UNIVERSITY OF CALIFORNIA
SANTA CRUZ

NON-ZERO-SUM, ADVERSARIAL DETECTION GAMES IN NETWORK SECURITY

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

APPLIED MATHEMATICS AND STATISTICS

by

Braden Cooper Soper

December 2015

The Dissertation of Braden Cooper Soper
is approved:

Associate Professor John Musacchio, Chair

Professor Hongyun Wang

Distinguished Professor Daniel Friedman

Tyrus Miller
Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

Abstract

Non-zero-sum, Adversarial Detection Games in Network Security

by

Braden Cooper Soper

In this dissertation we propose two novel non-zero-sum, adversarial detection games motivated by problems in network security. First we consider a local mean field, interdependent detection game between a network of defenders and a strategic attacker. Each defender chooses a detection threshold to test for the presence of a botnet infection, which can propagate between defenders if undetected. In order to avoid detection, the attacker balances stealth and aggression in his strategic utilization of the compromised network. We compare selfish, decentralized defenders to centrally planned defenders in order to examine the effects of network externalities on detection strategies. It is found that for fixed attack strategies, decentralized defenders choose thresholds that are either too low or too high than is socially optimal. When the attacker is strategic and the defenders are homogeneous, we prove the existence of a pure Nash equilibrium in both decentralized and centralized games. Through numerical approximations of the equilibria, we find that decentralized defenders can outperform a central planner in such games. It is observed that pure Nash equilibria often fail to exist when defenders are heterogeneous in their cost functions. In this case sufficient conditions are given to guarantee a Stackelberg equilibria.

Next a two-player, non-zero-sum, sequential detection game based on Wald's SPRT is presented. A defender seeks to sequentially detect the presence of an attacker via the drift of a stochastic process. The detection process is complicated by the attacker's ability to strategically choose the drift of the observed stochastic

process. We prove the existence of pure Nash equilibria and give sufficient conditions for the existence of Stackelberg equilibria with the defender as leader. It is shown that both low false positive costs and high prior probabilities of intrusion lead to an infinite number of Nash equilibria in which the defender makes no observations. Conversely both high false positive costs and low prior probabilities of intrusion lead to a finite number of non-trivial Nash equilibria. Through numerical examples we see that it is possible for the defender to do better using a Stackelberg equilibrium strategy than a Nash equilibrium strategy.

To my family.

Acknowledgments

There are many people I wish to thank for helping me throughout my graduate school career. First, a special thanks to my advisor John Musacchio for all the help and guidance throughout the research process. Thank you to my other dissertation committee members, Hongyun Wang and Daniel Friedman, for valuable comments and suggestions on this dissertation. Thank you to all the Applied Mathematics and Statistics faculty and graduate students for making graduate school a wonderful and memorable experience. In particular thank you to Marc Mangel for creating invaluable opportunities which enriched my graduate school experience immeasurably. Thank you to E.J. Dick and Alec MacCall of the National Marine Fisheries Service at the Southwest Fisheries Science Center for research and funding opportunities. Thank you to Celeste Matarazzo, Jim Brase, Dan Merl and Ryan Prenger at (or formerly at) Lawrence Livermore National Laboratory for summer research opportunities and for inspiring me to follow this line of research. Much love and infinite thanks to all my family: Mike and Pam Soper; Kate, Tristan, Lucas and Juna Grell; Stephen, Nenita, Olive and Charlotte Soper. Thanks to Rod and Elma Pinlac for knowing how to make a deal and providing me with computational resources. Thanks to the entire Kaijin family for giving me a second home. Finally and most importantly, thank you to my amazing wife Maria Theresa Pinlac, whose daily love, encouragement and support made this journey possible. Thank you all.

Chapter 1

Introduction

This dissertation is concerned with the development and analysis of novel non-zero-sum, *adversarial detection* games with applications to network security. More precisely we model encounters between strategic adversaries, a defender (or defenders) and a single network attacker, as non-zero-sum, intrusion detection games. The dissertation is composed of two parts, each of which focusses on a particular class of non-zero-sum, adversarial detection games which have not been considered previously.

In the first part we propose a class of games we call *interdependent detection games*. These games are related to the class of games known as *interdependent security games* (IDS) introduced by Heal and Kunreuther [42]. More specifically our model is motivated by the local mean field, security investment games introduced by Lelarge and Bolot [47,48]. The main difference being that we consider network effects on epidemic *detection*, as opposed to the network effects on security investments as is most common in the IDS games literature. As our main motivating example we model the strategic interactions between the decentralized, legitimate users of a compromised network (defenders) and an adversary (bot master) who illegitimately steals resources from the compromised computers.

In the second part of this dissertation we analyze a non-zero-sum, *sequential detection game*. The model is based on a continuous time version of Wald's Sequential Probability Ratio Test (SPRT) [78]. For this game we focus on a two-player game with a single defender and a single attacker.

We outline the two games, their applications to network security and our main contributions below.

1.1 Interdependent Detection Games

An important development in the study of information security was the realization that many of the difficulties with securing modern information systems are not simply due to technological problems, but arise due to misaligned incentives. One of the earliest observations of this fact was by Varian [75] in the context of the prevalence and destructiveness of Distributed Denial of Service (DDoS) attacks. A DDoS attack is often executed by a malicious individual, known as a *bot master* or *bot herder*, using a *botnet*, a large number of compromised computers under his control [25]. The bot master uses the botnet to send a large number of access requests to a single computer or server connected to the internet. If the number of requests is large enough they can overwhelm the targeted computer to the point of making it inoperable, effectively shutting it down. While there are no doubt technical failures which contribute to the prevalence of such attacks, it was pointed out by Varian that a misalignment of incentives make the DDoS problem much more difficult to deal with. In particular the owners of the compromised computers are not the intended targets of the DDoS attack, nor are they held responsible for the consequences of the attack. As such they do not have the correct incentives to invest in the high quality security measures that might prevent their computers from being compromised in the first place.

Anderson [10] elaborated on these observations and outlined more explicitly the need to incorporate microeconomic theory into the study of information/cyber security. Since then a growing body of literature has developed around the study of the economics of information security [11]. Due to the highly interdependent nature of modern information networks, a class of games known as *interdependent security games* (IDS) have become a rich research area within the scope of the economics of information security. IDS games were first introduced by Heal and Kunreuther [42] in a broader context not necessarily limited to information security. In fact one of their prime examples is that of airline security [43]. The main idea behind IDS games is that they extend the standard attacker-defender dynamics that is typical in security games to the case where there are multiple defenders who are subjected to interdependent security risks. In the case of airline security, the fact that a single piece of checked luggage passes through multiple airports when connecting flights are involved, results in a situation where the security level at each airport depends on the security efforts of *all* airports which can be reached by connecting flights. Hence the security risk posed by a bomb in a piece of luggage is an interdependent security risk for the airports. The resulting IDS problem can be modeled and analyzed with tools from game theory to better understand the motivations and expected behavior of the "players" involved.

Many important extensions to the IDS game framework have been proposed. Most IDS models assume fixed or randomized attack strategies, however some recent works have extended the IDS framework to include a *strategic attacker*. Chan et al. [27] considers a strategic attacker coupled with an IDS game. The model also considers the case that security investment reduces not only risk but also risk transfer. They call their more general model an *interdependent defense* (IDD) game. Bachrach et al. [14] considers an IDS game in which an attacker

chooses a single target in the network with the objective of destroying as many of the players in the network. Acemoglu et al. [1] considers a similar game but generalizes the results over various network topologies.

Another important extension to the IDS game framework was to consider propagating security threats. Note that the airline security problem mentioned above is one of risk transfer, not risk propagation. Extending security interdependence to include network propagation allows one to consider large scale epidemic threats such as computer or biological viruses. In a propagation IDS game individuals can be in one of a discrete number of states. In its simplest form the states correspond to *susceptible* and *infected*. A vast literature exists on modeling the dynamics and control of epidemic processes [12, 36]. The distinction here is that in propagation based IDS games an epidemic model is coupled with a security investment model at the individual level. Thus the problem becomes a multi-player game as opposed to an optimal control problem.

Building on the the mathematical epidemiology literature many propagation IDS games have been consider. [54] considered a N-intertwinded SIS model coupled with a noncooperative IDS game. In [44] the author utilizes evolutionary game models to study the effects of security and insurance investments on contagion in network. [73] uses a mean field approach to study the effects of learning about epidemic rates on security investment by selfish agents. The model in [1] also considers contagion as a security threat.

One of the primary difficulties in dealing with propagation based IDS games is that the models get incredibly complex [45]. Epidemic models themselves are highly non-linear, and coupling a large population game with it can result in intractable models. For this reason many of the propagation based IDS games rely on mean field approximation models such as in [44, 73].

A mean field approach with a very distinct character was introduced by Lelarge and Bolot [47, 48]. They called their model a *local mean field* model and it allows one to model the local correlation structure of the underlying network while also explicitly modeling a single agent's strategic behavior. By utilizing ideas from the statistical physics literature and combinatorial optimization, the authors study the network effects on security investments by selfish agents facing a propagating security threat. By focussing on large networks and considering the limiting case that the number of agents approaches infinity, one obtains a tractable model able to deal with propagating security threats and selfish individual behavior over large networks.

When one agent's exposure to security risks depends on the actions of others, then investment in security is not the only strategic decision that will be affected by the presence of network effects. In particular, if an agent is interested in detecting the presence of a security threat, then his chosen detection strategy should depend on the detection strategies of others. While detection is not explicitly an economic action, it can be understood in the context of the IDS literature. In particular, approaching the detection problem as a Bayesian decision theoretic problem [22], the defending agent can be given a utility function which determines his optimal detection strategy. So instead of optimal investment strategies we are interested in optimal detection strategies. In this context interdependent detection problems are subject to similar forces that arise in networked information systems in general and interdependent security problems in particular.

We propose a novel network *interdependent detection game* which builds on the probability based models for interdependent security initiated by Heal and Kunreuther [42] and the propagation based models of Lelarge and Bolot [47, 48]. We utilize the local mean field framework of Lelarge and Bolot to model the local

correlation structure of large computer networks. Our model diverges from the existing IDS framework in that we do not explicitly consider security investments. Instead we model epidemic *detection* by interdependent agents subjected to the same propagating security threat. In theory it is possible to model both security investment and infection detection, but we do not consider security investment as it has been studied extensively in the IDS literature. We are instead interested in the network effects on the epidemic detection process.

We use botnets as our motivating example of an interdependent detection game. A unique feature of the botnet security threat is the highly distributed, decentralized network of victims whose compromised computers make up the botnet. These victims are not always the intended target of botnet attacks, yet they play a critical role in the functionality of the botnet. We propose a game between the operator of a botnet (*bot master*) and the decentralized, legitimate users of the compromised network (*defenders*).

Computers become infected either directly by the bot master or indirectly from neighboring computers via self-propagating malware. If an infection is successful, then the bot master gains control of the compromised computer and can use it for his own nefarious purposes. If the bot master is too aggressive in his utilization of the compromised computers, then it is more likely that the defenders will detect the infections. For example, if the bot master is continually using the compromised computers to send large volumes of spam, then there may be a noticeable degradation of each computer's performance due to excessive bandwidth consumption from the spamming. Some defenders may then decide to patch or replace their computers, thus ridding themselves of the infection and reducing the overall size of the botnet. Thus he must balance stealth and aggression in his strategic utilization of the botnet.

It is assumed that each defender makes a single observation of some performance metric which is corrupted by noise. Larger observed values are indicative of infection. Thus we model the strategy of each defender as the choice of a binary classification threshold which is used to classify their states as either *infected* or *susceptible*. The defenders of the network must decide how vigilant they will be in trying to detect the presence of a botnet infection while balancing the costs of false alarms and missed detections.

We begin our study by examining a game in which the defenders in the network are statistically homogeneous and we look for symmetric threshold detection strategies. The bot master is explicitly modeled as a strategic agent pitted against the legitimate users (or defenders) of a computer network targeted to become a botnet. Two scenarios are considered, one in which the defenders are decentralized and selfish, and one in which the defenders are coordinated by a central planner. In this way we are able to examine the efficiency of equilibrium defense strategies. We find that there are cases in which defenders tend to be overly vigilant and cases in which defenders tend to be overly indifferent compared to the social optimum.

Our main theoretical results are proving the existence of a symmetric, pure Nash equilibrium in the game with homogeneous defenders in both the decentralized and centralized game. We then numerically approximate the equilibria in order to study their qualitative features. We find that under certain parameter regimes it can be socially optimal to have a higher infection rate throughout the network. Furthermore we find that when the bot master is strategic the decentralized defenders can do better than a central planner. While this may seem contradictory it is a result of the fact that both the central planner and the decentralized defenders are playing a game with the strategic attacker. As such, the attacker plays different strategies at the corresponding equilibria. Since the

defenders' payoffs depend on the attacker's strategy, it can happen that decentralized equilibrium strategies result in smaller expected costs than centralized equilibrium strategies.

We then extend the model to defenders which are heterogeneous in their cost functions. For fixed attack strategies we prove the existence of decentralized population best response functions. These are population strategies mapping defender type to defender strategy from which no individual has an incentive to unilaterally deviate. This is similar to the idea of a *fulfilled expectations equilibrium* used in the economics literature [40]. Through numerical examples we observe two counter intuitive results which demonstrate network effects on interdependent detection games. 1. When the population tends to value missed detection over false alarms, selfish defenders wind up being too indifferent. 2. When the population tends to value false alarms over missed detections, selfish defenders wind up being too vigilant. This happens because selfish agents do not account for the externality of altering infection rates when they unilaterally change strategies. This result is analogous to the problem of *free-riding* which can occur in games with network externalities [40,76]. In the context of IDS games this means agents tend to under invest in security. Thus we see that free-riding in interdependent detection games is distinct from free-riding in IDS games. Free-riding in interdependent detection games can go in two directions: under-vigilant and over-vigilant.

In contrast to the homogeneous game, it is observed that pure Nash equilibria often fail to exist with heterogeneous defenders. The difficulty in establishing the existence of Nash equilibria in the heterogeneous game stems from the fact that when decentralized defenders play a population best response equilibrium strategy, we often find that the resulting expected utility function of the attacker is not quasi-concave. As a result we cannot guarantee the continuity of

the attacker's best response function, a condition typically needed for equilibrium existence proofs in continuous strategy games.

When Nash equilibria do not exist other solution concepts may be of interest. In this case a Stackelberg solution concept is adopted, and sufficient conditions are given to guarantee such an equilibrium with the attacker as leader. Numerical examples are given in order to examine the effects of heterogeneity and network effects on defender strategies. Unfortunately, the same problems that plague Nash equilibria also affect the Stackelberg equilibria with defenders as leaders. We discuss some of these difficulties and show possible ways to circumvent them.

Finally we point out the non-zero-sum nature of our interdependent detection game. Many network security games focus on an attacker's strategy in compromising a network, such as attack intensity [33], frequency [29] and location [23]. A unique feature of botnets is that the initial network intrusion is only the beginning of a potentially long-term network control and command objective. Many bot masters rent their botnets to customers wishing to perpetrate illicit activities (spam, DDoS attacks, click-fraud, etc.) [49]. Thus a bot master must maintain a sufficiently large network of infected computers to maintain a profitable enterprise. In the process scores of legitimate computer users become unwitting participants in cyber-criminal activities [71]. It should be clear that such an interaction between the legitimate computer users and the strategic bot master is not a zero-sum game. The bot master is not interested in inflicting maximal damage to the network. On the contrary, he is interested in the long term viability of the network to maximize his profits. Thus we are interested in the bot master's strategic utilization of the compromised network, an important component of the botnet security threat that has not been considered in theoretical models.

In summary our main contributions to the fields of adversarial detection games

are as follows:

- To the best of our knowledge we present the first *interdependent detection game*. The majority of the IDS literature has focussed on security investment as a measure of security effort. We propose studying adversarial threat detection in an interdependent environment.

- While strategic adversaries have been modeled in IDS games, most of theses papers have focussed on *zero-sum* interactions. That is, the loss incurred by the defenders is the direct gain by the attacker. To our knowledge we are the first to consider the case of *non-zero-sum* interactions between attacker and defenders in the IDS literature. Specifically we consider the case that the attacker is not interested in destroying the defenders but in remaining undetected while utilizing stolen resources.

- To the best of our knowledge non-cooperative distributed detection has not been considered in the literature. Distributed detection usually refers to interdependent sensors cooperatively trying to detect the presence of a signal. As such distributed detection can be modeled as a cooperative (or coalition) game and is usually approached as a design problem [4–6]. In our model the sensors (defenders) are non-cooperative and are only interested in detecting the presence of the infection individually.

1.2 Sequential Detection Games

The local mean field model we use to study interdependent detection is probabilistic in nature and does not explicitly consider the dynamic nature of intrusion detection. The focus of the model is on the interdependent nature of the defenders while dynamic considerations are ignored. In the second part of the dissertation

we shift focus away from the interdependent nature of intrusion detection and focus on dynamic, real-time adversarial detection. Specifically, in contrast to the previous model where defending agents made a single observation of their state on which they based their decision, we now consider the case where a single agent makes multiple observations over time, and must make a real-time detection decision.

As such we consider the defender to be a type of intrusion detection system [7]. Such systems have become an integral components in securing modern computer networks. In its simplest form intrusion detection can be thought of as a hypothesis testing problem. The null hypothesis being that you are not compromised and the alternative hypothesis being that you are compromised. While fixed sample size hypothesis testing can be effective, the dynamic nature of computer networks and the need for real-time detection suggests *sequential hypothesis testing* may be more suitable [80]. That is one does not fix the number of observations ahead of time. Instead after each observation a decision is made to either make another observation or to stop making observations and accept or reject the null hypothesis. Such a procedure is useful when there is a fixed cost $c > 0$ per observation.

For example let $X_1, X_2, X_3, ...$ be i.i.d. observations from a parametric distribution with density $f(\cdot|\theta)$ where θ is a real valued parameter. We wish to test the following simple hypothesis test:

$$H_0 : \theta = \theta_0,$$
$$H_1 : \theta = \theta_1,$$

for some specific values $\theta_0 < \theta_1$. We then seek a *sequential decision rule* (τ, δ) where $\tau \in \{1, 2, 3, ...\}$ is a *stopping time* and $\delta \in \{0, 1\}$ a decision rule. Both τ and δ are random variables such that at time τ you stop making observations

and choose hypothesis δ. If there is a prior probability $\pi \in [0,1]$ that the null hypothesis is correct and costs $\alpha > 0$ and $\beta > 0$ associated with false positives and false negatives, respectively, then the *Bayes risk* of the sequential decision rule (τ, δ) is

$$E_\pi[c\tau + \alpha \mathbb{1}_{\{\delta=1,\theta=\theta_0\}} + \beta \mathbb{1}_{\{\delta=0,\theta=\theta_1\}}]. \tag{1.1}$$

One then seeks the sequential decision rule which minimizes the Bayes risk.

In the classic work of Wald [78, 80] it was shown that optimal sequential decision rules can be found in the class of Sequential Probability Ratio Tests (SPRT). An SPRT has two values $A \in [0,1]$ and $B \geq 1$ associated with it and is defined in terms of the *likelihood ratio process*,

$$L_n = \frac{\prod_{i=1}^n f(X_i|\theta_1)}{\prod_{i=1}^n f(X_i|\theta_0)}.$$

We write SPRT(A, B) for the SPRT associated with values A, B. The SPRT(A, B) has the decision rule

$$\delta_{A,B} = \begin{cases} 0 & \text{if } L_n \leq A, \\ 1 & \text{if }, L_n \geq B. \end{cases}$$

The SPRT(A, B) has the stopping time

$$\tau_{A,B} = \inf\{n > 0 : L_n \notin (A, B)\}.$$

The optimality of the SPRT(A, B) can be considered from both a frequentist and Bayesian perspective [64]. We will be concerned primarily with the Bayesian approach and thus only state the relevant results. Namely that for any prior π there will exist values (A^*, B^*) such that the SPRT(A^*, B^*) minimizes the Bayes risk (1.1). While originally developed in the context of industrial quality control,

the method is a classical example of real-time anomaly detection and has many applications, including in network security [7].

Because we will be interested in the detection of a strategic adversary, it is natural to once again utilize game theory to better understand the incentives of both defender (detector) and attacker (intruder). While the importance of game theory in developing robust intrusion detection systems has been recognized [4–6], less attention has been paid to applying game theory to sequential detection problems. For the most part applications of game theory to sequential hypothesis testing have typically been restricted to robust minimax solutions [13, 22, 80], which assumes a zero-sum game between *observer* and *nature*. Given the vast array of security threats and strategic adversaries in the cyber domain, one potential shortcoming of the minimax approach is the fact that many non-cooperative, strategic encounters may not be zero-sum. If a defending agent has information about the type of adversary, such as the attacker's payoff function, then the defending agent may be able to leverage this information to find superior sequential detection tests.

Motivated by these considerations we propose and analyze a two-player, non-zero-sum, sequential detection game between a *defender* agent and an *attacker* agent. We present the model as an abstract attacker-defender game, but botnets (see Chapter 2) and electricity theft in the smart grid [26, 68] are two areas of application that we have in mind.

The defender is in charge of protecting a secured resource. This could be a single computer, a network of computers or some other cyber-physical infrastructure. The defender's objective is to sequentially detect whether or not his secured resource has been compromised by the attacker. The defender makes noisy observations of the system's state which we model by a stochastic process Z_t. We assume that whether or not the system is compromised can be discerned through

the drift of the process Z_t. For example, the observed process could be cumulative bandwidth usage, CPU load or energy consumption. It is his objective to do so in such a way that minimizes a payoff function which takes into account the expected observation time and both type I and type II detection errors. As such the defender's optimal sequential test is a version of Wald's Sequential Probability Ratio Test (SPRT) [80].

The attacker is interested in bypassing the defender's security in order to establish long-term, unrestricted access to the resources available on the system. The attacker's objective is not necessarily to destroy or damage the defender's system, but to utilize system resources. The attacker must then balance how aggressive he should be in utilizing resources of the compromised system and how stealthy he should be in order to avoid detection. The more the attacker utilizes the system, the more utility he obtains. However, this also increases the drift of the observed stochastic process, thus increasing the probability of detection.

The main theoretical result is a proof of the existence of pure Nash equilibria in the special case that the attacker does not discount future expected utility. Furthermore we give conditions for the existence of Stackelberg equilibria with the defender as leader in the special case that the defender's strategy is restricted to Wald's SPRT. Numerical examples are given to explore the qualitative features of the equilibria. It is observed that both low false positive costs for the defender and high prior probabilities of intrusion by the attacker lead to an infinite number of Nash equilibria in which the defender immediately classifies his system as compromised and the attacker receives no utility. Conversely, we see that both high false positive costs for the defender and low prior probabilities of intrusion by the attacker lead to a finite number of non-trivial Nash equilibria. Finally we see that it is possible for the defender to improve his outcome under the Stackelberg

equilibrium strategy in relation to the Nash equilibrium strategy.

Previous examples of sequential detection games have largely been restricted to discrete-time, zero-sum games. As mentioned above, minimax sequential detection assumes the form of a zero-sum game between observer and nature. Minimax sequential detection was an attempt to develop more robust sequential statistical tests [13] rather than explicitly address interference by a strategic adversary. Nevertheless, minimax sequential detection lends itself to an adversarial framework and has been used in game-theoretic settings. Such an approach was taken in [60] and [59] with applications to detecting access layer misbehavior in wireless networks. A discrete-time, non-zero-sum, network security classification game involving Wald's SPRT can be found in [15]. This work was largely numerical as the discrete-time SPRT in an adversarial setting is not amenable to analysis due to the "overshoot" problem [65]. To our knowledge these are some of the only attempts to apply game theoretic reasoning to sequential detection. A similar fixed sample size detection game dealing with electricity theft in the smart grid can be found in [26].

In summary our main contributions to the field of adversarial sequential detection are the following:

- We present the first non-zero-sum, continuous-time, sequential detection game.
- We provide the first analytical equilibrium results for non-zero-sum, sequential detection games based on Wald's SPRT.
- To the best of our knowledge we are the first to use results from the theory of optimal stopping and free-boundary problems [56, 64] to construct and analyze sequential detection games.

Chapter 2

A Local Mean Field Botnet Detection Game

As one of the major security threats to users of the internet, botnets exemplify the difficulties of network security: They are highly distributed, interconnected and complex. Furthermore a single botnet can contain thousands of computers, making scores of legitimate computer users unwitting participants in cybercriminal activities [71]. Though the research community has taken an interest in the botnet phenomenon, theoretical models of botnets are nascent.

A promising game-theoretic approach to modeling botnets is the local mean field model of Lelarge and Bolot [47, 48]. This approach contains many appealing features for dealing with the complexities of the botnet phenomenon. However, the strategic attacker, or *bot master*, is not explicitly considered as a strategic agent in the game. Furthermore the network of agents are deciding whether or not to invest in security rather than dealing directly with the threat of a botnet. We propose a novel security game which extends the Lelarge-Bolot model by explicitly modeling the bot master as a strategic agent pitted against the legitimate users of a computer network targeted to become a botnet. In our game the bot master

must consider the tradeoffs between stealth and aggressiveness in utilizing his botnet. The legitimate users of the network (*defenders*) act as intrusion detection systems and must consider the tradeoffs between botnet infections, false alarms (false positives) and missed detections (false negatives).

Standard network security games dealing with an *attacker-defender* dynamic often focus on the attacker's strategy in initially compromising a network. Such strategies include attack intensity [33], frequency [29] and location [23]. One unique feature of botnets is that the initial network intrusion is only the beginning of a potentially long-term network control and command objective. Many bot masters rent out the services of their botnets to perpetrate illicit activities (spam, DDoS attacks, click-fraud, etc.) on the behalf of paying customers [49]. Thus a bot master must maintain a sufficiently large network of infected computers in order to make a profit. To our knowledge the strategic behavior in maintaining a botnet has not been considered in theoretical models.

We address this shortcoming in our botnet game by having the main strategic variable of the bot master be the degree to which he utilizes his botnet. The bot master tries to gain control of each computer directly with a fixed probability of success. In addition compromised computers are capable of propagating the vulnerability to other computers in the network. If an infection is successful then the bot master gains control of the compromised computer and can use it for his own nefarious purposes. If the bot master is too aggressive in his use of the compromised computers, then it is more likely that the legitimate users will detect the infection. For example, if the bot master is continually using the compromised computers to send large volumes of spam, there may be a noticeable degradation in computer performance due to excessive bandwidth consumption from the spamming. Some defenders may then decide it is time to patch or replace

their computers, thus ridding themselves of the infection and reducing the overall size of the botnet.

In many network security games the defender is thought to be the intended target of an attack. In the case of botnets this might be the security administrator of a server under a DDoS attack or a network spam filter being bombarded by a spambot. Missing from such models are the legitimate computer users whose compromised computers make up the botnet. Understanding the strategic interactions between the legitimate users and the bot master should contribute to a better understanding of the botnet threat.

For simplicity we consider homogeneous defenders and look for pure, symmetric, mutual best responses among the legitimate users. We find sufficient conditions for the existence of a pure, symmetric Nash equilibrium in a game involving decentralized, selfish defenders as well as in a game with a central planner. Network effects are explored numerically. We find that the average number of neighbors, infection transmission probability and cost of raising an alarm all determine qualitatively distinct regimes of Nash equilibria.

Previous work on network security games have considered similar issues. In [77] the incentives of ad-networks and ISPs to invest in detecting botnets are studied. Interconnected agents, network externalities and security investments have been considered in [8, 38, 39, 42, 48, 55]. Botnet dynamics have been considered in [21] and [28], while botnet economics are considered in [62] and [49]. Two-player resource control games applied to network security were considered in [57, 74].

2.1 A Two-Player Botnet Detection Game

To motivate the large-population botnet detection game, we begin with a security game between a bot master (attacker) and the owner of a single targeted

computer (defender). Bot masters are able to compromise computers by exploiting flaws in the software and hardware of networked systems. A game-theoretic model of software/hardware manufacturers and their incentives to invest in reducing software system failures is presented in [37]. The authors classify software failures into two categories: *security failures* (failures caused by malicious and unauthorized access to a user's system) and *reliability failures* (failures which are not security failures). Two observations on which the authors base their game are 1) the source of both security and reliability failures is the same (software bugs), and 2) it is too costly for users to distinguish between the two types of failures. We incorporate these points into our game by modeling the defender's inability to reliably detect the presence of a botnet infection by observing some level of performance degradation in his computer. The performance degradation is due to either natural performance variability (reliability failure), or it is due to the cumulative effects of both natural performance variability *and* security issues related to a botnet infection (security failure).

We assume the bot master tries to infect the targeted computer as aggressively as possible and has an overall probability of success p. The bot master infects computers in order to illegitimately utilize available computational resources, i.e. CPU time, RAM, bandwidth, electrical power, etc. Specifically we wish to model how *aggressive* the bot master should be in utilizing these resources. In what follows we will not model a particular resource, instead we consider a general measurable resource $\mathcal{R}$ taking values on $\mathbb{R}^+$. The strategic variable for the bot master is $A \in \mathbb{R}^+$, a direct measurement of the amount of resource $\mathcal{R}$ the bot master uses. We will often refer to the strategy A as the bot master's *aggressiveness*, since larger values of A correspond to a more aggressive utilization of the resource $\mathcal{R}$.

We model the natural performance variability associated with resource $\mathcal{R}$ as a random variable S with support $\mathbb{R}^+$. We denote the cumulative distribution function of S by $F_S(x)$ and the probability density function by $f_S(x)$. Let $\chi \sim$ Bernoulli(p) where $\chi = 1$ if the direct infection by the bot master is successful and $\chi = 0$ otherwise. We assume χ and S are independent. The observable performance degradation is then modeled by the random variable $Z = A\chi + S \in \mathbb{R}^+$.

The defender is thought to be a typical user of a computer connected to the internet. Aware that there are potential security threats the defender must decide how vigilant to be in detecting such threats. As in [37] we assume the defender is unable to reliably distinguish between security failures and reliability failures. The defender must then consider the potential costs from both false alarms (false positives) and missed detections (false negatives).

Given the observation Z the defender wishes to determine whether or not his computer is infected. Assuming the distributions of S and χ are known, this becomes a simple hypothesis testing problem with hypotheses $H_0 : Z = S$ and $H_1 : Z = A + S$. Higher observed values of Z indicate a higher likelihood of infection (H_1), thus we take the strategic variable of the defender to be a threshold T which takes values in the support of S, namely $\mathbb{R}^+$. The strategy T determines a threshold classifier for the defender: If $Z \geq T$ the defender decides his computer is infected and takes appropriate measures to remediate the potential infection. If $Z < T$ then the defender decides his computer is not infected and takes no action.[1]

[1]More generally the defender is free to choose any decision rule mapping observations in $\mathbb{R}^+$ to the set $\{0, 1\}$. If $\mathcal{C}$ is the set of all such decision rules we assume the defender will choose a $g \in \mathcal{C}$ that minimizes his expected posterior loss, i.e. he chooses a Bayes decision rule. Under the conditions of our model there exists an optimal threshold classifier T^* which is a Bayes decision rule.

Define the detection indicator random variable $D = \mathbb{1}_{\{Z \geq T\}}$, i.e. $D = 1$ if $Z \geq T$ and $D = 0$ otherwise. We can then define the following indicator random variables corresponding to the possible outcomes of the defender's observation and decision:

$$
\begin{aligned}
W_{\mathrm{FP}} &= D(1-\chi), \\
W_{\mathrm{FN}} &= (1-D)\chi, \\
W_{\mathrm{TP}} &= D\chi, \\
W_{\mathrm{TN}} &= (1-D)(1-\chi).
\end{aligned}
$$

Thus W_{FP} indicates a false positive, W_{FN} a false negative, W_{TP} a true positive, and W_{TN} a true negative. Taking expectations we have the following:

$$
\begin{aligned}
E[W_{\mathrm{FP}}] &= [1 - F_S(T)](1-p), \\
E[W_{\mathrm{FN}}] &= F_S(T-A)p, \\
E[W_{\mathrm{TP}}] &= [1 - F_S(T-A)]p, \\
E[W_{\mathrm{TN}}] &= F_S(T)(1-p).
\end{aligned}
$$

We associate a cost with each of the possible detection outcomes: We let $c \geq 0$ be the cost of a false positive, $v \geq 0$ the cost of a false negative and $k \geq 0$ the cost of a true positive. The defender incurs no cost for correctly classifying his state as *not infected*, i.e. for true negatives. The cost v includes the cost associated with future lost resources when the detection is missed, hence it depends on how aggressively the bot master utilizes the compromised computer. Thus $v = v(A)$ is a function of the bot master strategy A. We assume it is differentiable, non-decreasing and unbounded in A. The cost c, on the other hand, is only experienced

under the case that the defender is not infected. Thus it does not depend on the strategy of the bot master and we assume it is a fixed constant. We interpret the constant $k \geq 0$ as the fixed cost associated with remediating a potential infection such as reinstalling an operating system, updating software or purchasing a new computer. In other words it is the cost of raising an alarm no matter the infection state. Given this interpretation we necessarily have $c \geq k$, the cost of a false alarm is greater than or equal to the cost of raising an alarm. We furthermore will assume $v(A) \geq k$ for all A, the cost of a missed detection is greater than the cost of raising an alarm. Thus both type I and type II errors are more costly than raising an alarm.

For each $(A, T) \in \mathbb{R}^+ \times \mathbb{R}^+$ the defender's cost $C(A, T)$ is as follows:

$$C(A, T) \triangleq cW_{\text{FP}} + v(A)W_{\text{FN}} + kW_{\text{TP}}.$$

The expected cost of the defender is now

$$E[C(A, T)] = c[1 - F_S(T)](1 - p) + [k + (v(A) - k)F_S(T - A)]p.$$

A successful attack is the same as a false negative and has indicator random variable W_{FN}. Define the function $g(A)$ to be the utility gained by the bot master from a successful attack given the bot master strategy A. We assume $g(A) \in C^2$ with $\frac{dg}{dA} > 0$, $\frac{d^2g}{dA^2} \leq 0$ and $g(0) = 0$. We define the attacker utility as

$$U(A, T) = g(A)W_{\text{FN}}(A, T).$$

The expected utility is then

$$E[U(A,T)] = g(A)F_S(T-A)p.$$

The best response correspondences for the defender and attacker, respectively, are

$$\sigma_1(A) = \arg\min_T E[C(A,T)],$$
$$\sigma_2(T) = \arg\max_A E[U(A,T)].$$

Throughout the remainder of the paper we will need several assumptions regarding the random variable S that will be pertinent to the proofs of many of our results. We summarize them here for reference. Because the defender is ultimately performing a hypothesis test (i.e. *infected* or *not infected*) we begin by defining the likelihood ratio function $L(A,T)$ and its limiting values with regards to his threshold strategy.

$$L(A,T) \triangleq \frac{f_S(T-A)}{f_S(T)} \tag{2.1}$$

$$L_0(A) \triangleq \lim_{T \downarrow A} L(A,T) \tag{2.2}$$

$$L_\infty(A) \triangleq \lim_{T \to \infty} L(A,T) \tag{2.3}$$

Assumption 1. *The following are sufficient conditions on the random variable S needed in the remainder of the paper.*

1. *The support of S is $\mathbb{R}^+$.*

2. *$F_S(\cdot) \in C^2$.*

3. $f_S(x) > 0$ *for all finite* $x > 0$.

4. $\frac{d}{dx}\left[\frac{F_S(x)}{f_S(x)}\right] > -1$.

5. $\frac{\partial L}{\partial T} \geq 0$.

6. $\frac{\partial L_\infty}{\partial A} > 0$.

7. $\frac{\partial L}{\partial T} > 0 \implies L_0(A) \equiv 0$. $\frac{\partial L}{\partial T} \equiv 0 \implies \frac{\partial L_0}{\partial A} > 0$.

The prototypical distribution satisfying the above conditions is $S \sim \text{gamma}(\alpha, \beta)$ with $\alpha \geq 1$ and $\beta > 0$. Note that $\alpha = 1$ gives $S \sim \exp(\beta)$, which is the case where $L(A, T)$ is constant in T while $\alpha > 1$ is the case $L(A, T)$ is strictly monotonically increasing in T. With this in mind we restrict our proofs to two distinct cases: either $\frac{\partial L}{\partial T} > 0$ or $\frac{\partial L}{\partial T} \equiv 0$ for all $T > A$. The proofs hold with slight modifications for the more general condition that $L(A, T)$ is non-decreasing in T. Note that if $L(A, T)$ were *decreasing* in T then the implied threshold classifier is not necessarily a Bayes decision rule, hence it would not necessarily be an optimal strategy. As such we do not consider such cases at this time.

