.[9] Consequently, the response of honest users is given by Eq. (4.4), when the set $\mathcal{V}$ has

[9]Unlike attacks, faults are unintentional, and therefore, no user tries to take advantage to get more profit.

a fixed demand. The next result shows that, given our previous considerations, attackers consume more energy than honest users.

Lemma 1. *Let $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$ be the demand without and with an attack (λ, γ), respectively. If $\|\boldsymbol{\zeta}_{\mathcal{V}}\| = \|\boldsymbol{\xi}_{\mathcal{V}}\| = Q_v$, then $\|\boldsymbol{\zeta}_{\mathcal{A}}\| \leq \|\boldsymbol{\xi}_{\mathcal{A}}\|$.*

Proof. Let us denote by $\boldsymbol{\zeta}$ and $\boldsymbol{\xi}$ the demand of normal and a fraudster users, respectively. From Eq. (4.4) we know that the demand without an attack satisfies

$$\dot{v}_i(\zeta_i) - 2\beta \|\boldsymbol{\zeta}_{\mathcal{A}}\| - b = 2\beta \|\boldsymbol{\zeta}_{\mathcal{V}}\|, \tag{4.14}$$

and the demand with an attack (λ, γ) satisfies Eq. (4.11)

$$\dot{v}_i(\xi_i) - 2\beta \|\boldsymbol{\xi}_{\mathcal{A}}\| + \left(1 - \frac{1}{\lambda}\right) \beta \|\boldsymbol{\xi}_{\mathcal{V}}\| - b = 2\beta \|\boldsymbol{\xi}_{\mathcal{V}}\|, \tag{4.15}$$

for all $i \in \mathcal{A}$. Now, let us define $\|\boldsymbol{\zeta}_{\mathcal{V}}\| = \|\boldsymbol{\xi}_{\mathcal{V}}\| = Q_v$ to observe the reaction of normal users and attackers to a given demand of victims Q_v. Let us assume by contradiction that $\|\boldsymbol{\zeta}_{\mathcal{A}}\| > \|\boldsymbol{\xi}_{\mathcal{A}}\|$, hence, there exists some $i \in \mathcal{A}$ such that $\zeta_i > \xi_i$. From Assumption 1 we know that $\dot{v}_i(\zeta_i) < \dot{v}_i\xi_i$. Therefore, Eq. (4.14) and Eq. (4.15) lead to

$$2\beta \left(\|\boldsymbol{\xi}_{\mathcal{A}}\| - \|\boldsymbol{\zeta}_{\mathcal{A}}\|\right) > \left(1 - \frac{1}{\lambda}\right) \beta \|\boldsymbol{\xi}_{\mathcal{V}}\|.$$

However, the previous expression implies that $0 > \|\boldsymbol{\xi}_{\mathcal{V}}\|$, which is not possible from the demand constraints. This contradiction leads to $\|\boldsymbol{\zeta}_{\mathcal{A}}\| \leq \|\boldsymbol{\xi}_{\mathcal{A}}\|$. □

4.4.2 Case with Full Information (DLC)

From Eq. (4.11) we can extract the following relationship:

$$\lambda = \frac{\beta||\boldsymbol{x}_{\mathcal{V}}||}{\dot{v}_i(x_i) - 2\beta||\boldsymbol{x}_{\mathcal{A}}|| - \beta||\boldsymbol{x}_{\mathcal{V}}|| - b}, \tag{4.16}$$

If the utility company knows the valuation function of users, then it can use the previous equation to determine the value of λ. Note that $\lambda = 1$ indicates normal behavior, while $\lambda > 1$ suggests an attack. For instance, if we replace the normal demand evaluated in $\boldsymbol{\zeta}$ (see Eq. (4.14)) into Eq. (4.16) we obtain $\lambda = \beta Q_v / \beta Q_v = 1$, which indicates normal behavior. Observe that if $||\boldsymbol{x}_{\mathcal{V}}|| = 0$, then the estimation of λ is equal to zero.

4.4.3 Case with Asymmetric Information (DP)

Naive Detection

It might be reasonable to raise alarms when the total demand falls below some threshold ϵ, determined using historic consumption data. However, an attacker can bypass the detection mechanism choosing attacks that satisfy $\|\boldsymbol{x}\| \leq (1-\epsilon)\|\boldsymbol{\mu}\|$. Also, this detection mechanism does not distinguish between faults and attacks because any demand beyond the threshold raises alarms.

We can design better detection mechanisms by considering the characteristics of the attacks.

Improved Detection

Since the utility ignores the valuation function of users, it cannot determine if the observed demand corresponds to a normal behavior (as in Section 4.4.2). Instead, with a sample of the demand without attacks, in this case $\boldsymbol{\mu}$, the utility can try to estimate whether the observed demand matches the expected consequences of a given attack (λ, γ). The following result provides a relation between the demand of attackers and victims

Proposition 3 (Proposition 3 in (Barreto and Cárdenas, 2015a))**.** *Let $\boldsymbol{\mu}$ be a solution to Eq.* (4.3) *and $\boldsymbol{x}$ the demand with an attack (λ, γ) that solves Eq.* (4.6)*. Then $\|\boldsymbol{x}_{\mathcal{A}}\| \geq \Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\|, \lambda, \gamma)$ and $\|\boldsymbol{x}_{\mathcal{V}}\| \leq \Lambda(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{A}}\|, \lambda, \gamma)$, where*

$$\Omega(\boldsymbol{\mu}, ||\boldsymbol{x}_{\mathcal{V}}||, \lambda, \gamma) = \frac{2}{(1+\lambda)} \left(||\boldsymbol{\mu}|| - ||\boldsymbol{x}_{\mathcal{V}}||\right),$$

and

$$\Lambda(\boldsymbol{\mu}, ||\boldsymbol{x}_{\mathcal{A}}||, \lambda, \gamma) = \frac{2\lambda}{(1+\lambda)} \left(||\boldsymbol{\mu}|| - ||\boldsymbol{x}_{\mathcal{A}}||\right).$$

Proof. From Eq. (4.11) and Eq. (4.12) we get

$$||\boldsymbol{x}_{\mathcal{A}}|| = \frac{1}{\beta(1+\lambda)} \left(\dot{v}_j(x_j) - 2\beta||\boldsymbol{x}_{\mathcal{V}}|| - b\right), \tag{4.17}$$

and

$$||\boldsymbol{x}_{\mathcal{V}}|| = \frac{\lambda}{\beta(1+\lambda)} \left(\dot{v}_i(x_i) - 2\beta||\boldsymbol{x}_{\mathcal{A}}|| - b + \nu_j\right), \tag{4.18}$$

where $\nu_j \geq 0$. The valuation of each user $v_i(\cdot)$ is a concave function, thus the marginal valuation $\dot{v}_i(\cdot)$ is decreasing and non-negative, therefore, from Proposition 2 we get

$$\dot{v}_i(x_i) \leq \dot{v}_i(\mu_i), \qquad \dot{v}_j(x_j) \geq \dot{v}_j(\mu_j).$$

The previous equations can be used along Eq. (4.4) to extract the following inequalities:

$$\dot{v}_i(x_i) \leq \dot{v}_i(\mu_i) \leq 2\beta||\boldsymbol{\mu}|| + b, \tag{4.19}$$

$$\dot{v}_j(x_j) \geq \dot{v}_j(\mu_j) \geq 2\beta||\boldsymbol{\mu}|| + b, \tag{4.20}$$

Here we can replace Eq. (4.19) and Eq. (4.20) in Eq. (4.17) and Eq. (4.18), respectively, resulting

$$||\boldsymbol{x}_{\mathcal{A}}|| \geq \frac{2}{\beta(1+\lambda)} \left(||\boldsymbol{\mu}|| - ||\boldsymbol{x}_{\mathcal{V}}||\right) = \Omega(\boldsymbol{\mu}, ||\boldsymbol{x}_{\mathcal{V}}||, \lambda, \gamma)$$

and

$$||\boldsymbol{x}_{\mathcal{V}}|| \leq \frac{2\lambda}{(1+\lambda)} \left(||\boldsymbol{\mu}|| - ||\boldsymbol{x}_{\mathcal{A}}||\right) = \Lambda(\boldsymbol{\mu}, ||\boldsymbol{x}_{\mathcal{A}}||, \lambda, \gamma).$$

□

With the previous boundaries we can estimate the demand of attackers $\|\boldsymbol{x}_{\mathcal{A}}\|$ using the demand of victims $\|\boldsymbol{x}_{\mathcal{V}}\|$. Hence, for some demand $\boldsymbol{q}$, if $\|\boldsymbol{q}_{\mathcal{A}}\|$ exceeds $\Omega(\|\boldsymbol{\mu}\|, \|\boldsymbol{q}_{\mathcal{V}}\|, \lambda, \gamma)$, then we can conclude that the system is under attack.

Now, let us investigate if this detection scheme allows us to differentiate attacks from faults. On one hand, we can distinguish faults from attacks if the response to faults $\|\boldsymbol{\zeta}_{\mathcal{V}}\|$ lies below the estimated demand, that is, if

$$\|\boldsymbol{x}_{\mathcal{A}}\| \geq \Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\|, \lambda, \gamma) \geq \|\boldsymbol{\zeta}_{\mathcal{A}}\|,$$

However, we cannot distinguish faults if the honest response of users exceeds or equals the the estimated demand of attackers, that is, if

$$\|\boldsymbol{x}_{\mathcal{A}}\| \geq \|\boldsymbol{\zeta}_{\mathcal{A}}\| \geq \Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\|, \lambda, \gamma).$$

This previous situation occurs when $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$, as proved in the following result.

Lemma 2. *If $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$ we cannot differentiate attacks from faults, because the demand of users in the set $\mathcal{A}$ is the same in both cases, that is, $\|\boldsymbol{x}_{\mathcal{A}}\| = \|\boldsymbol{\zeta}_{\mathcal{A}}\|$.*

Proof. First, if $\|\boldsymbol{x}_{\mathcal{V}}\| = \|\boldsymbol{\zeta}_{\mathcal{V}}\|$, then we can use the same procedure as in Proposition 2 to show that

$$x_i \geq \zeta_i \geq \mu_i, \tag{4.21}$$

for $i \in \mathcal{P}$. Moreover, Proposition 3 leads to

$$\|\boldsymbol{x}_{\mathcal{A}}\| \geq \Omega(\boldsymbol{\mu}, \boldsymbol{\zeta}, \|\boldsymbol{x}_{\mathcal{V}}\|, \lambda, \gamma),$$

which, using the fact that $\|\boldsymbol{x}_{\mathcal{V}}\| = \|\boldsymbol{\zeta}_{\mathcal{V}}\|$, can be rewritten as

$$\|\boldsymbol{\zeta}_{\mathcal{A}}\| \geq \|\boldsymbol{x}_{\mathcal{A}}\| - \frac{\lambda - 1}{2\lambda} \|\boldsymbol{x}_{\mathcal{V}}\| . \tag{4.22}$$

From Eq. (4.21) and Eq. (4.22) we obtain

$$\|\boldsymbol{x}_{\mathcal{A}}\| \geq \|\boldsymbol{\zeta}_{\mathcal{A}}\| \geq \|\boldsymbol{x}_{\mathcal{A}}\| - \frac{\lambda - 1}{2\lambda} \|\boldsymbol{x}_{\mathcal{V}}\| .$$

However, if $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$, then we conclude that $\|\boldsymbol{x}_{\mathcal{A}}\| = \|\boldsymbol{\zeta}_{\mathcal{A}}\|$, that is, we cannot distinguish the demand behavior of attackers and honest users. □

Moreover, the proposed detection scheme has some drawbacks. On one hand, the utility ignores the attack intensity λ, hence, it can fail to estimate the attackers' minimum demand. Observe that the demand estimation decreases with the intensity λ, that is,

$$\Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\| , \lambda - \epsilon, \gamma) \geq \Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\| , \lambda, \gamma),$$

for some $\epsilon \geq 0$. Therefore, if the utility uses $\tilde{\lambda} < \lambda$ to estimate the boundary, it can fail to detect attacks (e.g., have false negatives). On the other hand, if $\tilde{\lambda} > \lambda$, then it can classify faults as attacks (e.g., have false positives). The following example illustrates the ideas of the proposed detection scheme.

Example 4. *Fig. 4.7 shows the demand of both attackers and victims with attacks with constant intensity $\lambda = 1.7$ and γ taking values in the interval $[0.01, 1]$. Observe that the demand of victims decreases as γ increases, reaching the minimum demand $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$ for $\gamma \geq 0.05$. In other words, a set $\mathcal{A}$ with five or more attackers forces the remainder of the population to consume their minimum demand. Fig. 4.7 also shows the aggregate demand that users from the set $\mathcal{A}$ have when a fault, rather than a attack, reduces the demand of victims to $\|\boldsymbol{x}_{\mathcal{V}}\|$. Here we confirm the result from Lemma 1, which show that the demand resulting from a fault $\|\boldsymbol{\zeta}_{\mathcal{A}}\|$ is strictly lower than the demand with an attack $\|\boldsymbol{x}_{\mathcal{A}}\|$, except when $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$.*

Fig. 4.7 also shows the estimated demand of attackers with an attack $(\tilde{\lambda}, \gamma)$, defined by the function $\Omega(\boldsymbol{\mu}, \|\boldsymbol{x}_{\mathcal{V}}\| , \tilde{\lambda}, \gamma)$, where the estimated impact of the attack is $\tilde{\lambda} = 2$. Observe that the our estimation fails to distinguish attacks from faults when $\|\boldsymbol{x}_{\mathcal{V}}\|$ is equal or close to zero (see Lemma 2). Also, since $\Omega(\boldsymbol{\mu}, ||\boldsymbol{x}_{\mathcal{V}}||, \tilde{\lambda}, \gamma)$ is linear with respect to $||\boldsymbol{x}_{\mathcal{V}}||$, estimations with $\tilde{\lambda} > \lambda$ will move the boundary downwards, increasing the number of faults classified as attacks (false positives). On the other hand, estimating the boundary with $\tilde{\lambda} < \lambda$ will move it upwards, which can result in attacks classified as faults (false negative).

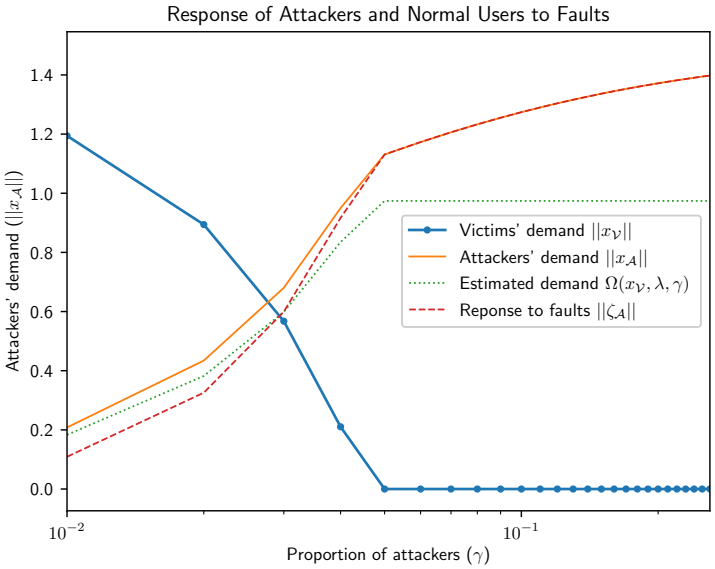

Figure 4.7: Total demand of attackers $\|\boldsymbol{x}_{\mathcal{A}}\|$ and honest users $\|\boldsymbol{\mu}_{\mathcal{A}}\|$ when $\|\boldsymbol{x}_{\mathcal{V}}\| = \|\boldsymbol{\zeta}_{\mathcal{V}}\|$ for $\lambda = 1.7$ and different values of γ. The demand of attackers is higher than the demand of honest users, except when $\|\boldsymbol{x}_{\mathcal{V}}\| = 0$.

4.5 Design of Penalties

In this section we design of penalties to make attacks unprofitable. Even though some attacks make their author untraceable, the utility company can impose penalties on all users that benefit from the attack. Although this strategy penalizes agents who involuntarily profited from the attack, this strategy can prevent rational agents from launching attacks. Below we analyze how to design penalties $\Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma)$ for each user $i \in \mathcal{P}$, having into account the asymmetric information of the utility.

4.5.1 Penalties with Full Information (DLC)

Intuitively, the penalties should be equal to the losses caused by the attack, i.e., the attackers should held accountable for the losses of the population (this is similar to the Clark pivot mechanism (Nisan et al., 2007)). Therefore, we design the penalties to satisfy

$$\sum_{i \in \mathcal{A}} \Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) = \Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq \Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma), \tag{4.23}$$

where the inequality follows from Proposition 1. In particular, we can penalize each attacker with an amount proportional to the benefit that they received with the attack, therefore, we

can select the penalties as

$$\Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) = \Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \frac{U_i(\boldsymbol{x}) - U_i(\boldsymbol{\mu})}{\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma)},$$

which satisfy Eq. (4.23). With the previous incentives we guarantee a negative profit for the profit attackers, thus

$$U_i(\boldsymbol{x}) - U_i(\boldsymbol{\mu}) - \Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \leq 0.$$

This scheme is desirable to the central planner, because it saves the expenses for repairing the damage caused to victims. Furthermore, the penalties exceed the profit earned by the attackers, making attacks unprofitable. Also, an utility with full information of the consumer preferences can compute the penalty functions $\Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma)$.

4.5.2 Penalties with Asymmetric Information (DP)

With asymmetric information we cannot calculate the losses caused by attacks; however, we can estimate the losses using the concavity of the utility functions. In particular, we can estimate an upper bound of the losses suffered by the j^{th} victim as

$$U_j(\boldsymbol{\mu}) - U_j(\boldsymbol{x}) \leq \nabla U_j(\boldsymbol{x})(\boldsymbol{\mu} - \boldsymbol{x}) = \sum\nolimits_{h \in \mathcal{P}} \frac{\partial}{\partial q_h} U_j(\boldsymbol{q}) \Big|_{\boldsymbol{q}=\boldsymbol{x}} (\mu_h - x_h),$$

where the marginal utility with respect to q_i is equal to

$$\frac{\partial}{\partial q_j} U_j(\boldsymbol{q}) = \dot{v}_j(q_j) - p(||\boldsymbol{q}||) - \beta q_j,$$
$$\frac{\partial}{\partial q_i} U_j(\boldsymbol{q}) = -\beta q_j.$$

However, the marginal valuation $\dot{v}_j(x_j)$ is unknown and we cannot estimate an upper bound, but the optimality condition in Eq. (4.12) provides a lower bound. Therefore, in this case the penalties must underestimate the losses of users, but we can guarantee that this estimation exceeds the benefit of attackers.

Consider the following lower bound on the losses victims

$$U_j(\boldsymbol{\mu}) - U_j(\boldsymbol{x}) \geq \nabla U_j(\boldsymbol{\mu})(\boldsymbol{\mu} - \boldsymbol{x}).$$

Summing over all $j \in \mathcal{V}$ we obtain the following boundary for the total losses

$$\Xi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \geq \sum_{j \in \mathcal{V}} \nabla U_j(\boldsymbol{\mu})(\boldsymbol{\mu} - \boldsymbol{x}) = \beta \left\|\boldsymbol{x}\right\| \left\|\boldsymbol{\mu}_{\mathcal{V}}\right\| - \beta \left\|\boldsymbol{\mu}\right\| \left\|\boldsymbol{x}_{\mathcal{V}}\right\|. \tag{4.24}$$

On the other hand, the attackers' profit has the following upper bound

$$U_i(\boldsymbol{x}) - U_i(\boldsymbol{\mu}) \leq \nabla U_i(\boldsymbol{\mu})(\boldsymbol{x} - \boldsymbol{\mu}) \tag{4.25}$$

Hence, the total benefit of attackers has the following upper bound

$$\Psi(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \leq \sum_{i \in \mathcal{A}} \nabla U_i(\boldsymbol{\mu})(\boldsymbol{x} - \boldsymbol{\mu}) = \beta \left\|\boldsymbol{\mu}\right\| \left\|\boldsymbol{x}_{\mathcal{A}}\right\| - \beta \left\|\boldsymbol{x}\right\| \left\|\boldsymbol{\mu}_{\mathcal{A}}\right\|. \tag{4.26}$$

Observe that the estimated losses in Eq. (4.24) are equal to the estimated benefits of attackers in Eq. (4.26). Therefore we design penalties of the form

$$\Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) = \nabla U_i(\boldsymbol{\mu})(\boldsymbol{x} - \boldsymbol{\mu}), \tag{4.27}$$

which make attacks unprofitable for the attackers (once the attack is detected). From Eq. (4.25) the profit of an attacker that is penalized with Eq. (4.27) is negative

$$U_i(\boldsymbol{x}) - U_i(\boldsymbol{\mu}) - \Phi_i(\boldsymbol{\mu}, \boldsymbol{x}, \gamma) \leq 0.$$

Fig. 4.8 shows an example of the effect of penalties on the profit of attackers. In this case the penalties are insufficient to compensate the damage caused to victims.