The following proposition characterizes the pure Nash equilibrium in the two-player game for the case that S satisfies the conditions in Assumption 1. We state the proposition without proof as it is a special case of Theorem 1. Recall a *pure Nash equilibrium* is a strategy pair $(A^*, T^*) \in \mathbb{R}^+ \times \mathbb{R}^+$ satisfying $T^* \in \sigma_1(A^*)$ and $A^* \in \sigma_2(T^*)$.

Proposition 1. *Suppose the random variable* S *satisfies Assumption 1. Then there exists a pure Nash equilibrium* $(A^*, T^*) \in \mathbb{R}^+ \times \mathbb{R}^+$ *in the two-player botnet detection game with* $T^* > A^* > v^{-1}(k)$.

2.2 A Large Population Botnet Detection Game

We now extend the two-player botnet detection game to a game with a large number of defenders in a network. In [48] network externalities in a security investment game between a large number of interconnected agents are studied. Their local mean field model focuses on the asymptotic properties of a security investment game in the limit as the number of agents grows indefinitely. We extend the model introduced in [48] by explicitly modeling the bot master as a strategic agent while allowing the legitimate users of the network to detect and remove infections. Following [48] we consider a sequence of rooted Erdős-Rényi random graphs $G(n, \lambda/n)$ and look for equilibria on the limiting graph as $n \to \infty$. By *rooted* we mean for each n the graph $G(n, \lambda/n)$ has a designated root node $v_{\varnothing}^{(n)}$ chosen uniformly at random from the set of n nodes. Because rooted Erdős-Rényi graphs converge to a rooted Galton-Watson Poisson Branching Process, denoted by $\mathcal{T}(\lambda)$, in the sense of local weak convergence [3], we restrict our analysis to $\mathcal{T}(\lambda)$. Results obtained on $\mathcal{T}(\lambda)$ can then be extended to equivalent results for the limiting process on $G(n, \lambda/n)$ as $n \to \infty$ using arguments similar to those found in [48]. The relevant convergence results for our model can be found in appendix A.8. For details on the objective method, local weak convergence and random distributional equations readers are referred to [3] and [2].

2.2.1 Agents, Costs and Utilities

We associate a unique defender a_i with each $i \in \mathcal{T}(\lambda)$. For simplicity we assume all defenders are homogeneous. In particular they have the same cost functions and each defender is equally likely to occupy any place in the network. Furthermore defenders are assumed to not know their exact locations in the network, but only have statistical knowledge of the network $\mathcal{T}(\lambda)$.

The bot master, denoted by b, does not occupy a node in the network $\mathcal{T}(\lambda)$, but instead wishes to gain control of the network by compromising defender nodes in $\mathcal{T}(\lambda)$. He does so by attempting to directly infect each node in $\mathcal{T}(\lambda)$ with some fixed probability of success. Furthermore infected nodes are capable of infecting neighboring nodes in $\mathcal{T}(\lambda)$. Let $X_i = 1$ if defender a_i is infected directly by the bot master b or indirectly via contagion from a neighboring defender in $\mathcal{T}(\lambda)$, and $X_i = 0$ otherwise.

Let $A \in \mathbb{R}^+$ be the strategic aggressiveness of the bot master. We again assume defenders use a threshold decision rule to detect infections. Let $T_i \in \mathbb{R}^+$ be the strategic threshold of defender a_i, and $Z_i \in \mathbb{R}^+$ the observation made by defender a_i. If $Z_i \geq T_i$ then the defender concludes his system has been compromised and takes measures to remediate the problem. If $Z_i < T_i$ then no infection is detected. In our model we assume that if the defender detects the infection, then he stops the infection and does not pass it on to his neighbors. The detection indicator random variable for defender a_i is $D_i = \mathbb{1}_{\{Z_i \geq T_i\}}$.

Let S_i denote the natural performance variability observed by defender a_i. We assume the S_i for all $i \in \mathcal{T}(\lambda)$ are i.i.d. copies of the random variable S which satisfy the conditions in Assumption 1. We then have $Z_i = AX_i + S_i$.

Because the observations Z_i are unreliable, defenders will have to balance the costs of false alarms (false positives) and missed detections (false negatives) in addition to the fixed cost of infection. We define indicator random variables for the outcomes of defender a_i's observation:

$$W^i_{\text{FP}} = D_i(1 - X_i),$$
$$W^i_{\text{FN}} = (1 - D_i)X_i,$$
$$W^i_{\text{TP}} = D_i X_i,$$

We associate the same costs for all defenders as in the two-player game, thus each defender a_i experiences a cost

$$C_i = cW^i_{\text{FP}} + v(A)W^i_{\text{FN}} + kW^i_{\text{TP}},$$

where $c \geq k \geq 0$ are constants and $v(A) \in C^1$ is non-decreasing and non-negative.

Due to the homogeneity of defenders in $\mathcal{T}(\lambda)$ our analysis we will focus on symmetric equilibria among the networked defenders. Let T denote a network-wide threshold played by all defenders. If defender a_i unilaterally deviates from the population threshold, i.e. $T_i \neq T$, then the expected cost of defender a_i, denoted by $\bar{C}_i$, will be a function of the strategy profile $(A, T, T_i) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$. Specifically

$$\bar{C}_i(A, T, T_i) = cE[W^i_{\text{FP}}] + v(A)E[W^i_{\text{FN}}] + kE[W^i_{\text{TP}}].$$

The best response correspondence for defender a_i is then

$$\sigma_i(A, T) = \arg\min_{T_i}\{\bar{C}_i(A, T, T_i)\}.$$

The bot master maximizes his expected utility by maximizing the cumulative computational resources stolen from across the compromised network. His utility will then depend on the fraction of defenders that are infected and miss the detection, as well as the degree to which he utilizes these infected computers. In the limit of a large population the expected proportion of agents experiencing a missed detection is equal to the probability of a false negative for a defender chosen uniformly at random from the network. Let defender $a_\varnothing$ be chosen uniformly at random from $\mathcal{T}(\lambda)$. Assuming a symmetric, network-wide threshold T, the bot

master's expected utility is

$$U(A,T) = g(A)E[W_{\mathrm{FN}}^{\varnothing}]. \tag{2.4}$$

As before we assume $g(A) \in C^2$ with $\frac{dg}{dA} > 0$, $\frac{d^2g}{dA^2} \leq 0$ and $g(0) = 0$. The bet response correspondence for the bot master is then

$$\sigma_b(T) = \arg\max_A \{g(A)E[W_{\mathrm{FN}}^{\varnothing}]\}.$$

2.2.2 Epidemic and Detection Processes

Our model on $\mathcal{T}(\lambda)$ is characterized by the following stochastic processes following [48]. Let $\chi_i \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(p)$ indicate a direct infection of defender a_i by the bot master b, and let $B_{kj} \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(q)$ indicate possible contagion between defenders a_k and a_j for all $k \neq j \in \mathcal{T}(\lambda)$. Furthermore we assume $B_{kj} = B_{jk}$ for all $k \neq j \in \mathcal{T}(\lambda)$. To take advantage of the structure of $\mathcal{T}(\lambda)$ we introduce the following Recursive Tree Process (RTP) [3]: For each $i \in \mathcal{T}(\lambda)$ let $\tilde{X}_i = 1$ if the infection reaches defender a_i either from a direct descendant in the rooted tree $\mathcal{T}(\lambda)$ or directly from the bot master, and $\tilde{X}_i = 0$ otherwise. Let $\tilde{D}_i$ be the indicator random variable indicating whether defender a_i detects such an infection. The defining equations for $\tilde{X}_i$ and $\tilde{D}_i$ are

$$\tilde{X}_i = 1 - (1 - \chi_i) \prod_{k \to i} (1 - B_{ki}(1 - \tilde{D}_k)\tilde{X}_k), \tag{2.5}$$

$$\tilde{D}_i = \mathbb{1}_{\{T_i \leq S_i + \tilde{X}_i A\}}. \tag{2.6}$$

Here $k \to i$ denotes that defender a_k is a direct descendant of defender a_i in the rooted tree. We also define the detection outcome indicator random variables

associated with the process $\{\tilde{X}_i\}_{i\in\mathcal{T}(\lambda)}$:

$$\tilde{W}^i_{FP} = \tilde{D}_i(1-\tilde{X}_i), \tag{2.7}$$

$$\tilde{W}^i_{FN} = (1-\tilde{D}_i)\tilde{X}_i, \tag{2.8}$$

$$\tilde{W}^i_{TP} = \tilde{D}_i\tilde{X}_i. \tag{2.9}$$

The introduction of the processes $\tilde{X}_i$ and $\tilde{D}_i$ is done to take advantage of the structure of $\mathcal{T}(\lambda)$. Note that $\tilde{X}_i$ and $\tilde{X}_j$ are independent of one another when a_i and a_j are the same distance away from the root node, while $\tilde{X}_i$ depends only on the children of a_i. Because of this structure we will be able to find analytical solutions for the distributions on $\tilde{X}_i$ and $\tilde{D}_i$. Moreover if $a_\varnothing$ is a *root defender* associated with a distinguished root node of $\mathcal{T}(\lambda)$, the distributions of $\tilde{X}_\varnothing$ and $\tilde{D}_\varnothing$ will corresponds exactly with the distributions of $X_\varnothing$ and $D_\varnothing$, respectively. As such we will be able to analytically derive the expected cost function of a root defender $a_\varnothing$, allowing a mean field analysis of the game.

As mentioned we focus our analysis on finding symmetric strategies for defenders in the network. Letting $T \in \mathbb{R}^+$ be a network-wide threshold we look for invariant processes satisfying (2.5)-(2.6). The fundamental Recursive Distributional Equations (RDEs) [2] which define the invariant process on $\mathcal{T}(\lambda)$ are as follows:

$$\tilde{X} \overset{d}{=} 1-(1-\chi)\prod_{k=1}^{N}(1-B_k(1-\tilde{D}_k)\tilde{X}_k), \tag{2.10}$$

$$\tilde{D}_k = \mathbb{1}_{\{T\le S_k+\tilde{X}_k A\}}. \tag{2.11}$$

The random variables $\chi \sim \text{Bernoulli}(p)$, $S_k \overset{\text{i.i.d.}}{\sim} F_S$, $B_k \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(q)$, and $N \sim \text{Poisson}(\lambda)$ are random variables independent of everything in the model.

The random variables $\tilde{X}$ and $\tilde{X}_k$, $k = 1, 2, ..., N$ are i.i.d. copies satisfying (2.10). Provided a distribution exists that satisfies (2.10) we can define the invariant decision outcome random variable

$$\tilde{D} \stackrel{d}{=} \mathbb{1}_{\{T \leq S + \tilde{X}A\}}, \tag{2.12}$$

as well as the invariant detection outcome indicator random variables:

$$\tilde{W}_{\text{FP}} \stackrel{d}{=} \tilde{D}(1 - \tilde{X}), \tag{2.13}$$

$$\tilde{W}_{\text{FN}} \stackrel{d}{=} (1 - \tilde{D})\tilde{X}, \tag{2.14}$$

$$\tilde{W}_{\text{TP}} \stackrel{d}{=} \tilde{D}\tilde{X}. \tag{2.15}$$

Equations (2.10)-(2.15) are the fundamental distributional equations describing all invariant solutions to equations (2.5)-(2.9). Because the propagation of the epidemic depends on missed detections, the distribution for $\tilde{W}_{\text{FN}}$ is of fundamental importance for finding solutions to the above distributional equations. The following result is analogous to Proposition 2 in [48].

Proposition 2. *For fixed $T \geq A \geq 0$, $0 < p \leq 1$ and $0 < q \leq 1$ the distributional equation for $\tilde{W}_{FN}$ has a unique solution: $\mathbb{P}(\tilde{W}_{FN} = 1) = 1 - \mathbb{P}(\tilde{W}_{FN} = 0) = h$, where $h = h(A, T, p, q, \lambda, F_S(\cdot))$ is the unique solution in $[0, 1]$ of the fixed point equation*

$$h = F_S(T - A)[1 - (1 - p)e^{-\lambda q h}]. \tag{2.16}$$

Proof. Let $h = \mathbb{P}(\tilde{W}_{\text{FN}} = 1)$. Since $P(\tilde{W}_{\text{FN}} = 1 | \tilde{X} = 0) = 0$ we have $h = \mathbb{P}(\tilde{W}_{\text{FN}} = 1 | \tilde{X} = 1)P(\tilde{X} = 1)$. Conditioned on $\tilde{X} = 1$ the distributional equation reduces to $\tilde{W}_{\text{FN}} \stackrel{d}{=} \mathbb{1}_{\{T > S + A\}}$, giving us $\mathbb{P}(\tilde{W}_{\text{FN}} = 1 | \tilde{X} = 1) = \mathbb{P}(T > S + A) = F_S(T - A)$.

The distributional equation for $\tilde{X}$ satisfies

$$\begin{aligned}
\mathbb{P}(\tilde{X}=0) &= \mathbb{P}((1-\chi)\prod_{k=1}^{N}(1-B_k(1-\tilde{D}_k)\tilde{X}_k)=1) \\
&= (1-p)\sum_{n=0}^{\infty}\left(\mathbb{P}\left(B_1\tilde{W}_{\mathrm{FN}}=0\right)\right)^n\mathbb{P}(N=n) \\
&= (1-p)\sum_{n=0}^{\infty}(1-qh)^n\frac{\lambda^n e^{-\lambda}}{n!} \\
&= (1-p)e^{-\lambda qh}.
\end{aligned}$$

Thus $\mathbb{P}(\tilde{X}=1) = 1-(1-p)e^{-\lambda qh}$, giving us (2.16).

Let $f(x,T,A) = F_S(T-A)[1-(1-p)e^{-\lambda qx}]$. Then f is continuous, increasing and concave in x. Since $f(0,T,A)>0$ and $f(1,T,A)\leq 1$ there must be a unique fixed point of f which depends on $A,T,p,q,\lambda,F_S(\cdot)$. □

With Proposition 2 the distributions of the remaining detection indicator random variables can be obtained in a similar manner.

Corollary 1. *Let $S\sim F_S(\cdot)$ with $T\geq A\geq 0$, $0<p\leq 1$ and $0<q\leq 1$. Then the distributional equations for $\tilde{W}_{FP}$ and $\tilde{W}_{TP}$ have unique solutions which depend on the distribution of $\tilde{W}_{FN}$. In particular if $E[\tilde{W}_{FN}]=h(A,T)$ then*

$$\begin{aligned}
E[\tilde{W}_{FP}] &= [1-F_S(T)](1-p)e^{-\lambda qh(A,T)}, \\
E[\tilde{W}_{TP}] &= [1-F_S(T-A)][1-(1-p)e^{-\lambda qh(A,T)}].
\end{aligned}$$

It is important to keep in mind that h depends on all parameters and choice variables of the model. In particular we will be interested in $h(A,T)$. We will often suppress this dependence in the notation for the sake of brevity. By the implicit function theorem h is differentiable in A and T provided that $f_S(T-A)\neq 0$. A

direct computation of the derivative of h with respect to A gives

$$\frac{\partial h}{\partial A} = -\frac{f_S(T-A)[1-(1-p)e^{-\lambda q h}]}{1-F_S(T-A)\lambda q(1-p)e^{-\lambda q h}}. \tag{2.17}$$

The following lemma guarantees the existence of the derivatives of h. It will also be useful in further analysis.

Lemma 1. *Define $\theta(A,T) = 1-\lambda q F_S(T-A)(1-p)e^{-\lambda q h}$ where h is defined as in Proposition 2. For any $\lambda q > 0$, $0 \le p < 1$ and $T \ge A \ge 0$, we have $0 < \theta(A,T) \le 1$.*

Proof. See Appendix A.1. □

From (2.17) we have $\frac{\partial h}{\partial A} \le 0$. Notice that $\frac{\partial h}{\partial T} = -\frac{\partial h}{\partial A}$, thus the same analysis above gives us $\frac{\partial h}{\partial T} \ge 0$. Also note that the dependence on A and T appears only in the form $T-A$ as arguments in $F_S(\cdot)$. Thus $h(A,T) = 0$ for all (A,T) pair with $A \ge T$. Furthermore for any finite A the limiting value of $h(A,T)$ as $T \to \infty$ is the same since for fixed finite A we have

$$\begin{aligned}\lim_{T\to\infty} h(A,T) &= \lim_{T\to\infty} F_S(T-A)[1-(1-p)e^{-\lambda q h(A,T)}] \\ &= 1-(1-p)e^{-\lambda q \lim_{T\to\infty} h(A,T)}.\end{aligned}$$

We will denote this limiting value by $h_\infty \triangleq \lim_{T\to\infty} h(A,T)$ and observe that it satisfies the fixed point equation

$$h_\infty = 1-(1-p)e^{-\lambda q h_\infty}. \tag{2.18}$$

Note that in this limit we recover the probability of infection in [48].

2.2.3 Equilibrium Analysis

In this section we derive the best response correspondences for the decentralized network of defenders and the bot master. We then state and prove the existence of a pure Nash equilibria in the decentralized botnet detection game.

Defender Population Best Response

To determine a decentralized network's best response, we investigate what happens when a single selfish defender deviates from a population threshold T that all other defenders in the network are playing. Because the root of the network $\mathcal{T}(\lambda)$ is chosen uniformly at random, we consider the associated root defender $a_\varnothing$ to be a "typical" deviant defender. Thus for a fixed strategy A and a fixed symmetric network-wide strategy T, we can define a deviant defender's strategy $T_\varnothing$, expected cost function $\bar{C}_\varnothing(A, T, T_\varnothing)$ and best response correspondence $\sigma_\varnothing(A, T)$. For a fixed strategy A we are interested in finding an equilibrium network strategy T^* such that $T^* \in \sigma_\varnothing(A, T^*)$. That is, if the entire network is playing the strategy T^*, then a single deviant node has no incentive to deviate from that threshold.

From the distributional equations (2.10)-(3.4) it is clear that if the root defender $a_\varnothing$ changes his threshold $T_\varnothing \neq T$, this will change his detection outcome, but it will not change his probability of infection. For the root we need to introduce new distributional equations:

$$\begin{aligned}
\tilde{D}_\varnothing &\stackrel{d}{=} \mathbb{1}_{\left\{S+\tilde{X}A \geq T_\varnothing\right\}}, \\
\tilde{W}^\varnothing_{\mathrm{FP}} &\stackrel{d}{=} \tilde{D}_\varnothing(1-\tilde{X}), \\
\tilde{W}^\varnothing_{\mathrm{FN}} &\stackrel{d}{=} (1-\tilde{D}_\varnothing)\tilde{X}, \\
\tilde{W}^\varnothing_{\mathrm{TP}} &\stackrel{d}{=} \tilde{D}_\varnothing\tilde{X}.
\end{aligned}$$

As mentioned, due to the tree structure of $\mathcal{T}(\lambda)$ the distributions for $\tilde{X}_\varnothing$ and $\tilde{D}_\varnothing$ are equal to the distributions for $X_\varnothing$ and $D_\varnothing$. We thus arrive at the following proposition whose proof is analogous to Proposition 2 and Corollary 1.

Proposition 3. *Let $S \sim F_S(\cdot)$ with $T, T_\varnothing \geq A \geq 0$, $0 < p \leq 1$ and $0 < q \leq 1$. Then the distributional equations for $W^\varnothing_{FP}, W^\varnothing_{FN}$ and $W^\varnothing_{TP}$ have unique solutions which depend on the distribution of W_{FN}. If $E[W_{FN}] = h(A,T)$ then the distributions are given by the following:*

$$\begin{aligned}
E[W^\varnothing_{FP}] &= [1 - F_S(T_\varnothing)](1-p)e^{-\lambda q h(A,T)}, \\
E[W^\varnothing_{FN}] &= F_S(T_\varnothing - A)[1 - (1-p)e^{-\lambda q h(A,T)}], \\
E[W^\varnothing_{TP}] &= [1 - F_S(T_\varnothing - A)][1 - (1-p)e^{-\lambda q h(A,T)}].
\end{aligned}$$

With the above proposition we can determine a root defender's expected cost function:

$$\begin{aligned}
\bar{C_\varnothing}(A,T,T_\varnothing) = c[1 - F_S(T_\varnothing)](1-p)e^{-\lambda q h} \\
+ (k + (v(A) - k)F_S(T_\varnothing - A))[1 - (1-p)e^{-\lambda q h}].
\end{aligned}$$

To simplify the notation we introduce the function $\ell(A) \equiv v(A) - k$. If $v(A)$ is constant on any interval, all subsequent results hold with slight modification. Thus without loss of generality we assume $\ell(A)$ to be differentiable and strictly monotonically increasing. We furthermore assume $\ell(0) = 0$. This gives us

$$\begin{aligned}
\bar{C_\varnothing}(A,T,T_\varnothing) = c[1 - F_S(T_\varnothing)](1-p)e^{-\lambda q h} \\
+ (k + \ell(A)F_S(T_\varnothing - A))[1 - (1-p)e^{-\lambda q h}]. \quad (2.19)
\end{aligned}$$

It is important to notice that h depends on A and T only, not on $T_\varnothing$. In particular

we have $\frac{\partial h}{\partial T_\varnothing} \equiv 0$. Thus the best response for a deviant root defender is

$$\sigma_\varnothing(A,T) = \arg\min_{T_\varnothing}\{c[1-F_S(T_\varnothing)](1-p)e^{-\lambda q h} + \ell(A)F_S(T_\varnothing - A)[1-(1-p)e^{-\lambda q h}]\}.$$

Taking the first derivative we then have

$$\frac{\partial \bar{C}_\varnothing}{\partial T_\varnothing} = -c f_S(T_\varnothing)(1-p)e^{-\lambda q h} + \ell(A) f_S(T_\varnothing - A)[1-(1-p)e^{-\lambda q h}]. \tag{2.20}$$

Since $f_S(T_\varnothing - A) = 0$ for all $T_\varnothing < A$ we have $\frac{\partial C}{\partial T_\varnothing} < 0$ for all $T_\varnothing \in (0, A)$. Thus any global minima will be in the interval $[A, \infty)$. Recall the likelihood ratio function:

$$L(A, T_\varnothing) \triangleq \frac{f_S(T_\varnothing - A)}{f_S(T_\varnothing)}.$$

Using (2.20) we define the following function:

$$V(A,T) \triangleq \frac{c}{\ell(A)} \frac{(1-p)e^{-\lambda q h}}{1-(1-p)e^{-\lambda q h}}. \tag{2.21}$$

From (2.20) we see that the single function $V(A,T) - L(A, T_\varnothing)$ has the same sign as $\frac{\partial \bar{C}_\varnothing}{\partial T_\varnothing}$. As such we have the relations

$$\frac{\partial \bar{C}_\varnothing}{\partial T_\varnothing} \overset{<}{\underset{\geq}{=}} 0 \iff L(A, T_\varnothing) \overset{<}{\underset{\geq}{=}} V(A,T). \tag{2.22}$$

Note that $L(A, T_\varnothing)$ depends on A and $T_\varnothing$ only, while $V(A,T)$ depends on A and T only. Thus for fixed A and T the functional form of $L(A, T_\varnothing)$ alone will determine the optimal response of a deviant defender. In particular if $L(A, T_\varnothing)$ is non-decreasing then $\bar{C}_\varnothing(A, T, T_\varnothing)$ will be quasi-convex in $T_\varnothing$. However as the network threshold T varies so too will the optimal response of a deviant defender.

To better understand how the functions $V(A,T)$ and $L(A,T_\varnothing)$ affect a deviant defender's best response, consider the following:

$$\frac{\partial V}{\partial T} = -\frac{c}{\ell(A)} \frac{\lambda q(1-p)e^{-\lambda q h}\frac{\partial h}{\partial T}}{[1-(1-p)e^{-\lambda q h}]^2} \leq 0.$$

Thus $V(A,T)$ is a non-increasing function of T and strictly monotonically decreasing for $T > A > 0$. We wish to show there exists a unique, pure, symmetric Nash equilibrium among defenders in the network in response to the bot master aggressiveness A, i.e. there exists a unique T^* such that $T^* \in \sigma_\varnothing(A,T^*)$. It turns out the monotonicity of $L(A,T_\varnothing)$ is sufficient to guarantee this.

Proposition 4. *For fixed $A \in \mathbb{R}^+$ if $L(A,T_\varnothing)$ is non-decreasing then there exists a unique network threshold $T^* \in \mathbb{R}^+$ such that $T^* \in \sigma_\varnothing(A,T^*)$, i.e. there is a unique, pure, symmetric Nash equilibrium among defenders in the network in response to the bot master strategy A.*

Proof. Fix $A \in \mathbb{R}^+$. Since $V(A,T)$ is strictly monotonically decreasing for $T > A$ and $L(A,T_\varnothing)$ is non-decreasing, then three possibilities exist: 1) there exists a unique value $\tilde{T} \in [A,\infty)$ such that $L(A,\tilde{T}) = V(A,\tilde{T})$, 2) $L(A,T_\varnothing) < V(A,T)$ for all $T_\varnothing, T \geq A$, and 3) $L(A,T_\varnothing) > V(A,T)$ for all $T_\varnothing, T \geq A$. Suppose the first case is true. Then there exists some values $\epsilon_1, \epsilon_2 \geq 0$ such that $L(A,T_\varnothing) < V(A,\tilde{T})$ for $T_\varnothing < \tilde{T} - \epsilon_1$, $L(A,T_\varnothing) = V(A,\tilde{T})$ for $\tilde{T} - \epsilon_1 \leq T_\varnothing \leq \tilde{T} + \epsilon_2$, and $L(A,T_\varnothing) > V(A,\tilde{T})$ for $T_\varnothing > \tilde{T} + \epsilon_2$. By (2.22) we have $\sigma_\varnothing(A,\tilde{T}) = [\tilde{T} - \epsilon_1, \tilde{T} + \epsilon_2]$. Clearly $\tilde{T} \in \sigma_\varnothing(A,\tilde{T})$. Furthermore, by the uniqueness of $\tilde{T}$ satisfying $L(A,\tilde{T}) = V(A,\tilde{T})$, it is the only value satisfying $\tilde{T} \in \sigma_\varnothing(A,\tilde{T})$.

Now suppose $L(A,T_\varnothing) < V(A,T)$ for all $T_\varnothing, T > A$. Then (2.22) implies that $\bar{C_\varnothing}(A,T,T_\varnothing)$ is monotonically decreasing in $T_\varnothing$ for all $T, T_\varnothing > A$. Thus $\sigma_\varnothing(A,T) = +\infty$ for all $A > 0$ and $T > A$. In particular $\sigma_\varnothing(A,\infty) = \infty$ and

$T^* = \infty$.

Finally suppose $L(A, T_\varnothing) > V(A, T)$ for all $T_\varnothing, T > A$. Now (2.22) implies that $\bar{C_\varnothing}(A, T, T_\varnothing)$ is monotonically increasing in $T_\varnothing$ for all $T, T_\varnothing > A$. Thus $\sigma_\varnothing(A, T) = A$ for all $A > 0$ and $T > A$. In particular $\sigma_\varnothing(A, A) = A$ and $T^* = A$. □

We can now define a network population best response correspondence $\sigma_p : \mathbb{R}^+ \to \mathbb{R}^+$ which maps a strategy A to the symmetric, mutual best response among defenders in the network for which no individual has an incentive to unilaterally deviate. Proposition 4 gives sufficient conditions under which $\sigma_p(A)$ is a single valued function, which can then be written as follows:

$$\sigma_p(A) = \begin{cases} A & \text{if } L(A,T) > V(A,T) \text{ for all } T, \\ +\infty & \text{if } L(A,T) < V(A,T) \text{ for all } T, \\ T^* & \text{o.w., where } T^* \text{ is the unique solution to } L(A,T^*) = V(A,T^*). \end{cases}$$

Bot Master Best Response

Because we are interested in finding pure Nash equilibria we wish to find under which conditions $U(A, T)$ is strictly quasi-concave and $\sigma_b(T)$ is single valued. Using (2.16) and (2.17) a first order optimality condition for a strategy $A^* \in \mathbb{R}^+$ can be expressed as

$$\frac{g(A^*)}{g'(A^*)} = \frac{F_S(T - A^*)}{f_S(T - A^*)}\theta(A^*, T). \tag{2.23}$$

Notice we have used the definition for $\theta(A, T)$ from Lemma 1. The following proposition gives a sufficient condition for the strict quasi-concavity of the bot master's expected utility function.

Proposition 5. *For $T > 0$, if there exists a unique $A^* \in \mathbb{R}^+$ satisfying (2.23), then $U(A,T)$ is strictly quasi-concave with a maximum at A^* and $\sigma_b(T) = A^*$. If $T = 0$ then $U(A,T) = 0$ for all $A \in \mathbb{R}^+$ and $\sigma_b(0) \equiv \mathbb{R}^+$.*

Proof. Let $T > 0$. Since $g(0) = 0$ we have $U(0,T) = 0$. Furthermore $U(A,T) = 0$ for all $A \geq T$ and $U(A,T) > 0$ for $A \in (0,T)$. Since $h(A,T)$ is a differentiable function in A, so too is $U(A,T)$. Thus by Rolle's Theorem there exists at least one A^* in the open interval $(0,T)$ such that $\frac{\partial U}{\partial A}\Big|_{A=A^*} = 0$. In addition since $U(0,T) = U(T,T) = 0$ and $U(A,T) > 0$ for $A \in (0,T)$, there must be at least one global maximum in the open interval $(0,T)$. Clearly if (2.23) has a unique solution A^* we must have $U(A^*,T)$ as a global maximum. Consequently A^* is the unique optimal response to the strategy T giving $\sigma_b(T) = A^*$. Now let $T = 0$. Since $h(A,T) = 0$ when $A \geq T$ we get $U(A,T) = 0$ for all A. Consequently $A \in \sigma_b(0)$ for all $A \in \mathbb{R}^+$. □

We now state a sufficient condition on $F_S(\cdot)$ to guarantee the strict quasiconcavity of the bot master's expected utility function for $T > 0$.

Proposition 6. *For fixed $T > 0$, if $\frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\right] < 1$ for all $A < T$ then (2.23) has a unique solution $A^* \in (0,T)$.*

Proof. We begin by establishing the following facts:

$$\frac{d}{dA}\left[\frac{g(A)}{g'(A)}\right] \geq 1, \tag{2.24}$$

$$\frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\theta(A,T)\right] < 1. \tag{2.25}$$

Directly differentiating gives $\frac{d}{dA}\left[\frac{g(A)}{g'(A)}\right] = 1 - \frac{g(A)g''(A)}{g(A)^2} \geq 1$. Differentiating $\frac{F_S(T-A)}{f_S(T-A)}\theta(A,T)$

we obtain

$$\frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\theta(A,T)\right] = \frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\right]\theta(A,T) + \left[\frac{1-\lambda q F_S(T-A)}{\theta(A,T)}\right](1-\theta(A,T)).$$

By assumption $\frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\right] < 1$ and by Lemma 1 we have $0 < \theta(A,T) \leq 1$. It remains to be shown $\left[\frac{1-\lambda q F_S(T-A)}{\theta(A,T)}\right] \leq 1$. If $F_S(T-A) \geq \frac{1}{\lambda q}$ then this is clearly the case. On the other hand if $F_S(T-A) < \frac{1}{\lambda q}$ then

$$1-\lambda q F_S(T-A) < 1 - F_S(T-A)(1-p)\lambda q e^{-\lambda q h} = \theta(A,T),$$

giving us $\left[\frac{1-\lambda q F_S(T-A)}{\theta(A,T)}\right] < 1$ from which we obtain $\frac{\partial}{\partial A}\left[\frac{F_S(T-A)}{f_S(T-A)}\theta(A,T)\right] < 1$. Since $\frac{g(0)}{g'(0)} = 0$ and $\frac{F_S(T)}{f_S(T)}\theta(0,T) > 0$, properties (2.24) and (2.25) guarantee that (2.23) has a unique solution. By Proposition 5 the result follows. □

Nash Equilibrium

We are now ready to prove the existence of a pure Nash equilibria in the large-population botnet detection game between the network of decentralized defenders and the bot master. As we have seen the function $V(A,T) = \frac{c}{\ell(A)}\frac{(1-p)e^{-\lambda q h(A,T)}}{1-(1-p)e^{-\lambda q h(A,T)}}$ played an important role in determining the response functions of the defenders. In the proceeding analysis we will need the limiting values of this function. For ease of exposition we define the following functions:

$$V_0(A) \triangleq \lim_{T \downarrow A} V(A,T) = \frac{c}{\ell(A)}\frac{1-p}{p}, \tag{2.26}$$

$$V_\infty(A) \triangleq \lim_{T \to \infty} V(A,T) = \frac{c}{\ell(A)}\frac{(1-p)e^{-\lambda q h_\infty}}{1-(1-p)e^{-\lambda q h_\infty}}, \tag{2.27}$$

where h_∞ is the unique solution to (2.18). Note that by the monotonicity of $\ell(A)$ we obtain $\frac{dV_0}{dA}, \frac{dV_\infty}{dA} < 0$. Furthermore the monotonicity of $V(A,T)$ in T gives us $V_\infty(A) < V_0(A)$ for all $A \in (0,\infty)$.

Finally, to establish our result we will need the following lemmas which give us important properties of the best response correspondences $\sigma_b(T)$ and $\sigma_p(A)$. For the next lemma we define the values A_0 and A_∞ as solutions to the following equations, provided unique, non-negative solutions exist:

$$L_\infty(A_\infty) = V_\infty(A_\infty), \tag{2.28}$$

$$L_0(A_0) = V_0(A_0). \tag{2.29}$$

Given our assumption that $\ell(0) = 0$ and $\lim_{A\to\infty} \ell(A) = +\infty$ we see from (2.27) that $\lim_{A\downarrow 0} V_\infty(A) = +\infty$ and $\lim_{A\to\infty} V_\infty(A) = +\infty$. By the monotonicity of $\ell(A)$ we obtain the monotonicity of $V_\infty(A)$. Given the monotonicity of $L_\infty(A)$ we are guaranteed for (2.28) to have a unique, non-negative solution. Under the assumptions of our model (2.29) is guaranteed to have a unique, non-negative solution only when $\frac{\partial L}{\partial T} \equiv 0$. When $\frac{\partial L}{\partial T} > 0$ we have $L_0(A) \equiv 0 < V_0(A)$ for all finite A, in which case $A_0 = +\infty$ and is not needed in the subsequent proofs.

Lemma 2. *Given the expected cost function $\bar{C}_\varnothing(A,T,T_\varnothing)$ in (2.19), the following properties of $\sigma_p(A)$ hold.*

1. *For $A \geq 0$, $\sigma_p(A) \geq A$.*
2. *For $0 \leq A \leq A_\infty$, $\sigma_p(A) = \infty$.*
3. *For $A > A_\infty$, $\sigma_p(A)$ is continuously differentiable with $\lim_{A\downarrow A_\infty} \sigma_p(A) = \infty$.*
4. *$\limsup_{A\to\infty} \sigma_p(A) - A = 0$. If $\frac{\partial L}{\partial T} \equiv 0$ then for $A \geq A_0$, $\sigma_p(A) = A$.*

Proof. See appendix A.2 □

Lemma 3. *Given the expected utility function $U(A,T)$ in (2.4) with $F_S(\cdot)$ satisfying the properties of Lemma 6, the following properties of $\sigma_b(T)$ hold.*

1. *For $T > 0$, $0 < \sigma_b(T) < T$.*

2. *For $T > 0$, $\sigma_b(T)$ is continuously differentiable.*

3. $\limsup_{T\to\infty} \sigma_b(T) = \infty$ *with*
   $\limsup_{T\to\infty} (T - \sigma_b(T)) > 0$.

4. *For all $A \in (0,\infty)$ there exists a finite $\tilde{T} > 0$ such that $\sigma_b(\tilde{T}) = A$.*

Proof. See Appendix A.3 □

We now state and prove the existence of a pure Nash equilibrium in the large-population botnet game on $\mathcal{T}(\lambda)$.