The penalties can fail when unintentional faults cause deviations from the expected behavior. Recall from Section 4.4.2 that faults can change the normal electricity demand from $\|\boldsymbol{\mu}\|$ to $\|\boldsymbol{\zeta}\|$, where $\|\boldsymbol{\mu}_{\mathcal{V}}\| > \|\boldsymbol{\zeta}_{\mathcal{V}}\|$. Since the reaction to any user is to increase his demand (even if it is honest), we have $\|\boldsymbol{\mu}_{\mathcal{A}}\| < \|\boldsymbol{\zeta}_{\mathcal{A}}\|$. In this case, the users that belong to $\mathcal{A}$ can receive penalties, even if $\|\boldsymbol{\zeta}_{\mathcal{A}}\|$ satisfies the normal behavior stated in Eq. (4.4).

4.6 Conclusions

In this work we introduced an attack model for DR programs that considers strategic adversaries who pursue different goals: either profit or damage the system. In particular, the fraudster attacker has specific goals, which balance both profit and anonymity (see Eq. (4.6)). On the other hand, the malicious attacker attempts to cause peaks in demand. We showed that, although the attackers have different objectives, they can use the same techniques to achieve their goals (in both centralized and decentralized systems).

We found that the structure of decentralized systems negatively affect both the attacker and the defender. On one hand, the fraudster can get less profit in decentralized systems (see Fig. 4.5). On the other hand, although the population has less losses in decentralized

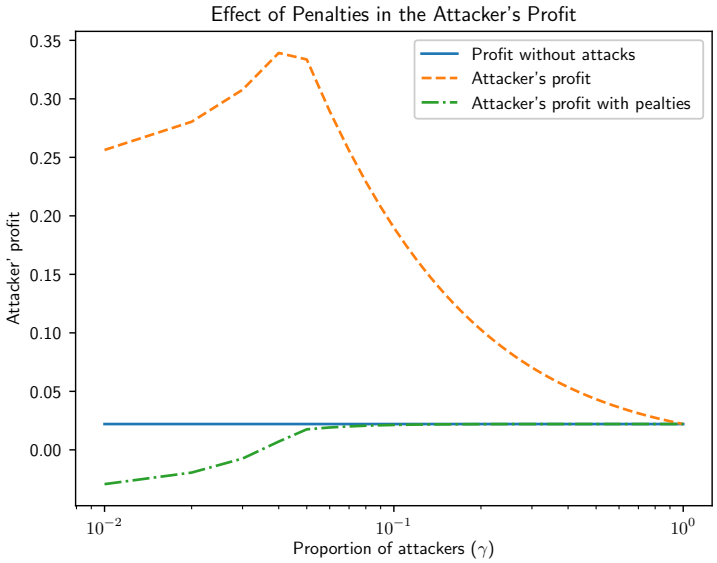

Figure 4.8: The design of penalties on the attacker's profit make it unprofitable to launch attacks, even with asymmetric information.

systems (see Fig. 4.6), the defender has more difficulties detecting and distinguishing attacks from accidental failures. Furthermore, the penalties on attackers cannot fully compensate the victims of the attacks.

Also, decentralized systems can allow users to verify the legitimacy of the incoming information. For instance, since all users receive the same signals, users can check the legitimacy of the received messages with multiple sources. In contrast, in centralized systems users receive private signals (which depend on their private preferences) and they have limited information to check whether the commands truly maximize the social goal.

4.6.1 Future Directions

It would be interesting to explore other protection schemes, such as *security by deception*. For example, the defender would detect attacks and diminish their impact installing honeypots and intrusion detection systems (IDS).

Another interesting direction consists in investigating how additional information about the user's demand (e.g., response of honest users to faults or the demand with real attacks) can improve the detection scheme.

CHAPTER 5

OPTIMAL SECURITY INVESTMENTS IN A PREVENTION AND DETECTION GAME

Authors – Carlos Barreto, Alvaro A. Cardenas, and Alain Bensoussan

The Computer Science Department, EC 31

The University of Texas at Dallas

800 West Campbell Road

Richardson, Texas 75080-3021

Corresponding author: Carlos Barreto.

The content of this chapter and Appendix A are reprinted with permission from: C. Barreto, A. A Cárdenas, and A. Bensoussan, "Optimal security investments in a prevention and detection game," In Proceedings of the Hot Topics in Science of Security: Symposium and Bootcamp, ©2017 ACM, Inc. https://doi.org/10.1145/3055305.3055314

5.1 Introduction

In the last decade cyber-attacks have become prevalent and the motivations behind the attacks more diversified. Attackers may be driven by monetary rewards (Cárdenas et al., 2009), by activism (Olson, 2013), or to cause sabotage (Lee et al., 2016; Barreto et al., 2014). To defend against these attacks, firms need to invest in cyber-security technologies; however, they have limited security resources and should solve the nontrivial problem of allocating them. For instance, if firms use all their resources in prevention technologies (e.g., best practices in access control, security policy enforcement, sandoboxing, frequent software updates, etc.), then they can fail to discover security breaches in their systems. Part of their efforts should include security monitoring, attack detection, identification, and mitigation (Li et al., 2011; Bejtlich, 2013). Attackers also face a similar resource allocation problem, between looking for new vulnerabilities, and the effort required in exploiting them (Axelrod and Iliev, 2014).

Incentives and the economics of information security is an area that has been explored significantly for over a decade (Anderson, 2001). In addition to studying security investments, the literature in security economics has also studied how to prevent attacks decreasing the adversary's incentives (Manshaei et al., 2013). For example, in our recent work on physical attacks against the Colombian power infrastructure, we found that an electric power transmission company redesigned electric tower repair contracts to minimize the incentives that service companies had to launch attacks against the power grid (Barreto and Cárdenas, 2016).

In this chapter we discuss how to balance investments in prevention and detection technologies, as a response to attackers who also have to manage investments to find vulnerabilities and exploiting them. Since attribution is a hard problem in cyber-attacks, we assume that the defender cannot penalize directly the attackers. Moreover, we analyze how limited resources and asymmetric information affect the investments of the defender.

We model the security of a firm's infrastructure using a Markov decision process (i.e., a *stochastic game* (Shapley, 1953; Puterman, 2014; Hernández-Lerma and Lasserre, 2012; Bensoussan, 2011)). With this model we formulate the interaction between an attacker and the firm (defender) as a repeated game in which the firm and the attacker seek to protect and compromise a system, respectively. We find that the defender cannot prevent the attacker from launching attacks once a vulnerability is discovered by the attacker. However, investing in prevention deters the interest of the attacker in the system, because it becomes more costly for the attacker to search vulnerabilities. On the other hand, we observe that the defender

prioritizes prevention over detection in cases with asymmetric information. Nonetheless, asymmetries in information do not have a significant impact in cases with low resources.

Our work is closely related to (Shiva et al., 2010; Alpcan and Basar, 2006), which model the interaction between an attacker and a defender with stochastic games. In these cases the defender has imperfect information about the state (i.e., security) of the system, due to imperfect sensors. This work departs from these works in the following aspects. First, we consider a non-zero sum game with nonlinear cost functions. Second, we give explicit solutions to the payoff functions of each player and consider scenarios with limited resources. Third, in our analysis of asymmetric information we adopt a pessimistic posture assuming that the defender cannot measure the state of the system.

The chapter is structured as follows: Section 5.2 introduces the stochastic game that models the interaction between an attacker and a defender. Sections 5.3 and 5.4 show the optimal strategies of both attacker and defender. In Section 5.5 we show some examples of the defender's strategy with information asymmetries and limited resources. We finalize in Section 5.6 with a discussion of the insights that can help firms deciding how to allocate their resources.

5.2 Attacker-Defender Model

We assume that the defender of a critical infrastructure (a nation state or a firm) can invest resources in *prevention* and *detection* strategies. These strategies have different purposes in the system: on one hand, prevention involves actions to minimize the vulnerability of systems; for instance, frequent software updates or secure software development can reduce the attack surface of the system. On the other hand, detection deals with identifying ongoing attacks (e.g., through a honeypot or an intrusion detection system) and responding to attacks. Thus, the defender has to choose the level of prevention and detection, denoted v_p and v_d respectively, where v_p, $v_d \in [0, 1]$.

We assume that the adversary attacks the system to get profit. In this case the attacker invests its resources in two activities: looking for vulnerabilities and exploiting them. The adversary must find vulnerabilities (either by itself or by purchasing them) before launching attacks, which have multiple intensity levels. The intensity of the attack determines both the profit and the probability that the defender discovers the attack. Intuitively, attacks of low intensity (or stealthy attacks) have a low probability of detection (Axelrod and Iliev, 2014). For instance, a malware tailored specifically to attack a firm has less chances of detected by the security community. However, if a malicious group uses the same vulnerabilities to attack

multiple firms (and get more profits from the attacks), then they also increase the risk of detection and can be associated as the same group behind a variety of attacks. Therefore attackers must act cautiously, because upon detection the defender can identify and correct the vulnerabilities of the system, which reduces the effectiveness of future attacks. Also, their actions can reveal their identity and goals. In summary, we assume that the attacker chooses the effort to look for vulnerabilities (or hack) the system and the intensity of the attacks, denoted by v_h and v_a, respectively, with $v_h,\, v_a \in [0, 1]$.[1]

In this case, we assume that the system (critical infrastructure) has two states, namely a compromised state S_0 and a state free of attackers S_1, where $S = \{S_0, S_1\}$ is the set of states. The actions of both players determine the transition between states, as illustrated in Fig. 5.1. Here, $\pi(v_a, v_d)$ denotes the probability of detecting (and stopping) an attack made with intensity v_a, when the detection scheme has degree v_d. On the other hand, $\delta(v_h, v_p)$ denotes the probability of discovering a vulnerability when the attacker makes an effort v_h to hack the system and the defender prevents vulnerabilities with effort v_p. The transition probabilities are defined as

$$\pi(v_a, v_d) = \mathbb{P}(x_{k+1} = S_1 | x_k = S_0, v_a, v_d),$$

and

$$\delta(v_h, v_p) = \mathbb{P}(x_{k+1} = S_0 | x_k = S_1, v_h, v_p),$$

where x_k is the state of the system at time k.

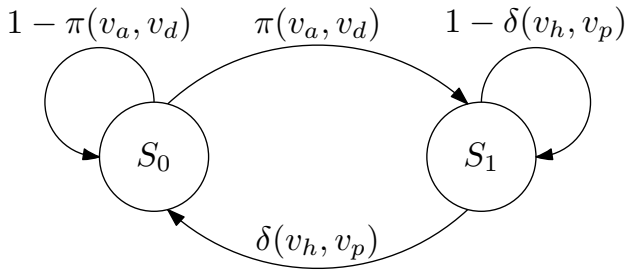

Figure 5.1: Markov process that describes the system's state transitions between a compromised state S_0 and a secure state S_1.

We consider typical decreasing marginal returns in the transition probabilities action. In particular, we assume that $\pi(v_a, v_d)$ is a continuous, increasing, and concave function with respect to v_a and v_d. That is, the probability of detection increases with the degree of the

[1]The effort of the agents determine their investments (or expenditures) to protect or attack the system.

attack and with the quality of the detector. We assume that the defender cannot detect attacks when it makes no effort to detect them ($v_d = 0$) or when the adversary does not execute attacks, that is, $\pi(0, v_d) = \pi(v_a, 0) = 0$. Also, we assume that $\pi(1, 1) = 1$.

On the other hand, $\delta(v_h, v_p)$ is continuous, increasing, and concave with respect to v_h and continuous, decreasing, and convex with respect to v_p. Thus, the probability of discovering a vulnerability increases with the effort of the attacker and decreases with the investment in prevention. We assume that the attacker fails to find vulnerabilities when the system has maximal prevention ($v_p = 1$) or with minimum effort of the attacker ($v_h = 0$), that is, $\delta(v_h, 1) = \delta(0, v_p) = 0$. Furthermore, we assume that $\delta(1, 0) = 1$.

5.2.1 Attacker

At state S_0 (vulnerable state), the attacker knows a vulnerability that allows attacks on the system. The benefit of the attack depends on its intensity v_a, and is represented by a continuous, increasing, and convex function $g_a : [0, 1] \to \mathbb{R}$, which satisfies $g_a(0) = 0$. Exploiting the vulnerability costs C_0 during each time period that the attack lasts i.e., C_0 is the attack operational overhead or the base cost of maintaining the infrastructure to carry out the attack. We assume that the defender patches the vulnerabilities (making them useless in the future) once he discovers the attack, hence, the system jumps to the secure state S_1. In the state S_1 (secure state) the attacker has to invest an amount C_v to discover another useful vulnerability. Here we assume that the attacker finds new vulnerabilities with probability $\delta(v_h, v_p)$. In summary, the attacker's cost with an strategy $v_A = (v_a, v_h)$ at state $x \in S$ is[2]

$$l_A(x, v_A) = \begin{cases} -g_a(v_a) + C_0 \mathbb{1}_{v_a > 0} & \text{if } x = S_0, \\ C_v \mathbb{1}_{v_h > 0} & \text{if } x = S_1. \end{cases}$$

5.2.2 Defender

The defender can implement a protection strategy $v_D = (v_d, v_p)$ that prevents and detects attacks with cost $C_p(v_p)$ and $C_d(v_d)$, respectively. The security scheme of the defender affects only its own cost and the transition probabilities of the system. The cost function of the defender is defined as

$$l_D(x, v_A, v_D) = \begin{cases} g_d(v_a) + C_p(v_p) + C_d(v_d) & \text{if } x = S_0, \\ C_p(v_p) + C_d(v_d) & \text{if } x = S_1, \end{cases}$$

[2]Here we assume that the attacker cannot be penalized. Thus, the cost of an attack only depends on the implementation of the attack. Furthermore, we assume that the attacker cannot discover new vulnerabilities in state S_0.

were $g_d(v_a)$ is the loss caused by an attack of intensity v_a. We assume that the cost functions g_d, C_d, and C_p are convex increasing.

Remark 2. *In this model we assume that the defender stops attacks once he detects them; however, the defender might be unable to fix all the vulnerabilities that allowed attacks. In particular, IoT devices have a short lifespan and fabricators do not have incentives to offer maintenance (e.g., invest in security patches or repairs) due to high costs (Leverett et al., 2017). Here we assume that the defender would stop attacks by restoring the system's components to their initial configuration, to undo adverse alterations. On the other hand, fixing known vulnerabilities becomes part of the efforts in prevention (v_p), and $C_p(\cdot)$ reflects such costs. Therefore, low investment in prevention allows the attacker to reuse vulnerabilities in its attacks.*

5.3 Optimal Attack Policy

Let us assume that the attacker chooses an optimal attack policy V_A that minimizes the cost of the attack, given some protection strategy $v_D = (v_d, v_p)$. We consider an infinite horizon decision problem in which the attacker wants to find the policy $V_A = \{v_{A,n}\}$ (sequence of actions where $v_{A,n} = (v_{a,n}, v_{h,n})$) that minimizes the cost functional

$$J^A(x_0, V_A, v_D) = l_A(x_0, v_{A,0}) + \beta \mathbb{E}_{x_0}^{V_A, v_D}\{l_A(x_1, v_{A,1}) + \dots + \beta \mathbb{E}_{x_{n-1}}^{V_A, v_D}\{l_A(x_n, v_{A,n}) + \dots\}\},$$

where x_0 is the initial state, $\beta \in [0, 1)$ is a discount factor and $x_k \in S$ and $v_{A,k}$ are the state and the attack strategy at time $k = 0, 1, \dots$, respectively. We can rewrite the cost functional as

$$J^A(x_0, V_A, v_D) = l_A(x_0, v_{A,0}) + \beta \mathbb{E}_{x_0}^{V_A, v_D}\{J^A(x_1, V_A, v_D)\}.$$

The discounted cost is bounded because the gain of the attacker is bounded by

$$\min\{0, C_0 - g_a(1)\} \leq l_A(x, v_A) \leq \max\{C_v, C_0\}.$$

Furthermore, since we have an infinite horizon, the control policy is stationary, that is, $v_{A,n} = v_{A,n+1} = v_A$ for every $n \geq 0$. Hence, we can define the attack policy as $V_A = v_A$, where $v_A = (v_a, v_h)$. Moreover, using the optimality principle (Hernández-Lerma and Lasserre, 2012; Bensoussan, 2011) we can prove that the minimum cost is equal to

$$u^A(x_0) = \inf_{v_A} J^A(x_0, v_A, v_D) = \inf_{v_A} \left\{ l_A(x_0, v_A) + \beta \mathbb{E}_{x_0}^{v_A, v_D}\{u^A(x_1)\} \right\}. \tag{5.1}$$

From the state transition in Fig. 5.1 we define the expected value of a function $\varphi(\cdot)$ (when the current state is x_0 and the strategies are v_A and v_D) as

$$\mathbb{E}_{x_0}^{v_A,v_D}\{\varphi(x_1)\} = \mathbb{E}\{\varphi(x_1)|x_0,v_A,v_D\} = \{\varphi(S_0) + (\varphi(S_1) - \varphi(S_0))\pi(v_a,v_d)\}\,\mathbb{1}_{x_0=S_0} + \{\varphi(S_1) + (\varphi(S_0) - \varphi(S_1))\delta(v_h,v_p)\}\,\mathbb{1}_{x_0=S_1}.$$

Although the attack strategy depends on the state, due to the form of the system we can simplify the attack strategy as

$$v_A(S_0) = (v_a, 0),$$
$$v_A(S_1) = (0, v_h).$$

Consequently, the value function in Eq. (5.1) evaluated at each state is

$$u^A(S_0) = \inf_{v_a\in[0,1]}\{\Psi(S_0,v_a,v_d)\} \tag{5.2}$$

and

$$u^A(S_1) = \inf_{v_h\in[0,1]}\{\Psi(S_1,v_h,v_p)\}, \tag{5.3}$$

where

$$\Psi(S_0,v_a,v_d) = -g_a(v_a) + C_0\mathbb{1}_{v_a>0} + \beta u^A(S_0) + \beta\pi(v_a,v_d)(u^A(S_1) - u^A(S_0)) \tag{5.4}$$

and

$$\Psi(S_1,v_h,v_p) = C_v\mathbb{1}_{v_h>0} + \beta u^A(S_1) + \beta\delta(v_h,v_p)(u^A(S_0) - u^A(S_1)). \tag{5.5}$$

The following result shows the strategy that minimizes the cost functions in Eq. (5.2) and (5.3).

Theorem 1. *The optimal strategy of the attacker is*

1. $v_a = 0$ *and* $v_h = 0$ *if* $K > 0$,
2. $v_a = 1$ *and* $v_h = 0$ *if* $K < 0$ *and* $B > 0$,
3. $v_a = 1$ *and* $v_h = 1$ *if* $K < 0$ *and* $B < 0$,

where $K = C_0 - g_a(1)$ *and* $B = C_v + \beta\frac{C_0 - g_a(1)}{1-\beta(1-\pi(1,v_d))}\delta(1,v_p)$.

The proof of this theorem is in Appendix A.1.

According to Theorem 1 the attacker makes binary decisions, i.e., whether attack or search vulnerabilities with maximum effort. In particular, the adversary attacks the system if he obtains a positive profit ($K < 0$), and this decision is independent of the defender's actions. However, the defense strategy can avoid attacks in the long term, by increasing the expected cost ($B > 0$) of discovering vulnerabilities. The following example shows the conditions in which hacking the system is unprofitable, leading to $v_h = 0$.

Example 5. *Here we construct the transition probabilities of the Markov process in Fig. 5.1 multiplying two functions that depend on the action of the attacker or the defender. In this way, we can specify the characteristics of the transition probabilities. Let us define the functions*

$$f_1(v, c) = \frac{e^c - e^{c(1-v)}}{e^c - 1}$$

and

$$f_2(v, c) = \frac{e^{c(1-v)} - 1}{e^c - 1}.$$

Observe that $f_1(\cdot, c)$ is concave and $f_2(\cdot, c)$ is convex, with $v \in [0, 1]$ and $c > 0$. The parameter c characterizes the change of the probability with respect to the variable v. Now we construct $\pi(\cdot)$ multiplying two concave functions and $\delta(\cdot)$ multiplying a concave and a convex function. Specifically,

$$\pi(v_a, v_d) = f_1(v_a, k_\pi) f_1(v_d, k_\pi)$$

and

$$\delta(v_h, v_p) = f_1(v_h, k_\delta) f_2(v_p, k_\delta).$$

In the experiments we select $k_\pi = k_\delta = 1$. On the other hand, we assume that the attacker has a linear profit function $g_a(v_a) \geq 4v_a$. Also, we select $C_0 = 1$, $C_v = 2$, and $\beta = 0.75$. With these parameters the optimal attack effort is $v_a = 1$, because $C_0 - g_a(1) < 0$.

Fig. 5.2 shows the optimal attack for different defense strategies and different values of $g_a(1)$ (the maximum benefit of the attacker). Here defense actions in the region below the line lead to hacks ($v_h = 1$). Observe that the region in which the attacker searches for vulnerabilities ($v_h = 1$) grows with the profit of the attack. Also, the defender can deter hacks ($v_h = 0$) on the system through investments in prevention ($v_p \geq 0.83$). On the contrary, investments in detection cannot discourage hacks by themselves.

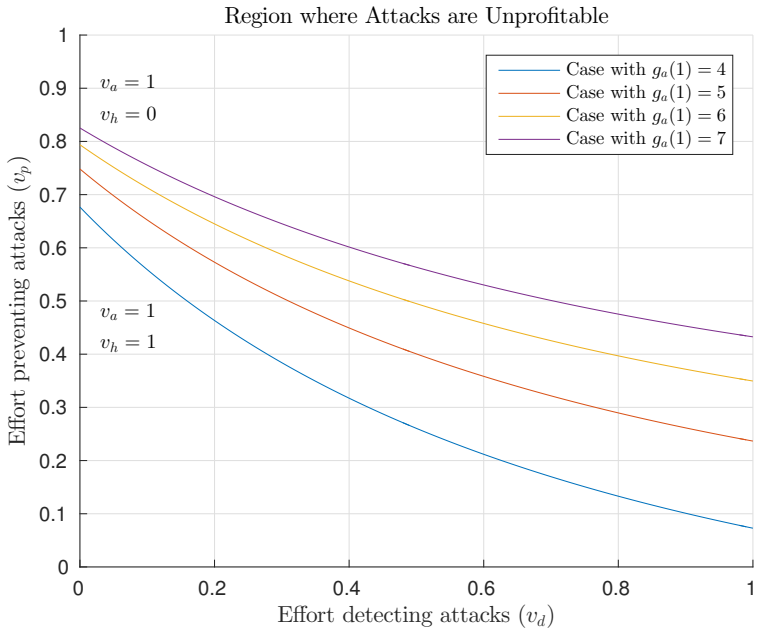

Figure 5.2: Optimal attack strategy as a function of the defense strategy $v_D = (v_d, v_p)$ and the maximum profit of the attacker $g_a(1)$. The region below the line, which lead to hacks ($v_h = 1$), grows with the attack's profit.

5.4 Optimal Defense Strategy

Here we consider two scenarios in which the defender either observes or ignores the state of the system, that is, whether the system is secure or vulnerable. Note that the defender can choose a different strategy in each state only if the state is observable. Otherwise, the defender should implement the same strategy in every state. In Section 5.5 we analyze the impact of the information about the state in the decisions of the defender.

5.4.1 Full Information

In this case we assume that the system's state is observed by the defender; however, the defender doesn't know the exact vulnerabilities. For instance, Stuxnet caused losses to an uranium enrichment plant, but the defender didn't realize the source of the failures until the worm was discovered (Zetter, 2014). Hence the defender in a compromised state S_1 cannot secure the system (jump to S_0) without investing in detection.

The defender chooses an optimal protection policy $V_D = \{v_{D,n}\}$ that minimizes the following cost functional, given an attack strategy $v_A = (v_a, v_h)$:

$$J^D(x_0, v_A, V_D) = l_D(x_0, v_{A,0}, v_{D,0}) + \beta \mathbb{E}_{x_0}^{v_A, V_D}\{J^D(x_1, v_A, V_D)\}.$$

Similar to the attacker's case, this cost functional is finite, since the cost function is bounded by $g_d(v_a) \leq l_D(x, v_A, v_D) \leq g_d(1) + C_p(1) + C_d(1)$. Furthermore, the defense policy is stationary, that is, $v_{D,n} = v_{D,n+1} = v_D$ for every $n \geq 0$. Therefore, the defense policy is $V_D = v_D$, where $v_D = (v_d, v_p)$. Henceforth we replace the policy V_D by the strategy v_D. Thus, the optimal solution satisfies

$$u^D(x_0) = \inf_{v_D} J^D(x_0, v_A, v_D) = \inf_{v_D} \left\{ l_D(x_0, v_A, v_D) + \beta \mathbb{E}_{x_0}^{v_A, v_D} \{u^D(x_1)\} \right\}.$$

The value function evaluated at each state is a follows

$$u^D(S_0) = \inf_{v_D \in [0,1]^2} \left\{ \Psi^D(S_0, v_a, v_D) \right\}$$

and

$$u^D(S_1) = \inf_{v_D \in [0,1]^2} \left\{ \Psi^D(S_1, v_h, v_D) \right\}.$$

where

$$\Psi^D(S_0, v_a, v_D) = g_d(v_a) + C_p(v_p) + C_d(v_d) + \beta u^D(S_0) + \beta \pi(v_a, v_d)(u^D(S_1) - u^D(S_0))$$

and

$$\Psi^D(S_1, v_h, v_D) = C_p(v_p) + C_d(v_d) + \beta u^D(S_1) + \beta \delta(v_h, v_p)(u^D(S_0) - u^D(S_1)).$$

In this case, it is not profitable to invest in detection (or prevention) when the system is in the secure (or vulnerable) state, respectively. Therefore, we can express the defense strategy as

$$\begin{aligned} v_D(S_0) &= (0, v_p), \\ v_D(S_1) &= (v_d, 0). \end{aligned} \tag{5.6}$$

The following result shows the defender's cost function with the strategy in Eq. (5.6).

Theorem 2. *The defender's discounted cost function is equal to*

$$J^D(S_0, v_A, v_D(S_0)) = \frac{Q(v_d)}{1-\beta} + \frac{\beta}{1-\beta} \frac{\pi(v_a, v_d)(W(v_p) - Q(v_d))}{1 + \beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)}$$

and

$$J^D(S_1, v_A, v_D(S_1)) = \frac{W(v_p)}{1-\beta} + \frac{\beta}{1-\beta} \frac{\delta(v_h, v_p)(Q(v_d) - W(v_p))}{1 + \beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)},$$

where $v_D(S_0) = (0, v_p)$ *and* $v_D(S_1) = (v_d, 0)$, $Q(v_d) = g_d(v_a) + C_d(v_d)$, *and* $W(v_d) = C_p(v_p)$.

The proof of this Theorem is in Appendix A.2.

The optimal defense strategy $v_D(x)$ minimizes the cost in every state simultaneously. However, we simplify the formulation finding an approximate solution that minimizes the total cost in all the states:[3]

$$\begin{aligned} \underset{v_D(S_1), v_D(S_0)}{\text{Minimize}} \quad & J^D(S_0, v_A, v_D(S_0)) + J^D(S_1, v_A, v_D(S_1)) \\ \text{subject to} \quad & \\ & v_D(S_0) = (0, v_p), \\ & v_D(S_1) = (v_d, 0), \\ & v_p, v_d \in [0, 1]. \end{aligned}$$

Note that if the cost functions $C_d(\cdot)$ and $C_p(\cdot)$ are identical, then the actions v_d and v_p are also identical.

5.4.2 Asymmetric Information

In this section we consider the case in which the defender cannot observe directly the state of the system. Instead, we assume that the defender has some belief about the initial state and knows the transition probabilities $\pi(\cdot)$ and $\delta(\cdot)$. The uncertainty of the initial state is captured with a probability distribution of the state:

$$\mathbb{P}(x_0 = S_0) = p \quad \text{and} \quad \mathbb{P}(x_0 = S_1) = 1 - p.$$

With the initial belief of the state and the transition probabilities we can compute the probability that the system will be at some state at time $n \geq 0$.

Moreover, we assume that the defender implements the same strategy in every state, that is, $v_D(S_0) = v_D(S_1) = v_D$. Thus, the defender aims to minimize the expected discounted costs. At time $n > 0$, the expected cost is

$$\hat{J}_n^D(v_A, v_D) = \mathbb{P}(x_n = S_0) l_D(S_0, v_A, v_D) + \mathbb{P}(x_n = S_1) l_D(S_1, v_A, v_D) + \beta \hat{J}_{n-1}^D(v_A, v_D) \quad (5.7)$$

with the initial expected cost

$$\begin{aligned} \hat{J}_0^D(v_A, v_D) &= \mathbb{P}(x_0 = S_0) l_D(S_0, v_A, v_D) + \mathbb{P}(x_0 = S_1) l_D(S_1, v_A, v_D) \\ &= g_d(v_a)\mathbb{P}(x_0 = S_0) + C_d(v_d) + C_p(v_p). \end{aligned}$$

The following result shows the defender's infinite horizon cost

[3] In the experiments we find that the approximate solution is almost equal to the optimal solution.

Theorem 3. *If* $\mathbb{P}(x_0 = S_0) = 1/2$*, then the defender's cost function with partial information is*

$$\hat{J}^D(v_A, v_D) = \frac{g_d(v_a)}{1-\beta}\gamma(v_A, v_D) + \frac{C_d(v_d) + C_p(v_p)}{1-\beta}$$

where

$$\gamma(v_A, v_D) = \begin{cases} \frac{1}{1-\beta}\frac{\delta}{\pi+\delta} & \text{if } 0 < \pi + \delta < 2 \\ \frac{1}{2}\frac{1}{1-\beta} & \text{otherwise} \end{cases}$$

and $\delta = \delta(v_h, v_p)$ *and* $\pi = \pi(v_a, v_d)$.

The proof of this Theorem is in Appendix A.3.

To guarantee convergence of the cost function we assume that $\mathbb{P}(x_0 = S_0) = 1/2$. Also, observe that the cost function has a discontinuity at $\pi(v_a, v_d) = \delta(v_h, v_p) = 0$.

The optimal defense strategy v_D is the solution to the following optimization problem

$$\begin{aligned} \underset{v_D}{\text{Minimize}} \quad & \hat{J}^D(v_A, v_D) \\ \text{subject to} \quad & \\ & v_D = (v_d, v_p), \\ & v_p, v_d \in [0, 1]. \end{aligned}$$

5.5 Attack-Defense Game

We have solved the optimal investment allocation problem for attackers and for defenders (with full and partial information), independently from each other. We now take into account that the investment choices of one party will be used by the other to change their own strategy. In particular, we model the investment choices between attackers and defenders as a simultaneous game between them (with full and partial information).

5.5.1 Defender with Full Information

We consider a simultaneous game to describe the interactions between attacker and defender. In this game, the defender selects the security scheme $v_D = (v_d, v_p)$ and the attacker selects its attack strategy $v_A = (v_a, v_h)$. The cost function of the attacker and the defender are $J^A(x_0, v_A, v_D)$ and $J^D(x_0, v_A, v_D)$, respectively.

We define the defender's cost functions as

$$C_d(v_d) = k_d(e^{v_d \ln 2} - 1),$$

and

$$C_p(v_d) = k_p(e^{v_p \ln 2} - 1).$$

These cost functions satisfy $C_i(0) = 0$ and $C_i(1) = k_i$. The loss function is defined as $g_d(v_a) = k_l v_a$. In the simulations, unless stated otherwise, we use the following parameters: $k_l = 4$, $k_d = 4$, and $k_p = 3$.

Below we analyze the Nash equilibrium that arises from the interaction between the attacker and the defender as a function of the defender's losses ($g_d(\cdot)$), and the cost of detection ($c_d(\cdot)$) and prevention ($c_p(\cdot)$).

Equilibrium as a Function of $g_d(\cdot)$

Here we consider the defender's loss parameter k_l in the interval $[0, 15]$ and observe the optimal defense strategy when $v_a = 1$ and $v_h \in \{0, 1\}$. The simulations in Fig. 5.3 show that system has a pure Nash equilibrium when the defense strategy allows looking for new vulnerabilities,[4] that is, when $v_h = 1$. If the defense strategy manages to make unprofitable attacks ($v_h = 0$), then the system does not have an pure Nash equilibrium. In such case, if $v_h = 1$, then the response of the defender will force the attacker to make $v_h = 0$. However, if $v_h = 0$, the attacker will have zero investment in prevention ($v_p = 0$), which encourages the attacker to make $v_h = 1$, repeating the cycle again.

Equilibrium as a Function of $c_d(\cdot)$

In this case, we analyze the response of the players when $k_p = 3$ and k_d is in the interval $[0, 7.5]$. The simulation results in Fig. 5.4 show that if $v_h = 1$, then it is always optimal to invest in prevention, even when detection is free ($k_d = 0$). Also, there is no pure Nash equilibrium, because the defense strategy when $v_h = 1$ will make attacks unprofitable, while if $v_h = 0$ the response of the defender allows looking for vulnerabilities.

Equilibrium as a Function of $c_p(\cdot)$

We analyze the response of the players when $k_d = 3$ and k_p is in the interval $[0, 7.5]$. The simulations in Fig. 5.5 show that if $v_h = 1$, then investing in detection is optimal even when $k_p = 0$. Also, the system has two Nash equilibria: i) if the defense strategy allows looking new vulnerabilities; ii) if the investment in prevention is large and avoids looking for vulnerabilities. The Nash equilibrium when $v_h = 1$ ($v_h = 0$) is reached for large (small) values of k_p.

[4]Intuitively, there is a Nash equilibrium when the defender's strategy with $v_h = 1$ ($v_h = 0$) remains below (above) the attacker's decision boundary.

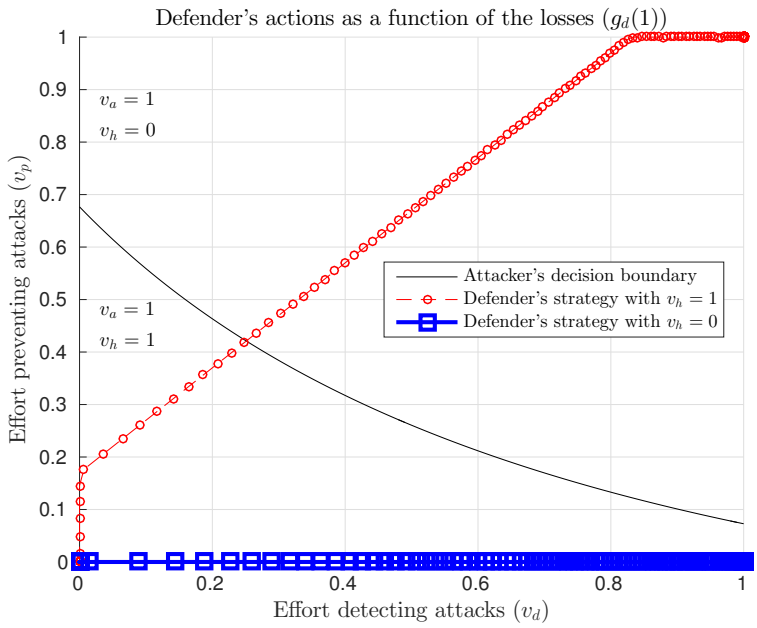

Figure 5.3: Change in the defender's strategy with the cost of losses, k_l, when $v_a = 1$ and $v_h = \{0, 1\}$. The system has a pure Nash equilibrium when the defense strategy allows looking for new vulnerabilities ($v_h = 1$).

Equilibrium with Budget Constraints

Let us consider a case in which the defender has limited resources to implement his strategy. Here, the total investment should be less or equal than some budget $E \geq 0$, that is,

$$C_d(v_d) + C_p(v_p) \leq E.$$

With this restriction, the defender's optimal strategy becomes

$$\begin{aligned} \underset{v_D(S_1), v_D(S_0)}{\text{Minimize}} \quad & J^D(S_0, v_A, v_D(S_0)) + J^D(S_1, v_A, v_D(S_1)) \\ \text{subject to} \quad & \\ & C_d(v_d) + C_p(v_p) \leq E, \\ & v_D(S_0) = (0, v_p), \\ & v_D(S_1) = (v_d, 0), \\ & v_d, v_p \in [0, 1]. \end{aligned}$$

Fig. 5.6 shows the optimal strategy of the defender as a function of the budget E, when $k_d = k_p = 3$, $k_l = 9$, and $E \in [0, k_p + k_d]$. Observe that with a limited budget (less than the necessary to implement an optimal strategy) the defender's optimal action prioritizes

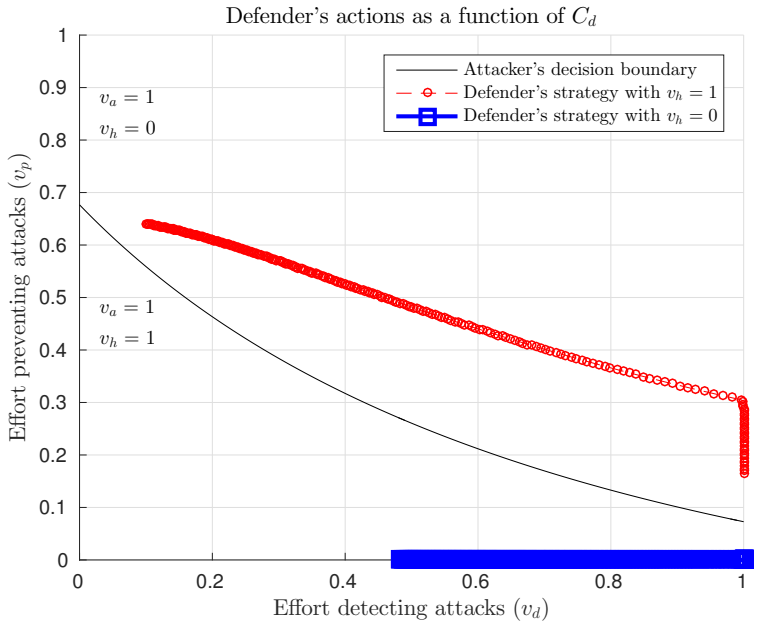

Figure 5.4: Change in the defender's strategy with the cost of detection, k_d, when $v_a = 1$ and $v_h = \{0, 1\}$. The system does not have a pure Nash equilibrium.

either detection or prevention, even though the cost function of both strategies is the same. In particular, with a low defense budget it is better to invest all (or at least most) of the resources in detection. As the budget increases it becomes better to invest more resources in prevention. Without budget constraints, the investments v_d and v_p are equal, because the costs of detection and prevention are equal.

5.5.2 Defender with Asymmetric Information

We now turn our attention to the case where the defender has asymmetric information. We use the same loss and cost functions we used in the last section for the case with full information.

Equilibrium as a Function of $g_d(\cdot)$

Fig. 5.7 shows the optimal action of the defender when $k_l \in [0, 15]$. When $v_h = 1$, the defender's best action is to either invest only in detection or invest in complete protection with some level of detection. Large investments in protection are optimal when attacks cause high losses. For lower losses, the best strategy is to invest only in detection (although the investment is relatively low).

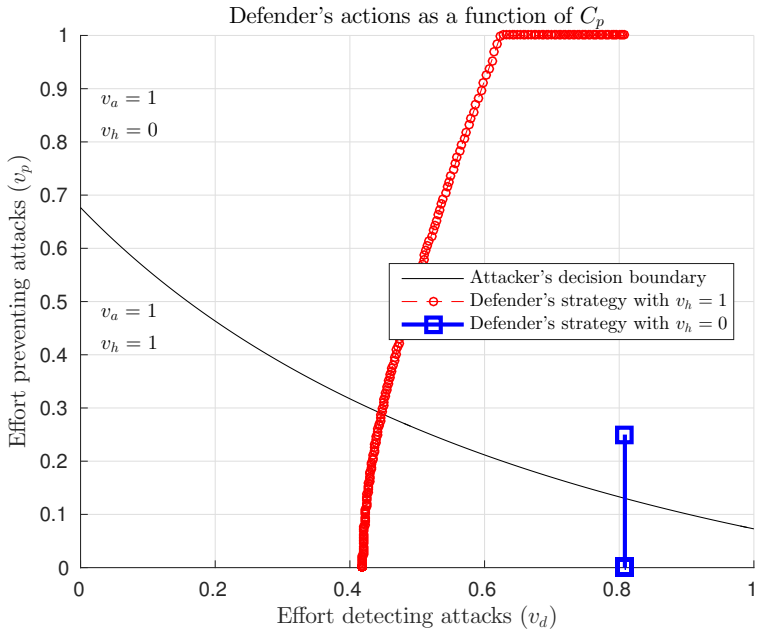

Figure 5.5: Change in the defender's strategy with the cost of prevention, k_p, when $v_a = 1$ and $v_h = \{0, 1\}$. The system has two Nash equilibria: i) if the defense strategy allows looking for new vulnerabilities; ii) if the investment in prevention is large and avoids searching vulnerabilities.