Theorem 1. *Let $S \sim F_S$ as in Assumption 1. Then in the large-population, botnet detection game on $\mathcal{T}(\lambda)$ with homogeneous defenders, there exists a pure Nash equilibrium $(A^*, T^*) \in \mathbb{R}^+ \times \mathbb{R}^+$ with all defenders playing the symmetric strategy T^*.*

Proof. First consider the case where $\frac{\partial L}{\partial T} \equiv 0$. Suppose $A \leq A_\infty$. It follows from Lemma 2 property 2 that for all $A \leq A_\infty$ we have $\sigma_p(A) = \infty$, hence $U(A, \sigma_p(A)) = \lim_{T\to\infty} Ah(A,T) = Ah_\infty$. But for any finite $M > 0$ we have $U(A+M, T^*) = (A+M)h_\infty > Ah_\infty = U(A, T^*)$ for all such A. It follows that A is not a best response and there are no pure Nash equilibrium with $A^* \in [0, A_\infty]$. On the other hand suppose $A \geq A_0$. By Lemma 2 property 4 $\sigma_p(A) = A$. Then we have $U(A, \sigma_p(A)) = U(A,A) = Ah(A,A) = 0$, and for sufficiently small $\epsilon > 0$ we have $U(A-\epsilon, A) = (A-\epsilon)h(A-\epsilon, A) > 0$ and the bot master will benefit from decreasing his aggressiveness . Clearly such an A is not a best response, and any strategy set (A^*, T^*) with $A^* \geq A_0$ is not a Nash equilibrium.

We thus restrict our attention to $A \in (A_\infty, A_0)$. By property 4 of Lemma 3 there exists a finite value T_∞ such that $\sigma_b(T_\infty) = A_\infty$ and a finite value $T_0 > 0$ such that $\sigma_b(T_0) = A_0$. By Lemma 2 we have $\sigma_p(\sigma_b(T_\infty)) = \infty > T_\infty$ and $\sigma_p(\sigma_b(T_0)) = \sigma_b(T_0) < T_0$. In other words when looking along the T axis at T_∞ the function $\sigma_p(\cdot)$ is above the function $\sigma_b(\cdot)$ while at T_0 the function $\sigma_p(\cdot)$ is below the function $\sigma_b(\cdot)$. By the continuity of both $\sigma_p(\cdot)$ and $\sigma_b(\cdot)$ the functions must cross at some point (A^*, T^*) giving us $\sigma_b(T^*) = A^*$ and $\sigma_p(A^*) = T^*$.

The proof for $\frac{\partial L}{\partial T} > 0$ is similar to the above with one exception. In this case we have $\sigma_p(A) > A$ for all A with $\lim_{A\to\infty} \sigma_p(A) = A$. Thus the continuity of $\sigma_p(A)$ and $\sigma_b(T)$ is not enough to guarantee the response functions cross.

Suppose $\sigma_b(\cdot)$ and $\sigma_p(\cdot)$ do not cross. From property 1 of Lemma 3 we have $\lim_{T\downarrow 0} \sigma_b(T) = 0$. Thus there must exist some finite T_∞ such that $0 < \sigma_b(T_\infty) < A_\infty$. Since $\sigma_p(A) = \infty$ for all $A \leq A_\infty$ we then have $\sigma_p(\sigma_b(T_\infty)) = \infty$. Thus $\sigma_p(\sigma_b(T_\infty)) > T_\infty$. By our assumption that $\sigma_b(\cdot)$ and $\sigma_p(\cdot)$ do not cross we must have $\sigma_p(\sigma_b(T)) > T$ for all $T > 0$. From Lemma 3 we have $T > \sigma_b(T)$ for all $T > 0$. Together this gives us the following:

$$\sigma_p(\sigma_b(T)) > T > \sigma_b(T). \tag{2.30}$$

From Lemmas 2 and 3 we have $\limsup_{T\to\infty} \sigma_b(T) = \infty$ and $\lim_{A\to\infty} \sigma_p(A) - A = 0$. It follows that

$$\limsup_{T\to\infty} [\sigma_p(\sigma_b(T)) - \sigma_b(T)] = \lim_{A\to\infty} [\sigma_p(A) - A] = 0.$$

Then by (2.30) $\limsup_{T\to\infty}[T - \sigma_b(T)] = 0$. But this contradicts $\limsup_{T\to\infty}[T - \sigma_b(T)] > 0$ from Lemma 3. Hence $\sigma_b(\cdot)$ and $\sigma_p(\cdot)$ must cross at least once. □

2.3 A Centralized, Large-Population Botnet Game

2.3.1 Centralized Expected Cost and Best Response

Another approach to the large-population botnet detection game is to consider the effects of a centralized planner on the game. Suppose there is a single player P_0 whose strategy space is $\mathbb{R}^+$ and chooses a threshold $T \in \mathbb{R}^+$ for *all* defenders to follow. We keep the same expected cost function and best response correspondence as for an individual defender. As such we define the centralized expected cost function and best response correspondence, respectively, as follows:

$$\begin{aligned}\bar{C}_c(A,T) \triangleq\ & c[1-F_S(T)](1-p)e^{-\lambda q h(A,T)} \\ & + (k+\ell(A)F_S(T-A))[1-(1-p)e^{-\lambda q h(A,T)}], \\ \sigma_c(A) \triangleq\ & \arg\min_T \{\bar{C}_c(A,T)\}.\end{aligned}$$

The difference is that now $\frac{\partial h}{\partial T} > 0$, since a central planner chooses a threshold for the entire population. As such the results from the decentralized case do not apply. In order to obtain equilibrium results for the centralized case we first need to establish the strict quasi-concavity of the central planner's expected utility in the symmetric threshold strategy T.

For ease of exposition we use the function $\rho(A,T)$ as defined in (A.1) in the proof of Lemma 2 in appendix A.2. With this new notation we can rewrite $h(A,T)$ and $\theta(A,T)$.

$$h(A,T) = F_S(T-A)(1-\rho(A,T)) \tag{2.31}$$

$$\rho(A,T) = (1-p)e^{-\lambda q h(A,T)} \tag{2.32}$$

$$\theta(A,T) = 1 - F_S(T-A)\lambda q \rho(A,T) \tag{2.33}$$

To prove the strict quasi-concavity of the central planner's expected utility we will need the function

$$h_\infty(\phi) = 1 - (1-p)e^{-\phi h_\infty(\phi)}$$

as well as the following technical lemmas.

Lemma 4. *Let $h_\infty(\phi)$ be defined as above. Then*

1. $\lim_{\phi\to 0} h_\infty(\phi) = p$.
2. $\lim_{\phi\to\infty} h_\infty(\phi) = 1$.
3. $\lim_{\phi\to 0} \phi e^{-\phi h_\infty(\phi)} = 0$.
4. $\lim_{\phi\to\infty} \phi e^{-\phi h_\infty(\phi)} = 0$.

Proof. The proofs follow directly from the definition of $h_\infty(\phi)$. □

Lemma 5. *For any $\phi \geq 0$, and $p \in [0,1]$ we have*

$$\phi e^{-\phi h_\infty(\phi)} \lessgtr \frac{1}{2(1-p)} \iff \phi e^{-\phi} \lessgtr \frac{e^{-\frac{1}{2}}}{2(1-p)}.$$

Proof. See appendix A.4. □

Lemma 6. *For any $\phi \geq 0$, and $p \in [0,1]$ we have*

$$1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} + (1-p)e^{-\phi h_\infty(\phi)} \geq 0.$$

Proof. See appendix A.5 □

Lemma 7. *For any $T \geq A$, $\lambda q > 0$ and $p \in [0,1]$, if $\frac{f_S(T-A)}{f_S(T)}$ is non-decreasing in T then*

$$1 - 2F_S(T-A)\lambda q\rho(A,T) + \rho(A,T) \geq 0.$$

Proof. See appendix A.6 □

We are now ready to prove the strict quasi-concavity of the Centralized Planner's expected cost function.

Proposition 7. *Fix $A > 0$, $\lambda q > 0$ and $p \in [0,1]$. If $\frac{f_S(T-A)}{f_S(T)}$ is non-decreasing in T, then the Central Planner's expected cost is strictly quasi-convex in T.*

Proof. Fix $A \geq 0$. First observe that $\frac{\partial C}{\partial T} \leq 0$ for $T \in [0, A]$. We thus consider $T > A$. Taking derivatives of (2.31) - (2.33) we obtain

$$\begin{aligned}
\frac{\partial h}{\partial T} &= \frac{f_S(T-A)(1-\rho)}{\theta}, \\
\frac{\partial \rho}{\partial T} &= -\frac{f_S(T-A)\lambda q\rho(1-\rho)}{\theta}, \\
\frac{\partial \theta}{\partial T} &= -\frac{f_S(T-A)\lambda q\rho}{\theta}\left[1 - F_S(T-A)\lambda q\right].
\end{aligned}$$

Taking first derivatives of $\bar{C}_c(A,T)$ we obtain

$$\begin{aligned}
\frac{\partial \bar{C}_c}{\partial T} = &-cf_S(T)\rho - c[1-F_S(T)]\lambda q\rho\frac{f_S(T-A)(1-\rho)}{\theta} \\
&+ \ell(A)f_S(T-A)(1-\rho) + (k + \ell(A)F_S(T-A))\lambda q\rho\frac{f_S(T-A)(1-\rho)}{\theta}.
\end{aligned}$$

By assumptions on $F_S(\cdot)$ we have $f_S(T-A)(1-\rho)>0$, and we can write

$$\frac{\partial \bar{C}_c}{\partial T}\frac{\theta}{\ell(A)f_S(T-A)\rho(1-\rho)} = -\frac{c}{\ell(A)}\frac{f_S(T)}{f_S(T-A)}\frac{\theta}{1-\rho} - \frac{c}{\ell(A)}[1-F_S(T)]\lambda q + \frac{1}{\rho} + \frac{k}{\ell(A)}\lambda q.$$

Define the function

$$M(A,T) \triangleq -\frac{c}{\ell(A)}\frac{f_S(T)}{f_S(T-A)}\frac{\theta}{1-\rho} - \frac{c}{\ell(A)}[1-F_S(T)]\lambda q + \frac{1}{\rho} + \frac{k}{\ell(A)}\lambda q. \quad (2.34)$$

Since $\frac{\theta}{\ell(A)f_S(T-A)(1-\rho)} > 0$ for all $T > A$ we have $\text{sign}\,(M) = \text{sign}\left(\frac{\partial \bar{C}_c}{\partial T}\right)$ for all $T > A$, so for fixed A the first order condition of optimality of $\bar{C}_c(A,T)$ in T is $M(A,T) = 0$. We proceed to show that $M(A,T)$ is strictly monotonically increasing for $T > A$, from which it follows that $\bar{C}_c(A,T)$ is strictly quasi-convex for $T > A$.

$\frac{\partial M}{\partial T} > 0$ if and only if

$$f_S(T-A)\lambda q\frac{1-\rho}{\theta\rho} + \frac{c}{\ell(A)}f_S(T)\lambda q > \frac{c}{\ell(A)}\frac{\partial}{\partial T}\left[\frac{f_S(T)}{f_S(T-A)}\right]\frac{\theta}{1-\rho} - \frac{c}{\ell(A)}\frac{f_S(T)\lambda q\rho}{\theta(1-\rho)}\left[2-F_S(T-A)\lambda q(1+\rho)\right],$$

if and only if

$$\frac{c}{\ell(A)}\frac{f_S(T)\lambda q}{\theta(1-\rho)}\left[1-2F_S(T-A)\lambda q\rho+\rho\right] > \frac{c}{\ell(A)}\frac{\partial}{\partial T}\left[\frac{f_S(T)}{f_S(T-A)}\right]\frac{\theta}{1-\rho} - f_S(T-A)\lambda q\frac{1-\rho}{\theta\rho}.$$

For notational convenience we introduce the following functions:

$$u(A,T) \triangleq \frac{c}{\ell(A)} \frac{f_S(T)\lambda q}{\theta(1-\rho)} \left[1 - 2F_S(T-A)\lambda q\rho + \rho\right],$$
$$v(A,T) \triangleq \frac{c}{\ell(A)} \frac{\partial}{\partial T} \left[\frac{f_S(T)}{f_S(T-A)}\right] \frac{\theta}{1-\rho} - f_S(T-A)\lambda q \frac{1-\rho}{\theta\rho}.$$

Thus $\frac{\partial M}{\partial T} > 0$ if and only if $u(A,T) > v(A,T)$. By assumption $\frac{\partial}{\partial T}\left[\frac{f_S(T)}{f_S(T-A)}\right] \leq 0$ for all $A, T \geq 0$ which gives us $v(A,T) \leq 0$. Furthermore it was assumed that $f_S(x) > 0$ for $x > 0$ which implies $v(A,T) < 0$ for $T > A$. On the other hand by Lemma 7 it follows that $u(A,T) \geq 0$. Hence $u(A,T) > v(A,T)$ and the result follows. □

2.3.2 Centralized Nash Equilibrium

We are now ready to establish the existence of a pure, symmetric (among defenders) Nash equilibrium for the centralized botnet game on $\mathcal{T}(\lambda)$. As we have seen the function $M(A,T)$ in (2.34) was crucial in determining the best responses for the central planner. In the proceeding analysis we will need the limiting values of this function. First note the following: $\lim_{T \downarrow A} \rho(A,T) = 1 - p$ and $\lim_{T \downarrow A} \theta(A,T) = 1$. For ease of exposition we use the value $\rho_\infty = (1-p)e^{\lambda q h_\infty}$ defined in (A.2) in the proof of Lemma 2. Similarly we define θ_∞ as the limiting value of $\theta(A,T)$ as $T \to \infty$:

$$\theta_\infty \triangleq \lim_{T \to \infty} \theta(A,T) = 1 - (1-p)\lambda q e^{-\lambda q h_\infty}.$$

We define the limiting values of the function $M(A,T)$ as follows:

$$M_0(A) \triangleq \lim_{T \downarrow A} M(A,T), \tag{2.35}$$

$$M_\infty(A) \triangleq \lim_{T \to \infty} M(A,T). \tag{2.36}$$

From the definition we have

$$M_\infty(A) = -\frac{c}{\ell(A)} \frac{1}{L_\infty(A)} \frac{\theta_\infty}{1-\rho_\infty} + \frac{1}{\rho_\infty} + \frac{k}{\ell(A)} \lambda q.$$

For the case $\frac{\partial L}{\partial T} > 0$ we have $\lim_{T \downarrow A} L(A,T) \equiv 0$ for all A. It follows that $M_0(A) = -\infty$ for all A. For the case $\frac{\partial L}{\partial T} \equiv 0$ we have

$$M_0(A) = -\frac{c}{\ell(A)} \frac{1}{L_0(A)} \frac{1}{p} - \frac{c}{\ell(A)} (1 - F_S(A)) \lambda q + \frac{1}{1-p} + \frac{k}{\ell(A)} \lambda q.$$

Similar to the decentralized case we use these functions to restrict the strategy space in which we look for equilibria. To do so we seek non-negative solutions to the following equations:

$$M_0(A) = 0, \tag{2.37}$$

$$M_\infty(A) = 0. \tag{2.38}$$

Unlike the decentralized case we cannot guarantee the strict monotonicity of (2.35) and (2.36), as such we cannot guarantee the existence or uniqueness of solutions to (2.37) and (2.38). If solutions to (2.37) and (2.38) exist we define A_0^{c} and A_∞^{c}

as follows:

$$A_0^{\mathrm{c}} \triangleq \inf\{x : M_0(x) = 0\}, \tag{2.39}$$

$$A_\infty^{\mathrm{c}} \triangleq \sup\{x : M_\infty(x) = 0\}. \tag{2.40}$$

Under the assumptions of our model equation (2.37) is guaranteed to have a non-negative solution only when $\frac{\partial L}{\partial T} \equiv 0$. When $\frac{\partial L}{\partial T} > 0$ we have $M_0(A) = -\infty$ for all A, in which case A_0^{c} is not defined, nor needed in subsequent proofs. Equation (2.38) is guaranteed to have a non-negative solution when $\lambda q < \frac{c}{k}\frac{\theta_\infty}{1-\rho_\infty}$. In what follows we assume (2.38) has a non-negative solution, however this is not a necessary assumption to arrive at the existence of a Nash equilibrium. The existence proof is nearly identical for the case that (2.38) has no solution. We focus on a single case for brevity.

To establish our result we will need the following lemmas, analogous to Lemmas 2 and 3, which give us important properties of the best response correspondences $\sigma_c(A)$ and $\sigma_b(T)$.

Lemma 8. *The following properties of $\sigma_c(A)$ hold.*

1. *For $A \geq 0$, $\sigma_c(A) \geq A$.*

2. *For $A > A_\infty^c$, $\sigma_c(A)$ is continuously differentiable with $\lim_{A \downarrow A_\infty^c} \sigma_c(A) = \infty$.*

3. *$\limsup_{A\to\infty} \sigma_c(A) - A = 0$. If $\frac{\partial L}{\partial T} \equiv 0$ then for $\limsup_{A \uparrow A_0^c} \sigma_c(A) - A = 0$.*

Proof. See Appendix A.7. □

Lemma 9. *The following properties of $\sigma_b(T)$ hold.*

1. *For $T > 0$, $0 < \sigma_b(T) < T$.*

2. *For $T > 0$, $\sigma_b(T)$ is continuously differentiable.*

3. $\limsup_{T\to\infty} \sigma_b(T) = \infty$ *with* $\limsup_{T\to\infty} (T - \sigma_b(T)) > 0$.

4. For all $A \in (0, \infty)$ *there exists a finite* $\tilde{T} > 0$ *such that* $\sigma_b(\tilde{T}) = A$.

Proof. The proof is unchanged from the decentralized case, as $\sigma_b(T)$ is independent of $\sigma_c(A)$. □

We are now ready to prove the existence of a pure Nash equilibrium in the centralized botnet detection game.

Theorem 2. *Let* $S \sim F_S$ *as in Assumption 1. Then in the centrally-planned, large-population, botnet detection game on* $\mathcal{T}(\lambda)$ *with homogeneous defenders, there exists a pure, symmetric (among defenders) Nash equilibrium.*

Proof. Given Lemmas 9 and 8 the proof is analogous to the decentralized case. □

2.4 Numerical Examples

In this section we examine numerical results to study network effects on defender threshold strategies. We will consider both the non-adversarial and adversarial setting. That is we consider two scenarios for the signal strength/aggressiveness A: fixed and strategic. When we consider A to be fixed we will be interested in computing the *Price of Anarchy* for the decentralized defenders in relation to the central planner. When we consider a strategic adversary and A is not fixed we will compute the pure, defender-symmetric Nash equilibria in the botnet detection game for both the decentralized and centralized games.

Recall that λ is the average number of neighbors in the underlying random graph while q is the probability of contagion between neighbors given the presence of the infection. Thus λq represents the defining parameter for the underlying

infection graph. Our parameter study involves numerically solving for Nash equilibria for the values $\lambda q \in [1, 15]$.

Recall that k is the cost of raising an alarm regardless of its outcome, c is the cost of a false alarm and $v(A)$ is the cost of a false negative given a bot master strategy of A. By necessity $c > k$. Thus we introduce the difference $r \triangleq c - k$ so that we can write $c = r + k$. Throughout we have assumed that $v(A) \geq k$. This means the cost of a missed detection is at least as great as the cost of raising an alarm, no matter the aggressiveness of the bot master. In our numerical examples we have made the further assumption that $v(A) = mA + k$ for some positive value m. We fix $k = 1$ and consider varying the values of r and m to understand the effects of the defender cost function on the best response strategies. Other parameter values used in the numerical examples are $p = 0.01$ and $S \sim \text{gamma}(2, 2)$.

We begin by pointing out why decentralized and centralized defenders arrive at different strategies in the first place. A decentralized, selfish individual will seek to minimize his own costs by unilaterally deviating from the current network strategy. Furthermore because the network population is large, a deviant decentralized defender assumes that changing his own strategy has a negligible effect on the network population as a whole. As such he assumes the probability of infection is fixed in relation to his own choice of strategy. He only has control over his own probability of a false alarm or missed detection conditioned on the probability of infection. The decentralized network threshold is then chosen via Nash equilibria: a threshold in which no individual has an incentive to deviate. On the other hand the central planner takes into account the fact that when he changes the thresholds across the network, this will change the unconditioned probability of false alarms and missed detections. Thus the central planner does a proper optimization which

takes the network effects into account. We then have two distinct methods for determining best responses to an attack value A: Nash equilibrium and differential calculus. There is no reason *a priori* why these two methods should arrive at the same solution, and as we will see below they almost surely do not. We numerically study these differences and their effects on the social welfare of the networked population.

2.4.1 Non-strategic attacker

We first examine the case of non-adversarial interdependent detection, i.e. the attack/signal strength A is fixed. This way we can explore network effects without the complications associated with equilibrium solutions between the attacker and defenders. Recall that for any fixed $A \geq 0$ there exists a unique symmetric, Nash equilibrium population response $T_d^*(A) = \sigma_p(A)$ as well as a unique centrally planned, symmetric population best response $T_c^*(A) = \sigma_c(A)$. We will be interested in the effects that the model parameters have on these best responses. In particular we will vary the network contagion potential parameter λq as well as the attack strength A. We will fix the cost of raising an alarm at $k = 1$ and consider various values for the detection costs r and m.

Threshold Comparison

Figure 2.1 shows us where in the λq-A plane $T_c^* > T_d^*$ and $T_c^* < T_d^*$ for four different combinations of r and m. The parameter r (cost of false alarm minus cost of raising an alarm) appears to determine the qualitative features of the plots in Figure 2.1: For "small" false alarm values ($r = 0.1$) we see $T_c^* < T_d^*$ for all computed values of $(\lambda q, A)$. For "large" false alarm values ($r = 5$) we see two distinct regimes, one in which $T_c^* < T_d^*$ and one in which $T_c^* < T_d^*$.

To understand why the parameter r controls this qualitative feature it is helpful to consider what a selfish, individual defender will want to do when the entire population plays the central planner's optimal threshold. First consider the case that r is small as in Figures 2.1a and 2.1c. In these cases r is small in relation to k which means the cost of a false alarm is relatively small in relation to the cost of raising an alarm. In this case $T_c^* < T_d^*$, i.e. the decentralized defenders are less vigilant (more tolerant of observed disturbances). When the cost of raising an alarm is high it stands to reason that a central planner will take this into account and raise fewer alarms. This is accomplished by choosing a relatively high threshold. However the central planner cannot choose too high a threshold since this will increase the missed detection rate and hence increase the rate of infection in the network, and higher infection rates will increase the number of alarms raised, which are costly! In contrast a selfish decentralized defender does not take the change in infection probability into account when choosing his threshold. Thus the temptation to choose high thresholds to avoid raising alarms is not tempered by a fear of increasing the infection rate which itself leads to more alarms. So if the network is playing the central planner's best response threshold, a selfish individual defender will want to unilaterally lower his own threshold to reduce alarms raised. This can be thought of as a free-rider-like effect. The cost of raising an alarm dominates the decision process and defenders are reluctant to raise alarms at all. Thus they tend to free-ride on the vigilance of the rest of the population. Since this line of reasoning is independent of λq or A we see the same phenomenon ($T_c^* < T_d^*$) throughout the λq-A plane.

The case when r is small as in Figures 2.1b and 2.1d is more interesting as we see the two distinct regimes. In these cases r is large relative to k which means that false alarms are costly compared to raising an alarm. Thus the central

planner counteracts this by choosing relatively higher threshold. This lowers the false alarm rate. But why do the decentralized agents settle on a *lower* threshold (which will result in more false alarms) at all when the cost of false alarms is high?

If A is sufficiently small and λq is sufficiently large, then the epidemic spreads more and infections are harder to detect. Thus the red regions in Figures 2.1b and 2.1d correspond to the cases where there is a high potential for contagion and the attacker is stealthy. This would suggest that the potential for missed detection is high. (High contagion potential + stealthy attacker = missed detection!) When considering a threshold to choose, one may be tempted to select a low threshold to mitigate this potentially high number of missed detections. But this is not the optimal strategy. Why not? Precisely for two reasons. The first reason is that the red regions appear only when r is high relative to k. That means the cost of a false alarm is high, and most of the cost is coming from the penalty for making a wrong decision and not the cost of raising the alarm. Thus a lower threshold will increase the number of false alarms, especially when we take the network effects into account: Lower thresholds mean more detections, which means fewer contagions which means a lower probability of infection. So you are raising more alarms even though the probability of infection has gone down. The net effect will be a high probability of false alarm (which are costly). The second factor is the stealthiness of the attacker. Not only does this make the attacker harder to detect, but it also makes missed detections less costly. For sufficiently small A, the decrease in missed detections that result from choosing a small threshold does not compensate for the increase in false alarms, because the low A means the missed detections were not costly to begin with. So by choosing low thresholds across the entire network you have increased the number of false alarms (which are costly) and decreased the number of missed detections (which are not costly)!

This reasoning shows us why the central planner does not choose a low threshold in this region. But why does a selfish individual not do the same? Why does a selfish individual lower his threshold in the face of costly false alarms? This seems counter intuitive. The reason is due to network externalities. An individual agent in a large network assumes (somewhat justifiably) that if he unilaterally changes his threshold then this will not affect the overall probability of infection in the network. Thus in the red region, if he lowers his threshold from the centrally planned threshold, it does not decrease the probability of infection, it only increases his probability of false alarms. Similarly the lower threshold decreases his probability of missed detection but it does not make infections less likely. Because λq is relatively large, as mentioned before, the epidemic is more likely to reach the individual, thus his reasoning leads him to conclude that he will benefit by lowering his threshold to counter the high probability of infection. But this reasoning, when applied to each selfish individual in the network puts a downward pressure on the network wide threshold relative to the central planner's optimal threshold. The Nash equilibrium response then winds up being lower than the central planner's optimal threshold, thus reducing the probability of infection across the network. This results in more false positives (which are costly). Thus by ignoring the network effect of choosing a lower threshold, a selfish individual overestimates the benefit from the reduction in missed detections.

Conversely we may ask, why doesn't the central planner lower his threshold to match the decentralized defenders? It is because he takes network effects into account. He sees that if he were to lower his threshold, then this would decrease the spread of the infection, which is a good thing, but this would also result in an increase in false alarms. And since false alarms are costly, the net effect would be that he is worse off than if he played a higher threshold.

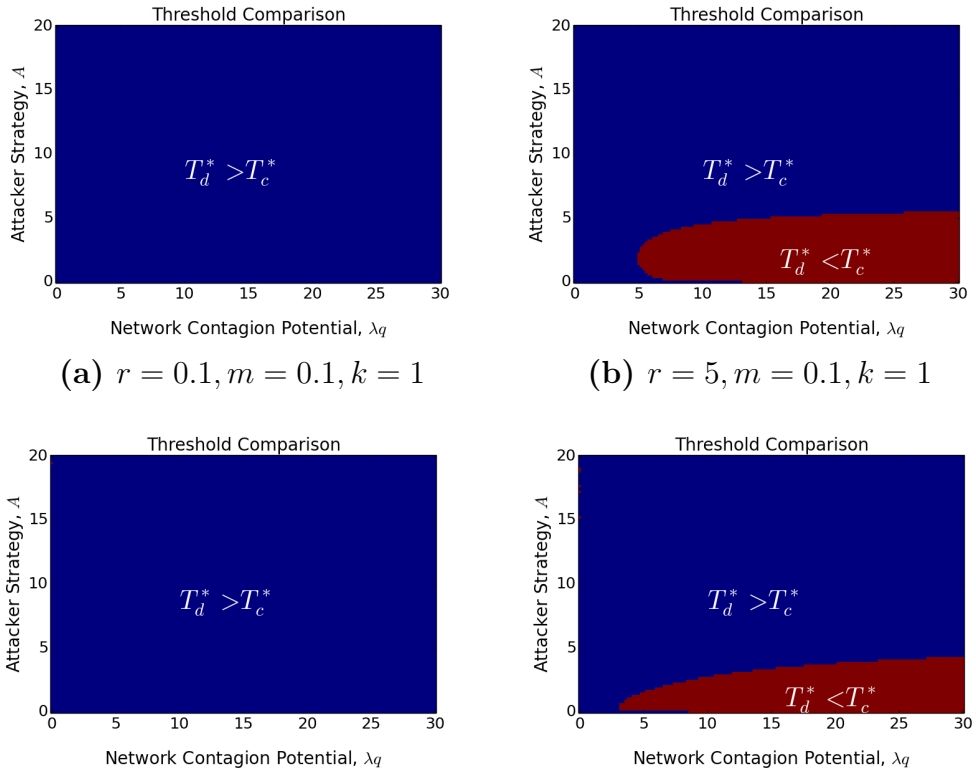

(a) $r = 0.1, m = 0.1, k = 1$ **(b)** $r = 5, m = 0.1, k = 1$

(c) $r = 0.1, m = 5, k = 1$ **(d)** $r = 5, m = 5, k = 1$

Figure 2.1: Threshold comparison. Red: $T_c^* > T_d^*$, Blue: $T_c^* < T_d^*$.

Note that we do not see the above phenomenon for low values of λq or large values of A. Why is this? Since the cost of a missed detection increase with A, when A is sufficiently large then the cost of a missed detection will dominate the cost of a false alarm and the central planner will take this into account, i.e. he lowers his threshold. This reduces the spread of the epidemic enough that selfish individual defenders will think they can free ride on the vigilance of the network. Similarly when λ is sufficiently small the probability of infection is low, and when this probability is small enough, selfish individual defenders will reason that they can reduce false alarm costs by raising their threshold from the central planner's optimal threshold. In both cases this is again similar to a free-rider-like effect in which selfish individuals free-ride off the vigilance of the network.

Price of Anarchy

For a fixed $(\lambda q, A)$ point we define the *Price of Anarchy* (PoA) as the ratio of decentralized cost to the centralized cost at their respective best response strategies:

$$PoA(\lambda q, A) \triangleq \frac{C(\lambda q, A, T_d^*(A))}{C(\lambda q, A, T_c^*(A))}.$$

Figure 2.6 shows the Price of Anarchy at each computed point in the λq-A plane. The parameter m (multiplicative component of missed detection cost) appears to determine the qualitative features of the plots in Figure 2.6. For "small" values, i.e. $m = 0.1$, we see a greater range in the Price of Anarchy. In particular the value of PoA takes on larger values and takes higher values over a larger range of the parameter space. For "large" values, i.e. $m = 5$ we see less variation in the Price of Anarchy.

First consider the case where m is "small", i.e. $m = 0.1$, shown in Figures 2.2a and 2.2b. We first observe that for values of A that are either very large or very small the PoA ≈ 1. When A is very large then there are two factors which lead to this observation: First, the infection is easier to detect with thresholds close to A, and second, the cost of a missed detection is relatively high. The first point leads to low false alarm costs. Combining this with the second point it can be seen that choosing a threshold close to A will minimize the expected cost. Thus a central planner will choose a threshold close to A and a selfish defender cannot free-ride on this vigilance without risking the missed detection costs. This can also be seen from the results of Lemmas 2 and 8: $T_c^*(A) - A \downarrow 0$ and $T_d^*(A) - A \downarrow 0$ as $A \to \infty$. Thus we have $|T_c^*(A) - T_d^*(A)| \downarrow 0$ as $A \to \infty$. As a consequence we will have $PoA(\lambda q, A) = \frac{C(\lambda q, A, T_d^*(A))}{C(\lambda q, A, T_c^*(A))} \to 1$ as $A \to \infty$. On the other hand when A is sufficiently small there is virtually no cost from missed detections. Thus expected costs are minimized by choosing large thresholds for both centralized

and decentralized defenders.

Now consider the case where again $m = 0.1$ but A is not too high or too low. When A is not too high a selfish individual defender will not suffer a great deal from missed detections (provided the rest of the network is playing a threshold that is not too high). Thus he will free ride and choose a higher than socially optimal threshold. If A is not too high and not too low and λq is relatively large, then this tendency to free ride winds up being more costly to the decentralized agents in terms of missed detections. Why? Because the high value of λq increases the spread of the infection so defenders are more likely to get infected. Since A is not too large, many of these infections will be harder to detect. Since A is not too small, the penalty for a missed detection is not trivial. The net effect is that the Price of Anarchy is greater in this regime. This explains the large hot spots in the graph for larger values of λq.

Suppose λq is large and A is not too high nor too low, so that the above reasoning does not apply. This means that the infection potential is high while the strength of the attack is high enough to generate a non-trivial cost of missed detection but low enough so that it is non-trivial to detect the presence of the infection. Now suppose m is low. This reduces the cost of missed detections for a given A when compared to high m. The central planner takes account of these factors and picks an optimal threshold $T_c^*(A, \lambda q)$. Because A is larger, selfish decentralized defenders tend to free-ride off of this strategy, i.e. they select higher thresholds than the central planner. Because m is low and because they do not take the network effects into account, they only consider reducing false alarms. They are free-riding on the vigilance of the network and raise fewer alarms. As a result decentralized defenders increase the spread of the infection in the network relative to the centrally planned network. Moreover the lower m leads

decentralized defenders to believing they can benefit from raising less alarms by raising their thresholds *more* than they would with higher m. So the lower missed detection cost induces them to free ride more (play higher thresholds) than they would were missed detection costs higher. These higher thresholds lead to higher infection rates. While this does in fact reduce false alarms it increases the missed detection cost more than they anticipated due to the higher infection rates.

On the other hand when m is larger decentralized defenders are not tempted to free ride as much. Thus when they raise their thresholds the subsequent increase in infection rates is not as severe and they do not deviate as far from the central planner's expected cost. This potentially explains why we see the large "hot spots" for the PoA when λq is large and A takes on medium values (Figures 2.2a and 2.2b).

The high POA that occur for low values of λq occur in all plots, thus we suspect that they do not depend on the parameters studied here. One possibility is that they are related to the epidemic threshold $\lambda q = 1$ where the infection graph has a fully connected large component. Because we have assumed an Erdős-Rényi random graph we know there exists an epidemic threshold at $\lambda q = 1$ [24]. Perhaps for a range of A values when the giant connected component appears the central planner is able to leverage the connectivity of the graph by controlling the spread of the infection in the network. After all this ability is his main advantage in selecting strategies. When there is no giant connected component he cannot utilize this ability.

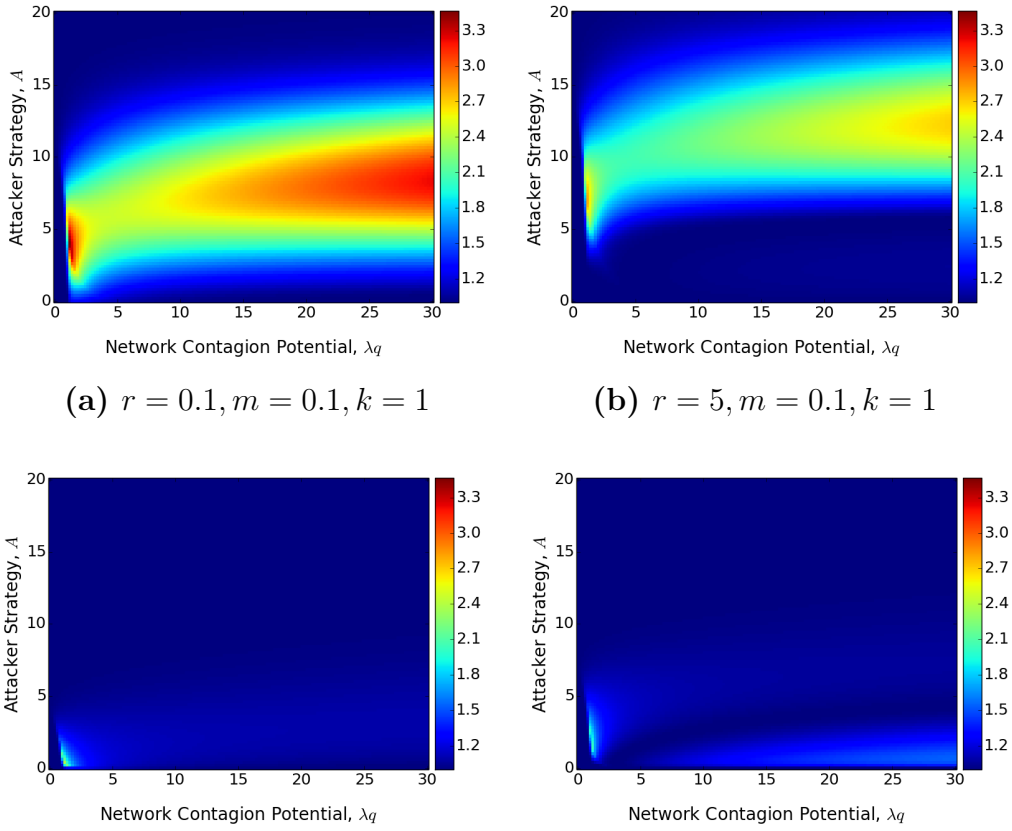

(a) $r = 0.1, m = 0.1, k = 1$ **(b)** $r = 5, m = 0.1, k = 1$

(c) $r = 0.1, m = 5, k = 1$ **(d)** $r = 5, m = 5, k = 1$

Figure 2.2: Price of Anarchy: $PoA(\lambda q, A) = C_d(A, T_d^*(A), \lambda q)/C_c(A, T_c^*(A), \lambda q)$.

2.4.2 Adversarial Interdependent Detection

Using the Nash equilibrium results from the previous section we can numerically estimate Nash equilibria for both decentralized and centralized games. Though we are not guaranteed of the uniqueness of a pure, symmetric Nash equilibrium in either the centralized or decentralized game, visual inspection of the best response functions in the strategy space suggests uniqueness of the equilibria in the examples considered. Figure 2.7 shows an example plot of the best response functions in the strategy space $\mathbb{R}^+ \times \mathbb{R}^+$ where the single crossings of the best response functions can be clearly seen. Given such a plot one can esti-

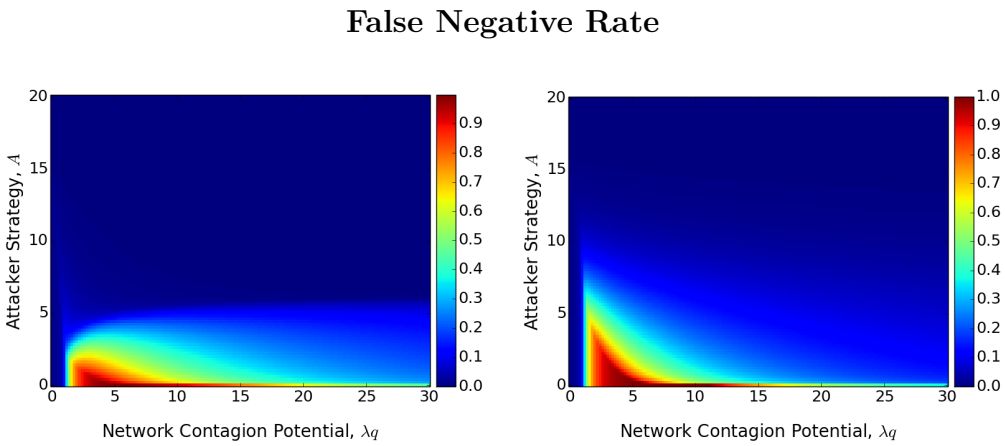

(a) Centralized **(b)** Decentralized

Figure 2.3: False negative rate $h(\lambda q, A)$ for centralized and decentralized defenders as a function of A and λq. This can be interpreted as the relative size of the botnet after defenders have performed their detection and removal processes. Parameters are $r = 5, m = 0.1, k = 1$.

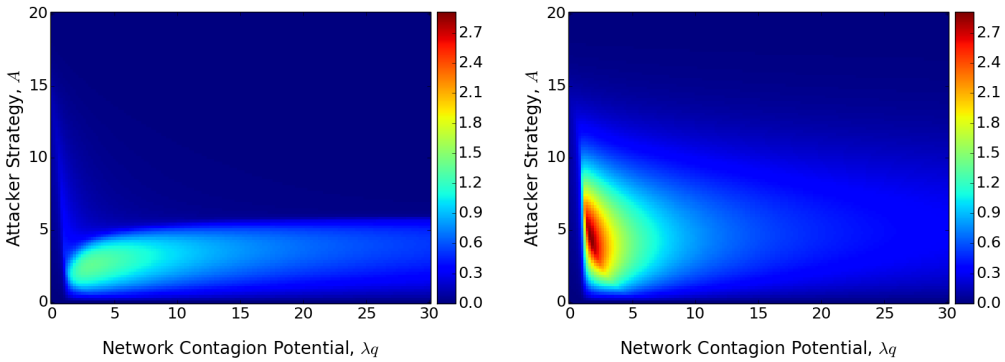

(a) Centralized **(b)** Decentralized

Figure 2.4: Bot master utility $U(A, T^*(A))$ for centralized and decentralized defenders as a function A and λq. Parameters are $r = 5, m = 0.1, k = 1$.

mate a bounded region in the strategy space where the Nash equilibrium exists. Given these bounds one can numerically solve for the equilibrium strategies by

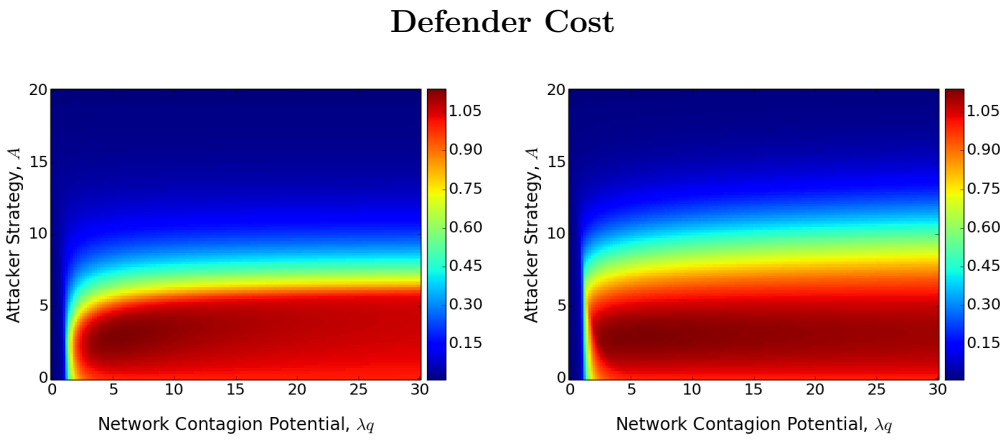

(a) Centralized **(b)** Decentralized

Figure 2.5: Defender cost $C(A, T^*(A))$ for centralized and decentralized defenders as a function A and λq. Parameters are $r = 5, m = 0.1, k = 1$.

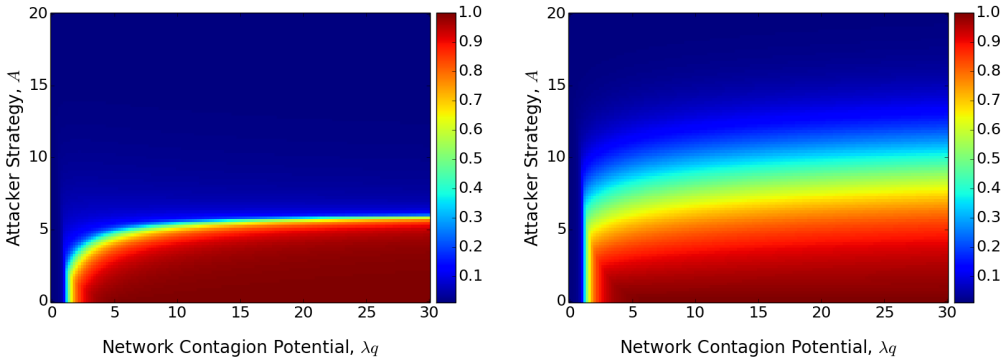

(a) Centralized **(b)** Decentralized

Figure 2.6: Probability of infection for centralized and decentralized defenders as a function A and λq. Parameters are $r = 5, m = 0.1, k = 1$.

minimizing a *regret* function $R(A, T)$ defined by

$$R(A, T) \triangleq (T - \sigma(A))^2 + (A - \sigma_a(T))^2,$$

where $\sigma(A)$ is either the decentralized or centralized best response function depending on which Nash equilibria you are solving for. This optimization problem

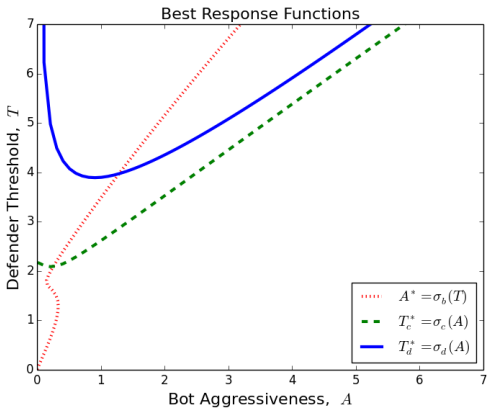

Figure 2.7: Best response functions in strategy space $\mathbb{R}^+ \times \mathbb{R}^+$. Parameters are $\lambda q = 5$, $k = 5$, $c = 1 + k$, $v(A) = A + k$, $p = 0.01$.

can be solved relatively efficiently with standard numerical solvers. The Nash equilibrium results in previous sections guarantee the existence of solutions for the numerical schemes.

For a given parameter set let (A_c^*, T_c^*) and (A_d^*, T_d^*) denote pure, symmetric Nash equilibria in the centralized and decentralized games, respectively. For the same parameter set let $\tilde{X}_c^*$ and $\tilde{X}_d^*$ denote indicator random variables for infection at equilibrium in the centralized and decentralized games, respectively. We denote the expected costs for a root defender at equilibrium by $\bar{C}_c^* = \bar{C}_c(A_c^*, T_c^*)$ and $\bar{C}_\varnothing^* = \bar{C}_\varnothing(A_d^*, T_d^*)$ in the centralized and decentralized botnet detection games respectively. Similarly we denote the bot master expected cost at equilibrium by $U_c^* = U(A_c^*, T_c^*)$ and $U_d^* = U(A_d^*, T_d^*)$. Finally we let p_c^* and p_d^* be the probability of infection at equilibrium, i.e. $P(\tilde{X}_c^* = 1) = p_c^*$ and $P(\tilde{X}_d^* = 1) = p_d^*$, in the centralized and decentralized games respectively.

We explore two parameter sets in solving for Nash equilibria in each game. In both cases we fix the cost of raising an alarm, $k = 1$, and set the cost of missed detections to be $v(A) = 5A + 1$, i.e. $m = 5$. We again vary the network contagion parameter $\lambda q \in (0, 15)$. Finally we consider two values for the cost of

a false alarm: $c = 0.1 + k = 1.1$ and $c = 5 + k = 6$. Recall this is equivalent to the parameter r taking on values 0.1 and 5 respectively. For comparison these parameter sets correspond to those used in Figures 2.1c and 2.1d.

When analyzing the defender strategies at equilibrium with an attacker, it is important to point out the main difference from the previous numerical results, which analyzed defender strategies for fixed A. Because the attacker is strategic, for any fixed set of parameters the decentralized defenders and central planner will select different thresholds. These different thresholds will result in the strategic attacker playing different attack strategies. This will be of importance in understanding some of the more counterintuitive results.

First consider the case that $r = 0.1$. As we saw from Figure 2.1c this parameter set is characterized by the decentralized defenders free-riding on the vigilance of the central planner. This results in the decentralized defenders choosing higher thresholds than the central planner for any fixed A. When we consider a strategic attacker and solve for the Nash equilibria we see that again, decentralized defenders choose higher thresholds than the central planner, as can be seen in Figure 2.9a. Consequently the attacker plays higher aggressiveness strategies against the decentralized defenders than against the central planner at equilibrium, as can be seen in Figure 2.9b. Moreover the attacker is able raise his aggressiveness in equilibrium in such a way that $T_d^* - A_d^* > T_c^* - A_c^*$, as can be seen in Figure 2.10a. As a result infection rates and missed detection rates go up (Figures 2.10b and 2.10d), as does the cost of a missed detection, since $A_d^* > A_c^* \implies v(A_d^*) > v(A_c^*)$. Because false positive costs are low the central planner is able to fair better while experiencing higher false positive rates at equilibrium. Furthermore since $v(A_d^*) > v(A_c^*)$ the missed detection costs the central planner experiences are lower as well. The overall effect is that the central planner is able to fair better in equilibrium for all

observed values of λq, as can be seen in Figure 2.9.

As in the non-adversarial setting, the case of higher false alarm costs is more interesting due to the different qualitative parameter regimes seen in Figure 2.1d. As we saw when $r = 5$ there are two distinct network externalities present among decentralized defenders. One is the same as when $r = 0.1$, namely that decentralized defenders free-ride on the vigilance of the central planner and choose higher thresholds. When the Nash equilibrium settles in this parameter regime we observe similar outcomes as described in the preceding paragraph: Decentralized defenders play higher thresholds at equilibrium compared to the central planner, and the attacker in turn is more aggressive in equilibrium against the decentralized defenders than against the central planner. This results in higher infection rates, missed detection rates, missed detection costs and overall expected costs for the decentralized defenders.

The second network externality occurs in the high-network-contagion-potential, stealthy-attacker parameter regime. In this regime the decentralized defenders choose higher thresholds than the central planner. When the Nash equilibria settle in this parameter regime we observe similar outcomes: Decentralized defenders choose lower thresholds than the central planner, as can be seen in Figure 2.11a and the attacker is less aggressive against the decentralized defenders than against the central planner at equilibrium, as can be seen in Figure 2.11b. Moreover the attacker lowers his aggressiveness at equilibrium in such a way that $T_d^* - A_d^* < T_c^* - A_c^*$, as can be seen in Figure 2.12a. As a result infection rates and missed detection rates go down (Figures 2.12b and 2.12d), as does the cost of a missed detection, since $A_d^* < A_c^* \implies v(A_d^*) < v(A_c^*)$.

When we examine the overall effect on the expected cost of the defenders in both the decentralized and centralized games, we see something very strik-

ing: For a certain range of λq we observe the decentralized defenders obtaining a *lower* expected cost at equilibrium compared to the central planner, i.e. $\bar{C}_{\varnothing}(A_d^*, T_d^*) < \bar{C}_c(A_c^*, T_c^*)$, as can be seen in Figure 2.11c. This may seem counter-intuitive because if both the central planner and the decentralized defenders faced the same aggressiveness A, the centralized planner would by definition have to do at least as well or better than the decentralized defenders. However, as mentioned earlier the equilibrium aggressiveness played against a centralized planner is not necessarily the same as the equilibrium aggressiveness played against decentralized defenders, i.e. $A_c^* \neq A_d^*$.

The following intuition explains how the difference between A_c^* and A_d^* tends to work against the central planner and for the decentralized defenders in this particular parameter regime. Consider the centralized equilibrium (A_c^*, T_c^*) and think about the natural dynamic of response and counter response between defenders and bot master that would follow if we replaced the centralized defenders with decentralized defenders. In this regime decentralized defenders play lower thresholds than a central planner prescribes, leading to a lower missed detection rate for the decentralized defenders. In response to lower thresholds, the bot master plays a lower aggressiveness than it would against a central planner, i.e. $A_d^* < A_c^*$. The lower A_d^* decreases the missed detection rate as well as the cost of a missed detection $v(A)$ that the decentralized defenders take on for playing the lower threshold T_d^*. The net effect is that the overall expected cost resulting from the decentralized defenders playing a lower threshold looks better when we take into account the accompanying change in the attacker's strategy.

Figure 2.8 shows iterated best response updates between the central planner and the bot master when starting at a decentralized Nash equilibrium profile. In figure 2.8a it can be seen that when false positive costs are relatively low ($r = 0.1$)

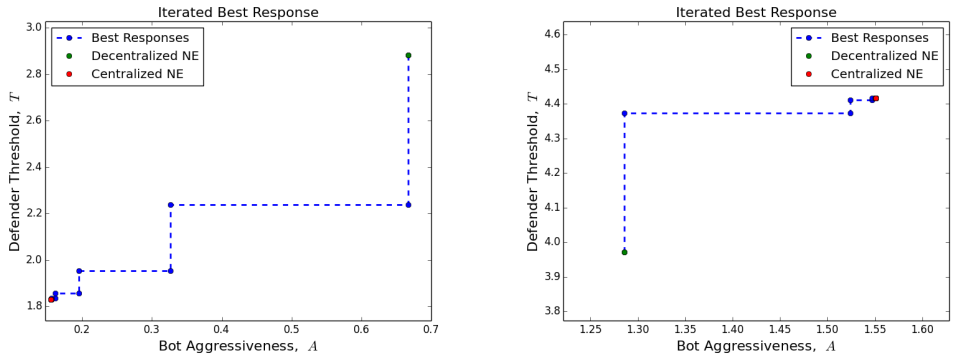

(a) $r = 0.1, m = 1, k = 1$ **(b)** $r = 5, m = 1, k = 1$

Figure 2.8: Iterated best response updates between central planner and bot master when starting at a decentralized Nash equilibrium profile.

the iterated best response updates converge to a centralized Nash equilibrium in which the central planner plays a lower threshold than decentralized defender play. On the other hand, in figure 2.8b it can be seen that when false positive costs are relatively high ($r = 5$) the iterated best response updates converge to a centralized Nash equilibrium in which the central planner plays a higher threshold than decentralized defenders play.

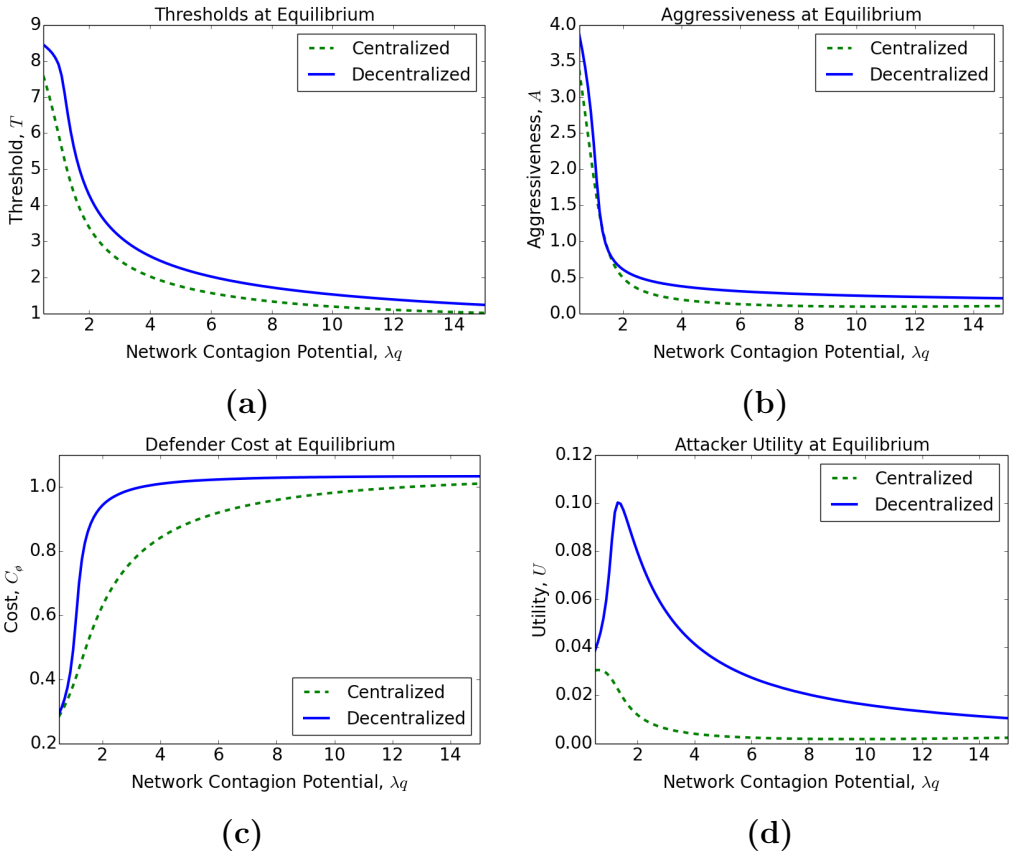

Figure 2.9: Strategies and payoffs at Nash equilibrium. Parameters are $r = 0.1, m = 5, k = 1$.

2.5 Conclusion

We have examined a novel interdependent detection game motivated by the botnet security threat. Both decentralized and centrally planned defenders were considered. Furthermore a single strategic adversarial attacker was included in the game. The existence of a pure, defender-symmetric Nash equilibrium was proved for both decentralized and centralized games. Network effects on defender strategy were explored via numerical approximation of the equilibria. It was seen that for fixed attack strategies decentralized defenders choose threshold strategies that either too high or too low. Furthermore when the attacker is strategic certain regimes of model parameters give rise to counterintuitive results, such as a central

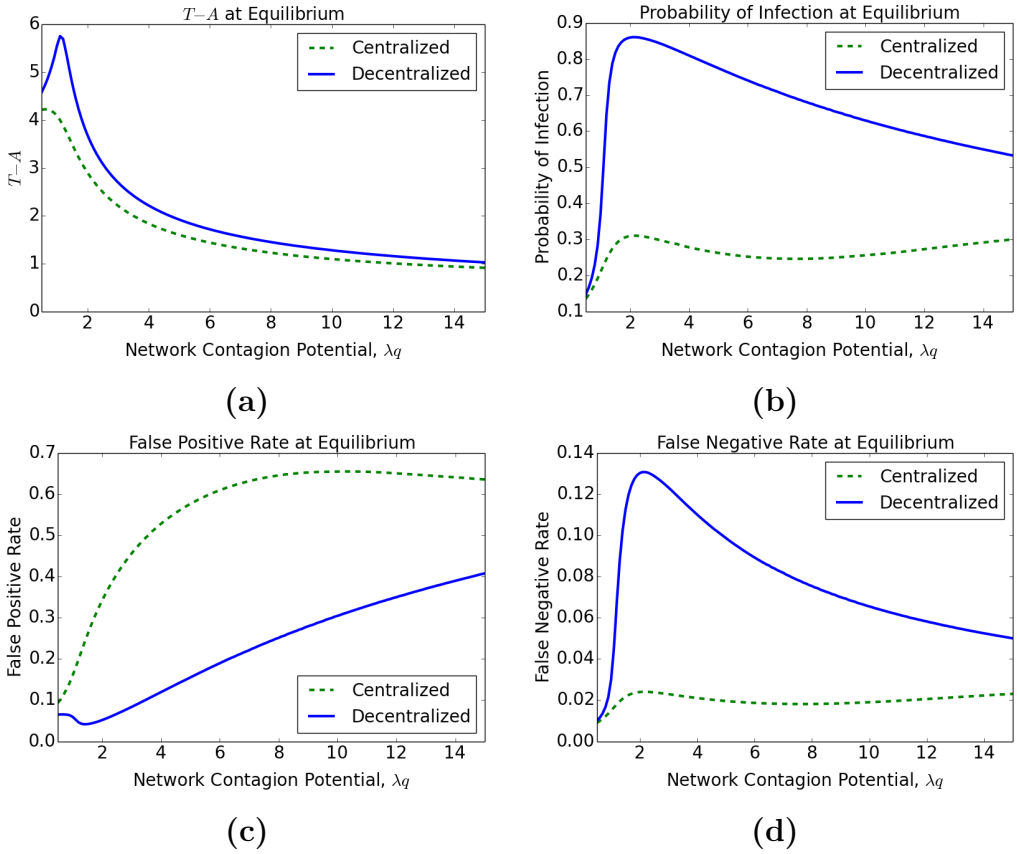

Figure 2.10: Strategy difference, probability of infection, and error probabilities at Nash equilibrium. Parameters are $r = 0.1, m = 5, k = 1$.

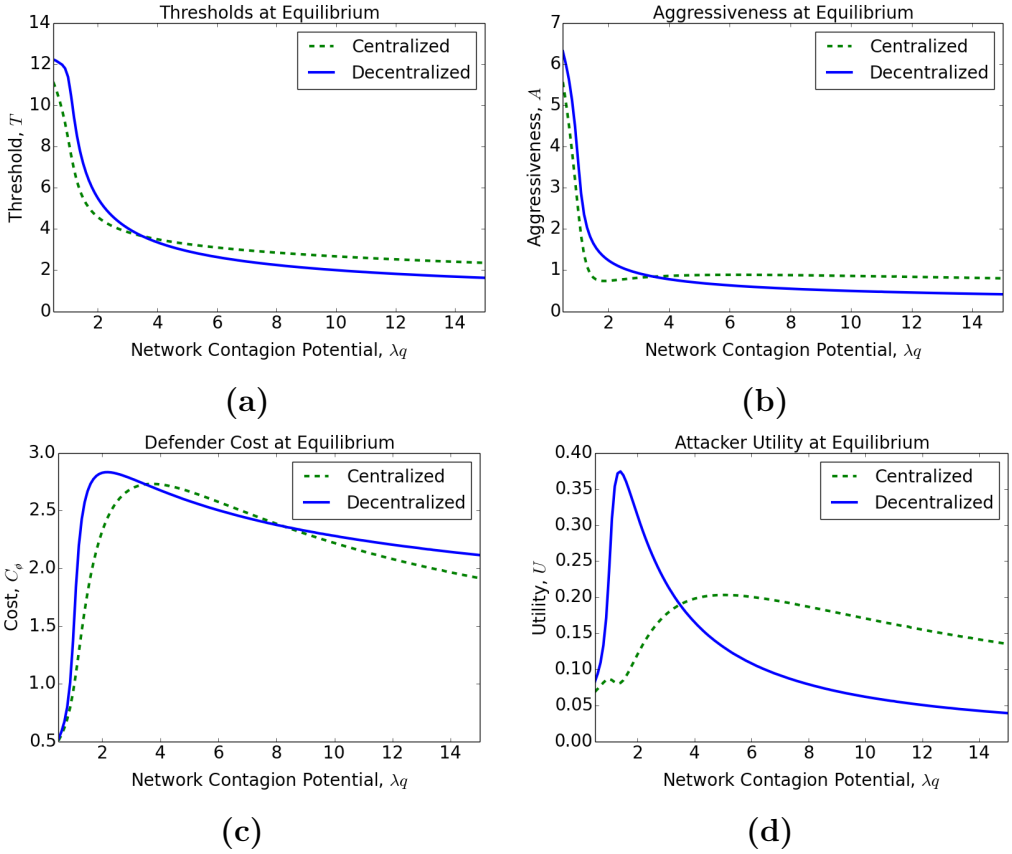

Figure 2.11: Strategies and payoffs at Nash equilibrium. Parameters are $r = 5, m = 5, k = 1$.

planner allowing higher infections rates than decentralized defenders, or decentralized defenders being able to outperform the central planner. Up to this point we have made simplifying assumptions in order to obtain a tractable model. In particular the assumption of defender homogeneity is restrictive, especially when dealing with social networks and internet topology. In the next chapter we extend the model to account for heterogeneous defenders.

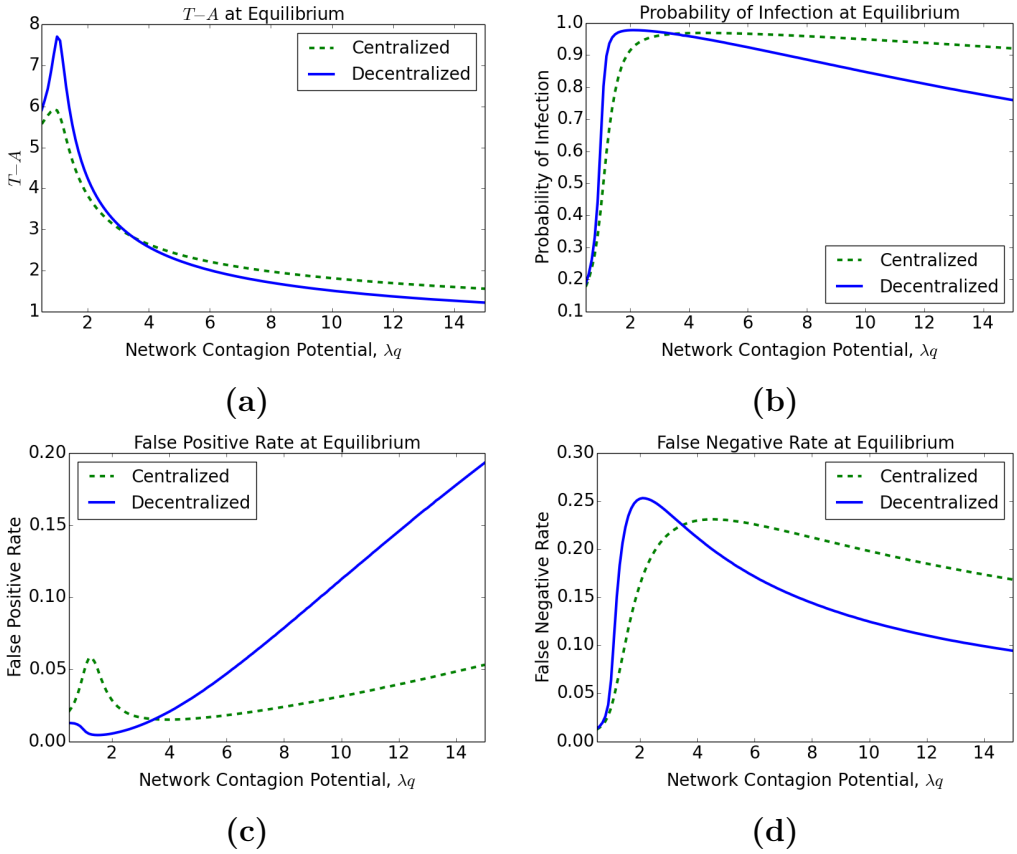

Figure 2.12: Strategy difference, probability of infection, and error probabilities at Nash equilibrium. Parameters are $r = 5, m = 5, k = 1$.

Chapter 3

A Heterogeneous Botnet Detection Game

3.1 Introduction

This chapter extends the results of the previous chapter by considering heterogeneous defenders. The specific heterogeneity is in the costs each defender attributes to false negatives and false positives when attempting to detect infections. We find that for any fixed bot master strategy there exists an equilibrium population response strategy in which no single defender has an incentive to deviate. However in many cases when a heterogeneous population responds with such an equilibrium strategy, the bot master will always have an incentive to deviate. In such cases there can be no pure Nash equilibria. However, if the bot master signals his strategy and the population responds optimally, then under some mild assumptions there will always exist an optimal signaling strategy for the bot master. As such we adopt a Stackelberg equilibrium concept for the heterogeneous botnet detection game with the bot master as leader. Sufficient conditions are

given to guarantee the existence of a Stackelberg equilibrium. We numerically compare these decentralized equilibrium strategies to a similar class of centrally planned strategies to explore network effects. The effects of agent heterogeneity on the equilibria are explored numerically as well.

3.2 The Botnet Detection Game

3.2.1 Agents

As in the homogeneous case we begin by considering a graph $G_n = \langle E, V \rangle$ with edge set E and vertex set V. We assume $|V| = n$ and with each vertex $i \in V$ we associate a computer m_i along with an independent defender d_i. With each edge $e_{ij} \in E$ we associate a network connection between computers m_i and m_j. Denote the malicious network attacker, or bot master, by b. The bot master does not exist in the graph G_n, but wishes to gain control of the m_i in order to utilize available computational resources. We assume the bot master is able to infiltrate the network by direct infection of some positive fraction of the m_i. Furthermore the infection can spread between the m_i via self-propagating malware. Once a machine m_i is infected (directly or indirectly) the bot master can use it for his own nefarious purposes. To model the epidemic process we will use a percolation model on the graph G_n. As such we assume the existence of an underlying probability space $(\Omega, \mathcal{F}, P)$ which all random variables in the model are defined on.

3.2.2 Strategic Variables

The bot master infects computers in order to illegitimately utilize available computational resources, for example CPU time, RAM, bandwidth or electrical

power. Specifically we wish to model how *aggressive* the bot master should be in utilizing these resources. We don't model a specific computational resource, but instead consider a general measurable resource $\mathcal{R}$ taking values on $\mathbb{R}^+$. The strategic variable for the bot master is $A \in \mathcal{A} \subseteq \mathbb{R}^+$, a measure of his aggressiveness in utilizing the resources of his botnet. We take A to be the directly observable amount of resource $\mathcal{R}$ the bot master uses.

The defenders are typical users of computers connected to the internet. Aware that there are potential security threats each defender must decide how vigilant he will be in detecting such threats. As in [37] we assume the defender is unable to reliably distinguish between security failures and natural variation in system performance. Thus defenders must consider the potential costs from both false positives and false negatives. We model the natural variation in system performance associated with m_i as a random variable S_i having support $\mathbb{R}^+$, cumulative distribution function $F_S(\cdot)$ and probability density function $f_S(\cdot)$. Let $X_i = 1$ if machine m_i is infected either directly by b or by contagion via another m_j ($e_{ij} \in E$) and $X_i = 0$ otherwise. Then the total observed variation in the performance of m_i is $Z_i = AX_i + S_i$.

Given the observation Z_i, the defender wishes to determine whether or not his computer is infected. Assuming the distributions of S_i and X_i are known this is a simple hypothesis testing problem with hypotheses $H_0 : X_i = 0$ and $H_1 : X_i = 1$. Higher observed values of Z_i indicate a higher likelihood of infection, thus we take the strategic variable of defender d_i to be a threshold T_i in a strategy space $\mathcal{T} \subseteq \mathbb{R}^+$. We interpret this as a measure of the defender's tolerance for software failures. If $Z_i \geq T_i$, the defender decides his computer is infected and takes appropriate measures to remediate the potential infection. If $Z_i < T_i$, then the defender takes no action. Define the detection indicator random variable

$D_i = \mathbb{1}_{\{Z_i \geq T_i\}}$. We can then define indicator random variables for the detection outcomes: $W^i_{FP} = (1 - X_i)D_i$ and $W^i_{FN} = X_i(1 - D_i)$ where W^i_{FP} indicates a false positive and W^i_{FN} a false negative for defender d_i.

3.2.3 Expected Cost/Utility and Best Responses

In the homogeneous game we assumed the cost of false positive c and false negative $v(A)$ were the same for all defenders in the network. In the heterogeneous game we drop this assumption. With each defender d_i we associate a cost of false positives, $c_i \geq 0$, and a cost of false negatives, $v_i \geq 0$. Cost c_i is the loss associated with a false positive. In this case the defender's computer is not infected, thus the loss incurred cannot depend on the strategy of the bot master. As such we assume c_i is constant. Cost v_i is the loss associated with false negatives. In this case the defender's computer is infected, and the loss incurred will include the resources stolen by the bot master. As such the loss will depend on the strategy of the bot master and we take $v_i(A)$ to be a non-decreasing function of A. Notice that neither c_i nor $v_i(A)$ depend on T_j for any $j \in V$. This is justified by the fact that a defender's tolerance for software failures will not alter the potential losses from an incorrect detection, it will only alter the probability of experiencing the loss.

In order to fully analyze the effects of defender heterogeneity in the game we make a further assumption regarding the function $v_i(A)$. In particular we assume that for each $i \in V$ we have $v_i(A) = \ell_i v(A)$ for some constant $\ell_i > 0$ and some common, non-decreasing function $v(A)$. As will become evident the cost ratio $\frac{c_i}{\ell_i}$ is critical in determining a defender's threshold strategy. In fact, if defenders d_i and d_j have identical cost ratios $\frac{c_i}{\ell_i} = \frac{c_j}{\ell_j}$, then their best response strategies will be the same. For this reason we treat the cost ratio $\frac{c_i}{\ell_i}$ as the type of the defender.

We assume all defenders are equally likely to occupy any place in the network, and there exists a distribution from which agent types are drawn i.i.d. Let θ_i be the random variable which returns the type of defender d_i, i.e. $\theta_i \equiv \frac{c_i}{\ell_i}$. We assume $\theta_i \in \Theta \subseteq \mathbb{R}^+$ for some locally compact set Θ. Denote the cumulative distribution function of θ_i by F_θ and its probability density function by f_θ.

It is the objective of each defender to minimize the expected costs associated with false positives and false negatives. Recall that W^i_{FP} indicates a false positive and W^i_{FN} a false negative for defender d_i. Thus the realized cost for defender d_i is

$$C_i = c_i W^i_{FP} + \ell_i v(A) W^i_{FN}.$$

Note that we have dropped the cost of raising and alarm k in this model that appears in the homogeneous game. This is done for simplicity. We could easily include it in this model if needed. The results will not change significantly.