On the other hand, if $v_h = 0$, the best action is to invest few resources in detection. The rationale is that although the attacker might not look for vulnerabilities, he can exploit them if he has the opportunity. Hence, it is necessary to guarantee some detection. The game has a Nash equilibrium when $v_h = 0$ and $v_p = 0$.

Equilibrium as a Function of $C_d(\cdot)$

Fig. 5.8 shows the defenders best action when $k_p = 3$ and $k_d \in [0, 7.5]$. In this case the defender invests only in detection, regardless of its cost. Nevertheless, the investment is larger when the cost is lower. Surprisingly, for large C_d the optimal strategy is not invest in security at all, that is ($v_d = 0$ and $v_p = 0$). The game has a Nash equilibrium when $v_h = 1$ and $v_p = 0$.

Equilibrium as a Function of $C_p(\cdot)$

Fig. 5.9 shows the defender's best action when $k_d = 3$ and $k_p \in [0, 7.5]$. In this case the defender has two optimal actions. When $v_h = 1$ the defender's best action is to either invest only in detection or invest in complete prevention with some level of detection. The defender

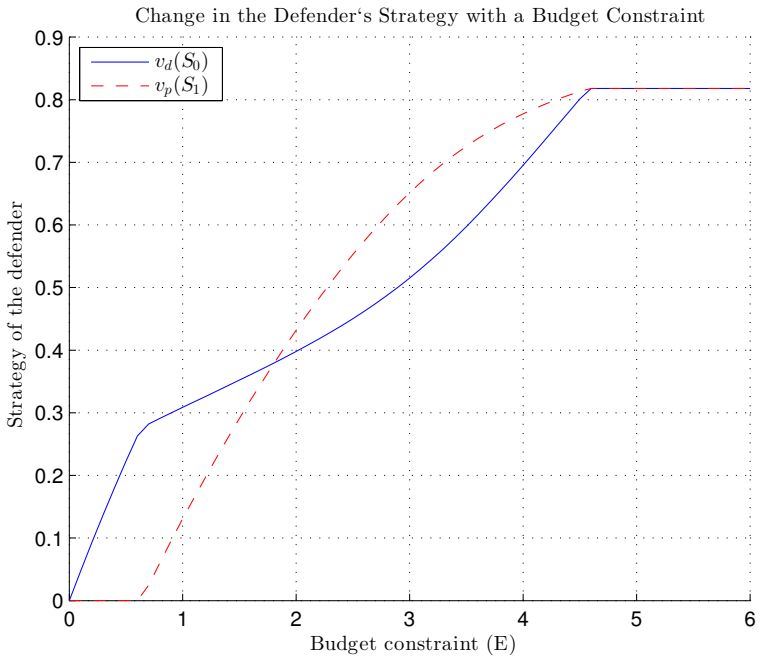

Figure 5.6: Change in the defender's strategy with the budget constraint E. For a small budget the best strategy is to prioritize detection over prevention (or vice versa).

chooses $v_p = 1$ when the cost of prevention is small. However, for large costs of prevention, the defender invests only in detection.

When $v_h = 0$, the defender invests in prevention if the cost is low, otherwise the investment is minimum; however, it invests roughly the same amount in detection, regardless of the cost $C_p(\cdot)$. The game has a Nash equilibrium when $v_h = 0$ and $v_p = 0$.

Equilibrium with a Budget Constraints

Similar to the case with full information, here we define a constraint in the expenses in security:

$$C_d(v_d) + C_p(v_p) \leq E.$$

Thus, the defender's optimization problem becomes

$$\begin{aligned} \underset{v_D}{\text{Minimize}} \quad & \hat{J}^D(v_A, v_D) \\ \text{subject to} \quad & \\ & C_d(v_d) + C_p(v_p) \leq E \\ & v_D = (v_d, v_p) \\ & v_d, v_p \in [0, 1] \end{aligned}$$

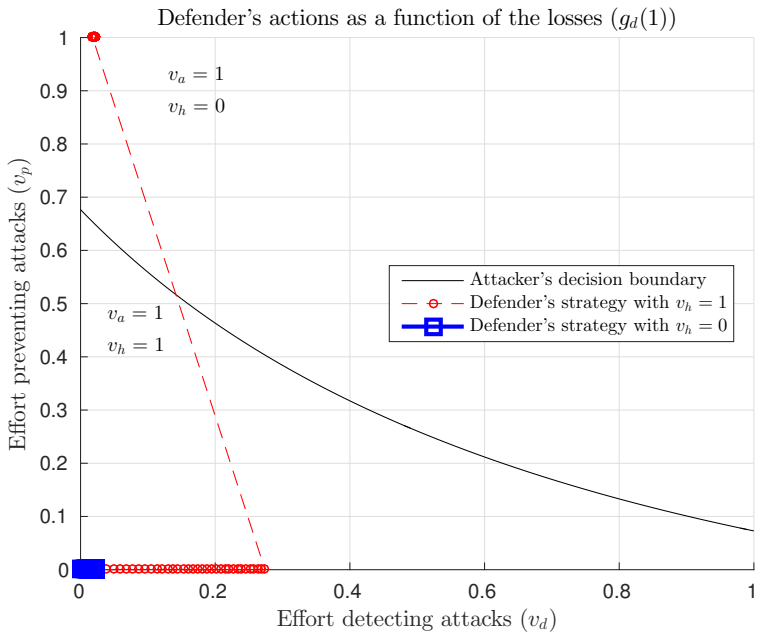

Figure 5.7: Change in the defender's strategy with the cost of losses, k_l, when $v_a = 1$ and $v_h = \{0, 1\}$. The system has a pure Nash equilibrium when the defense strategy allows searching new vulnerabilities ($v_h = 1$ and $v_p = 0$).

Fig. 5.10 shows the optimal defense when $k_d = k_p = 3$, $k_l = 9$, and $E \in [0, k_p + k_d]$. As in the case with full information, the optimal strategy with a low budget is to invest only in detection. However, as the budget increases, the investment in detection decreases and the investment in prevention increases. Due to the discontinuity of the cost function, for large budgets it is optimal to invest a small amount in detection.

5.6 Conclusions

In this work we find that the attacker's optimal strategy is binary, that is, either attack (or hack) with the maximum intensity or not at all. When the defender observes the system's state, we found two Nash equilibria, which allows or deters the attacker to discover vulnerabilities. Moreover, the optimal defense strategy has always a combination of both detection and prevention, even when the cost of prevention or detection is zero. Hence, no strategy alone minimizes the cost for the defender, although in some cases investing only in prevention can avoid hacks (see Fig. 5.2). In the experiments the investment in prevention and detection increases with the budget. We also find that with few resources the best strategy is to prioritize detection over prevention.

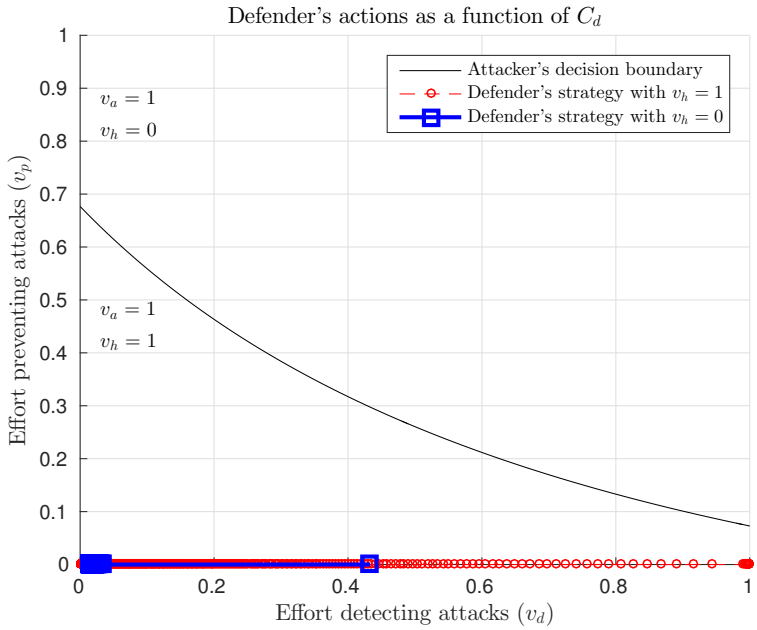

Figure 5.8: Change in the defender's strategy with the cost of detection, k_d, when $v_a = 1$ and $v_h = \{0, 1\}$. The game has a Nash equilibrium when $v_h = 1$ and and $v_p = 0$.

With limited information we found only one Nash equilibrium, in which the attacker has incentives to discover vulnerabilities, that is, the actions of the defender do not prevent attacks. Also, the defender tends to invest resources only in detection technologies. An exception occurs when either the cost of prevention is low or the losses are high; in such cases the defender tends to invest in maximum prevention. Moreover, similar to the case with full information, with a small budget the best strategy is to prioritize detection over prevention. In contrast, the investment in detection decreases for larger budgets.

5.6.1 Future Directions

In practice the defender usually ignores risk of security breaches, due to limited information to estimate the transition probabilities of the Markov process in Fig. 5.1. Therefore, it is necessary to develop techniques to deal with such uncertainties. In particular, the defenders could estimate the risk of breaches through security audits.

Also, it would be interesting to analyze how to deal with the negative effects of uncertainties. In particular, uncertainties about the system's state have a negative impact on the security investments. For example, in our experiments the defender with asymmetric information and limited budget invests less resources in detection (see Fig. 5.10). The government could design economic incentives to correct the negative behaviors that arise due to uncertainties.

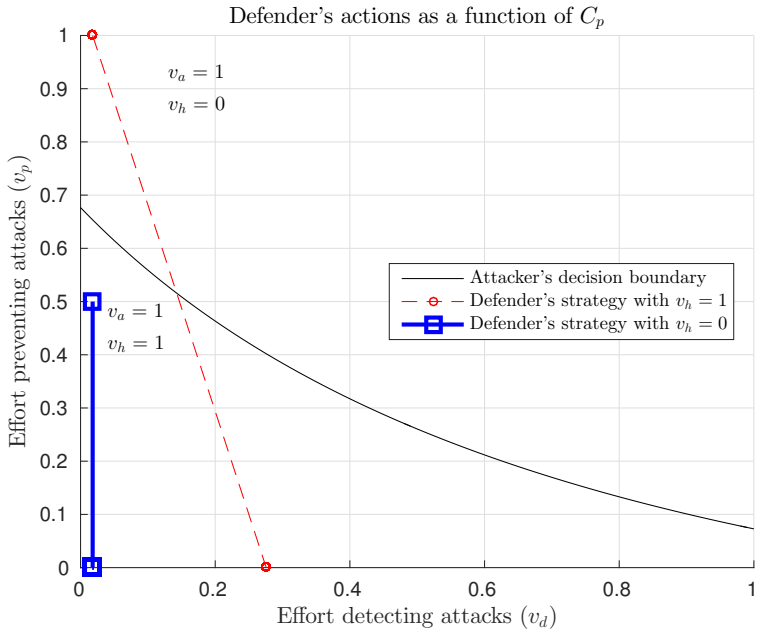

Figure 5.9: Change in the defender's strategy with the cost of prevention, k_p, when $v_a = 1$ and $v_h = \{0, 1\}$. The game has a Nash equilibrium when $v_p = 0$.

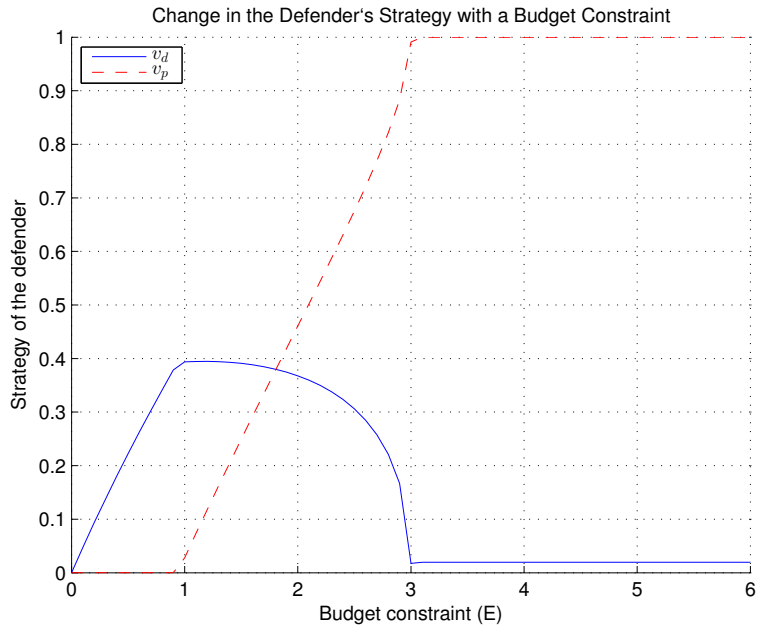

Figure 5.10: Change in the defender's strategy with the budget constraint E. For a small budget the best strategy is to prioritize detection over prevention. However, the investment in detection decrease for larger budgets.

CHAPTER 6

OPTIMAL INVESTMENTS WITH CYBER-INSURANCE

Authors – Carlos Barreto and Alvaro A. Cardenas

The Computer Science Department, EC 31

The University of Texas at Dallas

800 West Campbell Road

Richardson, Texas 75080-3021

Corresponding author: Carlos Barreto.

The content of this chapter is reprinted with permission from: C. Barreto and A. A. Cárdenas, "Optimal risk management in critical infrastructures against cyber-adversaries," In 2017 IEEE Conference on Control Technology and Applications (CCTA), ©2017 IEEE.

6.1 Introduction

When the adversaries pursue some economic benefit it is possible to design incentives to reduce the potential rewards that motivate their attacks. This idea follows from the seminal paper of Gary Becker on the economics of crime and further developments in information security (Becker, 1968; Schechter and Smith, 2003). For instance, (Barreto and Cárdenas, 2015a) studies a detection scheme and penalties to prevent false information attacks on smart grids. Another study (Barreto and Cárdenas, 2016) investigates how traditional contracts allowed electric tower repair companies to profit by sponsoring more attacks on electric transmission towers, and then shows a contract scheme that reduces the incentives of malicious repair companies. In such cases, knowing the adversary's goal is crucial to make the system a less attractive target for attackers.

However, incentives can fail to prevent attacks when he adversary pursues either a non-economic or a unknown goal.[1] In such cases the defender has to rely on risk management strategies to minimize the expected losses from attacks. Risk management seeks to either *mitigate* or *transfer* the risk of hazardous events. Risk mitigation includes detection and prevention of cyber attacks, while risk transfer includes mechanisms like cyber insurance,[2] where the defender transfers the risk to another party.

In this chapter we analyze how defenders can manage their investments efficiently to protect themselves against cyber-attacks. We propose a Markov decision process to model the system's exposure to attacks as a function of the investment in risk management strategies. In particular, we consider three risk management options: attack-detection (e.g., investing in intrusion detection systems), attack-prevention (e.g., investing in access control technologies), and cyber insurance. With this model we find the best investment strategy that reduces losses in the long term.

Our model has two key characteristics: First, we capture the fact that attack-attribution is a difficult problem assuming that the defender cannot penalize the attacker (adversaries often remain unknown or reside in countries where thy cannot be prosecuted). Second, the Markov process allows us to analyze the interaction among individuals, considering restrictions in their actions (Barreto and Cárdenas, 2017).

---

[1] Economic incentives can fail if the adversary pursues political interests, such as hacktivism, espionage, sabotage, terrorism, or war.

[2] Insurance against losses caused by cyber attacks.

Our work is closely related to (Preciado et al., 2014; Rasouli et al., 2014); however, we depart from them in that 1) we focus on the strategic aspect of the adversary, and 2) we consider the evolution of threats in time. Specifically, cyber-weapons become ineffective once the vulnerabilities are discovered and fixed by software vendors, forcing the attacker to find new ones. Moreover, in this work we extend our previous work (Barreto et al., 2017) by allowing the defender to invest in insurance.

The chapter is structured in the following way: Section 6.2 provides an introduction to the concept of insurance. Section 6.3 introduces the Markov decision process that models the exposure to cyber risks. Section 6.4 deals with the optimal investment in risk management with full and limited information. Section 6.5 shows an example of subsidies to incentivize the adoption of insurance and thus improve security investments in critical infrastructures. We finish the chapter with closing remarks in Section 6.6.

6.2 The Role of Insurance in Risk Management

Insurance is a mechanism in which an agent (insured) transfers its risk[3] to another party (insurer) (Bernstein, 1996). The insured pays periodically some premium P to the insurer and receives in exchange an indemnity I when an incident occurs causing losses L. The premium should be enough to cover claims and operation costs of the insurer; a common benchmark is the *fair premium principle*,[4] which defines the premium as the expected payments of claims ($P = \mathrm{E}[I]$) (Arrow, 1971; Mikosch, 2006). With a fair premium the accumulated payments equal the total claims (in the long run). One benefit of insurance is that the insured pays only a fraction of losses, rather than the total amount up front.

The insurer can reduce the variance of the average claims by collecting the risk of many customers in a pool, which allows a more precise calculation of the premium. Thus, insurance has a social nature and needs the participation of many individuals. In some cases this large participation arises organically, but in others, participation in insurance markets requires the government's intervention (e.g., governments often push regulations to improve the participation in car, earthquake, and health insurances).

In particular, due to the general underinvestment to protect critical infrastructures, the U.S. government is interested in promoting insurance, because successful insurance industries can incentivize firms to invest in mitigation of risk (when this leads to reductions in the

[3]Risk is the uncertainty associated with an outcome.

[4]Fair premiums do not account for the risk of the insurer and can lead to bankruptcy in practice.

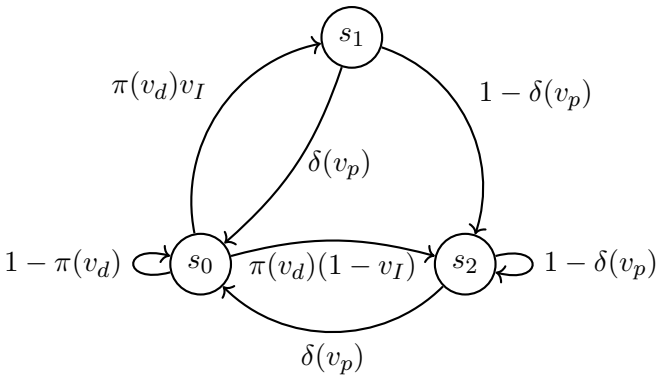

Figure 6.1: Markov process that describes changes in the security of the system. The system has a vulnerable state s_0 and two secure states s_1 and s_2, which differ in that s_1 occurs with insurance.

premium (Ehrlich and Becker, 1972)). Therefore, cyber insurance is expected to play a role improving the security of critical infrastructures and managing cyber risks of (potentially catastrophic) cyber attacks, as in the case of natural catastrophic events (Kunreuther, 2015).

6.3 Model

We assume that the defender can manage the risk of a system by investing in mitigation and transferring the risk. Mitigation involves investing in prevention (e.g., encryption and authentication) and detection (e.g., intrusion detection systems and audits) to reduce the success and damage of attacks. Moreover, we assume that there is a market of cyber insurance that offers coverage for cyber threats.

The defender can choose the degrees of investment in both prevention and detection, denoted by $v_p \in [0,1]$ and $v_d \in [0,1]$, respectively. In contrast, we assume that the insurance's cost is fixed and that the defender can decide only whether or not to purchase it. Thus, the decision to purchase insurance is denoted by $v_I \in \{0,1\}$. In summary, the action of the defender is a vector $v = (v_p, v_d, v_I) \in A$, where A is the set of possible actions, defined as $A = [0,1]^2 \times \{0,1\}$.

We consider a system with three states $S = \{s_0, s_1, s_2\}$, namely an insecure state s_0 and two secure states s_1 and s_2 (see Fig. (6.1)). In the insecure state s_0, the system is susceptible to attacks; e.g., an attacker knows a vulnerability of the system and is able to exploit it. The system cannot suffer attacks in any of the secure states; however, in s_1 the defender receives an indemnity from the insurer. Thus, when the defender purchases insurance, the system can

go from s_0 to s_1 and then to s_2 (the indemnity is paid only once) and without insurance the system goes from s_0 to s_2.