A *population strategy* is any function $T : \Theta \to \mathcal{T}$ which maps a defender type to an individual threshold strategy. Note that if a population strategy T is measurable (using standard σ-algebras generated by all Lebesgue measurable sets) then given $\theta \sim F_\theta$ the function $T(\theta)$ is a random variable describing the distribution of thresholds across the network. We assume that the only global information available to each agent in the game is statistical in nature. For example, each agent has knowledge of the distributions F_S and F_θ, but not the realized values of S_i and θ_i. Thus the probabilities that defender d_i assigns to the events W^i_{FP} and W^i_{FN} will depend on his own strategy T_i, the bot master strategy A and the distribution of defender strategies $T(\theta)$. The best response correspondence σ_i for

defender d_i is then

$$\sigma_i(A,T) = \arg\min_{T_i} \{c_i E[W^i_{FP}] + \ell_i v(A) E[W^i_{FN}]\}.$$

The bot master maximizes his expected utility by maximizing the cumulative computational resources stolen from the compromised network. His utility will then depend on the fraction of machines that are infected, say ζ, as well as the degree to which he utilizes the bots, which we measure by his aggressiveness, A. In the limit of a large population the expected proportion of infected machines is equal to the probability of a false negative of a defender chosen uniformly at random from the network. Letting $\varnothing$ be chosen uniformly at random from V we have $\zeta = E[W^{\varnothing}_{FN}]$. The bot master's expected utility is then $U(A,T) = g(A)\zeta = g(A)E[W^{\varnothing}_{FN}]$ and the set of bet responses, $\sigma_b(T)$, for the bot master is given by

$$\sigma_b(T) = \arg\max_{A} \{g(A)E[W^{\varnothing}_{FN}]\}.$$

3.3 Epidemic Process and Detection Model

In order to finish the construction of our game between the network of defending agents $\{d_i\}_{i\in V}$ and the bot master b, we need to define an epidemic process on G_n initiated by the bot master b. To model the epidemic process in the network we use a percolation model on the graph G_n, that is each $e_{ij} \in E$ admits contagion independently with probability $q \in (0,1]$. In order to initialize the epidemic we assume the bot master attempts to directly infect each m_i independently with probability of success $p \in (0,1]$. Following [48] and the previous chapter we can

define the following random variables for $i, j = 1, 2, 3, ..., n$:

$$\begin{aligned}
\chi_i &\sim \text{Bern}(p), \\
B_{ij} &\sim \text{Bern}(q), \\
X_i &= \mathbb{1}_{\{m_i \text{ is infected}\}}, \\
D_i &= \mathbb{1}_{\{d_i \text{ decides } m_i \text{ is infected}\}}.
\end{aligned}$$

The random variables χ_i indicate direct infection of defender d_i from the bot master b. The random variables B_{ij} indicate sufficient contact to admit contagion between d_i and d_j. We assume the χ_i and B_{ij} are independent of all random variables in the model and $B_{ij} = B_{ji}$ for all $i, j \in V$. As in the previous chapter we assume that if $D_i = 1$ then defender d_i takes immediate action and prevents the spread of any potential infection to his neighbors. The fundamental recursion which defines the epidemic/detection process on the graph G_n is then

$$X_i = 1 - (1 - \chi_i) \prod_{k \sim i} (1 - B_{ki}(1 - D_k)X_k), \tag{3.1}$$

where $k \sim i$ indicates $e_{ki} \in V$.

For large n this model on a general graph G_n is not amenable to analysis in the context of our game. In order to derive a tractable model we follow [48] and consider the asymptotic properties of *locally tree like* graphs. Specifically we restrict our attention to the limit of rooted Erdős-Rényi random graphs $G(n, \lambda/n)$ as $n \to \infty$. By *rooted* we mean that for each n the graph $G(n, \lambda/n)$ has a designated root vertex $v_\varnothing^{(n)}$ chosen uniformly at random from the set of n vertices. In the sense of *local weak convergence* [3] the limiting object of a sequence of rooted Erdős-Rényi random graphs is a rooted Galton-Watson Poisson Branching process, $T_\infty(\lambda)$.

The advantage of working with $T_\infty(\lambda)$ is that we can analytically derive the expected cost function of the *root defender* $d_\varnothing$ associated with the distinguished root vertex of $T_\infty(\lambda)$. Because the root is chosen uniformly at random we can treat $d_\varnothing$ as a randomly selected defender then perform a mean field analysis of the game. As such we restrict our analysis to $T_\infty(\lambda)$ and refer the readers to [3, 47, 48] for details on how to extend the subsequent results to $\lim_{n\to\infty} G(n, \lambda/n)$. Convergence results for the homogeneous game can be found in appendix A.8.

When working with $T_\infty(\lambda)$ we take advantage of the tree structure and define an alternative sequence of random variables, rather than the X_i above. If d_j is a direct descendant of d_i in $T_\infty(\lambda)$ we write $j \to i$. Let $\mathcal{O}_i = \{k \in T_\infty(\lambda) | k \to i\}$. For each $i \in T_\infty(\lambda)$ define $\tilde{X}_i$ and $\tilde{D}_i$ as follows:

$$\tilde{X}_i = \begin{cases} 1 & \text{if } m_i \text{ is infected by } b \\ & \text{or by some } m_j \text{ s.t. } j \in \mathcal{O}_i, \\ 0 & \text{otherwise,} \end{cases}$$

$$\tilde{D}_i = \mathbb{1}_{\{S_i + \tilde{X}_i A \geq T_i\}}.$$

The equations for the $\tilde{X}_i$ can then be expressed as

$$\tilde{X}_i = 1 - (1 - \chi_i) \prod_{k \to i} (1 - B_{ki}(1 - \tilde{D}_k)\tilde{X}_k). \tag{3.2}$$

We also introduce the alternate detection outcome indicator random variables $\tilde{W}^i_{\text{FP}} = (1 - \tilde{X}_i)\tilde{D}_i$ and $\tilde{W}^i_{\text{FN}} = \tilde{X}_i(1 - \tilde{D}_i)$.

The introduction of the processes $\tilde{X}_i$ and $\tilde{D}_i$ are done to take advantage of the structure of $T_\infty(\lambda)$. Note that $\tilde{X}_i$ and $\tilde{X}_j$ are independent random variables when d_i and d_j are the same distance away from the root node. Furthermore because $\tilde{X}_i$

and $\tilde{D}_i$ depend only on the children of d_i, the process $\{\tilde{X}_i\}_{i\in T_\infty(\lambda)}$ is a Recursive Tree Process [3]. As we will show this allows the distributions of the alternate processes $\{\tilde{X}_i\}_{i\in T_\infty(\lambda)}$, $\{\tilde{D}_i\}_{i\in T_\infty(\lambda)}$, $\{\tilde{W}^i_{\text{FP}}\}_{i\in T_\infty(\lambda)}$, and $\{\tilde{W}^i_{\text{FN}}\}_{i\in T_\infty(\lambda)}$ to be solved explicitly. Moreover, on $T_\infty(\lambda)$ the distributions of the random variables associated with a root defender and defined by (3.2) will coincide with the distributions of the random variables associated with a root defender and defined by (3.1), i.e. $(\tilde{X}_\varnothing, \tilde{D}_\varnothing, \tilde{W}^\varnothing_{\text{FP}}, \tilde{W}^\varnothing_{\text{FN}}) \stackrel{d}{=} (X_\varnothing, D_\varnothing, W^\varnothing_{\text{FP}}, W^\varnothing_{\text{FN}})$.

Because each root defender is chosen uniformly at random, and all agent's prior knowledge of the network is the same, we assume all other agent's would act similarly had they been chosen as the root defender. Specifically we assume each defender does not know his position in the graph nor who his neighbors are, but he does know the statistical properties of the network and the population. As such we expect there to be an *invariant process* [2] which solves (3.2). The fundamental Recursive Distributional Equation [2] which defines the invariant process on $T_\infty(\lambda)$ is

$$\tilde{X} \stackrel{d}{=} 1 - (1-\chi)\prod_{k=1}^{N}(1 - B_k(1-\tilde{D}_k)\tilde{X}_k), \tag{3.3}$$

$$\tilde{D}_k = \mathbb{1}_{\{T \le S + \tilde{X}_k A\}}. \tag{3.4}$$

$\tilde{X}$ and $\tilde{X}_k$ are i.i.d. random variables satisfying (3.3) while the random variables $\chi \sim \text{Bernoulli}(p)$, $S \sim \text{Gamma}(\alpha, \beta)$, $B_k \stackrel{iid}{\sim} \text{Bernoulli}(q)$ and $N \sim \text{Poisson}(\lambda)$ are independent of everything in the model. The random variable T is $\mathbb{R}^+$-valued with distribution function F_T, which is the distribution of threshold strategies across the network. The exact nature of this distribution will be addressed in the following section. Provided solutions exist for equations (3.3)-(3.4), the detection

outcome indicator random variables for the invariant process are

$$\tilde{W}_{\text{FP}} = (1 - \tilde{X})\tilde{D}, \tag{3.5}$$

$$\tilde{W}_{\text{FN}} = \tilde{X}(1 - \tilde{D}), \tag{3.6}$$

where $\tilde{D} = \mathbb{1}_{\{T \leq S + \tilde{X}A\}}$ is analogous to (3.4). We now find the distribution for the detection random variables. The proof follows from a slight generalization of Proposition 2 in Chapter 2.

Proposition 8. *Let $S \sim F_S(\cdot)$, $T \sim F_T(\cdot)$, $A \in \mathcal{A}$, $p \in (0, 1]$ and $q \in (0, 1]$. Then the distributions for $\tilde{W}_{FP}$ and $\tilde{W}_{FN}$ are given as follows:*

$$E[\tilde{W}_{FP}] = E[1 - F_S(T)](1 - p)e^{-\lambda q h},$$
$$E[\tilde{W}_{FN}] = h,$$

where $h = h(A, p, q, \lambda, F_S(\cdot), F_T(\cdot))$ is the unique solution in $[0, 1]$ of the fixed point equation

$$h = E[F_S(T - A)][1 - (1 - p)e^{-\lambda q h}].$$

From equations (3.3)-(3.6) it is clear that if the root defender $d_\varnothing$ alters his threshold $T_\varnothing$, this will change his decision rule, but it will not change what he observes. In other words the probabilities of a false positive and false negative will change, but the probability of the infection reaching him will not. So if the root defender unilaterally deviates from the population strategy T, then equations (3.3)-(3.6) are still valid for all defenders in the tree except the root. For the root

we need to introduce new equations:

$$
\begin{aligned}
D_{\varnothing} &= \mathbb{1}_{\{T_{\varnothing} < S + \tilde{X} A\}}, \\
W_{\mathrm{FP}}^{\varnothing} &= (1 - \tilde{X}) D_{\varnothing}, \\
W_{\mathrm{FN}}^{\varnothing} &= \tilde{X}(1 - D_{\varnothing}).
\end{aligned}
$$

The corresponding distributions are computed analogously as in Proposition 8.

Proposition 9. *Let $S \sim F_S(\cdot)$, $T \sim F_T(\cdot)$, $A \in \mathcal{A}$, $T_{\varnothing} \in \mathcal{T}$ with $T_{\varnothing} \geq A$, $p \in (0, 1]$ and $q \in (0, 1]$. Then the distributions for $W_{FN}^{\varnothing}$ and $W_{FP}^{\varnothing}$ have unique solutions which depend on the distribution of $\tilde{W}_{FN}$. If $E[\tilde{W}_{FN}] = h$ then the distributions are given by*

$$
\begin{aligned}
E[W_{FP}^{\varnothing}] &= [1 - F_S(T_{\varnothing})](1 - p) e^{-\lambda q h}, \\
E[W_{FN}^{\varnothing}] &= F_S(T_{\varnothing} - A)[1 - (1 - p) e^{-\lambda q h}].
\end{aligned}
$$

Summary of Notation:

- $A \in \mathcal{A}$: Strategy of bot master b.
- $T_i \in \mathcal{T}$: Strategy of defender d_i.
- $S_i \sim F_S$: Software reliability failure.
- $D_i = \mathbb{1}_{\{d_i \text{ decides } m_i \text{ is infected}\}}$.
- $p = P(\chi_i = 1)$: Probability b infects m_i.
- $q = P(B_{ij} = 1)$: Probability of contagion.
- $X_i = \mathbb{1}_{\{m_i \text{ is infected}\}}$.
- $X_i = 1 - (1 - \chi_i) \prod_{k \sim i} (1 - B_{ki}(1 - D_k)X_k)$.
- $Z_i = AX_i + S_i$: Observation made by d_i.
- $W_{FP}^i = \mathbb{1}_{\{d_i \text{ False Positive}\}}$.
- $W_{FN}^i = \mathbb{1}_{\{d_i \text{ False Negative}\}}$.
- c_i and v_i: Costs to d_i from FP and FN, respectively.
- θ_i: Type of defender d_i.
- g: Utility gained by b from FN.

3.4 Equilibrium Analysis

To analyze the above game we first look for population strategies which are mutual best responses for all defenders in the network. We call a mapping $\sigma_p : \mathcal{A} \times \Theta \to \mathcal{T}$ a *population best response correspondence* if for fixed $A \in \mathcal{A}$ the following relation holds:

$$\sigma_p(A, \theta_\varnothing) \in \sigma_\varnothing(A, \sigma_p(A, \cdot)). \tag{3.7}$$

That is, if the population strategy is $\sigma_p(A, \cdot)$, then a deviant root defender can do no better than to follow the population strategy and play $\sigma_p(A, \theta_\varnothing)$. Note that this defines a type of mean field Nash equilibrium among the defenders, i.e. a root defender has no incentive to deviate from the population strategy. If a population best response correspondence σ_p is found for each $A \in \mathcal{A}$, we can then look for equilibrium strategies between the defender population and the bot master.

3.4.1 Defender Equilibria

In this section we investigate the selfish behavior of heterogeneous defenders in response to a fixed bot master strategy $A \geq 0$. We show the existence of a population best response correspondence for each $A \in \mathcal{A}$. We begin by using the results of Sec. 3.3 to write down explicitly the expected cost and best response of a root defender. Let $T = T(A, \theta)$ denote a population response to the bot master strategy A. The realized cost experienced by root defender $d_\varnothing$ is

$$C_\varnothing = c_\varnothing W_{\text{FP}}^\varnothing + \ell_\varnothing v(A) W_{\text{FN}}^\varnothing.$$

The expected cost $\bar{C}_\varnothing$ is then

$$\bar{C}_\varnothing(T_\varnothing, T, A) = c_\varnothing[1 - F_S(T_\varnothing)](1-p)e^{-\lambda qh} + \ell_\varnothing v(A) F_S(T_\varnothing - A))[1 - (1-p)e^{-\lambda qh}],$$

where h is defined by the fixed point equation

$$h = \int_\Theta [F_S(T(A,\theta) - A)] dF_\theta [1 - (1-p)e^{-\lambda qh}]. \tag{3.8}$$

It is important to notice that h does not depend on $T_\varnothing$, but only T. In particular we have $\frac{\partial h}{\partial T_\varnothing} \equiv 0$. The best response for a deviant root defender is then

$$\sigma_\varnothing(A,T) = \arg\min_{T_\varnothing} \{c_\varnothing[1-F_S(T_\varnothing)](1-p)e^{-\lambda qh} + \ell_\varnothing v(A) F_S(T_\varnothing - A)[1-(1-p)e^{-\lambda qh}]\}.$$

Arguments similar to those in the homogeneous game guarantee the strict quasi-convexity of $\bar{C}_\varnothing$ in $T_\varnothing$. For ease of exposition we define the function

$$L(T_\varnothing, T, A) = v(A) \frac{f_S(T_\varnothing - A)}{f_S(T_\varnothing)} \frac{1 - (1-p)e^{-\lambda qh}}{(1-p)e^{-\lambda qh}}.$$

First order optimality implies the optimal response $T^*_\varnothing = \sigma_\varnothing(A,T)$ is determined by solutions to the equation $\frac{c_\varnothing}{\ell_\varnothing} = L(T_\varnothing, T, A)$. By the continuity and monotonicity of $\frac{f_S(T_\varnothing - A)}{f_S(T_\varnothing)}$ (Assumption 1), if a solution exists then it is unique, while if no solution exists, then for all $T_\varnothing$ we have one of the following: $\frac{c_\varnothing}{\ell_\varnothing} \gtreqless L(T_\varnothing, T, A)$. In the $<$ case $T^*_\varnothing = +\infty$, while in the $>$ case $T^*_\varnothing = A$. Thus for any $A \in \mathcal{A}$ and any $T \sim F_T$

we can define a best response function for a deviant defender:

$$\sigma_{\varnothing}(A,T) = \begin{cases} A & \text{if } \forall\ T_{\varnothing}, \frac{c_{\varnothing}}{\ell_{\varnothing}} < L(T_{\varnothing}, T, A), \\ T^*_{\varnothing} & \text{if } \exists\ T^*_{\varnothing} \text{ s.t. } \frac{c_{\varnothing}}{\ell_{\varnothing}} = L(T^*_{\varnothing}, T, A), \\ +\infty & \text{if } \forall\ T_{\varnothing}, \frac{c_{\varnothing}}{\ell_{\varnothing}} > L(T_{\varnothing}, T, A). \end{cases}$$

If defenders act selfishly, they will all follow a similar strategy. Thus we should look for population strategies of the above form that satisfy (3.7). The next proposition establishes, for a fixed A, the existence of such a population strategy.

Proposition 10. *Suppose the S_i have c.d.f. $F_S(\cdot) \in C^1$ and p.d.f. $f_S(\cdot)$ while $\theta_i \sim F_\theta(\cdot) \in C^1$ with support $\Theta \subseteq \mathbb{R}^+$. If for any $A \in \mathcal{A}$ the ratio $\frac{f_S(z-A)}{f_S(z)}$ is continuous and non-decreasing in $z \geq A$, then there exists a mapping $T^* : \mathcal{A} \times \Theta \to \mathcal{T}$ which for each $A \in \mathcal{A}$ is a population best response correspondence. Moreover any such mapping T^* is of the following form:*

$$T^*(A, \theta_i) = \begin{cases} A & \text{if } \theta_i < \kappa, \\ T_i & \text{if } \theta_i \in [\kappa, \omega], \\ +\infty & \text{if } \theta_i > \omega. \end{cases}$$

Here $\kappa = \lim_{t \downarrow A} L(t, T^*, A)$ and $\omega = \lim_{t \to +\infty} L(t, T^*, A)$ while T_i is a solution to $\theta_i = L(T_i, T^*, A)$.

Proof. Fix $A \in \mathcal{A}$. We suppress the dependence on A for notational clarity. The preceding discussion which lead to a deviant defender's best response function $\sigma_{\varnothing}$ makes clear the form of the desired equilibrium population strategy T^*. Let $x \in [0,1], t \in \bar{\mathbb{R}}^+, \theta \in \mathbb{R}^+$ and define the maps $\tilde{L} : \bar{\mathbb{R}}^+ \times [0,1] \to \mathbb{R}^+$ and $\sigma :$

$[0,1] \times \mathbb{R}^+ \to \bar{\mathbb{R}}^+$ as

$$\tilde{L}(z,x) = v(A)\frac{f_S(z-A)}{f_S(z)}\frac{1-(1-p)e^{-\lambda qx}}{(1-p)e^{-\lambda qx}},$$

and

$$\sigma(x,\theta) = \begin{cases} A & \text{if } \theta < \kappa(x), \\ t(x,\theta) & \text{if } \theta \in [\kappa(x),\omega(x)], \\ +\infty & \text{if } \theta > \omega(x). \end{cases}$$

where $t(\theta,x)$ is a solution to $\theta = \tilde{L}(t,x)$ while $\kappa(x) = \lim_{z\to\infty}\tilde{L}(z,x)$ and $\omega(x) = \lim_{z\downarrow A}\tilde{L}(z,x)$.

We prove the case when $\frac{f_S(z-A)}{f_S(z)}$ is strictly increasing in z. The proof can then be adapted to the more general non-decreasing case. The continuity and monotonicity of $\frac{f_S(z-A)}{f_S(z)}$ in z guarantees both the existence of $\kappa(x)$ and $\omega(x)$ and the existence and uniqueness of $t(x,\theta) \in \bar{\mathbb{R}}^+$ whenever $\theta \in [\kappa(x),\omega(x)]$. Define the map $G(x,t,\theta) = \tilde{L}(t,x) - \theta$. Then the implicit function theorem implies $t(x,\theta)$ is continuous and differentiable in both x and θ in open regions where $G(x,t,\theta) = 0$ has solutions, i.e. $\theta \in (\kappa(x),\omega(x))$. Since $F_S(\cdot)$ is differentiable it follows that $F_S(t(x,\theta)-A)$ is continuously differentiable in x and θ. Thus by Leibniz's rule

$$\begin{aligned} E[F_S(\sigma(x,\theta)-A)] &= \int_0^\infty F_S(\sigma(x,\theta)-A)dF_\theta \\ &= \int_{\kappa(x)}^{\omega(x)} F_S(t(x,\theta)-A)dF_\theta + \int_{\omega(x)}^\infty dF_\theta \\ &= \int_{\kappa(x)}^{\omega(x)} F_S(t(x,\theta)-A)dF_\theta + 1 - \int_0^{\omega(x)} dF_\theta \\ &= \int_{\kappa(x)}^{\omega(x)} F_S(t(x,\theta)-A)dF_\theta + 1 - F_\theta(\omega(x)) \end{aligned}$$

is differentiable in x. Now define $M(x) = E[F_S(t(x,\theta)-A)][1-(1-p)e^{-\lambda qx}]$. Since $M(x)$ is continuous and maps the unit interval into itself, Brouwer's fixed

point theorem guarantees the existence of a fixed point h satisfying $h = M(h)$.

Define $T^*(\theta) = \sigma(h, \theta)$. Then given a type distribution F_θ we can define the random variable $T^*(\theta)$. With Propositions 8 and 9 we can construct a botnet detection game with heterogeneous defenders as above and obtain $P(W_{\text{FN}}^{\varnothing} = 1) = h$. The analysis above regarding a deviant defender's optimal response to arbitrary population strategies guarantees that the threshold $T^*(\theta_\varnothing)$ will be optimal for the root defender. In other words $T^*(\theta_\varnothing) = \sigma_\varnothing(A, T^*)$. Since $A \in \mathcal{A}$ was arbitrary the mapping $T^*(A, \theta)$ is a population best response for each $A \in \mathcal{A}$. □

3.4.2 Example Population Best Response Function

We now look at a specific population best response function $T^*(A, \theta)$ and discuss how one can compute it. The first difficulty in determining the function $T^*(A, \theta)$ is that it is defined in reference to itself. This is a result of the self fulfilling nature of the strategy. If all agents assume the other agents are playing T^*, then they can do no better than to join them and follow the strategy T^*. This self referencing definition can be seen in the statement of Proposition 10 where the function $T^*(A, \theta)$ appears on the left hand side *and* the right side. Notice that the values κ and ω are both functionals of $T^*(A, \cdot)$.

This dependence comes in the form of the functional $h(A, T^*(A, \cdot))$, the false negative rate. Recall h is defined by the fixed point equation

$$h = \int_\Theta [F_S(T^*(A, \theta) - A)] dF_\theta [1 - (1 - p)e^{-\lambda q h}]. \tag{3.9}$$

The proof of Proposition 10 guarantees this equation has a solution. Furthermore if the likelihood ratio $\frac{f_S(z-A)}{F_S(z)}$ is monotonically increasing in z then we can solve for $T^*(A, \theta)$ more explicitly in terms of h. Plugging this back into the fixed point equation gives us an equation that can be solved numerically.

Let us look at a concrete example to understand this procedure. Suppose $S \sim \text{gamma}(\alpha, \beta)$ with $\alpha > 1$. Assume that the value h which solves (3.9) is known. This implies the values κ and ω are known too. Then using the definition of T^* from Proposition 10, for any $\theta \in [\kappa, \omega]$ we can write

$$\theta = v(A)\left(1 - \frac{A}{T^*(A,\theta)}\right)^{\alpha-1} e^{\beta A}\frac{1-(1-p)e^{-\lambda q h}}{(1-p)e^{-\lambda q h}}.$$

Solving for $T^*(A,\theta)$ we arrive at

$$T^*(A,\theta) = \frac{A}{1 - \left(\theta\frac{e^{-\beta A}}{v(A)}\frac{(1-p)e^{-\lambda q h}}{1-(1-p)e^{-\lambda q h}}\right)^{\frac{1}{\alpha-1}}}.$$

We then have an expression for $T^*(A,\theta)$ explicitly in terms of h. We can then plug this expression into (3.9) and we arrive at

$$h = \int_\Theta \left[F_S\left(\frac{A}{1 - \left(\theta\frac{e^{-\beta A}}{v(A)}\frac{(1-p)e^{-\lambda q h}}{1-(1-p)e^{-\lambda q h}}\right)^{\frac{1}{\alpha-1}}} - A\right)\right] dF_\theta[1-(1-p)e^{-\lambda q h}]. \quad (3.10)$$

While there is no hope of solving this equation for h, there is hope of solving it numerically. In this case care must be taken as the proof in Proposition 10 does not guarantee uniqueness. In certain cases uniqueness can be established depending on the distributions of S and θ.

Once a value for h is found that solves (3.10) we can explicitly write down the function $T^*(A,\theta)$. By solving the fixed point equation (3.10) we have "closed the loop" and can then numerically find the function $T^*(A,\theta)$. Figure 3.2 shows the best response function $T^*(A,\theta)$ for various values of A when $\theta \sim \text{uniform}(0.1, 10)$ and $S \sim \text{gamma}(2,2)$. Figure 3.1 shows corresponding histograms estimating the distribution of thresholds across the network. Notice that while the attack intensity A decreases more agents become indifferent and choose higher thresholds.

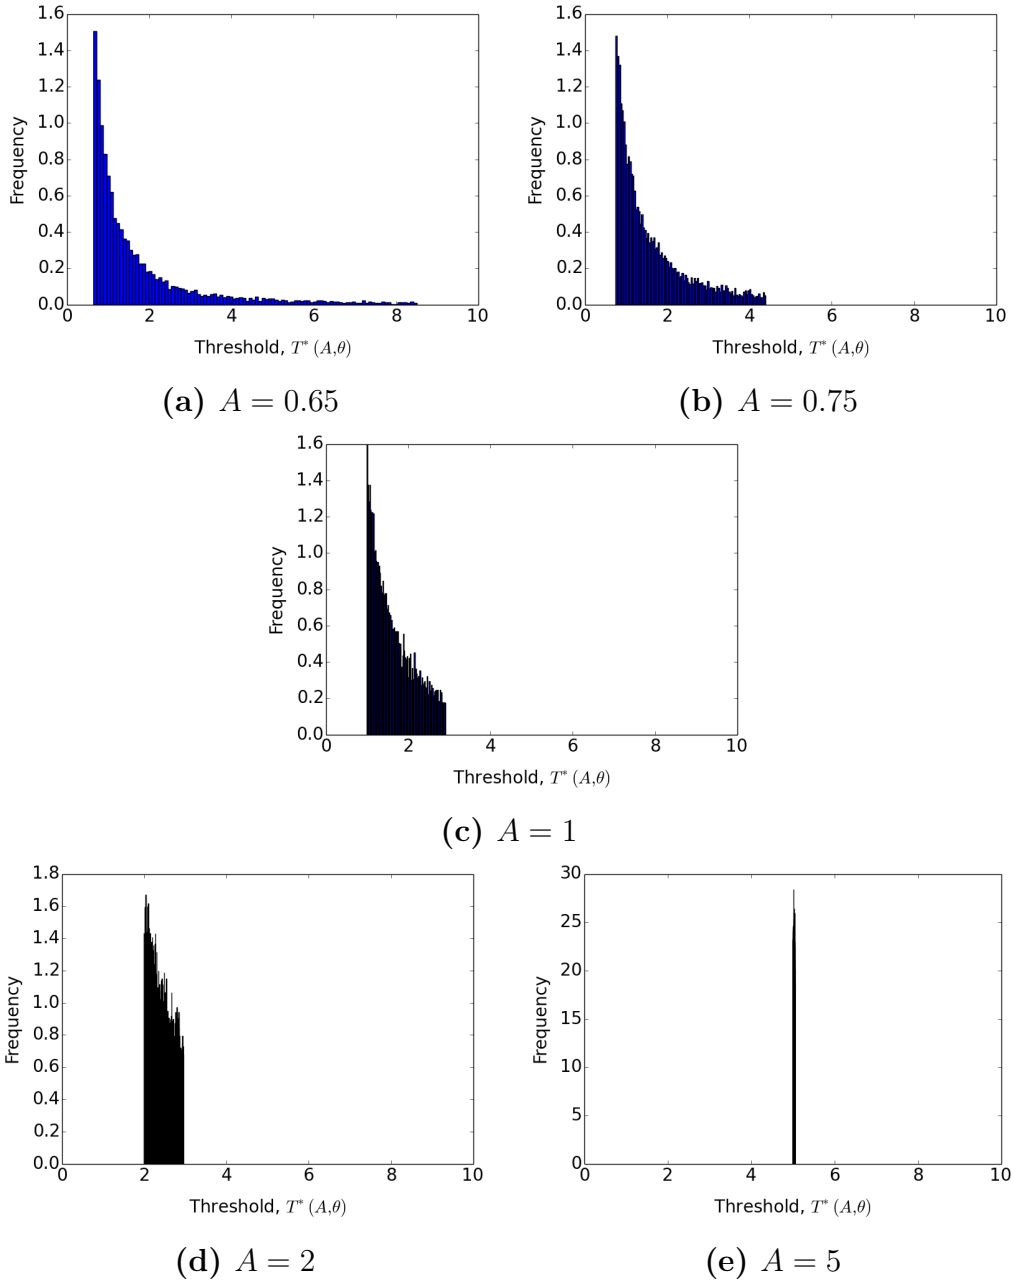

(a) $A = 0.65$ **(b)** $A = 0.75$

(c) $A = 1$

(d) $A = 2$ **(e)** $A = 5$

Figure 3.1: Histograms of $T^*(A, \theta)$ for 10,000 draws of $\theta \sim$ uniform$(0.1, 10)$, $S \sim$ gamma$(2, 2)$ and various values of A.

On the other hand when the attack intensity A increase more and more agents become vigilant and choose thresholds close to the attack value A.

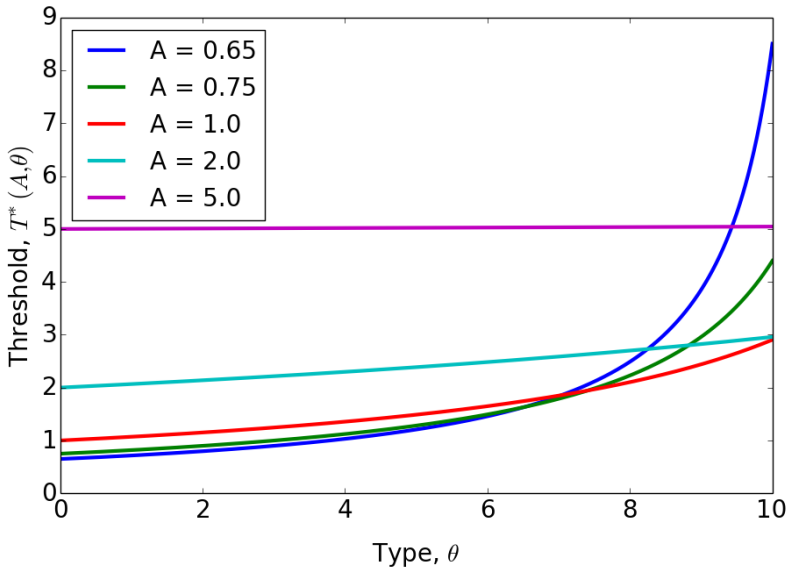

Figure 3.2: The best response function $T^*(A, \theta)$ for $\theta \sim \text{uniform}(0.1, 10)$, $S \sim \text{gamma}(2, 2)$ and various values of A.

3.4.3 Price of Anarchy

Just as we did for the homogeneous game, we can compute the Price of Anarchy for each attack strategy A. Instead of examining the relationship between PoA and the network contagion parameter λq we will examine varying the variance of the population. In this way we can examine the effect that population heterogeneity has on expected social cost. For simplicity we consider the case $S \sim \exp(\beta)$.

To complete such an analysis we must introduce a central planner as in the homogeneous case. We assume the central planner wishes to minimize the expected cost over the entire network of defenders. For each fixed strategy $A \in \mathcal{A}$ the central planner seeks a measurable function $T_c : \Theta \to \bar{\mathbb{R}}^+$ which minimizes the

functional

$$I(A,T_c) = \int_\theta c(\theta)[1-F_S(T_c(\theta))](1-p)e^{-\lambda q h(A,T_c)} + \ell(\theta)v(A)F_S(T_c(\theta)-A)[1-(1-p)e^{-\lambda q h(A,T_c)}]dF_\theta.$$

Because $h(A,T_c)$ is a functional of T_c itself, standard variational techniques are not available and a general optimal solution is difficult to find. Instead, we restrict the central planner to choose a function which is of the same functional form as that of the decentralized population best response T^*. In the case $S \sim \exp(\beta)$, for a fixed $A \in \mathcal{A}$ the function T^* is an unbounded step function which is uniquely determined by the location of the step. If we restrict the central planner to play functions of this form, we have a single parameter to optimize over. Specifically

$$T_c(\theta,x) = \begin{cases} A & \text{if } \theta < x, \\ +\infty & \text{if } \theta \geq x. \end{cases}$$

We can then define the *expected social cost* as

$$I(A,x) = \int_0^x c(\theta)[1-F_S(A)](1-p)e^{-\lambda q h(A,x)}dF_\theta + \int_x^\infty \ell(\theta)v(A)[1-(1-p)e^{-\lambda q h(A,x)}]dF_\theta.$$

Since $0 \leq I(A,x)$, there exists a value $\theta^*(A) \in \arg\min_{x\geq 0} I(A,x)$. We can then define the centralized expected social cost $I(A,\theta_c^*)$ and the decentralized expected

social cost $I_c(A, \theta_d^*)$ where

$$\theta_c^* \in \arg\min_{x \geq 0} I(A, x)$$
$$\theta_d^* = v(A) \exp(\beta A) \frac{1 - (1-p)e^{-\lambda q h(A, \theta_d^*)}}{(1-p)e^{-\lambda q h(A, \theta_d^*)}}.$$

Note that we have written $h(A, \theta_d^*)$ to stress the fact that the value h in this case depends on both A and θ_d^*. As stated before the proof of Proposition 10 guarantees us that such an h exists and it is not difficult to prove it is unique. The same methods as were used in the homogeneous case apply, with more algebra due to the complexity of the terms involved. As we saw in the previous section we have a means of computing this h and can thus compute θ_d^*.

Computing θ_c^* is less straight forward since we are not guaranteed that the function $I(A, x)$ is convex in x. However if we use a minimization algorithm which starts at the point θ_d^*, then even if we do not find a global minimum we are at least guaranteed to find a solution that is better than or equal to the decentralized solution. This will then provide us with a measure of efficiency of the decentralized best response strategy $T^*(A, \theta)$. We thus define the Price of Anarchy as

$$PoA \triangleq \frac{I(A, \theta_d^*)}{I(A, \theta_c^*)}.$$

We will assume $\theta \sim \text{gamma}(\phi, \gamma)$. We then fix $\phi = m\gamma$ for some $m > 0$ so that $E[\theta] = \frac{\phi}{\gamma} = \frac{m\gamma}{\gamma} = m$ and $\text{var}[\theta] = \frac{\phi}{\gamma^2} = \frac{m\gamma}{\gamma^2} = \frac{m}{\gamma}$. In this way we can fix the mean of the population's distribution and vary the variance. In particular we see that $\lim_{\gamma \to \infty} \text{var}[\theta] = 0$ and $\lim_{\gamma \downarrow 0} \text{var}[\theta] = \infty$. Letting $D = \frac{\text{var}[\theta]}{E[\theta]}$ denote the index of

dispersion we have the following interpretation of the parameter γ in this context:

$$\frac{1}{D} = \frac{E[\theta]}{\text{var}[\theta]} = \frac{m}{\frac{m}{\gamma}} = \gamma.$$

That is, γ is the inverse of the index of dispersion. Thus large values of γ correspond to low dispersion, i.e. the population is more homogeneous. Small values of γ correspond to high dispersion, i.e. the population is more heterogeneous.

Alternatively letting $SNR = \frac{E[\theta]}{\sqrt{\text{var}[\theta]}}$ be the signal-to-noise ratio of θ we have

$$SNR = \frac{E[\theta]}{\sqrt{\text{var}[\theta]}} = \frac{m}{\sqrt{\frac{m}{\gamma}}} = \sqrt{m\gamma}.$$

This interpretation will be helpful in characterizing the qualitative features of the Price or Anarchy. In particular we see two distinct regions of high PoA, one corresponding to $SNR < 1$ and the other to $SNR > 1$.

Figure 3.3 shows the Price of Anarchy for mean $m = 0.1, 0.5, 1$ and 10. To understand the qualitative features of these plots it helps to examine the case $m = 1$ more closely. Figure 3.4a shows this case on its own. It can be seen that there are two distinct regions where the PoA is relatively large. Relating γ to the SNR we see that these two regions roughly correspond to the cases $SNR < 1$ and the other to $SNR > 1$, or equivalently $\gamma < \frac{1}{m}$ and $\gamma > \frac{1}{m}$ respectively. The existence of these two regions can be explained by the following reasoning.