Fig. 6.1 shows the state transitions of the system.[5] The transition from a secure state (s_1 or s_2) to the insecure state s_0 occurs with probability $\delta(v_p)$, which increases with low investments in prevention. On the other hand, the transition probability from the insecure state s_0 to one of the secure states depends on both the investment in detection and insurance. In particular, the system jumps from s_0 to s_1 with probability $\pi(v_d)$ only if the defender pays a premium in advance (i.e., $v_I = 1$ in s_0); otherwise, the system jumps from s_0 to s_2 with probability $\pi(v_d)$. The probability of detecting an attack increases with investments in detection technologies.

Assumption 4. *The probability of detection $\delta : [0,1] \to [0,1]$ is a convex decreasing continuous function and satisfies $\delta(0) = 1$ and $\delta(1) = \varepsilon$.*

The probability of finding a vulnerability $\pi : [0,1] \to [0,1]$ is a concave increasing continuous function that satisfies $\pi(0) = \varepsilon$ and $\pi(1) = 1$.

We assume that new vulnerabilities are identified eventually even if the defender implements maximum protection, i.e., $\delta(1) = \varepsilon > 0$. Also, vulnerabilities can be discovered even if the attacker does not use them (Axelrod and Iliev, 2014), then, $\pi(0) = \varepsilon > 0$.

In this model the state transitions occur in discrete time instants $0, T, 2T, \ldots$, where T is the time period between transitions.

The cost of operating the system depends on the losses caused by attacks and the cost of the risk management strategy. We define the cost for the defender (in a single time period) as

$$l(x,v) = \begin{cases} L + C_p(v_p) + C_d(v_d) + P(x,v), \text{ if } x = s_0 \\ C_p(v_p) + C_d(v_d) + P(x,v) - I, \text{ if } x = s_1 \\ C_p(v_p) + C_d(v_d) + P(x,v), \text{ if } x = s_2, \end{cases}$$

where $x \in S$ is the state of the system, $v = (v_p, v_d, v_I)$ is the action of the defender, L is are the losses due to an attack, $C_p(v_p)$ is the cost of prevention, $C_d(v_d)$ is the cost of detection, I is the indemnity paid by the insurer, and $P(x,v)$ is a function that calculates the premium.

[5] We modify the model in (Barreto et al., 2017), which describes the interaction among an attacker and a defender, adding a state to consider the effect of cyber insurance. Also, we assume that the adversary attacks the system regardless of the defense strategy; therefore, in this case we do not model the attack strategy as a variable.

We assume that the defender seeks to maximize its *discounted utility* (Samuelson, 1937). In this case, we consider the following utility function for the defender:

$$H(x, v) = U(w - l(x, v)), \tag{6.1}$$

where w is some income and $U : \mathbb{R} \to \mathbb{R}$ is a concave increasing continuous function.[6]

6.3.1 Asymmetries in Information

Below we consider two situations: the defender either fully observes the state of the system, or has limited information about the state of the system. Although the former case is unrealistic in most scenarios, it provides a benchmark of the ideal defense strategy.

Full Information

In this case, the defender (and the insurer) observe the state of the system. In other words, the defender knows whether the attacker possess a cyber weapon that endangers the system, but the precise vulnerability of the system is unknown. In consequence, the defender can adjust its risk management strategy according to the current state.

Limited Information

In this case the defender does not know the true current state and therefore, he has to implement the same risk management strategy regardless of the state of the system. Nonetheless, we assume that the defender knows the characteristics of the system (e.g., transition probabilities), and therefore, he can estimate the probability that the system is in a particular state. The following result shows the stationary distribution of the Markov decision process when the defender implements a strategy v.

Lemma 3. *The stationary distribution* $\rho(v) = [\rho_0(v), \rho_1(v), \rho_2(v)]$, *where* $\rho_i(v) = \mathbb{P}[x = s_i|v]$, *of the Markov process depicted in Fig. 6.1 is equal to*

$$\rho(v) = \begin{cases} \rho_I(v), & \text{if } v_I = 1 \\ \rho_{NI}(v), & \text{if } v_I = 0, \end{cases}$$

[6] We make this assumptions because only risk averse users (who cave concave utility functions) purchase insurance.

where the stationary distribution with insurance is

$$\rho_I(v) = \left(\frac{\delta(v_p)}{\pi(v_d)+\delta(v_p)} \quad \frac{\pi(v_d)\delta(v_p)}{\pi(v_d)+\delta(v_p)} \quad \frac{\pi(v_d)(1-\delta(v_p))}{\pi(v_d)+\delta(v_p)} \right) \tag{6.2}$$

and the stationary distribution without insurance is

$$\rho_{NI}(v) = \left(\frac{\delta(v_p)}{\pi(v_d)+\delta(v_p)} \quad 0 \quad \frac{\pi(v_d)}{\pi(v_d)+\delta(v_p)} \right), \tag{6.3}$$

if $\pi(v_d) + \delta(v_p) \neq 0$.

Sketch proof. Let us begin with the case with insurance. Here the transition probability matrix of the system is

$$TP_I(v) = \begin{pmatrix} 1-\pi(v_d) & \pi(v_d) & 0 \\ \delta(v_p) & 0 & 1-\delta(v_p) \\ \delta(v_p) & 0 & 1-\delta(v_p) \end{pmatrix}.$$

It can be verified that Eq. (6.2) satisfies

$$\rho_I(v) = \rho_I(v) TP_I(v),$$

if $\pi(v_d) + \delta(v_p) \neq 0$. Now, the transition probability matrix of the system without insurance is

$$TP_{NI}(v) = \begin{pmatrix} 1-\pi(v_d) & 0 & \pi(v_d) \\ \delta(v_p) & 0 & 1-\delta(v_p) \\ \delta(v_p) & 0 & 1-\delta(v_p) \end{pmatrix}.$$

It can be verified that the stationary distribution $\rho_b(v)$ in Eq. (6.3) satisfies

$$\rho_{NI}(v) = \rho_{NI}(v) TP_{NI}(v)$$

if $\pi(v_d) + \delta(v_p) \neq 0$. □

6.3.2 Insurance

We assume that the defender can purchase insurance at any moment and that the premium is calculated based on the probability that an attack occurs during the next time period. Hence, based on the fair premium principle, the premium should be calculated using the probability of reaching the state s_1 in the next time period, that is,

$$P(x, v) = I \cdot \mathbb{P}[Claim|x, v] = I \cdot \mathbb{P}(x_{n+1} = s_1 | x_n = x, v),$$

where x is the current state and v is the strategy of the defender. Below we consider the cases with asymmetries in information.

Insurance with Full Information

Since the state of the system is known, we know that

$$\mathbb{P}(x_{n+1} = s_1 | x_n = x, v) = \begin{cases} \pi(v_d)\mathbb{1}_{v_I=1}, & \text{if } x_n = s_0 \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbb{1}_{v_I=1}$ is the indicator function. Hence,

$$P(x, v) = \begin{cases} I\,\pi(v_d) & \text{if } x = s_0 \text{ and } v_I = 1 \\ 0, & \text{otherwise.} \end{cases}$$

Thus, the defender pays a premium only in s_0.

Insurance with Limited Information

In this case, we assume that the defender can observe transitions from s_0 to the states s_1 and s_2, which occur when an attack is detected. Thus, the defender can make claims of losses to the insurer (although the system's state becomes uncertain in the following time periods). According to Lemma 3, the probability of a claim in the next time period is equal to

$$\mathbb{P}[Claim|v] = \mathbb{P}(x_{n+1} = s_1 | x_n = s_0, v)\mathbb{P}(x_n = s_0) = \rho_1(v).$$

Consequently, the fair premium for the case with limited information is

$$P(x, v) = P(v) = I \cdot \mathbb{P}(Claim|v) = I\rho_1(v). \tag{6.4}$$

In this case the premium does not depend on the state, that is, $P(v) = P(x, v)$.

6.3.3 Notes on the Model

We assume that only one transition occurs during each time period. For example, an attack during the k^{th} time period can be discovered only at the beginning of the $k + 1^{th}$ time period. Furthermore, we assume that losses are constant on every time period; however they accumulate while the attack remains undetected. This captures the intuition that attacks cause large or small losses depending on the duration of the attack.

We assume that the insurer gives indemnities when the insured shows precise evidence of a cyber attack, that is, when the system jumps from s_0 to s_1. This assumption agrees with Beck's definition of modern risks (or man-made risks) (Beck, 1992). According to Beck, people cannot determine their exposure to modern risks through simple inspection, but require the evaluation of an expert. According to this, cyber-insurance makes sense only when the defender invests in attack-detection mechanisms that allows him to make claims.

6.4 Optimal Defense Strategy

6.4.1 Full Information

We consider an infinite decision problem in which we want to find the control policy $V = \{v_0, v_1, v_2, \ldots, v_n, \ldots\}$ (where $v_k \in A$) that maximizes the performance criterion

$$J(x_0, V) = H(x_0, v_0) + \beta \mathbb{E}_{x_0}^{V} \Big\{ H(x_1, v_1) + \mathbb{E}_{x_1}^{V} \left\{ H(x_2, v_2) + \cdots + \beta \mathbb{E}_{x_{n-1}}^{V} \left\{ H(x_n, v_n) + \ldots \right\} \right\} \Big\},$$

where $\beta \in [0, 1)$ is a discount factor and x_k and v_k are the state and the defender's strategy at time $k = 0, 1, 2, \ldots$, respectively. Since we have an infinite horizon problem, the cost functional can be rewritten as

$$J(x_0, V) = H(x_0, v_0) + \beta \mathbb{E}_{x_0}^{V} \left\{ J(x_1, V) \right\}. \tag{6.5}$$

In this case, the losses of the defender are bounded by

$$\underline{l} \leq l(x, v) \leq \bar{l},$$

where $\underline{l} = C_p(0) + C_d(0) - I$ and $\bar{l} = L + C_p(1) + C_d(1) + I$. Thus,

$$w - \bar{l} \leq w - l(x, v) \leq w - \underline{l}.$$

Here we assume that $H(x, v)$ (see Eq. (6.1)) is positive and bounded for all $x \in S$ and $v \in A$, which implies that the cost functional in Eq. (6.5) is bounded.

In this infinite-horizon problem the optimal control policy is stationary, which means that $v_k = v_{k+1} = v$, for $k \geq 0$. Hence, the optimal policy has the form $V = \{v\}$. Now, using the optimality principle (Hernández-Lerma and Lasserre, 2012; Bensoussan, 2011) we can define the following value function:

$$u(x_0) = \sup_{v} J(x_0, v) = \sup_{v} \left\{ H(x_0, v) + \beta \mathbb{E}_{x_0}^{v} \left[u(x_1) \right] \right\}$$

The expectation of the value function $u(\cdot)$ is equal to

$$\mathbb{E}_{x_0}^{v}[u(x_1)] = \mathbb{E}[u(x_1)|x_0, v] = \left\{ u(s_0) + \pi(v_d) \left(u(s_1) v_I + u(s_2)(1 - v_I) - u(s_0) \right) \right\} \mathbb{1}_{x_0 = s_0} + \left\{ u(s_2) + \delta(v_p)(u(s_0) - u(s_2)) \right\} \mathbb{1}_{x_0 \neq s_0}.$$

Let us express the value function as

$$u(x) = \sup_{v} \{\Psi(x, v)\},$$

where

$$\Psi(x, v) = H(x, v) + \beta \mathbb{E}[u(x_{n+1})|x_n = x, v].$$

With full information the optimal strategy consists in choosing 1) prevention only in the secure states s_1 and s_2 and 2) detection and insurance only in state s_0.

6.4.2 Limited Information

Since the defender cannot observe the state of the system, the investment is independent of the vulnerabilities of the system. Hence, in the k^{th} the benefit of the defender is the expected benefit in the current time period plus the expected future earnings. The discounted benefit function becomes

$$\hat{J}_k(v) = \mathbb{E}[H(x_k, v)|v] + \beta \hat{J}_{k+1}(v),$$

where x_k is the state of the system in the k^{th} period. We know that the sequence $\{\hat{J}_k\}$ converges to some cost function $\hat{J}$ because $H(\cdot)$ is bounded. Thus,

$$\lim_{k\to\infty} \hat{J}_k(v) = \hat{J}(v) = \mathbb{E}[H(x, v)|v] + \beta \hat{J}(v).$$

From the previous equation we can rewrite the cost function as

$$\hat{J}(v) = \mathbb{E}[H(x, v)|v]/(1 - \beta),$$

which is equivalent to

$$\hat{J}(v) = \rho(v) \begin{pmatrix} M(v), & N(v), & G(v) \end{pmatrix}^{\top} /(1 - \beta),$$

where $\rho(v)$ is the stationary probability distribution for the case with asymmetric information (see Lemma 3), $M(v) = H(s_0, v)$, $N(v) = H(s_1, v)$, and $G(v) = H(s_2, v)$.

The next result shows sufficient conditions in which the defender with limited information about the state would (or would not) purchase insurance.

Theorem 4. *Let us consider an insurance with premium $\hat{P}$ and indemnity I. If the premium is higher than the fair premium, that is,*

$$\rho_1(v^1) I < \hat{P},$$

then the defender rejects insurance. Furthermore, if

$$\rho_1(v^0) I k_1(v^0)/k_2(v^0) > \hat{P},$$

then the defender purchases insurance, where $k_1(v) = \dot{U}(a(v)+I-\hat{P})$, $k_2(v) = \dot{U}(a(v)-L-\hat{P})$, $a(v) = w - c_p(v_p) - c_d(v_d)$, *and* v^1 *and* v^0 *are the optimal actions with and without insurance, respectively.*

Proof. Let $A_0 = [0,1]^2 \times \{0\}$ and $A_1 = [0,1]^2 \times \{1\}$ be the set of possible actions when the defender refuses ($v_I = 0$) or uses insurance ($v_I = 1$), respectively. Thus, we can express the problem of finding the optimal investment as

$$\max_{v \in A} \hat{J}(v) = \max\left\{\Gamma_0, \Gamma_1\right\}, \tag{6.6}$$

where

$$\Gamma_0 = \sup_{v \in A_0} \left\{\hat{J}(v)\right\} = \hat{J}(v^0)$$

and

$$\Gamma_1 = \sup_{v \in A_1} \left\{\hat{J}(v)\right\} = \hat{J}(v^1),$$

were $v^0 = (v_p^0, v_d^0, 0)$ and $v^1 = (v_p^1, v_d^1, 1)$ are the optimal strategies with and without insurance, respectively. From Eq. (6.6) we know that an individual would prefer (reject) insurance if $\Gamma_1 - \Gamma_0$ is greater (lower) than zero. Our purpose is to find boundaries of the difference $\Gamma_1 - \Gamma_0$ that allow us to determine the conditions to use or reject insurance.

Let us define the suboptimal actions $\tilde{v}^0 = (v_p^0, v_d^0, 1)$ and $\tilde{v}^1 = (v_p^1, v_d^1, 0)$. Now we can find the following boundaries of $\Gamma_1 - \Gamma_0$:

$$\hat{J}(\tilde{v}^0) - \hat{J}(v^0) \leq \Gamma_1 - \Gamma_0 \leq \hat{J}(v^1) - \hat{J}(\tilde{v}^1). \tag{6.7}$$

Let us consider first the upper bound in Eq. (6.7), which is equivalent to

$$\hat{J}(v^1) - \hat{J}(\tilde{v}^1) = \frac{\rho(v^1)}{1-\beta} \begin{pmatrix} M(v^1) - M(\tilde{v}^1) \\ N(v^1) - G(\tilde{v}^1) \\ G(v^1) - G(\tilde{v}^1) \end{pmatrix}. \tag{6.8}$$

Since $U(\cdot)$ is concave, the inequality

$$U(\mu - \sigma) \leq U(\mu) - \sigma \dot{U}(\mu)$$

is satisfied for every real numbers μ and σ. With the previous inequality we can find the following upper bound of Eq. (6.8)

$$\hat{J}(v^1) - \hat{J}(\tilde{v}^1) \leq \frac{\rho(v^1)}{1-\beta} \begin{pmatrix} -\hat{P}\dot{U}(a(v^1) - L) \\ (I - \hat{P})\dot{U}(a(v^1)) \\ -\hat{P}\dot{U}(a(v^1)) \end{pmatrix}, \tag{6.9}$$

where $a(v) = w - c_p(v_p) - c_d(v_d)$ and $a(v^1) = a(\tilde{v}^1)$. Since $U(\cdot)$ is concave increasing, we know that $\dot{U}(\cdot)$ is positive decreasing; hence, $\dot{U}(a(v^1) - L) \geq \dot{U}(a(v^1))$. Thus, we can replace $\dot{U}(a(v^1) - L)$ by $\dot{U}(a(v^1))$ in Eq. (6.9) to obtain the following upper bound

$$\hat{J}(v^1) - \hat{J}(\tilde{v}^1) \leq \frac{\dot{U}(a(v^1))}{1-\beta}(\rho_1(v^1)I - \hat{P}).$$

Thus, if $\rho_1(v^1)I < \hat{P}$, it is not convenient for the defender to purchase insurance. In other words, the defender does not accept an insurance that charges more than the fair premium (notice that $\rho_1(v^1)I$ is equal to the fair premium in Eq. (6.4)).

Let us consider now the lower bound of Eq. (6.7):

$$\hat{J}(\tilde{v}^0) - \hat{J}(v^0) = \frac{\rho(v^0)}{1-\beta} \begin{pmatrix} M(\tilde{v}^0) - M(v^0) \\ N(\tilde{v}^0) - G(v^0) \\ G(\tilde{v}^0) - G(v^0) \end{pmatrix}. \tag{6.10}$$

We use the following approximation of the utility function:

$$U(\mu - \sigma) \geq U(\mu) - \sigma\dot{U}(\mu - \sigma). \tag{6.11}$$

Thus, applying Eq. (6.11) in Eq. (6.10) we obtain the following lower bound

$$\hat{J}(\tilde{v}^0) - \hat{J}(v^0) \geq \frac{\rho(v^0)}{1-\beta} \begin{pmatrix} -\hat{P}\dot{U}(a(v^0) - L - \hat{P}) \\ (I - \hat{P})\dot{U}(a(v^0) + I - \hat{P}) \\ -\hat{P}\dot{U}(a(v^0) - \hat{P}) \end{pmatrix}.$$

We can replace $-\dot{U}(a(v^0) - \hat{P})$ and $-\dot{U}(a(v^0) + I - \hat{P})$ with $-\dot{U}(a(v^0) - L - \hat{P})$ to obtain the lower bound

$$\hat{J}(\tilde{v}^0) - \hat{J}(v^0) \geq \frac{1}{1-\beta}\left(\rho_1(v^0)Ik_1(v^0) - Pk_2(v^0)\right),$$

where $k_1(v^0) = \dot{U}(a(v^0) + I - \hat{P})$ and $k_2(v^0) = \dot{U}(a(v^0) - L - \hat{P})$. Hence, the defender will purchase insurance if $\rho_1(v^0)Ik_1(v^0) - \hat{P}k_2(v^0) > 0$. □

Since $\dot{U}(\cdot)$ is convex decreasing, then k_1 is lower than k_2. Thus, the premium must be strictly larger than the fair premium to guarantee that the defender purchases insurance. Hence, it is necessary to provide subsidies to improve the adoption of insurance (analyzed in Section 6.5). Below we illustrate the optimal investment in protection with an example.

Example 6. *We define the transition probabilities as*

$$\pi(v_d) = \left(\frac{e^1 - e^{(1-v_d)}}{e^1 - 1} + \varepsilon\right)\frac{1}{1+\varepsilon}$$

and

$$\delta(v_p) = \left(\frac{e^{(1-v_p)} - 1}{e^1 - 1} + \varepsilon\right)\frac{1}{1+\varepsilon}.$$

Thus, $\varepsilon \leq \pi(z), \delta(z) \leq 1$, *for* $z \in [0,1]$. *On the other hand, we use the following cost functions*

$$c_d(v_d) = k_d f_c(v_d) \quad \textit{and} \quad c_p(v_p) = k_p f_c(v_p),$$

where $f_c(\cdot)$ *is a concave function defined as*

$$f_c(z) = e^{z\log 2} - 1.$$

This function satisfies $f_c(0) = 0$ *and* $f_c(1) = 1$. *Furthermore, we define* $U(z) = \log z$. *In the simulations we select* $k_d = k_p = 3$, $\beta = 0.75$, $\varepsilon = 0.1$, $L = 6$, *and* $w = 746.5$. *In this case we define an insurance with indemnity* $I = 3$ *and premium*

$$P(v) = \eta I \rho_1(v), \tag{6.12}$$

with $0.5 \leq \eta \leq 2$. *We introduce the parameter* η *to observe the response of the defender to policies with costs above and below the fair premium, which corresponds to* $\eta = 1$.