When $\gamma >> \frac{1}{m} = 1$ the variance in the population is relatively low. That means most types are concentrated around the mean $m = 1$, i.e. $\theta_i = \frac{c_i}{\ell_i} \approx 1 \implies c_i \approx \ell_i$. Because $v(A) = A$, when $A < 1$ we have $c_i > \ell_i v(A)$. In words, most of the agents value false alarms more than missed detections when both the population variance and attack strength are low. Notice from Figure 3.4b that in this regime we have $\theta_d^* > \theta_c^*$, which means more decentralized agents are vigilant ($T_i = A$) and fewer

are indifferent ($T_i = \infty$) compared to the centralized defenders.

If most decentralized defenders value false alarms over missed detections, why are too many of them vigilant? Recall that vigilant defenders will raise more alarms, and more alarms raised will result in more false alarms. To understand what is going on here, it is necessary to consider what strategies the central planner prescribes. Because we are in a regime where most agents value false alarms more than missed detections, the central planner prescribes more defenders to be indifferent as a means of reducing false alarms across the network. Prescribing more indifferent defenders means fewer alarms raised and hence less false alarms. But there is a secondary effect of increasing the infection rate across the network: The missed detection rate increases while false alarm rate decreases.

The central planner is able to account for this in his choice of optimal strategies. The decentralized defenders do not consider this externality on their choice of threshold strategies. That is because the population best response which determines the strategies of decentralized defenders does not take into account the externality of altered infection rates that is associated with changes to individual strategies. As a result, when the population is playing a centralized best response strategy, there exists a subset of the population of which each member believes he can unilaterally decrease his threshold (from $T_i = \infty$ to $T_i = A$) in such a way that reduces his over all cost. The problem with this reasoning is that it does not take into account the associated change in infection rates across the network. As a result infection rates go down as the deviant agents become more vigilant, thus raising the number of false alarms experienced across the network! Notice that the larger γ is for $A < 1$ the PoA increase. That is because more defenders are concentrated around $\frac{c_i}{\ell_i} = 1$ and hence we have more defenders who value false alarms more than missed detection ($c_i > \ell_i A$). As such the central planner

prescribes more defenders to be indifferent ($T_i = \infty$). This leads to a high rate of infection and more of the decentralized defenders will incorrectly reason that they can decrease their thresholds and decrease their missed detection rates. This leads to a relatively higher PoA.

A similar explanation can explain the high PoA when $\gamma << 1$. In this case we have a high variance. Because $\theta \sim \text{gamma}(\gamma, \gamma)$ as $\gamma \downarrow 0$ most of the population is concentrated close to $\theta_i \approx 0 \implies c_i << \ell_i$. This means for a larger range of A values we have $c_i < \ell_i A$, i.e. most agents value missed detection over false alarms. Thus by similar reasoning as before (but in reverse) we can see why $\theta_d^* < \theta_c^*$ with high PoA for a larger range of A values when $\gamma << 1$. We again arrive at a somewhat counter intuitive result that when more agents value missed detection over false alarms, by ignoring the externality of altered infection rates associated with individual strategy changes, decentralized agents are not vigilant enough!

Finally we observe that changing the mean of the population changes which of the two regimes described above dominate to numerical results. As can be seen in Figure 3.3, for $E[\theta] < 1$ we tend to see the region corresponding to $\gamma << 1$, whereas when $E[\theta] > 1$ we tend to see the region corresponding to $\gamma >> 1$. We conjecture that the two regions described above, which roughly correspond to "low" variance and "high" variance, are differentiated by $\gamma < \frac{1}{m}$ and $\gamma > \frac{1}{m}$ respectively. Recall this corresponds to the signal-to-noise ratio $SNR < 1$ and $SNR > 1$ respectively.

3.5 Game Equilibria with Strategic Bot master

Given the above population best response T^*, it is natural to look for pure Nash equilibria as were found in the homogeneous botnet game. Letting $\mathcal{P}_{\Theta \to \mathcal{T}}$ be the class of all population strategies, a pure Nash equilibrium is any pair

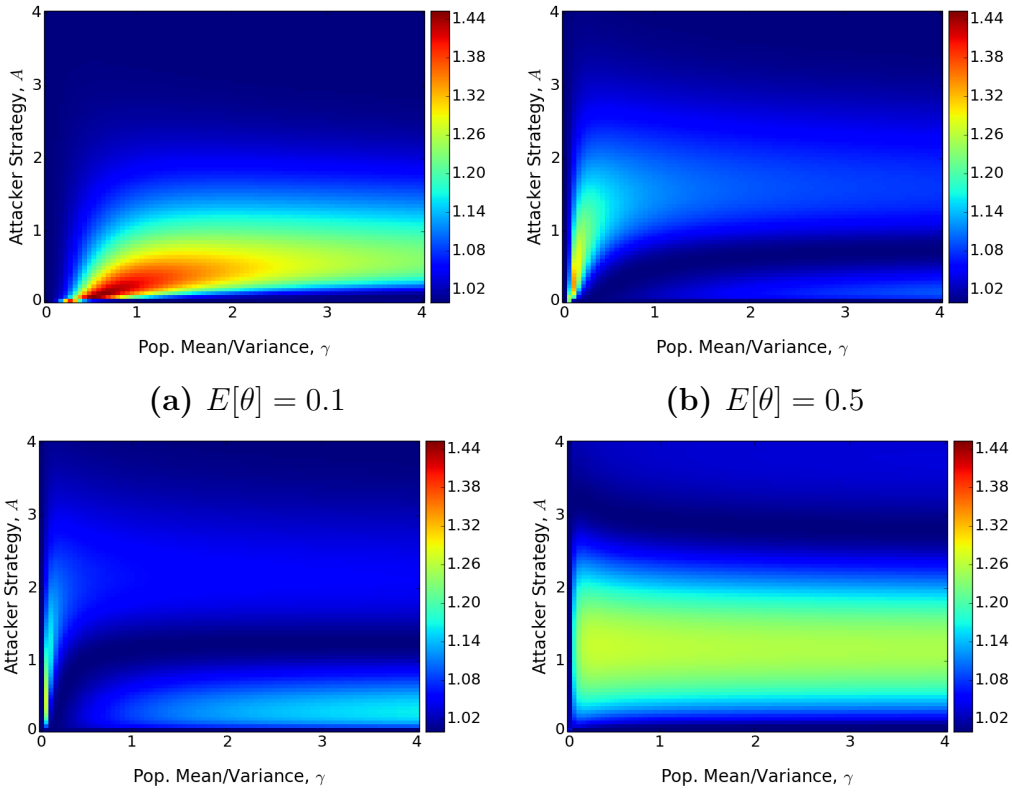

(a) $E[\theta] = 0.1$ **(b)** $E[\theta] = 0.5$

(c) $E[\theta] = 1$ **(d)** $E[\theta] = 10$

Figure 3.3: Price of Anarchy (PoA) in the heterogeneous botnet detection game.

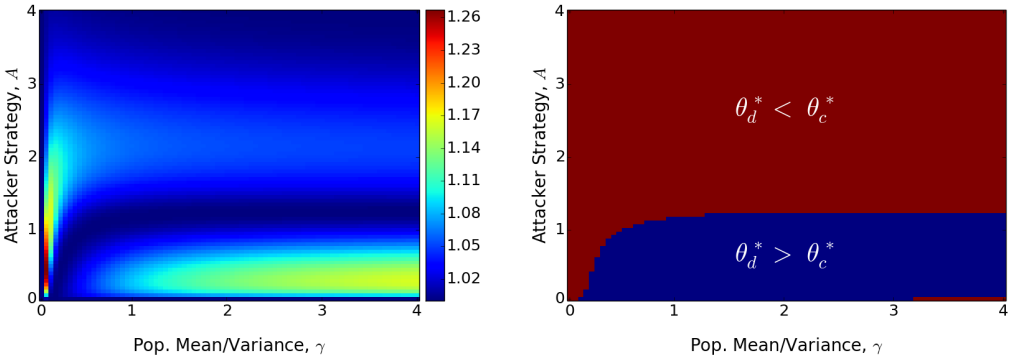

(a) Price of Anarchy **(b)** Red: $\theta_d^* < \theta_c^*$. Blue: $\theta_d^* > \theta_c^*$.

Figure 3.4: Price of Anarchy (PoA) and strategy comparison in the heterogeneous botnet detection game: $S \sim \exp(1)$, $\theta \sim \text{gamma}(\gamma, \gamma)$, $E[\theta] = 1$, $\text{var}(\theta) = \frac{1}{\gamma}$.

$(A^*, T^*) \in \mathcal{A} \times \mathcal{P}_{\Theta \to \mathcal{T}}$ such that $A^* \in \sigma_b(T^*)$ and $T^*(\theta_\varnothing) \in \sigma_\varnothing(A^*, T^*(\cdot))$. However what we find is that in many cases no such Nash equilibrium exists.

For example, suppose $\Theta = \mathbb{R}^+$ and $S \sim \exp(\beta)$. Assume that such a Nash equilibrium (A^*, T^*) does exist. Then T^* must be a population best response. Let κ and ω be as in Proposition 10 for this A^* and T^*. Let h^* be a value that satisfies (3.8) with $A = A^*$ and $T(\theta) = T^*(\theta)$. Note that $\kappa = \omega$ when $S \sim \exp(\beta)$.

First note that if $A^* = 0$ we have $\kappa = \omega = 0$ and $T^*(\theta_i) = +\infty$ for all types θ_i. However $h^* > 0$, so for any $A > 0$ we have $U(A, h^*) = Ah^* > A^*h^* = 0$. Thus the attacker can increase his expected utility by deviating from A^*. It follows that (A^*, T^*) is not a Nash equilibrium.

Now suppose $A^* > 0$. Since $S \sim \exp(\beta)$ the population best response T^* must have the form

$$T^*(\theta_i) = \begin{cases} A^* & \text{if } \theta_i \leq \kappa, \\ +\infty & \text{if } \theta_i > \kappa. \end{cases}$$

For $A \in [0, A^*)$ we can define $h(A)$ as a solution to

$$h = \int_\Theta F_S(T^*(\theta) - A) dF_\theta \left(1 - (1-p)e^{-\lambda q h}\right).$$

Using arguments similar to those in the homogeneous game we see that $\frac{\partial h}{\partial A} < 0$ and the bot master's utility function is strictly quasi-concave for $A \in [0, A^*]$. Thus there exists a value $\tilde{A} \in (0, A^*)$ such that $U(\tilde{A}, h) > U(A^*, h^*)$. Again it follows that (A^*, T^*) is not a Nash equilibrium. Thus by contradiction no such Nash equilibrium can exist.

In the above scenario if a heterogeneous population responds in equilibrium to a particular strategy A^*, then the attacker will always have an incentive to deviate from this attack strategy. One may then look for mixed attack strategies that lead

to Nash equilibria. That is we seek a distribution F_A over the strategy space $\mathcal{A}$ which the bot master can use to randomize his strategy. The problem with this approach is the fact that the defender type space is infinite and the defenders are decentralized. As a result if the attacker chooses a mixed strategy F_A, then he can only ever make a set of the population having measure zero indifferent to their strategies.

As an example consider the above case with $S \sim \exp(\beta)$ and Θ a locally compact subset of the real line. Suppose the bot master plays a mixed strategy F_A over the set of strategies $\mathcal{A}$. Assume further that the network of defenders is playing a population strategy $T(\theta)$. Now consider what a single decentralized defender will do in response the the strategy profile (F_A, T). Using arguments similar to those above, the probability of a false negative can be shown to be the solution, h, of the fixed point equation

$$h = \int_{\mathcal{A}} \int_{\Theta} F_S(T(\theta) - A) dF_\Theta dF_A [1 - (1-p)e^{-\lambda q h}].$$

With this it can be seen that a deviant root defender's expected cost function is as follows:

$$\begin{aligned}\bar{C}_\varnothing(T_\varnothing, T, F_A) = c_\varnothing [1 - F_S(T_\varnothing)](1-p)e^{-\lambda q h} \\ + \ell_\varnothing \int_{\mathcal{A}} v(A) F_S(T_\varnothing - A) dF_A [1 - (1-p)e^{-\lambda q h}].\end{aligned}$$

We then have the best response for a deviant root agent as

$$\begin{aligned}\sigma_\varnothing(F_A, T) = \arg\min_{T_\varnothing} \{c_\varnothing [1 - F_S(T_\varnothing)](1-p)e^{-\lambda q h} \\ + \ell_\varnothing \int_{\mathcal{A}} v(A) F_S(T_\varnothing - A) dF_A [1 - (1-p)e^{-\lambda q h}]\}.\end{aligned}$$

Because of the decentralized nature of the defender population, an individual selfish agent will always have an incentive to play a pure strategy. This is a result of the strict quasi convexity of his expected cost function regardless of the strategies of the population and the bot master. In order for the population strategy to be an equilibrium best response strategy we require that

$$T^*(\theta_i) = \begin{cases} \min \mathcal{A} & \text{if } \theta_i < \kappa, \\ +\infty & \text{if } \theta_i > \kappa \end{cases}$$

where

$$\kappa = \int_{\mathcal{A}} v(A) e^{\beta A} dF_A \frac{1-(1-p)e^{-\lambda q h}}{(1-p)e^{-\lambda q h}}.$$

Thus the only subset of the population that will be indifferent to their choice of strategy will be all defenders with type $\theta_i = \kappa$. But this is a set of measure zero and thus will have no impact on the equilibrium in the game. As such there will be no mixed strategy Nash equilibria in the heterogeneous botnet detection game.

3.5.1 Stackelberg Equilibria

Because there are cases where pure and mixed Nash equilibria do not exist, we may wish to consider alternative equilibrium solution concepts. In this section we examine Stackelberg equilibrium as a solution concept for the heterogeneous botnet detection game.

A Stackelberg equilibrium is one in which there is an explicit order to the game and one player is designated as a *leader* and the other players are *followers*. Moreover there is an implicit assumption that the leader has the authority or ability to guarantee that the followers do in fact follow. We will consider both attacker-as-leader and defenders-as-leader Stackelberg games. Numerical exam-

ples for attacker-as-leader Stackelberg games are given.

Defenders as Leader

When considering Stackelberg games with the defenders as leader we should take care to make explicit what we mean. When we consider the central planner, then it is clear that there is a single player who moves first. But when we consider decentralized defenders it is not clear what is meant by leader since there are an infinite number of self interested defenders in this case. Moreover we have only defined decentralized population strategies in terms of population best response functions. Thus we need an initial value for the decentralized defenders to "best respond" to.

For example let $t_0 \geq 0$. Then we can define the best response function $T^*(\theta, t_0)$ in the same way as in the previous section. By restricting the decentralized defenders to play best response type functions, we have essentially reduced the strategy space of the decentralized defenders to a single parameter with the mapping

$$t_0 \mapsto T^*(\cdot, t_0).$$

Note that in general this mapping is not unique. We will restrict attention to the cases in which it is unique to avoid complications. Further analysis is required for cases in which the mapping is not unique. Assuming $T^*(\cdot, t_0)$ can be identifies by $t_0 \geq 0$, we can then ask, how will the attacker best respond to this strategy? We have already defined the attacker's utility function:

$$U(A, T^*(\cdot, t_0)) = g(A)h(A, T^*(\cdot, t_0))$$

where h satisfies

$$h = \int_{\Theta} F_S(T^*(\theta, t_0) - A) dF_\theta \left(1 - (1-p)e^{-\lambda q h}\right).$$

One of the main difficulties with analyzing equilibria with heterogeneous defenders is that the attacker's utility function is not guaranteed to be quasi-concave, and as a result we cannot guarantee the continuity of the best response correspondence when it is single valued. This creates theoretical difficulties for proving existence of Nash equilibria and also numerical difficulties for computing approximate Stackelberg equilibria. Nevertheless by the continuity of $U(A, T^*(\cdot, t_0))$ in A we can define the attacker's best response correspondence as

$$\sigma_a(t_0) = \arg\max_A U(A, T^*(\cdot, t_0)).$$

Thus if the attacker best responds to the strategy t_0, the missed detection probability becomes $h(\sigma_a(t_0), T^*(\cdot, t_0))$, where h satisfies

$$h = \int_{\Theta} F_S(T^*(\theta, t_0) - \sigma_a(t_0)) dF_\theta \left(1 - (1-p)e^{-\lambda q h}\right). \tag{3.11}$$

To simplify notation we will define for each t_0 the value $h^*(t_0)$ to be the value satisfying (3.11).

Once the best response of the attacker is taken into consideration we need a mechanism for the decentralized defenders to agree upon a particular strategy t_0 to use. One option is to simply minimize the expected social cost as the central planner does. However one may object to the fact that decentralized defenders will not be able to agree on this strategy, as it may result in a large percentage of the population playing strategies which they are not happy with, i.e. they want

to unilaterally deviate from. An alternative mechanism by which to choose a t_0 value is to minimize a function which measures the degree to which the population wishes to deviate from the resulting strategy profile once the attacker has chosen a best response. For example define the functions $D_m : t_0 \mapsto \mathbb{R}^+$ as follows:

$$D_m(t_0) \triangleq \int_\Theta |F_S(T^*(\theta, t_0)) - F_S(T^*(\theta, \sigma_a(t_0)))|^m d\theta.$$

For $m > 0$ this function is a metric on $\mathbb{R}$ which measures the *regret* across the population of defenders. Since D_m is a metric we have $D_m(t_0) \geq 0$ for all t_0. Note that if there existed a value t_0^* such that $\sigma_a(t_0^*) = t_0^*$ then the strategy profile $(t_0^*, T^*(\cdot, t_0^*))$ is a Nash equilibrium and we have $D_m(t_0^*) = 0$. Absent a Nash equilibrium it stands to reason that the population would be interested in minimizing the regret across the network, as this will minimize the degree to which they want to unilaterally deviate from the resulting strategy profile. Because $D_m(t_0)$ is bounded below, if it were continuous in t_0 then we could guarantee the existence of a global minimum. We could then define a *Decentralized Stackelberg equilibrium* as any strategy t_s that satisfies

$$t_s \in \arg\min_{t_0} D_m(t_0).$$

However we are not guaranteed that $D_m(t_0)$ is continuous in t_0, because we are not guaranteed that $\sigma_a(t_0)$ is continuous in t_0. Further consideration of the attacker's best response correspondence is needed. One way to deal with this is to consider only sufficiently large t_0. Notice that as $t_0 \to +\infty$ we have $T^*(\theta, t_0) - t_0 \to 0$ for all θ. Thus for large enough t_0 the population will be highly concentrated around a single strategy and we can approximate the attacker's best response correspondence as if it is responding to the single strategy t_0. In this case we are

guaranteed for $\sigma_a(t_0)$ to be single valued and continuous. Let

$$t_m \triangleq \inf\{x \geq 0 : \sigma_a(t_0) \in \mathcal{C} \text{ for all } t_0 \in (x, \infty)\}.$$

Then $\sigma_a(t_0) \in \mathcal{C}$ for all $t_0 > t_m$. then we can define the *Decentralized Stackelberg equilibrium* as any strategy t_s that satisfies

$$t_s \in \arg\min_{t_0 \geq t_m} D_m(t_0).$$

If

$$t_m \in \{x \geq 0 : \sigma_a(t_0) \in \mathcal{C} \text{ for all } t_0 \in (x, \infty)\}$$

then we are guaranteed for such an equilibrium to exist.

Alternatively one might consider minimizing the *surprise* of the decentralized agents. Define the surprise as

$$S_m(t_0) \triangleq |t_0 - \sigma_a(t_0)|^m$$

This is a metric when $\sigma_a(t_0)$ is single valued and $S_m(t_0) = 0$ implies $(t_0, T^*(, t_0))$ is a Nash equilibrium.

Unfortunately, because of the attacker's discontinuous and possibly set valued best response correspondence, it is very difficult to guarantee that such a Stackelberg equilibria will exist under any of the above definitions. The best one could hope for is to numerically search for ϵ-Stackelberg equilibria given one of the above mechanisms for coordinating the decentralized defenders. When defenders best respond we always have $T^*(\theta, t_0) \geq t_0$ for all θ. Thus if we require that the defenders choose a strategy that forces the attacker to play at most the minimum threshold played by all defenders, i.e. $\sigma_a(t_0) \leq t_0$, we can define decentralized

Stackelberg equilibria given any of the mechanisms discussed here. This could also give us a way of numerically computing the optimal response of the attacker more efficiently. If we can restrict our attention to search between 0 and t_0 we can use standard optimization tools, as the attacker's utility function is guaranteed to be quasi-concave in this region.

Finally we note that one can always trivially define a centralized Stackelberg equilibrium by forcing the centralized planner to prescribe a single strategy for the entire network of defenders. This of course reduces to the homogeneous game considered in the previous chapter. But since decentralized heterogeneous defenders will never settle on a single strategy as a population best response, it is not clear how to define a decentralized Stackelberg strategy which can be used to compare to this centralized Stackelberg strategy.

Attacker as Leader

Given the situation in which no pure or mixed Nash equilibria exist, the attacker may consider signaling his intended strategy knowing that the networked population will respond with an optimal equilibrium strategy. The bot master would then seek a strategy

$$A^* = \arg\max_A U(A, T^*(A, \cdot)),$$

where $T^*(A, \cdot)$ is the population best response to strategy A. The next theorem states the existence of such an optimal signaling strategy, i.e. there exists a Stackelberg equilibrium. For simplicity we restrict our attention to the case $\frac{f_S(z-A)}{f_S(z)}$ is constant in z, i.e. $S \sim \exp(\beta)$. Furthermore we require $\Theta = \mathbb{R}^+$ and F_θ to be

continuously differentiable with density f_θ such that

$$\frac{(1 - F_\theta(x))^2}{f_\theta(x)} = O(x). \tag{3.12}$$

Many parametric distributions on $\mathbb{R}^+$ satisfy (3.12) such as Gamma, Log-Normal and Pareto distributions.

Theorem 3. *There exists a Stackelberg equilibrium in the heterogeneous population botnet detection game when $S \sim \exp(\beta)$ and F_θ satisfies (3.12).*

Proof. To determine the existence of a Stackelberg equilibrium we consider the heterogeneous population best response as a function of the attacker strategy A. In particular the optimal population response T^* depends on A so we write $T^*(A, \theta)$. We write h^* for a value satisfying

$$\begin{aligned} h^* &= \int_\Theta F_S(T^*(A,\theta) - A) dF_\theta [1 - (1-p)e^{-\lambda q h^*}], \\ &= (1 - F_\theta(L(A))) \left[1 - (1-p)e^{-\lambda q h^*}\right], \end{aligned}$$

where $L(A) = v(A)e^{\beta A \frac{1-(1-p)e^{-\lambda q h^*(A)}}{(1-p)e^{-\lambda q h^*(A)}}}$. We then have the expected utility of the bot master as $U(A, T^*(A, \cdot)) = g(A)h^*(A)$.

If $A = 0$ we have $U(0, T^*(0, \cdot)) = 0$. On the other hand $\lim_{A\to\infty} L(A) = +\infty$ which implies $\lim_{A\to\infty} h^*(A) = 0$. Furthermore the properties of $g(\cdot)$ imply $g(A) \leq A$ which give us $g(A)h^*(A) \leq Ah^*(A)$ for all A. For notational convenience define $\rho = (1-p)e^{-\lambda q h^*(A)}$. We then have

$$0 \leq \lim_{A\to\infty} g(A)h^* \leq \lim_{A\to\infty} Ah^* = \lim_{A\to\infty} \frac{A}{1/h^*} = \lim_{A\to\infty} -\frac{h^{*2}}{\frac{\partial h^*}{\partial A}},$$

where the last equality follows from l'Hopital's rule. After computing $\frac{\partial h^*}{\partial A}$ the last

term becomes

$$\lim_{A\to\infty} \frac{(1-F_\theta(L(A)))^2(1-\rho)[1+\lambda q f_\theta(L(A))L(A)]}{f_\theta(L(A))\beta e^{\beta A\frac{1-\rho}{\rho}}[k'(A)/\beta+v(A)]}$$
$$\leq \lim_{A\to\infty} \frac{(1-F_\theta(L(A)))^2(1-\rho)[1+\lambda q f_\theta(L(A))L(A)]}{\beta f_\theta(L(A))L(A)}.$$

Notice that if $\lim_{x\to\infty} x f_\theta(x) > 0$ then the last limit is exactly zero and we are done. However if $\lim_{x\to\infty} x f_\theta(x) = 0$ then condition (3.12) guarantees the limit goes to zero. Since $\lim_{A\to\infty} U(A, T^*(A,\cdot)) = \lim_{A\to 0} U(A, T^*(A,\cdot)) = 0$ and $U(A, T^*(A,\cdot)) < \infty$, there must exist a global maximum for some $A^* \in (0,\infty)$. The strategy profile $(A^*, T^*(A^*,\cdot))$ is a Stackelberg equilibrium. □

3.5.2 Attacker as Leader Numerical Results and Discussion

Similar arguments as above guarantee the existence of a Stackelberg equilibrium in a game between the bot master and a centrally planned network. Note that the central planner and decentralized defenders do not necessarily choose the same step location θ^* in response to a given A. Thus we let $\theta_c(A)$ and $\theta_d(A)$ be the optimal step locations chosen by the central planner and decentralized defenders, respectively, in response to the strategy A. The Stackelberg equilibrium is uniquely determined by the bot master strategy A^* and the location of the step θ^*. Thus we let (A_c^*, θ_c^*) and (A_d^*, θ_d^*) denote the Stackelberg equilibria in the centralized network and decentralized network, respectively.

We numerically investigate the Stackelberg equilibria of the following model: $S_i \sim \exp(1)$, $\theta_i = \epsilon_i + \frac{\gamma-1}{\gamma}$, with $\epsilon_i \sim \exp(\gamma), \gamma \geq 1, c_i = \sin(\arctan(\theta_i)), \ell_i = \cos(\arctan(\theta_i))$, $g(A) = v(A) = A$, $p = 0.1$, and $\lambda q = 5$. Notice that $E[\theta_i] = 1$ and $Var[\theta_i] = \frac{1}{\gamma^2}$. Thus as $\gamma \to \infty$ we have $Var[\theta_i] \to 0$. We compare the effects

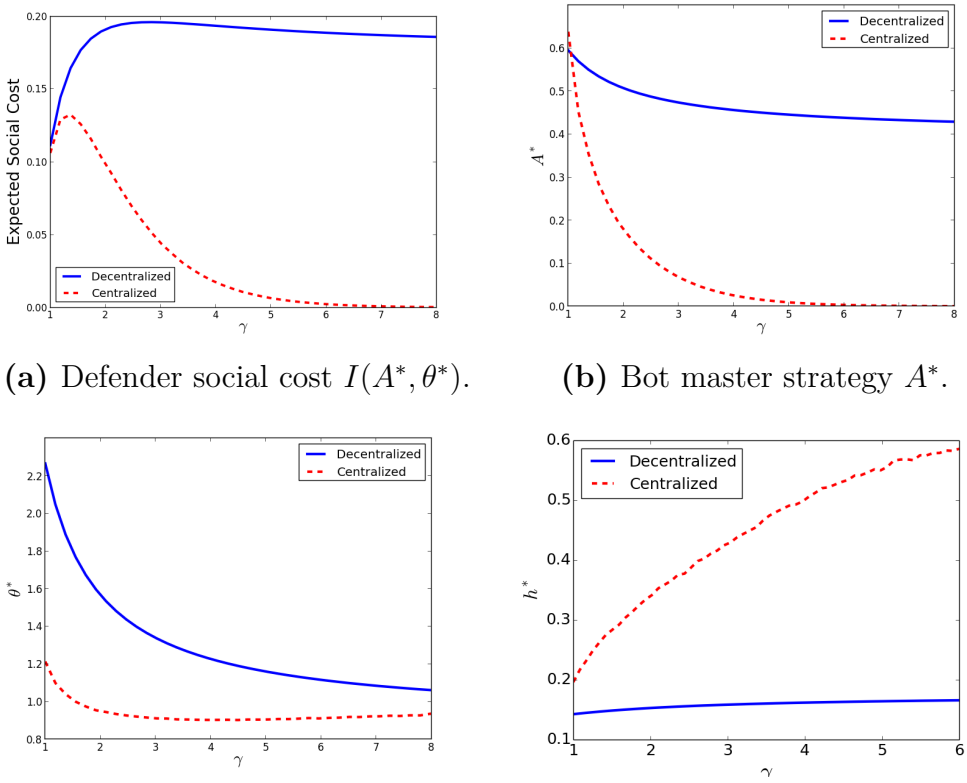

(a) Defender social cost $I(A^*, \theta^*)$.

(b) Bot master strategy A^*.

(c) Defender step location θ^*.

(d) Missed detection rate h^*.

Figure 3.5: Values at Stackelberg equilibrium with attacker as leader for varying values of γ. Larger values of γ correspond to less heterogeneity in the population.

of increasing the relative heterogeneity by decreasing $\gamma \downarrow 1$. Figures 3.5a, 3.5b, 3.5c and 3.5d show expected social cost, A^*, θ^* and h^*, respectively, for both the centralized and decentralized games. We note the following qualitative features from the numerical results:

- The equilibrium strategies θ_c^*, θ_d^*, A_c^* and A_d^* tend to decrease in γ. Thus increasing population heterogeneity results in the bot master stealing more resources from fewer defenders.

- For larger γ (lower heterogeneity) the central planner prescribes the minimum threshold to more defenders, i.e. $\theta_c(A) > \theta_d(A)$. In return the bot

master is forced to play lower values of A at equilibrium against the central planner. The result is $\theta_c^* < \theta_d^*$ and $A_c^* < A_d^*$, which has the effect of more false negatives and less false positives for the central planner. However the lower $v(A_c^*)$ offsets the costs associated with more false negatives. The net effect is that a centrally planned network admits a larger, less aggressive botnet and has lower expected social cost than the decentralized defenders, i.e. $I(A_c^*, \theta_c^*) < I(A_d^*, \theta_d^*)$.

- For $\gamma \approx 1$ (higher heterogeneity) the central planner prescribes the minimum threshold to fewer defenders, i.e. $\theta_c(A) < \theta_d(A)$. In return the bot master plays higher values of A at equilibrium against the central planner. The result is $\theta_c^* < \theta_d^*$ and $A_c^* > A_d^*$. The lower θ_c^* again results in more false negatives and fewer false positives for the central planner. The higher A_c^*, on the other hand, reduces the false negative rate enough to counter the increase in $v(A_c^*)$. The net effect is that a centrally planned network admits a larger, more aggressive botnet than the decentralized defenders, but still achieves a lower expected social cost, i.e. $I(A_c^*, \theta_c^*) < I(A_d^*, \theta_d^*)$. This may seem counterintuitive, but the lower false positive rate is enough to guarantee a lower expected social cost for the centrally planned network. We speculate that if defenders incur a fixed cost for raising an alarm this effect may diminish.

A possible criticism of the current model is that network attackers are often able to respond to the strategies of network defenders with relative ease. Thus it is not clear that a bot master would have an incentive to communicate his strategy if he can easily adapt to new defense strategies. While this is a valid point, such asymmetric advantages are usually more relevant for the initial intrusion of the network, which is not the focus of our model. After a network is compromised a

strategic decision must be made on how to utilize the compromised network. It is not clear that the network attacker can rely solely on technical skills in this capacity.

There exist several avenues for future research. Investigating the effects of different heterogeneities, such as heterogeneity in the resource $\mathcal{R}$, is of interest. In this case the bot master may then adopt different strategies A_i for each type of resource R_i. One could also explore the effects of different distributions placed on types, F_θ. For example, a Pareto distribution exhibits qualitatively different equilibria than observed in the present model while other distributions may allow for non-unique population best response strategies.

Chapter 4

A Two-Player Adversarial Sequential Detection Game

4.1 Introduction

Intrusion detection systems (IDS) have become an integral component in securing modern computer networks. The most basic underlying statistical model used by an IDS is a simple hypothesis testing procedure. While fixed sample size hypothesis testing can be effective, the dynamic nature of computer networks and the need for real time detection suggests sequential hypothesis testing may be more suitable. As the classic work by Wald [80] showed, if the observation process is costly, sequential hypothesis testing can result in a significant decrease in costs, both in terms of observation time and detection error rates.

While the importance of game theory in developing robust intrusion detection systems has been recognized [4–6], less attention has been paid to applying game theory to sequential detection problems. For the most part applications of game theory to sequential hypothesis testing have typically been restricted to robust

minimax solutions [13, 80], which assumes a zero-sum game between *observer* and *nature*. Given the vast array of security threats and strategic adversaries in the cyber domain, one potential shortcoming of the minimax approach is the fact that many non-cooperative, strategic encounters may not be zero-sum. If a defending agent has information about the type of adversary, such as the attacker's payoff function, then the defending agent may be able to leverage this information to find superior sequential detection tests.

This paper examines a two-player, non-zero-sum, sequential detection game between a *defender* agent and an *attacker* agent. The game is motivated by problems arising in the cyber-security domain. Botnets [67] or electricity theft in the smart grid [26] are examples of such scenarios we have in mind. The defender's objective is to sequentially detect whether or not his secured cyber infrastructure has been compromised by the attacker. It is his objective to do so in such a way that minimizes a payoff function which takes into account the expected observation time and both type I and type II detection errors. As such the defender's optimal sequential test is a version of Wald's Sequential Probability Ratio Test (SPRT) [80]. The attacker is interested in bypassing the defender's security in order to establish long-term, unrestricted access to the resources available on the system. The attacker's objective is not necessarily to destroy or damage the defender's system, but to utilize system resources. The attacker must then balance how aggressive he should be in utilizing resources of the compromised system and how stealthy he should be in order to avoid detection.

The main theoretical result is a proof of the existence of pure Nash equilibria in the special case that the attacker does not discount future expected utility. Furthermore we give conditions for the existence of Stackelberg equilibria with the defender as leader in the special case that the defender's strategy is restricted

to Wald's SPRT. Numerical examples are given to explore the qualitative features of the equilibria. It is observed that both low false positive costs for the defender and high prior probabilities of intrusion by the attacker lead to an infinite number of Nash equilibria in which the defender immediately classifies his system as compromised and the attacker receives no utility. Conversely we see that both high false positive costs for the defender and low prior probabilities of intrusion by the attacker lead to non-trivial Nash equilibria. Finally we see that it is possible for the defender to improve his outcome under the Stackelberg equilibrium strategy in relation to the Nash equilibrium strategy.

Previous examples of sequential detection games have largely been restricted to discrete-time, zero-sum games. As mentioned above minimax sequential detection assumes the form of a zero-sum game between observer and nature. Minimax sequential detection was an attempt to develop more robust sequential statistical tests [13] rather than explicitly address interference by a strategic adversary. Nevertheless, minimax sequential detection lends itself to an adversarial framework and has been used in game-theoretic settings. Such an approach was taken in [60] and [59] with applications to detecting access layer misbehavior in wireless networks. A discrete-time, non-zero-sum, network security classification game involving Wald's SPRT can be found in [15]. This work was largely numerical as the intractability of the discrete-time SPRT in an adversarial setting is not amenable to analysis. To our knowledge these are some of the only attempts to apply game theoretic reasoning to sequential detection. A similar fixed sample size detection game dealing with electricity theft in the smart grid can be found in [26].

4.2 Sequential Detection

The theory of sequential hypothesis testing was initiated by Wald [78]. The asymptotic analysis used to obtain approximate solutions to error probabilities and expected stopping times is related to detecting the drift of a brownian motion. The continuous sample paths of standard Brownian motion avoid the problem of *over shoot* encountered in the discrete-time case. In addition, when considering a large number of i.i.d. sequential observations, one can approximate the cumulative sum of the observations, appropriately scaled in time and space, by a Brownian motion. As such we will focus on the continuous-time case of sequentially detecting the drift of a Brownian motion. Furthermore our model focuses on the Bayesian point of view of this problem. We first give a brief overview of some standard results on the optimal sequential detection of the drift of a Brownian motion from the Bayesian point of view. For a detailed treatment the reader is referred to the texts [56, 58, 64].

Let $(\Omega, \mathcal{F}, P_\pi)$ be a probability space with $\pi \in [0,1]$. We assume there exists a random variable $\theta \in \{0,1\}$ such that

$$P_\pi(\theta = 1) = 1 - P_\pi(\theta = 0) = \pi.$$

Let Z_t be a stochastic process of the form

$$Z_t = \theta\mu t + W_t,$$

where $\mu \neq 0$ and W_t is a standard Brownian motion under P_π. Let P_i be the distribution of the observed process Z_t assuming $\theta = i$ for $i = 0, 1$. Then we can write $P_\pi = (1-\pi)P_0 + \pi P_1$. Similarly let E_i be the expectation operator under P_i for $i = 0, 1, \pi$.