Fig. 6.2 shows the optimal strategy of the defender as a function of η. *In this example, a defender with full information purchases insurance with a cost below the fair premium (*$\eta \leq 0.7$*). Furthermore, a defender with limited information accepts policies with a cost close but not superior to the fair premium (as predicted by the Theorem 4). This suggests that the insurance becomes more important as the uncertainty of the user increases, because the defender is willing to pay a higher premium. Moreover, with full information the defender invests most of its resources in detection. With limited information, when insurance is not used, the defender prioritizes prevention over detection; however, the opposite occurs when attacker has insurance.*

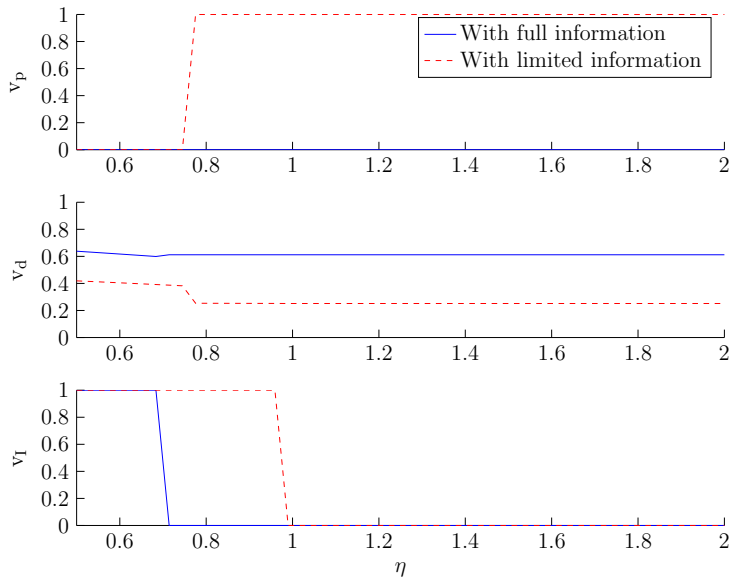

Figure 6.2: Optimal strategy with and without full information for different cost of the premium (see Eq. (6.12). With limited information the defender accepts higher premiums, which shows the importance of insurance in situations with uncertainty.

6.5 Subsidies on Indemnities

The results of the previous section show that the defender purchases insurance when the cost is lower than the fair premium. However, insurers often add a safety loading factor $\gamma > 0$ to the fair premium, so they charge

$$\tilde{P}(v) = (1+\gamma)\rho_1(v)I(v), \tag{6.13}$$

where $I(v)$ is the indemnity paid by the insurer. From Theorem 4 we know that the defender needs incentives to purchase an insurance policy with the premium in Eq. (6.13), since it is larger than the fair premium.

In this section, we show the effect of a subsidy $\Omega(v)$ granted to a defender (who has insurance) when an accident occurs. We assume that the subsidy is an indemnity that the defender receives from the government. Hence, the total indemnity of the defender becomes

$$\tilde{I}(v) = I(v) + \Omega(v). \tag{6.14}$$

With the premium in Eq. (6.13) and the indemnity in Eq. (6.14), the defender purchases insurance if (see Theorem 4)

$$\rho_1(v)\tilde{I}(v)k_1(v)/k_2(v) > \tilde{P}.$$

Assuming that $\dot{U}(z) = 1/z$, the previous expression is equivalent to

$$I(v) + \Omega(v) > (1+\gamma)I(v)\frac{a(v) + I(v) + \Omega(v) - \tilde{P}(v)}{a(v) - L - \tilde{P}(v)}.$$

Thus, the subsidy must satisfy:

$$\Omega(v) > I(v)\frac{(1+\gamma)r(v) - b(v)}{b(v) - (1+\gamma)I(v)}, \tag{6.15}$$

where

$$r(v) = a(v) + I(v) - \tilde{P}(v)$$

and

$$b(v) = a(v) - L - \tilde{P}(v).$$

Example 7. *In this experiment we use a loading factor $\gamma = 0.5$ and for simplicity we consider subsidies that satisfy Eq. (6.15) with equality. Moreover, we consider two alternatives for the insurance's indemnity $I(v)$, namely full and limited coverage. With full coverage the indemnity covers the losses accumulated during the attack. Thus, if the system stays in s_0 during m time periods, then the total losses are $L \cdot m$. For simplicity we define the indemnity with unlimited coverage as*

$$I(v) = L/\pi(v_d),$$

where $1/\pi(v_d)$ is the expected number of periods that the system remains in the state s_0 (i.e., the expected duration of an attack). On the other hand, policies with limited coverage offer a constant indemnity, in this case, we select $I(v) = 3$.

Fig. 6.3 shows the optimal defense strategy with subsidies. As we can see, the strategy is different with full and limited coverage. On one hand, subsidies with limited coverage improve the investment in detection and guarantee that the defender purchases insurance policies with a loading factor γ. On the other hand, when the insurer offers full coverage, the defender purchases insurance policies regardless of the cost (η); however, the defender also invests less resources in detection and prevention. This occurs because subsidies with full coverage create a problem of moral hazard, in which the defender does not has incentives invest in the protection of the system. Fig. 6.3b shows that the subsidies granted with full coverage are at least 10^{10} times the subsidies granted with limited coverage.

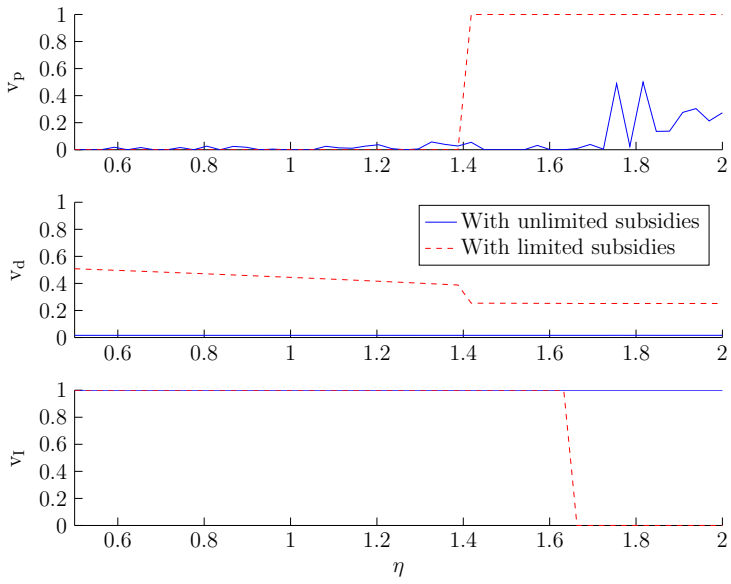

(a) Optimal defense strategy.

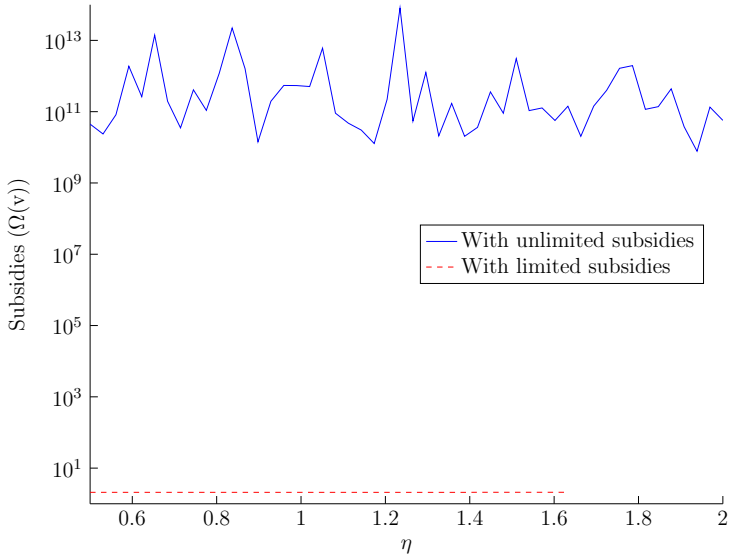

(b) Subsidies granted to the defender.

Figure 6.3: Defense strategy with subsidies and both full and limited coverage for different cost premiums. Full coverage improves the adoption of insurance; however, the defender looses incentives to invest in protection.

6.6 Conclusions

We model the dynamic interaction between a defender and an attacker as a Markov process and analyze the best defense strategy with asymmetric information. With limited information the defender accepts insurance policies close to the fair premium, contrary to the case with full information. Therefore, uncertainty makes insurance more attractive.

The introduction of insurance has the potential to improve the investment that firms make in security. In particular, we find that the investments in detection improves with low premiums, whereas investments in prevention increase when the premium exceeds a threshold.

Since the defender does not accept fair premiums, we propose a subsidy scheme to incentive the adoption of insurance. We find that policies with unlimited indemnities create perverse incentives that stimulate low investments in security. Therefore, incentives for insurance have to be designed with care not to have negative consequences. We observe an example in which policies with limited indemnities prevent the perverse incentives.

CHAPTER 7

CONCLUSION

In this research we investigate the effect that economic policies have in the security of critical infrastructures, in particular, the power grid. We refer to economic policies as mechanisms to allocate resources, such as contractual rules, markets, and investment strategies. We hypothesized that that these schemes affect the security of CPSs, since they impact the profit and difficulty to implement attacks.

We found that economic policies can affect the security of systems in many ways. First, in Chapter 3 we learned that parties involved in the protection of systems can leverage the anonymity of attacks to profit sponsoring attacks. Nevertheless, an appropriate design of the protection policies can prevent such situations. Second, in Chapter 4 we found that the structure of power markets can affect both attackers and defenders. In particular, with distributed systems we can increase the difficulty to implement attacks; however, it becomes more difficult to detect and penalize attacks. Third, in Chapter 5 we found that, the defender can prevent attacks event if he cannot penalize directly his adversaries. Furthermore, in Chapter 6 we found that uncertainties make insurance more attractive; however, the defender accepts insurance with cost lower than the fair premium. Also, insurance with subsidies and full coverage can create perverse incentives that reduce investments in security.

From the cases analyzed we conclude that it is possible to improve the security of systems through a careful design of economic policies. In particular, the the defender can prevent attacks having into account the motivation and limitations of attackers. For instance, the defender can prevent attacks by depleting the resources of its adversaries, without punishing them directly. The models analyzed in this work capture general security issues and we believe that our work can help to improve the security of CPSs from different types of adversaries. The following are some areas for further research, necessary to apply our ideas in practical scenarios.

7.1 Future Research

7.1.1 Optimal Design of Contracts

In the model of Chapter 3 a defender can design contractual policies minimizing the total costs. For example, the defender can choose the number of companies and the lotteries to assign contracts balancing the increasing costs of repair services and operational losses from

attacks. If having a large pool of contractors increases excessively the repair costs, then it would be convenient to have few companies and allow some attacks. Also, the contracts should contemplate the possibility that multiple contractors can cooperate to defraud the electric company.

7.1.2 Improve Detection of Attacks

In Chapter 4 we show how to detect attacks analyzing the market's equilibrium. It is also important to design schemes to detect malicious attacks, which cause peaks in demand. In such cases the detection would need real time monitoring, because the attackers leverage the physical dynamics of the system (rather than its equilibrium) to cause damage (Barreto et al., 2014, 2013). Moreover, it would be interesting to explore how to improve the detection of attacks using honeypots and intrusion detection systems.

7.1.3 Regulation on Security Protection

Uncertainties about the system's state have a negative impact on the security investments (see Section 5.5). For example, in our experiments a defender with asymmetric information invests less resources in detection (in comparison with the case with full information). A regulator would design economic incentives to mitigate these negative behaviors.

7.1.4 Risk Estimation

Most works on risk management assume that risks can be estimated precisely, which implies knowledge of the likelihood of attacks and their impact. However, estimating risks in practice could be unfeasible due to the lack of information and the complex interrelations among firms (Johnson et al., 2014). Furthermore, statistical properties change upon observation, because rational individuals react to the acquired information. Hence, we need verify the precision of risk estimations (or develop better risk measures[1]) and design risk management strategies that consider uncertainties in the risk (Danıelsson, 2002; Danielsson, 2008). It might be interesting to investigate how to use security audits to estimate the evolving risk of companies. Particularly, an insurer should find a balance between the frequency of the audits (which have some cost) and the accuracy of the risk estimations needed to offer policies with some degree of confidence.

[1](Danıelsson, 2002) shows that risk measures with *value at risk* (VaR) lack robustness and volatility, that is, forecasts are inaccurate and fluctuate between time periods.

APPENDIX

PROOFS ON THE OPTIMAL INVESTMENT IN PROTECTION

A.1 Attacker's Optimal Strategy

Before finding the attacker's optimal strategy we need the following results:

Lemma 4. *The attacker's cost functional is a contraction mapping.*

Proof. Let us define the operator

$$T(u)(x) = \inf_{v_A} \left\{ l_A(x, v_A) + \beta \mathbb{E}_x^{v_A, v_D} \{u(x')\} \right\}.$$

We can find an upper bound of $T(u)(x)$, called $\hat{T}(u)(x)$, using some suboptimal action $\hat{v}$:

$$T(u)(x) \leq \hat{T}(u)(x) = l_A(x, \hat{v}) + \beta \mathbb{E}_x^{\hat{v}, v_D} \{u(x')\}.$$

Now, let us consider two cost functions $u_1(\cdot)$ and $u_2(\cdot)$ with optimal actions $v_1(\cdot)$ and $v_2(\cdot)$, respectively. Now, the difference of the operator evaluated in each cost function is upper bounded by

$$|T(u_1)(x) - T(u_2)(x)| \leq |\hat{T}(u_1)(x) - T(u_2)(x)|.$$

If we select $\hat{v} = v_2$, then the previous expression with $x = S_0$ becomes

$$|T(u_1)(S_0) - T(u_2)(S_0)| \leq \beta |\pi(v_a, v_d)(u_1(S_1) - u_2(S_1)) + (1 - \pi(v_a, v_d))(u_1(S_0) - u_2(S_0))|,$$

which has the following upper bound

$$|T(u_1)(S_0) - T(u_2)(S_0)| \leq \beta \pi(v_a, v_d) |(u_1(S_1) - u_2(S_1)| + \beta(1 - \pi(v_a, v_d)) |u_1(S_0) - u_2(S_0)|$$

Finally, since $0 \leq \pi(v_a, v_d) \leq 1$, we have

$$|T(u_1)(S_0) - T(u_2)(S_0)| \leq \beta \max\{|u_1(S_0) - u_2(S_0)|, |u_1(S_1) - u_2(S_1)|\}.$$

Similarly,

$$|T(u_1)(S_1) - T(u_2)(S_1)| \leq \beta \max\{|u_1(S_0) - u_2(S_0)|, |u_1(S_1) - u_2(S_1)|\}.$$

Thus,

$$||T(u_1)(x) - T(u_2)(x)||_\infty \leq \beta ||u_1 - u_2||_\infty$$

Consequently, the cost functional is a contraction mapping. □

We use an iterative approximation of the attacker's value function $u_A(x)$ (see Eq. (5.1)) to find the attacker's optimal strategy. We define the approximation of the value function as

$$u_{n+1}(x) = \inf_{v_n \in [0,1]} \{ l_A(x, v_n) + \beta \mathbb{E}_x^{v_n, v_D} \{ u_n(x) \} \}, \tag{A.1}$$

with $u_n(x) \to u(x) = u_A(x)$, $v_n(x) \to v_A$ for all $x \in S$, where $u_0(x) = 0$ and $v_n(x) = (v_{a,n}(x), v_{h,n}(x))$. The approximations satisfy the following property:

$$0 \geq u_1(x) \geq u_2(x) \geq \ldots \geq u_A(x).$$

This iterative approximation is possible because $u_A(x)$ is a contraction mapping, which allows us to use the Banach fixed point theorem (Banach, 1922).

The following result shows that the attacker's value function is greater in the secure state.

Lemma 5. *The iterative cost function at time k satisfies*

$$u_k(S_0) \leq u_k(S_1)$$

Proof. Let us assume by contradiction that $u_n(S_0) > u_n(S_1)$. This implies that there exists some time k such that $u_k(S_0) \leq u_k(S_1)$ and

$$u_{k+1}(S_0) > u_{k+1}(S_1). \tag{A.2}$$

In particular, we can adjust $u_0(x)$ to guarantee that $u_k(S_0) = u_k(S_1)$. Thus, from Eqs. (5.2), (5.3), (5.4), (5.5), and (A.1) we know that

$$u_{k+1}(S_0) = \inf_v \{ C_0 \mathbb{1}_{v>0} - g_a(v) + \beta u_k(S_0) \} \leq \beta u_k(S_0)$$

and

$$u_{k+1}(S_1) = \inf_v \{ C_v \mathbb{1}_{v>0} + \beta u_k(S_1) \} = \beta u_k(S_1).$$

Hence, $u_{k+1}(S_0) \leq u_{k+1}(S_1)$, which contradicts our initial hypothesis in Eq. (A.2), which proves that $u_k(S_0) \leq u_k(S_1)$. □

The following result shows that if an attack is profitable at the k^{th} iteration, then it will remain profitable in the next iterations.

Lemma 6. *If $v_{a,k} > 0$ (or $v_{h,k} > 0$) for some time $k \geq 1$, then $v_{a,k'} > 0$ (or $v_{h,k'} > 0$) for all $k' > k$.*

Proof. Since $u_n \to u$ as $n \to \infty$, $u_n(x) \geq u(x)$, and $u_0(x) = 0$, we have

$$0 \geq u_{k-1} \geq u_k \geq u_{k+1}. \tag{A.3}$$

Let us assume that $v_{(a,h),k} > 0$. This implies that attacks are profitable, that is, $0 > u_k(x)$. Now, if we select in the next time period $v_{a,k+1} = 0$ or $v_{h,k+1} = 0$ results

$$u_{k+1}(x) = \beta u_k(x).$$

However, since $u_k(x)$ is negative we have $u_k(x) < u_{k+1}(x)$, which contradicts Eq. (A.3). □

In the following result we show that the attacker's optimal strategy consists in attacking with full intensity or not attacking at all.

Lemma 7. *The strategy of the attacker is always either 1 or 0.*

Proof. Let us show first that $v_{h,k} \in \{0, 1\}$. First, if $v_{h,k} > 0$ for some k, then from Eq. (A.1) the optimal cost is

$$u_k(S_1) = \min_v \left\{ C_v + \beta u_{k-1}(S_1) + \beta\delta(v, v_p)(u_{k-1}(S_0) - u_{k-1}(S_1)) \right\}.$$

From Lemma 5 we know that $u_k(S_0) \leq u_k(S_1)$, thus, the difference of $u_{k-1}(S_0)$ and $u_{k-1}(S_1)$ is negative. Since $\delta(\cdot)$ is increasing in v_h, then the optimal attack is $v_{h,k} = 1$.

Now we will show that $v_{a,k} \in \{0, 1\}$. Observe that if $v_{a,k} > 0$ for some k, then the optimal cost is

$$u_k(s_0) = \min_v \left\{ \Psi_{S_0,k}(v, v_d) \right\},$$

where

$$\Psi_{S_0,k}(v, v_d) = C_0 - g_a(v) + \beta u_{k-1}(S_1) + \beta(1 - \pi(v, v_d))u_{k-1}(S_0).$$

Since $g_a(\cdot)$ and $(1 - \pi(v, v_d))$ are convex with respect to v_a, then $\Psi_{S_0,k}(v, v_d)$ is concave with its minimum at either 0 or 1 (recall that $u_{k-1}(S_0) < 0$). Let us consider the value of $\Psi_{S_0,k}(v, v_d)$ at the extremes:

$$\Psi_{S_0,k}(0, v_d) = C_0 + \beta u_{k-1}(S_0)$$

and

$$\Psi_{S_0,k}(1, vd) = C_0 - g_a(1) + \beta u_k(S_1)$$

From Eq. (A.1) we can extract the following properties

$$u_k(S_0) \geq C_0 - g_a(1) + \beta u_{k-1}(S_0)$$

and

$$u_k(S_1) \leq \beta u_{k-1}(S_0).$$

Hence,

$$\Psi_{S_0,k}(0, vd) \geq C_0(1+\beta) - \beta g_a(1) + \beta^2 u_{k-1}(S_0)) \tag{A.4}$$

and

$$C_0 - g_a(1) + \beta^2 u_{k-1}(S_0) \geq \Psi_{S_0,k}(1, v_d) \tag{A.5}$$

From Eqs. (A.4) and (A.5) we know that $\Psi_{S_0,k}(0, v_d)$ is greater than $\Psi_{S_0,k}(1, v_d)$. Hence, the optimal strategy is $v_{a,k} = 1$, for every $k \geq 1$, since $C_0 \geq 0$. □

Now we are ready to show that optimal strategy of the attacker.