It is assumed that the process Z_t is observed in order to test the following hypothesis:

$$H_0 : Z_t = W_t, t \geq 0, \tag{4.1}$$

$$H_1 : Z_t = \mu t + W_t, t \geq 0. \tag{4.2}$$

The observer seeks a sequential decision rule (τ, δ), where $\tau \in \mathbb{R}$ is a stopping time and $\delta \in \{0, 1\}$ a final decision. Note both τ and δ are random variables under P_π which depend on the stochastic process Z_t. We define the cost associated with the observer's detection as

$$C(\tau, \delta) \triangleq \tau + \alpha \mathbb{1}_{\{\theta=0, \delta=1\}} + \beta \mathbb{1}_{\{\theta=1, \delta=0\}},$$

where $\alpha, \beta > 0$ are the costs associated with false positives and false negatives, respectively, and $\mathbb{1}_{\{.\}}$ is the indicator random variable defined on $(\Omega, \mathcal{F}, P_\pi)$. The observer wishes to minimize the overall expected value of this cost with respect to P_π. We thus define the *value function*

$$V(\pi) \triangleq \inf_{\tau, \delta} E_\pi \left[C(\tau, \delta) \right]. \tag{4.3}$$

As in the classical, discrete-time case considered by Wald, solving (4.3) involves reducing it to an optimal stopping problem [64]. Specifically (4.3) can be reduced to

$$V(\pi) = \inf_{\tau} E_\pi \left[\tau + \min \left\{ \alpha(1 - \pi_\tau), \beta \pi_\tau \right\} \right], \tag{4.4}$$

where π_t is the *posterior probability process*,

$$\pi_t \triangleq \left(\frac{\pi}{1-\pi}\Lambda_t\right) \Big/ \left(1 + \frac{\pi}{1-\pi}\Lambda_t\right),$$

and Λ_t is the *likelihood ratio process*,

$$\Lambda_t \triangleq \exp\left\{\mu\left(Z_t - \frac{1}{2}\mu t\right)\right\}.$$

The optimal decision δ^* for any stopping time τ is given by $\delta^* = 1$ if $\pi_\tau \geq \frac{\alpha}{\alpha+\beta}$ and $\delta^* = 0$ if $\pi_\tau \leq \frac{\alpha}{\alpha+\beta}$. All that is needed is to determine the optimal stopping time τ^* satisfying (4.4). By reducing the optimal stopping problem (4.4) to a free-boundary problem [56] one can find an explicit formula for $V(\pi)$ and the optimal stopping time τ^*. The following theorem gives us this main result due to Shiryaev [64]. For ease of notation we define the functions $\psi : [0,1] \to \mathbb{R}$ and $\nu : [0,1]^3 \to \mathbb{R}$ as follows:

$$\begin{aligned}
\psi(x) &\triangleq (1-2x)\log\left(\frac{x}{1-x}\right), \\
\nu(x,y,z) &\triangleq \frac{2}{z^2}(\psi(x) - \psi(y)) + \left(\beta - \frac{2}{z^2}\psi'(y)\right)(x-y) + \beta y.
\end{aligned}$$

Theorem 4 (Shiryaev [64])**.** *For $\pi \in (0,1)$ the value function $V(\pi)$ in (4.4) is given by*

$$V(\pi) = \begin{cases} \nu(\pi, \pi_\ell, \mu) & \text{if } \pi \in (\pi_\ell, \pi_u), \\ \min\{\alpha(1-\pi), \beta\pi\} & \text{if } \pi \notin (\pi_\ell, \pi_u), \end{cases}$$

where the values π_ℓ and π_u satisfy $\pi_\ell \in \left(0, \frac{\alpha}{\alpha+\beta}\right)$ and $\pi_u \in \left(\frac{\alpha}{\alpha+\beta}, 1\right)$ and are the

unique solutions to the following transcendental equations:

$$\nu(\pi_u, \pi_\ell, \mu) = \alpha(1 - \pi_u), \tag{4.5}$$

$$\left.\frac{\partial \nu(x, \pi_\ell, \mu)}{\partial x}\right|_{x=\pi_u} = -\alpha. \tag{4.6}$$

The optimal stopping time is

$$\tau^* = \inf\{t \geq 0 : \pi_t \notin (\pi_\ell, \pi_u)\}.$$

In what follows it will be convenient to express the stopping time in terms of the likelihood ratio process Λ_t. As such we define the values A and B as

$$A \triangleq \min\left\{\frac{1-\pi}{\pi}\frac{\pi_\ell}{1-\pi_\ell}, 1\right\}, \tag{4.7}$$

$$B \triangleq \max\left\{\frac{1-\pi}{\pi}\frac{\pi_u}{1-\pi_u}, 1\right\}. \tag{4.8}$$

The optimal stopping time is then

$$\tau^* = \inf\{t \geq 0 : \Lambda_t \notin (A, B)\}.$$

Note the restriction $A \leq 1 \leq B$. The case $A = 1$ corresponds to the observer immediately accepting H_0, i.e. $\tau^* = 0$ and $\delta^* = 0$ a.s. The case $B = 1$ corresponds to the defender immediately accepting H_1, i.e. $\tau^* = 0$ and $\delta^* = 1$ a.s.

4.3 Adversarial Sequential Detection

We now consider a two-player, non-zero-sum, sequential detection game motivated by problems arising in the cyber-security domain. A *defender* is in charge of protecting a secured cyber-system. This could be a single computer, a network

of computers or some other cyber-physical infrastructure like the smart grid. The defender's objective is to detect whether or not this system has been compromised by the attacker. The defender makes noisy observations of the system's state which we model by a stochastic process Z_t. We assume that whether or not the system is compromised can be discerned through the drift of the process Z_t. For example, the observed process could be cumulative bandwidth usage, CPU load or energy consumption.

An *attacker* is interested in infiltrating the defender's system in order to establish long-term, unrestricted access to the resources available on the system. The attacker's objective is not necessarily to destroy or damage the defender's system. In fact he may be interested in the long-term viability of the system's resources so that he may benefit by illicitly using them. The more the attacker utilizes the system the more utility he obtains. However this also increases the drift of the observed stochastic process, thus increasing the probability of detection. The attacker must then balance how aggressively he utilizes the resources of the compromised system and how stealthy he should be to avoid detection.

As in the previous section we assume $Z_t = \theta\mu t + W_t$ on the same probability space $(\Omega, \mathcal{F}, P_\pi)$ for some fixed $\pi \in [0, 1]$. In the context of our game the random variable θ has the interpretation of whether or not the attacker's attempts at bypassing the defender's security were successful. This is not a strategic variable for the attacker. It is assumed that θ is independent of W_t and all strategic actions. Instead we are interested in analyzing the attacker's strategic utilization of the compromised system given a successful intrusion. Since the attacker's aggressiveness in utilizing system resources is reflected in the drift of the stochastic process Z_t, we take the drift $\mu > 0$ to be the strategic variable of the attacker.

4.3.1 Defender Expected Cost and Best Response

In this section we focus on finding the defender's best response to an attacker's strategy $\mu > 0$. If μ is known by the defender we assume he seeks a sequential decision rule (τ, δ) that minimizes the expected cost associated with the observation and detection process. We furthermore assume the defender uses the same expected cost function as in the classical SPRT case. However in the context of our game we expect the cost of a false negative β to increase with μ. This follows from the interpretation of μ as a measure of stolen resources from the defender's system. The higher μ is, the more the defender loses. Thus we assume $\beta(\mu)$ is a monotonically increasing function with $\beta(0) = 0$. Furthermore we assume $\lim_{\mu\to\infty} \beta(\mu) = \infty$. On the other hand the cost of a false positive α should not depend on the value of μ since under H_0 the attacker fails to bypass the defender's security and is unable to utilize any resources. Thus we assume α is constant in μ. Also note that without loss of generality we are assuming a per unit time observation cost of 1.

The expected cost to the defender under the prior P_π is then

$$E_\pi \left[C(\tau, \delta)\right] = E_\pi[\tau] + \alpha P_0(\delta = 1)(1 - \pi) + \beta(\mu) P_1(\delta = 0)\pi.$$

The defender wishes to minimize this expected cost, thus the value function is defined as

$$V_\pi(\mu) = \inf_{\tau,\delta} E_\pi \left[C(\tau, \delta)\right]. \tag{4.9}$$

Note that we consider the value to be a function of the drift μ as opposed to the parameter π as in (4.3). This is because we assume $\pi \in [0, 1]$ is fixed throughout while the equilibrium analysis depends on variations in μ. Despite this difference,

(4.9) is equivalent to the classical sequential detection of the drift of a Brownian motion for fixed $\mu > 0$. By modifying Theorem 4 to include $\beta(\mu)$ we can define functions $\pi_u(\mu)$ and $\pi_\ell(\mu)$ as the unique solutions to (4.5) and (4.6). Then (4.7) and (4.8) give us functions $A(\mu)$ and $B(\mu)$. For each $\mu > 0$ we will refer to the associated sequential probability ratio test for testing H_0 versus H_1 as $\text{SPRT}(A(\mu), B(\mu))$. We then have the following.

Proposition 11. *If the attacker chooses a drift $\mu > 0$, then a best response for the defender is the SPRT$(A(\mu), B(\mu))$.*

4.3.2 Adversarial Sequential Detection Statistics

Suppose now that the defender plays the $\text{SPRT}(A(\mu), B(\mu))$ in response to the attacker strategy μ. In order to determine the attacker's best response to the strategy $\text{SPRT}(A(\mu), B(\mu))$ we will need to consider what happens when the defender keeps his statistical test fixed while the attacker unilaterally deviates from the value μ and chooses a value $\tilde{\mu}$. In this case we can no longer rely on standard results from Bayesian sequential detection theory to derive the SPRT statistics. Recall that the stopping time and error probabilities are determined by the likelihood ratio process Λ_t, which is a function of the observed data Z_t. The defender derives Λ_t from the hypotheses H_0 and H_1. However, when the attacker chooses a drift $\tilde{\mu} \neq \mu$ there is a corresponding hypothesis, $\tilde{H}_1$, conditional distribution $\tilde{P}_1$ and prior distribution $\tilde{P}_\pi$, which are all distinct from H_1, P_1 and P_π, respectively:

$$\begin{aligned} \tilde{H}_1 : Z_t &= \tilde{\mu}t + W_t, t \geq 0, \\ \tilde{P}_1(\cdot) &= \tilde{P}_\pi(\cdot|\theta = 1), \\ \tilde{P}_\pi &= (1-\pi)P_0 + \pi\tilde{P}_1. \end{aligned}$$

When the defender evaluates Λ_t at the observed data, it will no longer be a true likelihood ratio process under $\tilde{P}_\pi$, but a corrupted likelihood ratio process. In this case the expected stopping times and error probabilities under prior $\tilde{P}_\pi$ and hypotheses H_0 and $\tilde{H}_1$ will diverge from those anticipated by the defender under prior P_π and hypotheses H_0 and H_1.

In the next two propositions we state standard results regarding error rates and expected stopping times under the assumption that the defender chooses the $\text{SPRT}(A(\mu), B(\mu))$ for testing H_0 versus H_1, while the attacker chooses the drift $\tilde{\mu}$ with alternative hypothesis $\tilde{H}_1$.

We still assume $\theta \sim \text{Bern}(\pi)$ and W_t is a standard Brownian motion, but we define the following stochastic processes to account for the discrepancy in drifts:

$$\begin{aligned}
\tilde{Z}_t &\triangleq \tilde{\mu}\theta t + W_t, \\
\tilde{\Lambda}_t &\triangleq \exp\left\{\mu\left(\tilde{Z}_t - \frac{1}{2}\mu t\right)\right\}, \\
T &\triangleq \inf\{t \geq 0 : \tilde{\Lambda}_t \notin (A, B)\}, \\
T_A &\triangleq \inf\{t \geq 0 : \tilde{\Lambda}_t \leq A\}, \\
T_B &\triangleq \inf\{t \geq 0 : \tilde{\Lambda}_t \geq B\}.
\end{aligned}$$

Proposition 12. *Suppose given the prior P_π and hypotheses H_0 and H_1 a sequential detection procedure SPRT(A, B) is used to test the hypothesis H_0 versus H_1 as in (4.1) and (4.2), respectively. Given the process $\tilde{Z}_t$, the error probabilities under the prior $\tilde{P}_\pi$ and hypotheses H_0 and $\tilde{H}_1$, respectively, are*

$$\begin{aligned}
P_0(\delta = 1) &= \frac{1 - A}{B - A}, \\
\tilde{P}_1(\delta = 0) &= \frac{1 - B^{1-2\frac{\tilde{\mu}}{\mu}}}{A^{1-2\frac{\tilde{\mu}}{\mu}} - B^{1-2\frac{\tilde{\mu}}{\mu}}}.
\end{aligned}$$

Proof. Given the definitions above we have

$$\tilde{\Lambda}_t = \exp\left\{\mu\left((\tilde{\mu}\theta - \frac{1}{2}\mu)t + W_t\right)\right\}.$$

Under H_0 we have $\tilde{\Lambda}_t = \exp\left\{-\frac{1}{2}\mu^2 t + \mu W_t\right\}$ while under $\tilde{H}_1$ we have $\tilde{\Lambda}_t = \exp\left\{(\mu\tilde{\mu} - \frac{1}{2}\mu^2)t + \mu W_t\right\}$. Recall that a stochastic process of the form

$$\exp\left\{-\frac{1}{2}\xi^2 t + \xi W_t\right\}$$

is a martingale for any $\xi \in \mathbb{R}$. Setting $\xi = \mu$ we see that $\exp\left\{-\frac{1}{2}\mu^2 t + \mu W_t\right\}$ is a martingale. Since T is a bounded stopping time it follows from the optional stopping theorem that $E_0[\tilde{\Lambda}_T] = E\left[\exp\left\{\mu W_T - \frac{1}{2}\mu^2 T\right\}\right] = 1$.

Setting $\xi = 1 - 2\frac{\tilde{\mu}}{\mu}$ we have

$$\begin{aligned}\exp\left\{-\frac{1}{2}\xi^2 t + \xi W_t\right\} &= \exp\left\{-\frac{1}{2}(1 - 2\frac{\tilde{\mu}}{\mu})^2\mu^2 t + (1 - 2\frac{\tilde{\mu}}{\mu})\mu W_t\right\} \\ &= \left(\exp\left\{(\mu\tilde{\mu} - \frac{1}{2}\mu^2)t + \mu W_t\right\}\right)^{1-2\frac{\tilde{\mu}}{\mu}}.\end{aligned}$$

It follows that $\left(\exp\left\{(\mu\tilde{\mu} - \frac{1}{2}\mu^2)t + \mu W_t\right\}\right)^{1-2\frac{\tilde{\mu}}{\mu}}$ is a martingale, and the optional stopping theorem gives us

$$\tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}\right] = E\left[\left(\exp\left\{(\mu\tilde{\mu} - \frac{1}{2}\mu^2)t + \mu W_t\right\}\right)^{1-2\frac{\tilde{\mu}}{\mu}}\right] = 1.$$

We now have

$$E_0\left[\tilde{\Lambda}_T\right] = \tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}\right] = 1. \tag{4.10}$$

Furthermore we can write

$$E_0\left[\tilde{\Lambda}_T\right] = E_0\left[\tilde{\Lambda}_T|T_A < T_B\right] P_0(\delta = 0) + E_0\left[\tilde{\Lambda}_T|T_B < T_A\right] P_0(\delta = 1),$$
$$\tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}\right] = \tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}|T_A < T_B\right] \tilde{P}_1(\delta = 0) + \tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}|T_B < T_A\right] \tilde{P}_1(\delta = 1).$$

Also note the following:

$$E_0\left[\tilde{\Lambda}_T|T_A < T_B\right] = A,$$
$$E_0\left[\tilde{\Lambda}_T|T_B < T_A\right] = B,$$
$$\tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}|T_A < T_B\right] = A^{1-2\frac{\tilde{\mu}}{\mu}},$$
$$\tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}|T_B < T_A\right] = B^{1-2\frac{\tilde{\mu}}{\mu}}.$$

This then gives us

$$E_0\left[\tilde{\Lambda}_T\right] = AP_0(\delta = 0) + BP_0(\delta = 1),$$
$$\tilde{E}_1\left[\tilde{\Lambda}_T^{1-2\frac{\tilde{\mu}}{\mu}}\right] = A^{1-2\frac{\tilde{\mu}}{\mu}}\tilde{P}_1(\delta = 0) + B^{1-2\frac{\tilde{\mu}}{\mu}}\tilde{P}_1(\delta = 1).$$

Combing the above with (4.10) we have

$$AP_0(\delta = 0) + BP_0(\delta = 1) = 1, \tag{4.11}$$
$$A^{1-2\frac{\tilde{\mu}}{\mu}}\tilde{P}_1(\delta = 0) + B^{1-2\frac{\tilde{\mu}}{\mu}}\tilde{P}_1(\delta = 1) = 1. \tag{4.12}$$

Note also that

$$P_0(\delta = 0) + P_0(\delta = 1) = 1, \tag{4.13}$$
$$\tilde{P}_1(\delta = 0) + \tilde{P}_1(\delta = 1) = 1. \tag{4.14}$$

Using equations (4.11)-(4.14) to solve for $P_0(\delta = 1)$ and $\tilde{P}_1(\delta = 0)$ gives us the desired result. □

Proposition 13. *Suppose given the prior P_π a sequential detection procedure SPRT(A, B) is used to test the hypothesis H_0 versus H_1 as in (4.1) and (4.2), respectively. Given the process $\tilde{Z}_t$, the expected stopping times under the prior $\tilde{P}_\pi$ and hypotheses H_0 and $\tilde{H}_1$, respectively, are*

$$E_0[T] = -\frac{2}{\mu^2}\frac{(B-1)\log A + (1-A)\log B}{B-A},$$
$$\tilde{E}_1[T] = \frac{(B^{1-2\frac{\tilde{\mu}}{\mu}} - 1)\log A + (1 - A^{1-2\frac{\tilde{\mu}}{\mu}})\log B}{\mu(\tilde{\mu} - \frac{1}{2}\mu)(B^{1-2\frac{\tilde{\mu}}{\mu}} - A^{1-2\frac{\tilde{\mu}}{\mu}})}.$$

Proof. Given Proposition 12, the proof is analogous to the classical SPRT case. See, for example, [58]. □

4.3.3 Attacker Expected Utility and Best Response

In order to determine which alternate drift $\tilde{\mu}$ the attacker will choose in response to a given defender strategy SPRT$(A(\mu), B(\mu))$, we will need to determine an expected utility function for the attacker. Under H_0 the system is not compromised, in which case the attacker gets no utility. Under $\tilde{H}_1$ there are two possibilities. Either $\tilde{\Lambda}_T = A$, in which case the defender decides (incorrectly) that the system is not compromised, or $\tilde{\Lambda}_T = B$, in which case the defender decides (correctly) that the system is compromised. Let T_c be the amount of time the attacker controls the compromised system. Under H_0 we have $T_c = 0$, while under $\tilde{H}_1$ we have

$$T_c = \begin{cases} T & \text{if } \tilde{\Lambda}_T = B, \\ +\infty & \text{if } \tilde{\Lambda}_T = A. \end{cases}$$

It follows that $\tilde{E}_1[T_c] = +\infty$. For this reason we introduce a discount function e^{-rt} over an infinite time horizon. Let $f(\tilde{\mu})$ be the instantaneous utility gained by the attacker given the strategy $\tilde{\mu}$ and define the attacker's expected utility as

$$U_r(\tilde{\mu}, \mu, A, B) \triangleq \tilde{E}_1 \left[\int_0^{T_c} f(\tilde{\mu}) e^{-rt} dt \right].$$

For any realized time t we can write the utility of the attacker as $\frac{f(\tilde{\mu})}{r}(1 - e^{-rt})$. Thus we write the expected utility as

$$\begin{aligned} U_r(\tilde{\mu}, \mu, A, B) &= \tilde{E}_1 \left[\frac{f(\tilde{\mu})}{r} \left(1 - e^{-rT_c} \right) \right] \\ &= \frac{f(\tilde{\mu})}{r} \left(1 - \tilde{E}_1 \left[e^{-rT_c} \right] \right). \end{aligned}$$

In order to determine $\tilde{E}_1 \left[e^{-rT_c} \right]$, we define the following stopping times and some corresponding lemmas:

$$\begin{aligned} \tau_c^x &\triangleq \inf\{t \geq 0 : W_t = c - xt\}, \\ \tau_{a,b}^x &\triangleq \inf\{t \geq 0 : W_t \notin (a - xt, b - xt)\}. \end{aligned}$$

It is assumed throughout that $W_0 = 0$ with probability one.

Lemma 10. *For constants $x, c \in \mathbb{R}$ and $r > 0$*

$$\tilde{E}_1[e^{-r\tau_c^x}] = \begin{cases} e^{c(x + \sqrt{x^2 + 2r})} & \text{if } c \leq 0, \\ e^{-c(-x + \sqrt{x^2 + 2r})} & \text{if } c > 0. \end{cases}$$

Proof. (See [30] Chapter 7, Exercise 5.4 for a similar problem.) We use the fact that $\exp\left(\xi W_t - \frac{1}{2}\xi^2 t\right)$ is a martingale for any $\xi \in \mathbb{R}$. By the optional stopping

theorem we have

$$\tilde{E}_1\left[\exp\left(\xi W_{\tau_c^x} - \frac{1}{2}\xi^2\tau_c^x\right)\right] = \tilde{E}_1\left[\exp\left(\xi W_0 - \frac{1}{2}\xi^2 0\right)\right] = 1.$$

Furthermore $W_{\tau_c^x} = c - x\tau_c^x$ so

$$\begin{aligned}\tilde{E}_1\left[\exp\left(\xi W_{\tau_c^x} - \frac{1}{2}\xi^2\tau_c^x\right)\right] &= \tilde{E}_1\left[\exp\left(\xi(c - x\tau_c^x) - \frac{1}{2}\xi^2\tau_c^x\right)\right] \\ &= \tilde{E}_1\left[\exp\left(\xi c - (x\xi + \frac{1}{2}\xi^2)\tau_c^x\right)\right].\end{aligned}$$

Combing the results gives us

$$\tilde{E}_1\left[\exp\left(-(x\xi + \frac{1}{2}\xi^2)\tau_c^x\right)\right] = \exp\left(-\xi c\right).$$

Setting $r = x\xi + \frac{1}{2}\xi^2$ and solving for ξ gives $\xi = -x \pm \sqrt{x^2 + 2r}$. If $c > 0$ choose $\xi = -x + \sqrt{x^2 + 2r}$ and if $c < 0$ choose $\xi = -x - \sqrt{x^2 + 2r}$. The result follows.

□

Lemma 11. *If $a < 0 < b$, then*

$$\begin{aligned}\tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_a^x < \tau_b^x]\tilde{P}_1(\tau_a^x < \tau_b^x) &= \frac{e^{a(x+\sqrt{x^2+2r})}\left(1 - e^{-2b\sqrt{x^2+2r}}\right)}{1 - e^{-2(b-a)\sqrt{x^2+2r}}}, \\ \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_b^x < \tau_a^x]\tilde{P}_1(\tau_b^x < \tau_a^x) &= \frac{e^{b(x-\sqrt{x^2+2r})}\left(1 - e^{2a\sqrt{x^2+2r}}\right)}{1 - e^{-2(b-a)\sqrt{x^2+2r}}}.\end{aligned}$$

Proof. By the strong Markov property we have

$$\begin{aligned}
\tilde{E}_1\left[e^{-r\tau_a^x}\right] &= \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_a^x < \tau_b^x]\tilde{P}_1(\tau_a^x < \tau_b^x) \\
&\quad + \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_b^x < \tau_a^x]\tilde{P}_1(\tau_b^x < \tau_a^x)\tilde{E}_1[e^{-r\tau_{a-b}^x}], \\
\tilde{E}_1\left[e^{-r\tau_b^x}\right] &= \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_b^x < \tau_a^x]\tilde{P}_1(\tau_b^x < \tau_a^x) \\
&\quad + \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_a^x < \tau_b^x]\tilde{P}_1(\tau_a^x < \tau_b^x)\tilde{E}_1[e^{-r\tau_{b-a}^x}].
\end{aligned}$$

By Lemma 10 this gives

$$\begin{aligned}
e^{a(x+\sqrt{x^2+2r})} &= \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_a^x < \tau_b^x]\tilde{P}_1(\tau_a^x < \tau_b^x) \\
&\quad + \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_b^x < \tau_a^x]\tilde{P}_1(\tau_b^x < \tau_a^x)e^{(a-b)(x+\sqrt{x^2+2r})}, \\
e^{-b(-x+\sqrt{x^2+2r})} &= \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_b^x < \tau_a^x]\tilde{P}_1(\tau_b^x < \tau_a^x) \\
&\quad + \tilde{E}_1[e^{-r\tau_{a,b}^x}|\tau_a^x < \tau_b^x]\tilde{P}_1(\tau_a^x < \tau_b^x)e^{(a-b)(-x+\sqrt{x^2+2r})}.
\end{aligned}$$

Solving for the unknown terms gives us the desired result. □

Given the above lemmas we can obtain the attacker's discounted expected utility in closed form. Note that the following result is valid for any $\mu > 0, A \in [0,1], B \geq 1$, but in the equilibrium analysis to come later we will assume that $A(\mu)$ and $B(\mu)$ are the optimal SPRT choices associated with μ.

Proposition 14. *The attacker's discounted expected utility is*

$$U_r(\tilde{\mu}, \mu, A, B) = \frac{f(\tilde{\mu})}{r}\left[1 - \frac{B^{\frac{1}{\mu}(x+y)}\left(1 - A^{\frac{2}{\mu}y}\right)}{B^{\frac{2}{\mu}y} - A^{\frac{2}{\mu}y}}\right],$$

where $x = \tilde{\mu} - \frac{1}{2}\mu$ *and* $y = \sqrt{(\tilde{\mu} - \frac{1}{2}\mu)^2 + 2r}$.

Proof. By definition we have

$$U_r(\tilde{\mu}, \mu, A, B) = \frac{1}{r} f(\tilde{\mu})(1 - \tilde{E}_1\left[e^{-rT_c}\right]),$$

thus it is enough to determine $\tilde{E}_1\left[e^{-rT_c}\right]$. Because $T_c = +\infty$ whenever $T_A < T_B$ we have $\tilde{E}_1\left[e^{-rT_c}|T_A < T_B\right] = 0$ and we can write

$$\tilde{E}_1\left[e^{-rT_c}\right] = \tilde{E}_1\left[e^{-rT_B}|T_B < T_A\right]\tilde{P}_1(T_B < T_A).$$

For fixed μ, A, B using the change of variables $x = \tilde{\mu} - \frac{1}{2}\mu$ and setting $a = \frac{1}{\mu}\log A$ and $b = \frac{1}{\mu}\log B$, we see that $T_B \stackrel{d}{=} \tau_b^x$ and $T_A \stackrel{d}{=} \tau_a^x$. Thus

$$\tilde{E}_1\left[e^{-rT_c}\right] = \tilde{E}_1\left[e^{-r\tau_b^x}|\tau_b^x < \tau_a^x\right]P_1(\tau_b^x < \tau_a^x).$$

Using lemma 11 this gives us

$$\tilde{E}_1\left[e^{-rT_c}\right] = \frac{e^{b\left(x-\sqrt{x^2+2r}\right)}\left(1 - e^{2a\sqrt{x^2+2r}}\right)}{1 - e^{2(a-b)\sqrt{x^2+2r}}},$$

or equivalently

$$\tilde{E}_1\left[e^{-rT_c}\right] = \frac{B^{\frac{1}{\mu}\left(x+\sqrt{x^2+2r}\right)}\left(1 - A^{\frac{2}{\mu}\sqrt{x^2+2r}}\right)}{B^{\frac{2}{\mu}\sqrt{x^2+2r}} - A^{\frac{2}{\mu}\sqrt{x^2+2r}}}.$$

□

4.3.4 Approximate Utility Function

We assume that the intruder is interested in establishing long-term, unrestricted access to the defender's system. As such we will be interested in the case

that the discount factor r is small. At this time a general analysis of the utility function U_r for arbitrary $r > 0$ is unavailable due to its intractability. However if we consider the limit $r \downarrow 0$, then we obtain a tractable utility function which approximates the case $r \approx 0$.

Define the attacker's asymptotic expected utility function and asymptotic best response correspondence, respectively, as follows:

$$U_0(\tilde{\mu}, \mu, A, B) \triangleq \lim_{r \downarrow 0} r U_r(\tilde{\mu}, \mu, A, B),$$
$$\sigma_0(\mu, A, B) \triangleq \arg\max_{\tilde{\mu} \geq 0} U_0(\tilde{\mu}, \mu, A, B).$$

Proposition 15. *For $\tilde{\mu} \geq 0, \mu > 0, A \in [0,1]$ and $B \geq 1$ with $A < B$ we have*

$$U_0(\tilde{\mu}, \mu, A, B) = f(\tilde{\mu}) \tilde{P}_1(\delta = 0).$$

Proof. We again use the notation $x = \tilde{\mu} - \frac{1}{2}\mu$ and $y = \sqrt{x^2 + 2r}$ as in Proposition 14. We consider three cases. First assume $\tilde{\mu} > \frac{1}{2}\mu$, i.e. $x > 0$. Then

$$\begin{aligned}
\lim_{r \downarrow 0} \frac{B^{\frac{1}{\mu}(x+y)} \left(1 - A^{\frac{2}{\mu}y}\right)}{B^{\frac{2}{\mu}y} - A^{\frac{2}{\mu}y}} &= \frac{B^{\frac{1}{\mu}2x} \left(1 - A^{\frac{2}{\mu}x}\right)}{B^{\frac{2}{\mu}x} - A^{\frac{2}{\mu}x}} \\
&= \frac{B^{2\frac{\tilde{\mu}}{\mu} - 1} \left(1 - A^{2\frac{\tilde{\mu}}{\mu} - 1}\right)}{B^{2\frac{\tilde{\mu}}{\mu} - 1} - A^{2\frac{\tilde{\mu}}{\mu} - 1}} \\
&= \frac{1 - A^{1 - 2\frac{\tilde{\mu}}{\mu}}}{B^{1 - 2\frac{\tilde{\mu}}{\mu}} - A^{1 - 2\frac{\tilde{\mu}}{\mu}}}.
\end{aligned}$$

Now assume $\tilde{\mu} < \frac{1}{2}\mu$, i.e. $x < 0$. In this case $\lim_{r \downarrow 0} x + y = 0$ and $\lim_{r \downarrow 0} y = |x|$.

Thus we have

$$\lim_{r\downarrow 0} \frac{B^{\frac{1}{\mu}(x+y)}\left(1-A^{\frac{2}{\mu}y}\right)}{B^{\frac{2}{\mu}y}-A^{\frac{2}{\mu}y}} = \frac{1-A^{\frac{2}{\mu}|x|}}{B^{\frac{2}{\mu}|x|}-A^{\frac{2}{\mu}|x|}} = \frac{1-A^{1-2\frac{\tilde{\mu}}{\mu}}}{B^{1-2\frac{\tilde{\mu}}{\mu}}-A^{1-2\frac{\tilde{\mu}}{\mu}}}.$$

Finally we consider the case $\tilde{\mu} = \frac{1}{2}\mu$, where we have $x = 0$ and $y = \sqrt{2r}$:

$$\frac{B^{\frac{1}{\mu}(x+y)}\left(1-A^{\frac{2}{\mu}y}\right)}{B^{\frac{2}{\mu}y}-A^{\frac{2}{\mu}y}} = \frac{B^{\frac{1}{\mu}\sqrt{2r}}\left(1-A^{\frac{2}{\mu}\sqrt{2r}}\right)}{B^{\frac{2}{\mu}\sqrt{2r}}-A^{\frac{2}{\mu}\sqrt{2r}}}.$$

Taking limits we have

$$\lim_{r\downarrow 0} \frac{B^{\frac{1}{\mu}\sqrt{2r}}\left(1-A^{\frac{2}{\mu}\sqrt{2r}}\right)}{B^{\frac{2}{\mu}\sqrt{2r}}-A^{\frac{2}{\mu}\sqrt{2r}}} = -\frac{\log A}{\log B - \log A} = \lim_{\tilde{\mu}\to\frac{1}{2}\mu} \tilde{P}_1(\delta = 0).$$

From this analysis we see that $\lim_{r\downarrow 0} rU(\tilde{\mu}, \mu, r) = f(\tilde{\mu})\tilde{P}_1(\delta = 0)$. □

In order to prove the existence of pure Nash equilibria we will need the strict quasi-concavity of the attacker's expected utility function U_0 in his strategy $\tilde{\mu}$. The required analysis will be easier if the expected utility is written in terms of the natural exponential function, e^x. We thus introduce the following change of variables: $x = \tilde{\mu} - \frac{1}{2}\mu, a = \frac{1}{\mu}\log A$ and $b = \frac{1}{\mu}\log B$. We can then rewrite the probability of a false negative as $\tilde{P}_1(\delta = 0) = 1 - \frac{e^{2bx}\left(1-e^{2ax}\right)}{e^{2bx}-e^{2ax}}$. With this observation we define the function

$$g(x, a, b) \triangleq \frac{e^{2bx}\left(1-e^{2ax}\right)}{e^{2bx}-e^{2ax}}.$$

The following lemma regarding the function g will help us establish the strict quasi-concavity of the attacker's expected utility U_0 in his strategy $\tilde{\mu}$.

Lemma 12. *For $a < 0 < b$, $\frac{\partial}{\partial x}\left[\frac{1-g(x,a,b)}{\frac{\partial g}{\partial x}}\right] \leq 0$ for all $x \in \mathbb{R}$.*

Proof. Since

$$\frac{\partial g}{\partial x} = \frac{-2ae^{2ax}(1-e^{-2bx}) - 2be^{2(a-b)x}(1-e^{2ax})}{(1-e^{2(a-b)x})^2},$$

we have

$$\frac{1-g(x)}{\frac{\partial g}{\partial x}} = \frac{e^{2bx} - 1 - e^{2ax} + e^{2(a-b)x}}{-2ae^{2bx} + 2a - 2b + 2be^{2ax}}.$$

Define the functions

$$\begin{aligned} h(x) &\triangleq e^{2bx} - 1 - e^{2ax} + e^{2(a-b)x}, \\ \ell(x) &\triangleq -2ae^{2bx} + 2a - 2b + 2be^{2ax}, \\ F(x) &\triangleq \frac{h(x)}{\ell(x)}. \end{aligned}$$

Clearly $F(x) = \frac{1-g(x)}{\frac{\partial g}{\partial x}}$. We will show that $F'(x) < 0$ for all x. First note that

$$\begin{aligned} F'(x) &= \frac{h'(x)\ell(x) - h(x)\ell'(x)}{\ell(x)^2}, \\ h'(x) &= 2be^{2bx} - 2ae^{2ax} + 2(a-b)e^{2(a-b)x}, \\ \ell'(x) &= -4ab(e^{2bx} - e^{2ax}). \end{aligned}$$

One then arrives at

$$\frac{1}{8}F'(x)\ell(x)^2 e^{-2ax} = (b-a)^2\cosh(2bx) - b^2\cosh(2(b-a)x) + 2a(2b-a).$$

Define the right hand side as

$$R(x) \triangleq (b-a)^2 \cosh(2bx) - b^2 \cosh(2(b-a)x) + 2a(2b-a).$$

We then have

$$R''(x) = 4(b-a)^2 b^2 (\cosh(2bx) - \cosh(2(b-a)x)) \leq 0.$$

Furthermore $R(0) = 0$ and $R'(x) = 0$ if and only if $x = 0$. Since $R(x)$ is concave, $R(0) = 0$ is a global maximum giving us $R(x) \leq 0$ for all x. It follows that $F'(x) \leq 0$, establishing our result. □

Proposition 16. *Fix $\mu > 0, A \in [0,1]$ and $B \geq 1$. If $A < 1 < B$, then the intruder expected utility $U_0(\tilde{\mu}, \mu, A, B)$ is strictly quasi-concave in $\tilde{\mu}$. If $A = 1 < B$ then $U_0(\tilde{\mu}, \mu, A, B) = f(\tilde{\mu})$. If $A < 1 = B$ then $U_0(\tilde{\mu}, \mu, A, B) \equiv 0$.*

Proof. The first order condition (FOC) of optimality, $\frac{\partial U_0}{\partial \tilde{\mu}} = 0$, can be written $\frac{\partial f}{\partial \tilde{\mu}} \left(1 - g\left(\tilde{\mu} - \frac{1}{2}\mu, \frac{1}{\mu}\log A, \frac{1}{\mu}\log B\right)\right) = f(\tilde{\mu}) \frac{\partial g}{\partial \tilde{\mu}}$. Fixing $\mu > 0, A \in [0,1]$ and $B \geq 1$ and abusing notation we write $g(\tilde{\mu}) = g\left(\tilde{\mu} - \frac{1}{2}\mu, \frac{1}{\mu}\log A, \frac{1}{\mu}\log B\right)$.