Proof of Theorem 1. We can prove that the attacker's optimal cost functional is a contraction mapping. This allows us to use the Banach fixed point theorem to find iteratively the optimal cost

$$u_{n+1}(x) = \inf_{v_n \in [0,1]} \{l_A(x, v_n) + \beta \mathbb{E}_x^{v_n, v_D}\{u_n(x')\}\}$$

with $u_n(x) \to u_A(x)$ and $u_0(x) = 0$ for all $x \in S$.

We use this property to find the optimal strategy of the attacker. First let us examine the cost functional and the optimal strategy for $n = 1$:

$$u_1(S_0) = \inf_{v \in [0,1]} \{-g_a(v) + C_0 \mathbb{1}_{v>0}\},$$

Since $g_a(\cdot)$ is increasing we have

$$u_1(S_0) = \min\{0, K_1\},$$

where $K_1 = C_0 - g_a(1)$. On the other hand,

$$u_1(S_1) = \inf_{v \in [0,1]} \{C_v \mathbb{1}_{v>0}\} = 0.$$

In summary, the optimal strategies are

$$v_{a,1} = \begin{cases} 1 & \text{if } K_1 < 0 \\ 0 & \text{otherwise} \end{cases}; \quad v_{h,1} = 0.$$

Note that if $K_1 > 0$, then attacks are not profitable and the cost function converges, that is, $u_0(x) = u_1(x) = u_A(x)$. This means that $v_{a,1}$ and $v_{h,1}$ are the optimal strategies. Intuitively,

the attacker won't invest in vulnerabilities if the attacks are too expensive. Therefore, $v_a = 0$ and $v_h = 0$, which proves the first part of the theorem.

Now, if $K_1 < 0$, then the optimal action is $v_{a,1} = 1$ and $v_{h,1} = 0$. From Lemmas 6 and 7 we know that $v_{a,k} = 1$ for $k > 1$. We are interested in finding the conditions that make $v_a = 1$ and $v_h = 0$ the optimal strategy. From Lemma 6 we know that if $v_h = 0$, then $v_{h,k} = 0$ for all time $k \geq 1$. Observe that if $v_{h,k} = 0$, then $u_k(S_1) = 0$ for all $k \geq 1$. This fact can be used to compute $u_n(S_0)$. First, let us suppose that $v_a = 1$ and $v_h = 0$. Then, $u_1(S_0) = K_1$ and the optimal cost for $n = 2$ is

$$u_2(S_0) = K_1 + K_1\beta(1 - \pi(1, v_d)) = K_2.$$

For $n = 3$ we have

$$u_3(S_0) = K_1 + K_2\beta(1 - \pi(1, v_d)) = K_3.$$

We can generalize the cost as

$$u_n(S_0) = K_1 + K_{n-1}\beta(1 - \pi(1, v_d)) = K_n.$$

This expression can be rewritten as

$$u_n(S_0) = K_1 \sum_{i=0}^{n-1} \beta^i (1 - \pi(1, v_d))^i.$$

This cost function converges to

$$u(S_0) = \lim_{n\to\infty} u_n(S_0) = \frac{K_1}{1 - \beta(1 - \pi(1, v_d))}. \tag{A.6}$$

Now, let us consider the conditions to guarantee $v_{h,n} = 0$. The cost functional at time n is

$$u_n(S_1) = \inf_{v\in[0,1]} \{C_v \mathbb{1}_{v>0} + u_{n-1}(S_0)\beta\delta(v, v_p)\}.$$

Note that the objective function is convex in $(0, 1]$. Hence, the minimum cost is at some extreme value. Here, the optimal strategy is

$$v_{h,n} = \begin{cases} 1 & \text{if } C_v + u_{n-1}(S_0)\beta\delta(1, v_p) < 0, \\ 0 & \text{otherwise.} \end{cases}$$

Hence, hacking the system is unprofitable if

$$C_v + u_{n-1}(S_0)\beta\delta(1, v_p) > 0, \tag{A.7}$$

for all $n > 1$. Since $u_{n-1}(S_0) \geq u_n(S_0) \geq u_A(S_0)$, we know that Eq. (A.7) is satisfied for all $n > 1$ if it is satisfied when $n \to \infty$. Thus, we can use Eq. (A.6) to show that $v_h = 0$ if

$$C_v + \beta \frac{C_0 - g_a(1)}{1 - \beta(1 - \pi(1, v_d))} \delta(1, v_p) > 0. \tag{A.8}$$

Observe that if Eq. (A.8) is not satisfied, there is some k such that $v_{h,k} = 1$, which implies that $v_h = 1$, □

A.2 Defender's Cost Function with Full Information

Proof of Theorem 2. We can prove that the defender's cost functional is a contraction mapping using the same argument used to prove the Lemma 4. Hence, as in the attacker's case, we can use an iterative approximation of the cost function to extract some properties of the defender's optimal strategy. However, due to the complexity of the problem, we focus on finding the defender's cost function $J^D(x_0, v_A, v_D)$, rather than the defender's optimal strategy.

Let us consider iteration function at state S_0

$$u^D_{n+1}(S_0) = g_d(v_a) + \inf_{v_d, v_p} \left\{ C_p(v_p) + C_d(v_d) + \beta u^D_n(S_0) + \beta \pi(v_a, v_d)(u_n(S_1) - u^D_n(S_0)) \right\}.$$

Observe that in the state S_0 it is optimal to set $v_p = 0$ (i.e., $v_p(S_0) = 0$). On the other hand, at state S_1 we have

$$u^D_{n+1}(S_1) = \inf_{v_d, v_p} \left\{ C_p(v_p) + C_d(v_d) + \beta u^D_n(S_1) + \beta \delta(v_h, v_p)(u^D_n(S_0) - u^D_n(S_1)) \right\}.$$

In this case the optimal strategy is to set $v_d = 0$ (i.e., $v_d(S_1) = 0$).

We can use the fact that $v_p(S_0) = 0$ and $v_d(S_1) = 0$ to find the cost function $J^D(x, v_A, v_D)$ through iterations. Thus, assuming that $J^D_0(S_0, v_A, v_D) = 0$ and $J^D_0(S_1, v_A, v_D) = 0$ we have for $n = 1$

$$J^D_1(S_0, v_A, v_D) = g_d(v_a) + C_d(v_d) = Q(v_d)$$

and

$$J^D_1(S_1, v_A, v_D) = C_p(v_p) = W(v_p)$$

Now, for $n = 2$ we have

$$J^D_2(S_0, v_A, v_D) = Q(v_d)(1 + \beta) + \beta \pi(v_a, v_d)(W(v_p) - Q(v_d)).$$

and

$$J_2^D(S_1, v_A, v_D) = W(v_p)(1+\beta) + \beta\delta(v_h, v_p)(Q(v_d) - W(v_p)).$$

For $n = 3$ we have

$$J_3^D(S_0, v_A, v_D) = Q(v_d)(1+\beta+\beta^2) + \beta\pi(v_a, v_d)(W(v_p) - Q(v_d))r_3$$

and

$$J_3^D(S_1, v_A, v_D) = W(v_p)(1+\beta) + \beta\delta(v_h, v_p)(Q(v_d) - W(v_p))r_3,$$

where

$$r_3 = (1 + \beta + \beta(1 - \pi(v_a, v_d) - \delta(v_h, v_p))).$$

We can generalize the cost function in the nth iteration as

$$J_n^D(S_0, v_A, v_D) = Q(v_d)\sum_{i=0}^{n-1}\beta^i + \beta\pi(v_a, v_d)(W(v_p) - Q(v_d))r_n$$

and

$$J_n^D(S_1, v_A, v_D) = W(v_p)\sum_{i=0}^{n-1}\beta^i + \beta\delta(v_h, v_p)(Q(v_d) - W(v_p))r_n$$

where

$$r_n = \sum_{i=0}^{n-2}\beta^i + \beta r_{n-1}(1 - \pi(v_a, v_d) - \delta(v_h, v_p)), \tag{A.9}$$

with $r_1 = 0$ and $r_2 = 1$.

Recall that the sequence $J_n^D(x, v_A, v_D)$ converges to $J^D(x, v_A, v_D)$. This implies that the the factor r_n converges to an stationary value r, which can be found evaluating Eq. (A.9) when $n \to \infty$:

$$\lim_{n\to\infty} r_n = r = \frac{1}{1-\beta} + \beta r(1 - \pi(v_a, v_d) - \delta(v_h, v_p)).$$

From this expression we know that

$$r = \frac{1}{1-\beta}\frac{1}{1 + \beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)}$$

Now we are ready to find the cost function when $n \to \infty$:

$$J^D(S_0, v_A, v_D) = \frac{Q(v_d)}{1-\beta} + \frac{\beta}{1-\beta}\frac{\pi(v_a, v_d)(W(v_p) - Q(v_d))}{1 + \beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)}$$

and

$$J^D(S_1, v_A, v_D) = \frac{W(v_p)}{1-\beta} + \frac{\beta}{1-\beta}\frac{\delta(v_h, v_p)(Q(v_d) - W(v_p))}{1 + \beta(\pi(v_a, v_d) + \delta(v_h, v_p) - 1)}.$$

□

A.3 Defender's Cost Function with Asymmetric Information

Proof of Theorem 3. We can generalize the cost function in Eq. (5.7) as

$$\hat{J}_n^D(v_A, v_D) = g_d(v_a)\gamma_n + (C_d(v_d) + C_p(v_p)) \sum_{i=0}^{n} \beta^i,$$

where

$$\gamma_n = \mathbb{P}(x_n = S_0) + \beta\gamma_{n-1} \tag{A.10}$$

with $\gamma_0 = \mathbb{P}(x_0 = S_0)$. For simplicity let us define $\pi = \pi(v_a, v_d)$ and $\delta = \delta(v_h, v_p)$. Thus,

$$\mathbb{P}(x_1 = S_0) = \mathbb{P}(x_0 = S_0)(1 - \pi) + \mathbb{P}(x_0 = S_1)\delta,$$

and replacing $\mathbb{P}(x_0 = S_1) = 1 - \mathbb{P}(x_0 = S_0)$ we obtain

$$\mathbb{P}(x_1 = S_0) = \mathbb{P}(x_0 = S_0)(1 - \pi - \delta) + \delta = p(1 - \pi - \delta) + \delta.$$

We can generalize the probability that the system's state at time n is equal to S_0 as

$$\mathbb{P}(x_n = S_0) = p(1 - \pi - \delta)^n + \delta \sum_{i=0}^{n-1} (1 - \pi - \delta)^i.$$

If the cost function converges when $n \to \infty$, then the parameter γ_n converges to γ, that is,

$$\lim_{n\to\infty} \gamma_n = \gamma \tag{A.11}$$

Replacing (A.10) into (A.11) we obtain

$$\gamma = \lim_{n\to\infty} \mathbb{P}(x_n = S_0) + \beta\gamma.$$

This equation can be rewritten as

$$\gamma = \frac{1}{1-\beta} \lim_{n\to\infty} \mathbb{P}(x_n = S_0) = \frac{1}{1-\beta} \lim_{n\to\infty} \left\{ p(1 - \pi - \delta)^n + \delta \sum_{i=0}^{n-1} (1 - \pi - \delta)^i \right\}.$$

Observe that $0 \leq \pi + \delta \leq 2$. Here we need to consider three cases: i) if $\pi = 0$ and $\delta = 0$ then

$$\gamma = \frac{p}{1-\beta}.$$

ii) if $\pi = 1$ and $\delta = 1$, then

$$\gamma = \frac{1}{1-\beta} \lim_{n\to\infty} \left\{ p(-1)^n + \sum_{i=0}^{n-1} (-1)^i \right\}.$$

In this case γ converges only if $p = 1/2$, resulting

$$\gamma = \frac{p}{1-\beta}.$$

iii) if $0 < \pi + \delta < 2$ then

$$\gamma = \frac{1}{1-\beta} \lim_{n\to\infty} \left\{ p(1-\pi-\delta)^n + \delta \frac{1-(1-\pi-\delta)^n}{\pi+\delta} \right\}.$$

which is equal to

$$\gamma = \frac{1}{1-\beta} \frac{\delta}{\pi+\delta}.$$

Thus, if $p = 1/2$ then

$$\hat{J}^D(v_A, v_D) = g_d(v_a)\gamma + \frac{C_d(v_d) + C_p(v_p)}{1-\beta},$$

where

$$\gamma = \begin{cases} \frac{1}{1-\beta}\frac{\delta}{\pi+\delta}, & \text{if } 0 < \pi + \delta < 2, \\ \frac{p}{1-\beta}, & \text{otherwise,} \end{cases}$$

with $\pi = \pi(v_a, v_d)$ and $\delta = \delta(v_h, v_p)$. □

REFERENCES

Ablon, L., M. C. Libicki, and A. A. Golay (2014). Markets for cybercrime tools and stolen data: Hackers' bazaar. Technical report, Rand Corporation.

Alpcan, T. and T. Basar (2006). An intrusion detection game with limited observations. In *12th Int. Symp. on Dynamic Games and Applications, Sophia Antipolis, France*, Volume 26.

Anderson, R. (2001). Why information security is hard, an economic perspective. In *Proceedings 17th annual Computer Security Applications Conference, ACSAC 2001*, New York, NY, USA, pp. 358–365. ACM.

Arrow, K. J. (1971). Insurance, risk and resource allocation. *Essays in the theory of risk-bearing*, 134–143.

Axelrod, R. and R. Iliev (2014). Timing of cyber conflict. *Proceedings of the National Academy of Sciences 111*(4), 1298–1303.

Banach, S. (1922). Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fund. Math 3*(1), 133–181.

Barreto, C. (2014). Population dynamics toolbox (pdtoolbox). `https://github.com/carlobar/PDToolbox_matlab`. Accessed: April 10, 2018.

Barreto, C. and A. A. Cárdenas (2015a, Dec). Detecting fraud in demand response programs. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 5209–5214. IEEE.

Barreto, C. and A. A. Cárdenas (2015b, July). Incentives for demand-response programs with nonlinear, piece-wise continuous electricity cost functions. In *2015 American Control Conference (ACC)*, pp. 4327–4332. IEEE.

Barreto, C. and A. A. Cárdenas (2016). Perverse incentives in security contracts: A case study in the colombian power grid. In *the Annual Workshop on the Economics of Information Security (WEIS)*.

Barreto, C. and A. A. Cárdenas (2017, aug). Optimal risk management in critical infrastructures against cyber-adversaries. In *2017 1st IEEE Conference on Control Technology and Applications (CCTA)*. IEEE.

Barreto, C., A. A. Cardenas, and A. Bensoussan (2017, Apr). Optimal security investments in a prevention and detection game. In *Proceedings of the Hot Topics in Science of Security: Symposium and Bootcamp*, HoTSoS, New York, NY, USA, pp. 24–34. ACM.

Barreto, C., A. A. Cárdenas, and N. Quijano (2013). Controllability of dynamical systems: Threat models and reactive security. In *4th International Conference on Decision and Game Theory for Security - Volume 8252*, GameSec 2013, New York, NY, USA, pp. 45–64. Springer-Verlag New York, Inc.

Barreto, C., A. A. Cárdenas, N. Quijano, and E. Mojica-Nava (2014). CPS: Market Analysis of Attacks Against Demand Response in the Smart Grid. In *Proceedings of the 30th Annual Computer Security Applications Conference*, ACSAC '14, New York, NY, USA, pp. 136–145. ACM.

Barreto, C., J. Giraldo, A. A. Cárdenas, E. Mojica-Nava, and N. Quijano (2014, Nov). Control systems for the power grid and their resiliency to attacks. *IEEE Security Privacy 12*(6), 15–23.

Barreto, C., E. Mojica-Nava, and N. Quijano (2013, Dec). Design of mechanisms for demand response programs. In *52nd IEEE Conference on Decision and Control*, pp. 1828–1833. IEEE.

Batz, D., J. Brenton, D. Dunn, G. Williams, P. Clark, S. Elwart, E. Goff, B. Harrell, C. Hawk, M. Henrie, et al. (2011). Roadmap to achieve energy delivery systems cybersecurity. Technical report, Energy Sector Control Systems Working Group (ESCSWG), Washington, DC.

Beck, U. (1992). *Risk society: Towards a new modernity*, Volume 17. Sage.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy 76*(2), 169–217.

Bejtlich, R. (2013). *The practice of network security monitoring: understanding incident detection and response*. No Starch Press.

Bensoussan, A. (2011). *Dynamic programming and inventory control*, Volume 3 of *Studies in Probability, Optimization and Statistics*. IOS Press.

Bernstein, P. L. (1996). *Against the gods: The remarkable story of risk*. Wiley New York.

Biener, C., M. Eling, and J. H. Wirfs (2015). Insurability of cyber risk: an empirical analysis. *The Geneva Papers on Risk and Insurance Issues and Practice 40*(1), 131–158.

Böhme, R. and G. Schwartz (2010). Modeling cyber-insurance: Towards a unifying framework. In *the Annual Workshop on the Economics of Information Security (WEIS)*.

Boyd, S. and L. Vandenberghe (2004). *Convex optimization*. Cambridge university press.

Cambridge Centre for Risk Studies (2015). Business blackout: The insurance implications of a cyber attack on the us power grid. Technical report, Lloyd's.

Caracol radio (2009). Capturan a funcionarios de una empresa que atentaba contra las torres eléctricas de ISA. http://caracol.com.co/radio/2009/06/15/judicial/1245050340_829038.html. Accessed February 1, 2016.

Cárdenas, A. A., S. Radosavac, J. Grossklags, J. Chuang, and C. J. Hoofnagle (2009). An economic map of cybercrime. In *The 37th Research Conference on Communication, Information and Internet Policy (TPRC)*.

Caskey, J. and N. B. Ozel (2017). Earnings expectations and employee safety. *Journal of Accounting and Economics 63*(1), 121–141.

Chen, L., N. Li, S. H. Low, and J. C. Doyle (2010). Two market models for demand response in power networks. In *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 397–402. IEEE.

Cherepanov, A. (2017). Telebots are back: Supply-chain attacks against ukraine. https://www.welivesecurity.com/2017/06/30/telebots-back-supply-chain-attacks-against-ukraine/. Accessed October 10, 2017.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice 2*, 19–33.

Congreso de Colombia (1993). Ley 80 de 1993. http://www2.igac.gov.co/igac_web/UserFiles/File/web%202008%20/ley%2080-93.pdf. Accessed February 1, 2016.

Congreso de Colombia (2013). Decreto 1510 de 2013. http://www.alcaldiabogota.gov.co/sisjur/normas/Norma1.jsp?i=53776#163. Accessed February 1, 2016.

Danıelsson, J. (2002). The emperor has no clothes: Limits to risk modelling. *Journal of Banking & Finance 26*(7), 1273–1296.

Danielsson, J. (2008). Blame the models. *Journal of Financial Stability 4*(4), 321–328.

Ehrlich, I. and G. S. Becker (1972). Market insurance, self-insurance, and self-protection. *Journal of political Economy 80*(4), 623–648.

El Tiempo (2000). ELN retiene obreros de torres. http://www.eltiempo.com/archivo/documento/MAM-1305561. Accessed September 22, 2016.

Eling, M. and J. H. Wirfs (2016). Cyber risk: Too big to insure? risk transfer options for a mercurial risk class. Technical report, Institute of Insurance Economics.

Fahrioglu, M. and F. L. Alvarado (1998, jan). Designing cost effective demand management contracts using game theory. In *Proceedings of the IEEE Power Engineering Society Winter Meeting*, Volume 1, pp. 427–432. IEEE.

Finkle, J. (2017). Hackers halt plant operations in watershed cyber attack. https://www.reuters.com/article/us-cyber-infrastructure-attack/hackers-halt-plant-operations-in-watershed-cyber-attack-idUSKBN1E8271. Accessed April 16, 2018.

Fu, K. (2016). Infrastructure disruption: Internet of things security. http://docs.house.gov/meetings/IF/IF17/20161116/105418/HHRG-114-IF17-Wstate-FuK-20161116.pdf. Accessed October 24, 2017.

Gellings, C. W. (2009). *The smart grid: enabling energy efficiency and demand response*. The Fairmont Press, Inc.

Gettys, J. (2018). Mythology about security. https://gettys.wordpress.com/2018/04/09/mythology-about-security/. Accessed April 26, 2018.

Greenberg, A. (2015). Hackers remotely kill a jeep on the highway—with me in it. https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/. Accessed January 24, 2018.

Greenberg, A. (2017). Petya ransomware epidemic may be spillover from cyberwar. https://www.wired.com/story/petya-ransomware-ukraine. Accessed October 10, 2017.

Greene, K. (2006). Catching cyber criminals. https://www.technologyreview.com/s/405467/catching-cyber-criminals/. Accessed May 19, 2017.

Groves, T. (1973). Incentives in Teams. *Econometrica 41*, 617–631.

Hardin, G. (1968, December). The Tragedy of the Commons. *Science 162*(3859), 1243–1248.

Harrington, W. and R. Morgenstern (2007). Economic incentives versus command and control: What's the best approach for solving environmental problems? *Acid in the Environment*, 233–240.

Hernández-Lerma, O. and J. B. Lasserre (2012). *Discrete-time Markov control processes: basic optimality criteria*, Volume 30. Springer Science & Business Media.

Hofbauer, J. and K. Sigmund. (1998). *Evolutionary Games and Population Dynamics*. Cambridge UK: Cambridge University Press.

Huang, L., J. Walrand, and K. Ramchandran (2012). Optimal smart grid tariffs. In *Information Theory and Applications Workshop (ITA), 2012*, pp. 212–220.

Ibars, C., M. Navarro, and L. Giupponi (2010, October). Distributed demand management in smart grid with a congestion game. In *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 495–500.

Interconexión Eléctrica S.A. E.S.P. (ISA) (2018). Informes empresariales. http://www.isa.co/es/sala-de-prensa/Documents/Forms/AllItems.aspx?RootFolder=%2Fes%2Fsala-de-prensa%2FDocuments%2Fnuestra-compania%2Finformes-empresariales. Accessed April 10, 2018.

Johnson, B., A. Laszka, and J. Grossklags (2014). The complexity of estimating systematic risk in networks. In *2014 IEEE 27th Computer Security Foundations Symposium (CSF)*, pp. 325–336. IEEE.

Karpesky Lab (2014). Cybercrime, inc.: how profitable is the business? https://www.kaspersky.com/blog/cybercrime-inc-how-profitable-is-the-business/15034/. Accessed October 7, 2016.

Karush, W. (1939). Minima of functions of several variables with inequalities as side constraints. *M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago*.

Kelly, S., E. Leverett, E. J. Oughton, J. Copic, S. Thacker, R. Pant, L. Pryor, G. Kassara, T. Evan, S. J. Ruffle, M. Tuveson, A. W. Coburn, D. Ralph, and J. W. Hall (2016). Integrated infrastructure: Cyber resiliency in society, mapping the consequences of an interconnected digital economy. Technical report, Centre for Risk Studies, University of Cambridge.

Koppel, T. (2016). *Lights out: a cyberattack, a nation unprepared, surviving the aftermath*. Broadway Books.

Krebs, B. (2014). Target hackers broke in via hvac company. https://krebsonsecurity.com/2014/02/target-hackers-broke-in-via-hvac-company. Accessed May 18, 2017.

Krebs, B. (2016). Who makes the iot things under attack? https://krebsonsecurity.com/2016/10/who-makes-the-iot-things-under-attack/. Accessed October 7, 2016.

Krebs, B. (2017). Who is Anna-Senpai, the Mirai Worm Author? https://krebsonsecurity.com/2017/01/who-is-anna-senpai-the-mirai-worm-author/. Accessed May 19, 2017.

Kuhn, H. and A. Tucker (1951). Nonlinear programming. In *Proceedings of 2nd Berkeley Symposium*, pp. 481–492. University of California Press.

Kunreuther, H. (2015). The role of insurance in reducing losses from extreme events: The need for public–private partnerships. *The Geneva Papers on Risk and Insurance Issues and Practice 40*(4), 741–762.

Laffont, J. J. and J. Tirole (1993). *A theory of incentives in procurement and regulation*. MIT press.

Langner, R. and P. Pederson (2013). Bound to fail: Why cyber security risk cannot simply be "managed" away. Technical report, Brookings.

Laszka, A., M. Felegyhazi, and L. Buttyan (2014, August). A survey of interdependent information security games. *ACM Comput. Surv. 47*(2), 23:1–23:38.

Latham & Watkins (2014, April). Cyber insurance: A last line of defense when technology fails. Technical report, Latham & Watkins.

Lee, R. M., M. J. Assante, and T. Conway (2016, March). Analysis of the cyber attack on the ukrainian power grid. Technical report, SANS Industrial Control Systems.

Leverett, E., R. Clayton, and R. Anderson (2017). Standardisation and certification of the 'internet of things'. In *the Annual Workshop on the Economics of Information Security (WEIS)*.

Li, F., A. Lai, and D. Ddl (2011). Evidence of advanced persistent threat: A case study of malware for political espionage. In *2011 6th International Conference on Malicious and Unwanted Software*, pp. 102–109. IEEE.

Li, N., L. Chen, and S. H. Low (2011, jul). Optimal demand response based on utility maximization in power networks. In *2011 IEEE Power and Energy Society General Meeting*, pp. 1–8. IEEE.

Liu, Y., S. Hu, and T.-Y. Ho (2016). Leveraging strategic detection techniques for smart home pricing cyberattacks. *IEEE Transactions on Dependable and Secure Computing 13*(2), 220–235.

Liu, Y., M. K. Reiter, and P. Ning (2009). False data injection attacks against state estimation in electric power grids. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*, New York, NY, USA, pp. 21–32. ACM.

Liyan, J., R. J. Thomas, and L. Tong (2012, January). Impacts of malicious data on real-time price of electricity market operations. In *45th Hawaii International Conference on System Sciences*, pp. pp.1907–1914.

Manshaei, M. H., Q. Zhu, T. Alpcan, T. Bacşar, and J.-P. Hubaux (2013, July). Game theory meets network security and privacy. *ACM Comput. Surv. 45*(3), 25:1–25:39.

Mikosch, T. (2006). *Non-life insurance mathematics: an introduction with stochastic processes*. Universitext. Springer-Verlag Berlin Heidelberg.

Mojica-Nava, E., C. Barreto, and N. Quijano (2015, Nov). Population games methods for distributed control of microgrids. *IEEE Transactions on Smart Grid 6*(6), 2586–2595.

Mossburg, E., J. Gelinne, and H. Calzada (2016). Beneath the surface of a cyberattack: A deeper look at business impacts. Technical report, Deloitte.

Nakashima, E. (2017). U.s. military cyber operation to attack isis last year sparked heated debate over alerting allies. https://www.washingtonpost.com/world/national-security/us-military-cyber-operation-to-attack-isis-last-year-sparked-heated-debate-over-alerting-allies/2017/05/08/93a120a2-30d5-11e7-9dec-764dc781686f_story.html. Accessed April 25, 2018.

Negrete-Pincetic, M., F. Yoshida, and G. Gross (2009, June). Towards quantifying the impacts of cyber attacks in the competitive electricity market environment. In *2009 IEEE PowerTech*.

Newman, L. H. (2015). Medical devices are the next security nightmare. https://www.wired.com/2017/03/medical-devices-next-security-nightmare/. Accessed January 24, 2018.

Nisan, N., T. Roughgarden, É. Tardos, and V. V. Vazirani (2007). *Algorithmic Game Theory*. 32 Avenue of the Americas, New York, NY 10013-2473, USA: Cambridge University Press.

Olson, P. (2013). *We Are Anonymous*. Random House.

Pagliery, J. (2016). Regulate cybersecurity or expect a disaster, experts warn congress. CNN. http://money.cnn.com/2016/11/16/technology/cybersecurity-regulation-congress/index.html?iid=hp-toplead-dom. Accessed October 24, 2017.

Papadimitriou, C. H. (2001). Algorithms, games, and the internet. In *In STOC*, pp. 749–753. ACM Press.

Perlroth, N. (2017). Hackers Are Targeting Nuclear Facilities, Homeland Security Dept. and F.B.I. Say. https://www.nytimes.com/2017/07/06/technology/nuclear-plant-hack-report.html. Accessed October 16, 2017.

Petersen, K. L. (2008). Terrorism: When risk meets security. *Alternatives 33*(2), 173–190.

Ponemon Institute (2014). Critical infrastructure: Security preparedness and maturity. Technical report, Ponemon Institute.

Ponemon Institute (2015). 2015 global cyber impact report. Technical report, Ponemon Institute.

Ponemon Institute (2016). 2016 cost of data breach study: Global analysis. Technical report, Ponemon Institute.

Preciado, V. M., M. Zargham, C. Enyioha, A. Jadbabaie, and G. J. Pappas (2014, March). Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems 1*(1), 99–108.

Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Rasouli, M., E. Miehling, and D. Teneketzis (2014). A supervisory control approach to dynamic cyber-security. In *International Conference on Decision and Game Theory for Security*, pp. 99–117.

Risk Management Solutions, Inc. and Centre for Risk Studies, University of Cambridge (2016). Managing cyber insurance accumulation risk. Technical report, Risk Management Solutions, Inc. and Centre for Risk Studies, University of Cambridge.

Romanosky, S., L. Ablon, A. Kuehn, and T. Jones (2017). Content analysis of cyber insurance policies: How do carriers write policies and price cyber risk? Working paper, RAND Corporation.

Roozbehani, M., M. Dahleh, and S. Mitter (2010, October). Dynamic Pricing and Stabilization of Supply and Demand in Modern Electric Power Grids. In *First IEEE Smart Grid Commnunications Conference (SmartGridComm)*.

Roozbehani, M., M. A. Dahleh, and S. K. Mitter (2012). Volatility of power grids under real-time pricing. *Power Systems, IEEE Transactions on 27*(4), 1926–1940.

Samadi, P., R. Schober, and V. W. Wong (2011). Optimal energy consumption scheduling using mechanism design for the future smart grid. In *2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pp. 369–374. IEEE.

Samuelson, P. A. (1937). A note on measurement of utility. *The review of economic studies 4*(2), 155–161.

Sandholm, W. H. (2011, January). *Population Games and Evolutionary Dynamics (Economic Learning and Social Evolution)* (1 ed.). The MIT Press.

Schechter, S. and M. Smith (2003). How much security is enough to stop a thief? In *Computer Aided Verification*, pp. 122–137.

Schneier, B. (2017). The internet of things will upend our industry. *IEEE Security and Privacy 15*(2), 108–108.

Semana (2008). Negocio redondo. `http://www.semana.com/nacion/articulo/negocio-redondo/94315-3`. Accessed February 1, 2016.

Shapley, L. S. (1953). Stochastic games. *Proceedings of the national academy of sciences 39*(10), 1095–1100.

Shiva, S., S. Roy, H. Bedi, D. Dasgupta, and Q. Wu (2010). A stochastic game model with imperfect information in cyber security. In *International Conference on Cyber Warfare and Security*, pp. 308. Academic Conferences International Limited.

Shleifer, A. (1985). A theory of yardstick competition. *The RAND Journal of Economics*, 319–327.

Smith, B. (2017). The need for urgent collective action to keep people safe online: Lessons from last week's cyberattack. `https://blogs.microsoft.com/on-the-issues/2017/05/14/need-urgent-collective-action-keep-people-safe-online-lessons-last-weeks-cyberattack`. Accessed May 18, 2017.

Stark, H. (2017). A bipartisan bill to strengthen cybersecurity for the internet of things. `https://www.forbes.com/sites/haroldstark/2017/08/20/a-bipartisan-bill-to-strengthen-cybersecurity-for-the-internet-of-things/`. Accessed October 24, 2017.

Swiss Re (2017). Sigma explorer. `http://www.sigma-explorer.com/`. Accessed: 2017-05-25.

Tan, R., V. B. Krishna, D. K. Yau, and Z. Kalbarczyk (2013). Impact of integrity attacks on real-time pricing in smart grids. In *ACM Conference on Computer and Communications Security (CCS 2013)*.

Vardakas, J. S., N. Zorba, and C. V. Verikoukis (2015, Firstquarter). A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. *IEEE Communications Surveys Tutorials 17*(1), 152–178.

Vardi, M. Y. (2017). Cyber insecurity and cyber libertarianism. *Communications of the ACM 60*(5), 5–5.

Verizon (2017). 2017 data breach investigations report. Technical report, Verizon.

Vickrey, W. (1961). Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, 8–37.

Zetter, K. (2014, November). An unprecedented look at stuxnet, the world's first digital weapon. `http://www.wired.com/2014/11/countdown-to-zero-day-stuxnet`. Accessed: Mar 1, 2017.

Zetter, K. (2015). A cyberattack has caused confirmed physical damage for the second time ever. `http://www.wired.com/2015/01/german-steel-mill-hack-destruction/`. Accessed October 16, 2017.

Zetter, K. (2016, mar). Inside the cunning, unprecedented hack of ukraine's power grid. `http://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/`. Accessed October 16, 2017.

Zimmerman, R., C. E. Restrepo, N. Dooskin, J. Fraissinet, R. Hartwell, J. Miller, and W. Remington (2005). Diagnostic tools to estimate consequences of terrorism attacks against critical infrastructure. In *Proceedings of the U.S. Department of Homeland Security conference, Working Together: Research and Development Partnerships in Homeland Security, Boston, MA*.

BIOGRAPHICAL SKETCH

Carlos Barreto was born in Garagoa, Boyacá, Colombia. He received a Bachelor of Science degree in electrical engineering from Universidad Distrital Francisco José de Caldas in Bogotá, Colombia in 2011. In 2013 he received a Master of Science degree in electrical engineering with emphasis in control science from Universidad de los Andes in Bogotá, Colombia. In the fall of 2013, he entered the computer science graduate program at The University of Texas at Dallas, Richardson, TX, USA. Carlos received a Master of Science degree with emphasis in intelligent systems from The University of Texas at Dallas in 2017 and defended his PhD dissertation in the Spring of 2018. His PhD research focuses on the security of critical infrastructures, which unlike traditional information security, does not focus on protecting information, but protecting physical processes from cybernetic attacks. He uses game theory and control theory to build models of systems, analyze possible security issues, and devise protection schemes to handle cyber risks.

CURRICULUM VITAE

Carlos Barreto
April 12, 2018

CONTACT INFORMATION

E-mail: Carlos.BarretoSuarez@utdallas.edu
Website: `http://utdallas.edu/~carlos.barretosuarez`

EDUCATION

University of Texas at Dallas, Richardson, Texas Spring 2018

Doctor of Philosophy in Computer Science. *Emphasis:* Intelligent Systems

Relevant Coursework: Algorithms and Data Structures, Design and Analysis of Computer Algorithms, Machine Learning, Stochastic Dynamic Programming

Dissertation: Role of economic policies in the security of critical infrastructures

Universidad de los Andes, Bogotá, Colombia Aug. 2013

Master of Science in Electronic Engineering. *Emphasis:* Control

Relevant Coursework: Optimization, Non-linear systems, Game theory

Thesis: Design of Economic Incentives in Demand Response Programs

Universidad Distrital Francisco José de Caldas, Bogotá, Colombia Sept. 2011

Bachelor of Science with Honors in Electronic Engineering

Emphasis: Control, Telecommunications, Nanotechnology, Computational Intelligence

Thesis: Comparative Study over FPGA of Four Embedded Systems based on Soft-Cores and uCLinux

INTERESTS

Research

Control Theory, Game theory, Mechanism design (inverse game theory), Cyber physical systems security, Machine Learning, Algorithms, Optimization, Distributed resource allocation

AWARDS & HONORS

Young researchers 2012 grant, conferred by Colciencias, an institution equivalent to the National Science Foundation (NSF) in the U.S.

SKILLS

Operating Systems

GNU/Linux (Debian and Arch-Linux) and Windows

Computer

Command Interpreters, LaTeX, Matlab/Simulink, Octave, Mathematica, LabVIEW

Programming

Main: Python, Java, Matlab

Used in the past: C, C#, Assembler, VHDL, Verilog, PHP, JavaScript, HTML

Languages

Spanish (Native), English (Fluent)

PUBLICATIONS

Journals

E. Mojica-Nava, C. Barreto, N. Quijano, "Population Games Methods for Distributed Control of Microgrids," IEEE Transactions on Smart Grid, vol. 6, no. 6, pp. 2586-2595, Nov. 2015

C. Barreto, J. Giraldo, A. Cárdenas, E. Mojica-Nava, N. Quijano, "Control Systems for the Power Grid and Their Resiliency to Attacks," IEEE Security & Privacy, vol. 12, no. 6, pp. 15-23, Nov.-Dec. 2014

Conferences

C. Barreto and A. A. Cárdenas, "Optimal risk management in critical infrastructures against cyber-adversaries," in 2017 IEEE Conference on Control Technology and Applications (CCTA), pp. 2027-2032, Aug. 2017

C. Barreto, A. A Cárdenas, and A. Bensoussan. "Optimal security investments in a prevention and detection game," In Proceedings of the Hot Topics in Science of Security: Symposium and Bootcamp, pp. 24-34. Apr. 2017.

C. Barreto and A. A. Cárdenas, "Perverse incentives in security contracts: A case study in the Colombian power grid," in the Annual Workshop on the Economics of Information Security (WEIS), 2016.

C. Barreto, A. A. Cárdenas, "Detecting Fraud in Demand Response Programs," in Proceedings of IEEE 54th Conference on Decision and Control (CDC), pp. 5209-5214, Dec. 2015

C. Barreto, A. A. Cárdenas, "Incentives for Demand-Response Programs with Nonlinear, Piece-Wise Continuous Electricity Cost Functions," in Proceedings of the IEEE American Control Conference (ACC), pp. 4327 - 4332, July 2015

C. Barreto, A. A. Cárdenas, N. Quijano, E. Mojica-Nava, "CPS: Market Analysis of Attacks Against Demand Response in the Smart Grid," in Proceedings

of the 30th Annual Computer Security Applications Conference (ACSAC '14), pp. 136-145, Dec. 2014

J. Valente, C. Barreto, A. Cárdenas, "Cyber-Physical Systems Attestation," 2014 IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS), pp. 354-357, May 2014

C. Barreto, E. Mojica-Nava, and N. Quijano, "Design of mechanisms for demand response programs," in Proceedings of 2013 IEEE 52nd Annual Conference on Decision and Control (CDC), pp. 1828-1833, Dec. 2013

C. Barreto, A. A. Cárdenas, and N. Quijano, "Controllability of Dynamical Systems: Threat Models and Reactive Security," in Decision and Game Theory for Security, pp. 45-64, Nov. 2013

C. Barreto, E. Mojica-Nava, and N. Quijano, "A population dynamics model for opinion dynamics with prominent agents and incentives," in Proceedings of IEEE American Control Conference (ACC), pp. 4575-4580, June 2013

M. Melgarejo, A. Gaona and C. Barreto, "Adaptive fuzzy equalization based on neuron grouping for time-varying non-linear channels," Ingeniería y Universidad, vol. 15, no.2, pp.423 – 443, July 2011

Posters

C. Barreto and A. Cárdenas, "Market Analysis of Attacks Against Demand Response in the Smart Grid," at Graduate Summer School on Games and Contracts for Cyber-Physical Security, Los Angeles, 2015

INDIVIDUAL PROJECTS

PDTooldbox

PDToolbox contains a set of Matlab functions to implement evolutionary dynamics from game theory with multiple populations. This toolbox is designed to facilitate the implementation of any game with different evolutionary dynamics or revision protocols. Available at `https://github.com/carlobar/PDToolbox_matlab`

SERVICE

Reviewer

IEEE Conference on Decision and Control (CDC), IEEE Multiconference on Systems and Control, IEEE Conference on Control Technology and Applications (CCTA), IEEE Colombian Conference on Automatic Control (CCAC), American Control Conference (ACC), Conference on Decision and Game Theory for Security (GameSec), IEEE Pervasive Computing, IEEE Transactions on Smart Grid , IEEE Transactions on Information Forensics and Security, IEEE Transactions on Cybernetics, IEEE

Transactions on Automatic Control, IEEE Systems Journal, IEEE Control Systems Letters, IEEE Internet Computing, Computer Networks Journal, Dynamic Games and Applications, Applied Energy, IEEE Design & Test.

REFERENCES

Alvaro Cárdenas

Assistant Professor

Erik Jonsson School of Engineering and Computer Science

University of Texas at Dallas

alvaro.cardenas@utdallas.edu

Dr. Cárdenas is my PhD advisor

Nicanor Quijano

Associate Professor

Department of Electic and Electronic Engineering

Universidad de los Andes

nquijano@uniandes.edu.co

Dr. Quijano was my Master's advisor

Eduardo Mojica-Nava

Associate Professor

Department of Electrical and Electronics Engineering

Universidad Nacional de Colombia

eamojican@unal.edu.co

Dr. Mojica was my Master's co-advisor