First consider the case that $A < 1 < B$. Since $\frac{\partial x}{\partial \tilde{\mu}} = 1$ we have $\frac{\partial g}{\partial \tilde{\mu}} \equiv \frac{\partial g}{\partial x}$. Then by definition of the function g we have $\frac{\partial g}{\partial \tilde{\mu}} > 0$, while by assumption $\frac{\partial f}{\partial \tilde{\mu}} > 0$. The first order condition for optimality is then

$$\frac{1 - g(\tilde{\mu})}{\frac{\partial g}{\partial \tilde{\mu}}} = \frac{f(\tilde{\mu})}{\frac{\partial f}{\partial \tilde{\mu}}}. \tag{4.15}$$

By assumption the function $f(\tilde{\mu})$ is concave, thus giving us $\frac{\partial}{\partial \tilde{\mu}}\left[\frac{f(\tilde{\mu})}{\frac{\partial f}{\partial \tilde{\mu}}}\right] > 0$. By Lemma 12 we have $\frac{\partial}{\partial \tilde{\mu}}\left[\frac{1-g}{\frac{\partial g}{\partial \tilde{\mu}}}\right] \leq 0$. It follows that there is at most one solution satisfying the FOC.

Furthermore we see that $\lim_{\tilde{\mu}\downarrow 0} f(\tilde{\mu}) = 0$ and

$$\begin{aligned}\lim_{\tilde{\mu}\downarrow 0} \frac{1-g(\tilde{\mu})}{\frac{\partial g}{\partial \tilde{\mu}}} &= \lim_{x\downarrow -\mu/2} \frac{e^{2bx} - 1 - e^{2ax} + e^{2(a-b)x}}{-2ae^{2bx} + 2a - 2b + 2be^{2ax}} \\ &= \frac{\mu}{2} \frac{(B-A)(B-1)}{A\log A(B-1) - B\log B(A-1)} > 0.\end{aligned}$$

By monotonicity and continuity either $\frac{1-g(\tilde{\mu})}{\frac{\partial g}{\partial \tilde{\mu}}} > \frac{f(\tilde{\mu})}{\frac{\partial f}{\partial \tilde{\mu}}}$ for all $\tilde{\mu} \geq 0$ or there exists a unique $\tilde{\mu}$ satisfying (4.15). Furthermore

$$\lim_{\tilde{\mu}\to\infty} \frac{1-g(\tilde{\mu})}{\frac{\partial g}{\partial \tilde{\mu}}} = \lim_{x\to\infty} \frac{1-g(x,a,b)}{\frac{\partial g}{\partial x}} = -\frac{1}{2a} < \infty.$$

Since $f(\tilde{\mu}) \geq 0, f'(\tilde{\mu}) > 0$ and $f''(\tilde{\mu}) \leq 0$ we must have $\lim_{\tilde{\mu}\to\infty} \frac{f(\tilde{\mu})}{\frac{\partial f}{\partial \tilde{\mu}}} = \infty$. Thus we are guaranteed that there exists a unique $\tilde{\mu}$ satisfying (4.15), and the intruder expected utility $U_0(\tilde{\mu}, \mu, A, B)$ is strictly quasi-concave in $\tilde{\mu}$.

If $A = 1 < B$ then $\tilde{P}_1(\delta = 0) = 1$ for all $\tilde{\mu} \geq 0$ giving us $U_0(\tilde{\mu}, \mu, A, B) = f(\tilde{\mu})$. Finally if $A < 1 = B$ then $\tilde{P}_1(\delta = 0) = 0$ for all $\tilde{\mu} \geq 0$, giving us $U_0(\tilde{\mu}, \mu, A, B) \equiv 0$. □

In what follows we will assume that for a given $\mu > 0$ the defender plays the associated SPRT$(A(\mu), B(\mu))$. Because the values of $A(\mu)$ and $B(\mu)$ are uniquely determined by the optimality of the SPRT, we can consider the attacker's best response strategy to be a best response to the single variable μ. Abusing notation we will write the attacker's best response correspondence as $\sigma_0(\mu)$ where it is understood that $\mu \mapsto \text{SPRT}(A(\mu), B(\mu))$.

Proposition 17. *Fix $\mu > 0$ and assume the defender plays the SPRT$(A(\mu), B(\mu))$. If $A(\mu) < 1 < B(\mu)$ then the best response correspondence $\sigma_0(\mu)$ is single valued and differentiable at μ. If $A(\mu) = 1 < B(\mu)$ then $\sigma_0(\mu) = +\infty$. If $A(\mu) < 1 = B(\mu)$ then $\sigma_0(\mu) \equiv \mathbb{R}^+$.*

Proof. Define the function

$$M(u,v) \triangleq \frac{\frac{1-B(v)^{1-2\frac{u}{v}}}{A(v)^{1-2\frac{u}{v}}-B(v)^{1-2\frac{u}{v}}}}{\frac{\partial}{\partial\tilde{\mu}}\left[\frac{1-B(v)^{1-2\frac{\tilde{\mu}}{v}}}{A(v)^{1-2\frac{\tilde{\mu}}{v}}-B(v)^{1-2\frac{\tilde{\mu}}{v}}}\right]_{\tilde{\mu}=u}} + \frac{f(u)}{\frac{\partial f}{\partial\tilde{\mu}}\Big|_{\tilde{\mu}=u}}.$$

Then the attacker's first order condition of optimality can be written as $M(\tilde{\mu},\mu) = 0$. Given (4.5) and (4.6) the implicit function theorem implies that $\pi_u(\mu)$ and $\pi_\ell(\mu)$ are differentiable functions of μ. Then by (4.7) and (4.8), $A(\mu)$ and $B(\mu)$ are differentiable functions in μ in open neighborhoods where $A(\mu) < 1 < B(\mu)$. Furthermore $f(\tilde{\mu})$ is twice differentiable in $\tilde{\mu}$ by assumption. It follows that the function $M(u,v)$ is differentiable in both u and v.

Suppose for some $\mu_0 > 0$ the associated SPRT$(A(\mu_0), B(\mu_0))$ satisfies $A(\mu_0) < 1 < B(\mu_0)$. Furthermore suppose there exists a value $\tilde{\mu}_0$ such that the first order condition (4.15) is satisfied, i.e. $M(\tilde{\mu}_0, \mu_0) = 0$. By the implicit function theorem there exists a differentiable function $y : \mathbb{R} \to \mathbb{R}$ such that $y(\mu_0) = \tilde{\mu}_0$ and $M(y(\mu),\mu) = 0$ for all μ in some open neighborhood of μ_0 . By definition $\sigma_0(\mu)$ is the set of best responses to the strategy SPRT$(A(\mu), B(\mu))$. Since the pairs $(y(\mu),\mu)$ solve the FOC the strict quasi-concavity of U_0 implies $\sigma_0(\mu) = y(\mu)$ for all μ in some open neighborhood of μ_0. Thus $\sigma_0(\mu)$ is differentiable at μ_0.

From Proposition 16 we see that if $A = 1 < B$ then $U_0(\tilde{\mu},\mu,A,B) = f(\tilde{\mu})$, which is monotonically increasing in $\tilde{\mu}$. Thus $\sigma_0(\mu) = +\infty$. Again from Proposition 16 we see that $A < 1 = B$ implies $U_0(\tilde{\mu},\mu,A,B) \equiv 0$, which means all strategies are equally valid, i.e. $\sigma_0(\mu) \equiv \mathbb{R}^+$. □

4.4 Equilibrium Analysis

We now establish the existence of Nash equilibria in the adversarial sequential detection game in the limiting case $r \downarrow 0$. More specifically we will show the existence of a value μ^* such that

$$U_0(\mu, \mu^*, A(\mu^*), B(\mu^*)) \leq U_0(\mu^*, \mu^*, A(\mu^*), B(\mu^*))$$
$$\text{for all } \mu \geq 0.$$

Note that this implies the attacker has no incentive to unilaterally deviate from the strategy μ^*. Furthermore since the defender is playing a best response to the strategy μ^*, namely the SPRT$(A(\mu^*), B(\mu^*))$, he too has no incentive to unilaterally deviate from his strategy. As such the strategy profile $(\mu^*, \mu^*, A(\mu^*), B(\mu^*))$ is a Nash equilibrium. In our analysis we have restricted attention to the case where the defender plays an SPRT. However, the existence of a value μ^* implies the existence of a Nash equilibrium in the more general case in which the defender is free to choose any sequential decision rule. Since the value μ^* is non-random we refer to this as a pure Nash equilibirum.

For the Nash equilibrium existence proof, we will need to understand the asymptotic behavior of the functions $A(\mu)$ and $B(\mu)$, which we establish in the following lemma.

Lemma 13. *Let SPRT$(A(\mu), B(\mu))$ be the sequential probability ratio tetst associated with the drift $\mu > 0$. Then $\lim_{\mu \downarrow 0} B(\mu) = \infty$, $\lim_{\mu \downarrow 0} A(\mu) = 1$ and $\lim_{\mu \to \infty} A(\mu) = 0$. Furthermore there exists a value $\mu' > 0$ such that $\lim_{\mu \downarrow \mu'} A(\mu) = 1$ and $A(\mu) < 1$ for all $\mu > \mu'$.*

Proof. Recall from Theorem 4 that $\pi_\ell(\mu)$ and $\pi_u(\mu)$ are guaranteed to satisfy $0 < \pi_\ell(\mu) < \frac{\alpha}{\alpha + \beta(\mu)} < \pi_u(\mu) < 1$. Since $\beta(\mu)$ is monotonically increasing in μ with

$\beta(0) = 0$ and $\lim_{\mu\to\infty} \beta(\mu) = \infty$ we must have $\frac{\alpha}{\alpha+\beta(\mu)}$ monotonically decreasing with $\lim_{\mu\downarrow 0} \frac{\alpha}{\alpha+\beta(\mu)} = 1$ and $\lim_{\mu\to\infty} \frac{\alpha}{\alpha+\beta(\mu)} = 0$. With these limits the bounds on $\pi_\ell(\mu)$ and $\pi_u(\mu)$ immediately give us $\lim_{\mu\downarrow 0} \pi_u(\mu) = 1$ and $\lim_{\mu\to\infty} \pi_\ell(\mu) = 0$. Hence (4.7) and (4.8) give us $\lim_{\mu\downarrow 0} B(\mu) = \infty$ and $\lim_{\mu\to\infty} A(\mu) = 0$.

From Theorem 4 we have $\frac{\partial \nu(\pi,\pi_\ell,\mu)}{\partial \pi}\Big|_{\pi=\pi_u(\mu)} = \beta(\mu) + \frac{2}{\mu^2}(\psi'(\pi_u(\mu)) - \psi'(\pi_\ell(\mu)))$. Taking $\mu \downarrow 0$ on both sides of (4.6) we arrive at

$$\lim_{\mu\downarrow 0} \frac{2}{\mu^2}(\psi'(\pi_u) - \psi'(\pi_\ell)) = -\alpha. \tag{4.16}$$

Since $\lim_{\mu\downarrow 0} \frac{1}{\mu^2} = \infty$ we must have $\lim_{\mu\downarrow 0}(\psi'(\pi_u) - \psi'(\pi_\ell)) = 0$, otherwise (4.16) would not hold. Since $\psi'(x) = -2\log\frac{x}{1-x} + \frac{1-2x}{x}$ it is easily verified that the only possibility is $\lim_{\mu\downarrow 0} \pi_\ell(\mu) = 1$, for otherwise $\lim_{\mu\downarrow 0}(\psi'(\pi_u) - \psi'(\pi_\ell)) \neq 0$.

By the continuity of $\pi_\ell(\mu)$ along with its limiting values, there must exist a value $\mu' > 0$ such that $\lim_{\mu\downarrow\mu'} \pi_\ell(\mu) = \pi$ and $\pi_\ell(\mu) < \pi$ for all $\mu > \mu'$. Thus as $\mu \downarrow \mu'$ we have $A(\mu) \uparrow 1$ and $A(\mu) < 1$ for all $\mu > \mu'$. □

For the final lemma define the functions $\phi : \mathbb{R}\times\mathbb{R} \to \mathbb{R}, L : \mathbb{R} \to \mathbb{R}$ and $R : \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ as $\phi(x,y) \triangleq 2\frac{y}{x} - 1, L(x) \triangleq \frac{f(x)}{f'(x)}$ and

$$R(x,y) \triangleq -\frac{B(y)^{-\phi(x,y)}(B(y)^{\phi(x,y)} - 1)(B(y)^{\phi(x,y)} - A(y)^{\phi(x,y)})}{(B(y)^{\phi(x,y)} - 1)\log A(y) + (1 - A(y)^{\phi(x,y)})\log B(y)}.$$

The attacker's FOC can then be written as

$$L(\tilde{\mu}) = R(\tilde{\mu}, \mu). \tag{4.17}$$

The next lemma establishes the limiting values of these functions which will be needed in the proof of the main existence theorem.

Lemma 14. *For $\mu > 0$ such that $A(\mu) < 1 < B(\mu)$ we have the following:*

$$\begin{aligned}
\lim_{\tilde{\mu}\downarrow 0} L(\tilde{\mu}) &= 0, \\
\lim_{\tilde{\mu}\to+\infty} L(\tilde{\mu}) &= +\infty, \\
\lim_{\tilde{\mu}\to\frac{1}{2}\mu} R(\tilde{\mu},\mu) &= -\frac{\mu}{\log A(\mu)}, \\
\lim_{\tilde{\mu}\to\infty} R(\tilde{\mu},\mu) &= -\frac{1}{2}\frac{\mu}{\log A(\mu)}.
\end{aligned}$$

Proof. The limiting values of L are obtained directly from the assumptions on the function f, namely that $f(x) \geq 0, f'(x) > 0$ and $f''(x) \leq 0$. The limiting values of R are obtained via L'Hôpital's rule. □

Theorem 5. *There exists a pure Nash equilibrium in the non-zero-sum, sequential detection game in the limiting case $r \downarrow 0$.*

Proof. Suppose the defender is testing for the strategy μ and has chosen the associated optimal sequential test $\mathrm{SPRT}(A(\mu), B(\mu))$. Assuming μ is fixed with $A(\mu) < 1 < B(\mu)$, the monotonicity properties of L and R in $\tilde{\mu}$ give us bounds on the unique value $\tilde{\mu}$ witch satisfies (4.17). Define the values $\tilde{\mu}_{\frac{1}{2}}, \tilde{\mu}_\ell, \tilde{\mu}^*$ as the unique solutions to the following equations:

$$\begin{aligned}
L(\tilde{\mu}_{\frac{1}{2}}) &= \lim_{\tilde{\mu}\to\frac{1}{2}\mu} R(\tilde{\mu},\mu) = -\frac{\mu}{\log A(\mu)}, \\
L(\tilde{\mu}_\ell) &= \lim_{\tilde{\mu}\to\infty} R(\tilde{\mu},\mu) = -\frac{1}{2}\frac{\mu}{\log A(\mu)}, \\
L(\tilde{\mu}^*) &= R(\tilde{\mu}^*,\mu).
\end{aligned}$$

Given the assumptions on the function f, namely $f(x) \geq 0, f'(x) > 0, f''(x) \leq 0$, we must have $L'(x) > 0$. If $y = L(x)$ then there exists an inverse function L^{-1} such that $L^{-1}(y) = x$. As such we have $\tilde{\mu}_{\frac{1}{2}} = L^{-1}\left(-\frac{\mu}{\log A(\mu)}\right)$ and $\tilde{\mu}_\ell = L^{-1}\left(-\frac{1}{2}\frac{\mu}{\log A(\mu)}\right)$.

By the monotonicity of L and R in $\tilde{\mu}$ we have $\tilde{\mu}_\ell < \tilde{\mu}^*$, i.e. $L^{-1}\left(-\frac{1}{2}\frac{\mu}{\log A(\mu)}\right) < \tilde{\mu}^*$. For any μ satisfying $A(\mu) < 1 < B(\mu)$ the value $\tilde{\mu}^*$ is the unique best response of the attacker, i.e. $\sigma_0(\mu) = \tilde{\mu}^*$. Thus we have the lower bound $L^{-1}\left(-\frac{1}{2}\frac{\mu}{\log A(\mu)}\right) < \sigma_0(\mu)$.

We first consider what happens as $\mu \downarrow 0$. By Lemma 13 there exists a value $\mu' > 0$ such that as $\lim_{\mu \downarrow \mu'} A(\mu) = 1$ and $A(\mu) < 1$ for all $\mu > \mu'$. Thus we have $\lim_{\mu \downarrow \mu'} L^{-1}\left(-\frac{1}{2}\frac{\mu}{\log A(\mu)}\right) = +\infty$. Since this is a lower bound on $\sigma_0(\mu)$ it follows that $\lim_{\mu \downarrow \mu'} \sigma_0(\mu) = +\infty$.

We now consider what happens as $\mu \to \infty$. First observe that as $\mu \to \infty$ we have $\frac{\alpha}{\alpha+\beta(\mu)} \downarrow 0$. Thus we must have $\pi_\ell \downarrow 0$ as well as $A \downarrow 0$. Thus for large enough μ we must have $A(\mu) < 1$. Suppose at some point μ^* we have $B(\mu^*) = 1$. Then $\sigma_0(\mu^*) \equiv \mathbb{R}^+$ and we have $\mu^* \in \sigma_0(\mu^*)$, i.e. μ^* is a pure Nash equilibrium.

Suppose there are no values satisfying $B(\mu) = 1$. We consider two cases. First suppose $\sigma_0(\mu) \leq \frac{1}{2}\mu$ as $\mu \to \infty$. As we already showed $\lim_{\mu \downarrow \mu'} \sigma_0(\mu) = +\infty$ and $\sigma_0(\mu) < +\infty$ for all $\mu > \mu'$. We thus have that $\sigma_0(\mu) - \mu > 0$ for small enough μ while $\sigma_0(\mu) - \mu < 0$ for large enough μ. By the continuity of $\sigma_0(\mu)$ there must exist a value μ^* satisfying $\sigma_0(\mu^*) - \mu^* = 0$, i.e. there exists a Nash equilibrium.

Now suppose $\sigma_0(\mu) > \frac{1}{2}\mu$ as $\mu \to \infty$. For all μ satisfying this condition we must have $L(\frac{1}{2}\mu) < R(\frac{1}{2}\mu, \mu) = -\frac{\mu}{\log A(\mu)}$. Otherwise the monotonicity properties of L and R in $\tilde{\mu}$ would violate our assumption that $\sigma_0(\mu) > \frac{1}{2}\mu$. As such we can obtain a tighter upper bound on $\sigma_0(\mu)$. Specifically we must have $\sigma_0(\mu) < L^{-1}\left(-\frac{\mu}{\log A(\mu)}\right)$. Suppose $\limsup_{\mu\to\infty} -\frac{\mu}{\log A(\mu)} < \infty$. Then for large enough μ we will have $\sigma_0(\mu) < \mu$. Again by the continuity of $\sigma_0(\mu)$ there must exist a value μ^* satisfying $\sigma_0(\mu^*) - \mu^* = 0$, i.e. there exists a Nash equilibrium. Now suppose $\limsup_{\mu\to\infty} -\frac{\mu}{\log A(\mu)} = \infty$. Note that $L'(x) = 1 - \frac{f(x)f''(x)}{(f'(x))^2} \geq 1$ with equality if and only if $x = 0$. Thus the inverse function satisfies $\frac{d}{dy}\left[L^{-1}\right] = \frac{1}{L'(y)} \leq 1$ with

equality if and only if $y = L(0)$. As such the inverse function has a unique fixed point y_0 satisfying $y_0 = L^{-1}(y_0)$. Furthermore we must have $L^{-1}(y) < y$ for all $y > y_0$. It follows that $\sigma_0(\mu) < \mu$ as $\mu \to +\infty$. Again by continuity of $\sigma_0(\mu)$ there must exist a value μ^* satisfying $\sigma_0(\mu^*) - \mu^* = 0$, i.e. there exists a pure Nash equilibrium. This exhausts all possibilities and completes the proof. □

When facing a strategic attacker, there is no guarantee that he will choose the strategy μ the defender is testing for. The above theorem suggests that the defender can find a pure Nash equilibrium μ^* that satisfies $\mu^* \in \sigma_0(\mu^*)$. Assuming the defender is rational and strategic, playing such a strategy guarantees that the $\mathrm{SPRT}(A(\mu^*), B(\mu^*))$ is optimal for detecting the drift μ^* and insures that the attacker has no incentive to deviate from the strategy μ^*.

An alternative approach is for the defender to anticipate the best response of the attacker. The defender may then seek to find a strategy that is optimal given that the attacker will himself play a best response to the defender's strategy. In other words the defender may seek a *Stackelberg equilibrium* solution. We restrict the strategy space of the defender to be the set of all SPRTs. Thus we assume that for a given $\mu > 0$ the defender chooses the optimal $\mathrm{SPRT}(A(\mu), B(\mu))$ with respect to the measure P_π as in Theorem 4. In this way the defender's strategy is simply to choose a drift $\mu > 0$ to test against the null hypothesis of zero drift, and it is understood that $\mu \mapsto \mathrm{SPRT}(A(\mu), B(\mu))$. If the attacker plays a strategy $\tilde{\mu}$ in response, we can then define the *corrupted random cost* w.r.t. the measure $\tilde{P}_\pi$:

$$\tilde{C}(\mu, \tilde{\mu}) \triangleq T + \alpha \mathbb{1}_{\{\theta=0, \tilde{\Lambda}_T = B(\mu)\}} + \beta(\tilde{\mu}) \mathbb{1}_{\{\theta=1, \tilde{\Lambda}_T = A(\mu)\}}.$$

We then define the *corrupted value function* as

$$\tilde{V}_\pi(\mu, \tilde{\mu}) \triangleq \tilde{E}_\pi \left[\tilde{C}(\mu, \tilde{\mu}) \right].$$

Note that in the special case $\tilde{\mu} = \mu$ we have $\tilde{V}_\pi(\mu, \mu) = V_\pi(\mu)$. This follows from the optimality of the SPRT and the fact that $\tilde{P}_\pi = P_\pi$ whenever $\tilde{\mu} = \mu$. If the attacker best responds we have $\tilde{\mu} = \sigma_0(\mu)$. The *Stackelberg value* w.r.t. the measure $\tilde{P}_\pi$ is then $\tilde{V}_\pi^* \triangleq \inf_{\mu>0} \tilde{V}_\pi(\mu, \sigma_0(\mu))$. The Stackelberg equilibrium, if it exists, is any $\mu_s \in \arg\min_{\mu>0} \tilde{V}_\pi(\mu, \sigma_0(\mu))$.

In the next theorem we give sufficient conditions on $(A(\mu), B(\mu))$ to guarantee the existence of a Stackelberg equilbirium with the defender as leader under the restriction that he play only SPRTs. The first assumption states that for larger values of μ, the defender requires more evidence to accept H_0. The second assumption states that the defender will never accept H_1 without making observations, regardless of the value of μ.

Theorem 6. *For $\mu > 0$ suppose the defender restricts his sequential tests to the optimal SPRT$(A(\mu), B(\mu))$. If $A(\mu)$ is monotonically decreasing in μ and if there exists a constant $c > 1$ such that $B(\mu) > c$ for all $\mu > 0$, then there exists a Stackelberg equilibrium in the limiting case $r \downarrow 0$.*

Proof. Assume that $A(\mu)$ is monotonically decreasing in μ and that there exists a constant $c > 1$ such that $B(\mu) > 1$ for all $\mu > 0$. Let μ'' be the unique solution to $\pi_\ell(\mu'') = \pi$. By Proposition 17 we must have $\sigma_0(\mu) = +\infty$ for all $\mu \leq \mu''$ and $\sigma_0(\mu) < +\infty$ for all finite $\mu > \mu''$. Furthermore we have $\lim_{\mu \downarrow \mu''} \sigma_0(\mu) = +\infty$. We wish to show $\lim_{\mu \to +\infty} \sigma_0(\mu) = +\infty$. Since $B(\mu) > c > 1$ for all $\mu > 0$ we must $\lim_{\mu \to +\infty} B(\mu) > 1$. Thus

$$\lim_{\mu \to +\infty} P_1(\delta = 0) = \lim_{\mu \to +\infty} \frac{1 - B(\mu)}{A(\mu) - B(\mu)} = 1 - \lim_{\mu \to +\infty} \frac{1}{B(\mu)} \geq 1 - \frac{1}{c} > 0.$$

Suppose $\lim_{\mu\to+\infty}\sigma_0(\mu) = \mu^* < +\infty$. Then for any $\epsilon > 0$

$$\lim_{\mu\to+\infty}\frac{U_0(\sigma_0(\mu)+\epsilon,\mu)}{U_0(\sigma_0(\mu),\mu)} = \frac{f(\mu^*+\epsilon)}{f(\mu^*)} > 1,$$

where the equality follows from the continuity of f and the inequality follows from the monotonicity of f. But then for large enough μ we have $U_0(\sigma_0(\mu)+\epsilon,\mu) > U_0(\sigma_0(\mu)+\epsilon,\mu)$ which contradicts the optimality of $\sigma_0(\mu)$. Thus our assumption was incorrect and we can conclude that $\lim_{\mu\to+\infty}\sigma_0(\mu) = +\infty$.

We now have $\lim_{\mu\to+\infty}\sigma_0(\mu) = \lim_{\mu\downarrow\mu''}\sigma_0(\mu) = +\infty$ with $\sigma_0(\mu) < +\infty$ and differentiable for all $\mu > \mu''$ and $\sigma_0(\mu) = +\infty$ for all $\mu \leq \mu''$. Now consider the defender's expected corrupted value function $\tilde{V}_\pi(\mu,\tilde{\mu})$ evaluated at $(\mu,\sigma_0(\mu))$. Since $\tilde{V}_\pi(\mu,\sigma_0(\mu)) > 0$ and $\beta(x) \to +\infty$ as $x \to +\infty$, we must have $\lim_{\mu\downarrow\mu''}\tilde{V}_\pi(\mu,\sigma_0(\mu)) = \lim_{\mu\to+\infty}\tilde{V}_\pi(\mu,\sigma_0(\mu)) = +\infty$. The continuity of $\tilde{V}_\pi(\mu,\tilde{\mu})$ in μ follows from Propositions 12 and 13. Then the continuity of $\sigma_0(\mu)$ for $\mu > \mu''$ gives us the continuity of $\tilde{V}_\pi(\mu,\sigma_0(\mu))$ in $\mu > \mu''$. From this continuity there must exist a global minimum in the open interval $(\mu'',+\infty)$. Thus any value $\mu_s \in \arg\min_{\mu\in(\mu'',+\infty)}\tilde{V}_\pi(\mu,\sigma_0(\mu))$ is a Stackelberg equilibrium. □

4.5 Numerical Examples

Unless otherwise noted, the following values were used in the numerical results: $\alpha = 1, \beta(\mu) = \mu, f(\tilde{\mu}) = \tilde{\mu}$ and $\pi = 0.25, 0.6$. Figure 4.1 shows the attacker's best response correspondence $\sigma_0(\mu)$ for $\pi = 0.25$ and varying μ, as well as the identity line. Any point μ^* at which $\sigma_0(\mu)$ cross the identity line satisfies $\mu^* \in \sigma_0(\mu^*)$ and is thus a Nash equilibrium. In this particular example it appears that the equilibrium is unique. Figure 4.2 shows the attacker's best response correspondence $\sigma_0(\mu)$ for $\pi = 0.6$ and varying μ, as well as the identity line. Notice in this example the best

response correspondence is set valued on an interval $\mathcal{I}$. As discussed in Theorem 5 we have $\sigma_0(\mu) \equiv \mathbb{R}^+$ for all $\mu \in \mathcal{I}$, thus we are guaranteed that $\mu \in \sigma_0(\mu)$ for all $\mu \in \mathcal{I}$. The interval $\mathcal{I}$ corresponds to the set $\{\mu \geq 0 : B(\mu) = 1\}$. As such, we have an infinite number of Nash equilibria in which the defender makes no observations and immediately accepts hypothesis H_1 while the attacker gets zero utility.

To understand how the infinite Nash equilibria can arise, we numerically compute the Nash equilibria for various parameter values. In particular figure 4.3 shows Nash equilibria for varying false positive cost $\alpha \in [0, 6]$ and figure 4.4 shows Nash equilibria for varying prior $\pi \in [0, 1]$. Notice in figure 4.3 that for small false positive values we see infinite Nash equilibria. This can be explained by the following reasoning: When the defender's penalty for a false positive is sufficiently low, then the majority of his expected costs come from observation time and false negatives. As such for a large range of μ values he can minimize his expected costs by making *no* observations and making the probability of a false negative zero, i.e. immediately choose H_1. As we saw in figure 4.2 it is exactly this situation which leads to the attacker having set-valued best responses and consequently infinite Nash equilibria. As one would expect, as the cost of a false positive increases, this reasoning applies to a smaller and smaller range of μ, and for large enough false positive costs the above reasoning does not apply. We point out that for larger false positive costs, the Nash equilibria are relatively robust to increases in the cost of a false positive.

A similar line of reasoning applies to figure 4.4, but in reverse. Since π is the prior probability that the system is compromised, higher values of π lead to the probability of a false positive being lower and the probability of a false negative being higher. Thus for large values of π there is a range of μ values for which

the defender minimizes his expected cost by immediately classifying his system as compromised. An interesting observation is that a high probability of a successful intrusion by the attacker, i.e. high values of π, leads to Nash equilibria in which the attacker receives no utility. Thus what may at first appear to be an advantage for the attacker leads to a worst-case outcome for him at the Nash equilibria. It is in fact better for the attacker to have a lower probability of successful intrusion. Note that in contrast to varying false positive costs, Nash equilibria vary much more with variations in the prior π.

Using Theorem 5 we can numerically compute the values $A(\mu)$ and $B(\mu)$ in order to compare the functions $V_\pi(\mu)$ and $\tilde{V}_\pi(\mu, \sigma_0(\mu))$. Figure 4.5 shows these functions for the value $\pi = 0.25$ and varying μ. As one might expect, for values μ that are far away from the Nash equilibrium μ^* we have $V_\pi(\mu) < \tilde{V}_\pi(\mu, \sigma_0(\mu))$. In fact it appears that $\tilde{V}_\pi(\mu, \sigma_0(\mu)) - V_\pi(\mu) \to +\infty$ as $\mu \downarrow 0$ and $\mu \to +\infty$. Notice however that for some values near the Nash equilibrium the defender's outcome is *improved* by the attacker's best response. While this may appear to contradict the optimality of the SPRT test, it in fact does not. Recall that $\tilde{V}_\pi(\mu, \sigma_0(\mu))$ is computed w.r.t. the measure $\tilde{P}_\pi$ while $V_\pi(\mu)$ is computer w.r.t. the measure P_π. As such all SPRT optimality results do not apply. To guarantee optimality the defender would need to update his test to reflect the change in the distribution caused by the attacker, i.e. choose the test $\text{SPRT}(A(\sigma_0(\mu), B(\sigma_0(\mu))))$.

Figure 4.5 also appears to show the existence of a Stackelberg equilibrium. The function $\tilde{V}_\pi(\mu, \sigma_0(\mu))$ appears to have a unique minimum at a point μ_s which is distinct from μ^*, the Nash equilibrium. Thus by signaling an appropriately chosen sequential test to the attacker, the defender can improve his expected outcome relative to the Nash equilibrium.

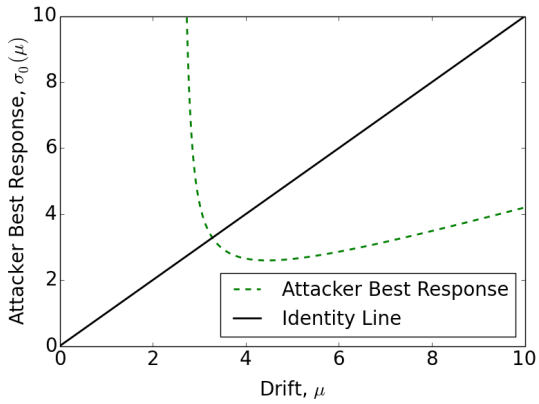

Figure 4.1: Attacker best response and the identity line for the case $\pi = 0.25$. The point at which they cross, μ^*, satisfies $\mu^* \in \sigma_0(\mu^*)$, i.e. μ^* is a Nash equilibrium.

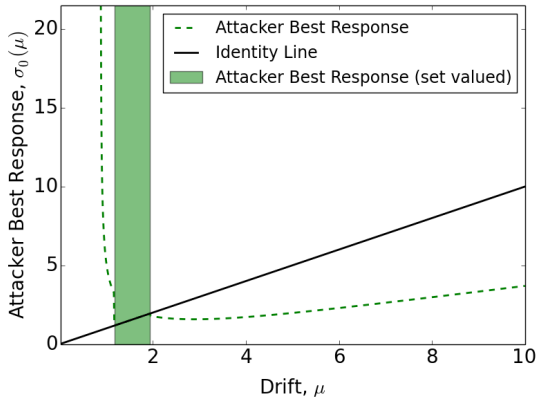

Figure 4.2: Attacker best response and the identity line for the case $\pi = 0.6$. All points at which they cross satisfy $\mu^* \in \sigma_0(\mu^*)$, i.e. μ^* is a Nash equilibrium.

4.6 Conclusion

In summary, we have presented a novel two-player, non-zero-sum, sequential detection game motivated by problems in the cyber-security domain. We proved that in the special case that the attacker's discount rate approaches zero the game admits Nash equilibria in which the defender plays Wald's sequential probability ratio test (SPRT). Furthermore we gave sufficient conditions for the existence of a Stackelberg equilibrium with the defender as leader. Through numerical

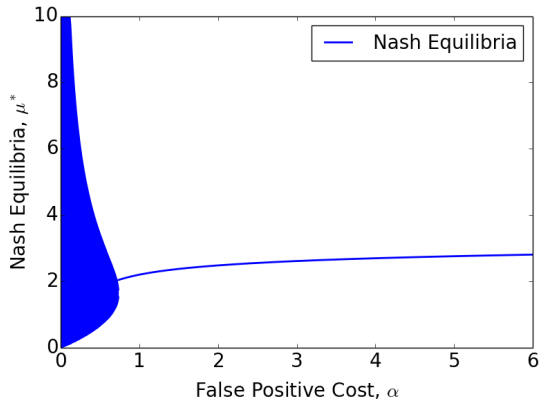

Figure 4.3: Nash equilibria for varying false positive cost, α, for the case $\pi = 0.5$.

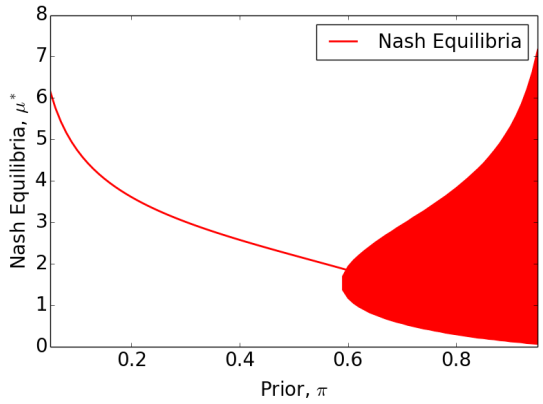

Figure 4.4: Nash equilibria for varying prior, π, for the case $\alpha = 1$.

examples we showed that it is possible for the defender to do better under the Stackelberg equilibrium than the Nash equilibrium. We also showed through numerical examples that both low false positive costs and high prior probabilities of intrusion lead to an infinite number of Nash equilibria in which the attacker receives no utility.

Several avenues for future research exist in adversarial sequential detection. First, we would like to consider the more general discount factor $r > 0$. Second, considering different payoff functions for various adversaries should give rise to qualitatively distinct equilibria. For example, in the context of network security a

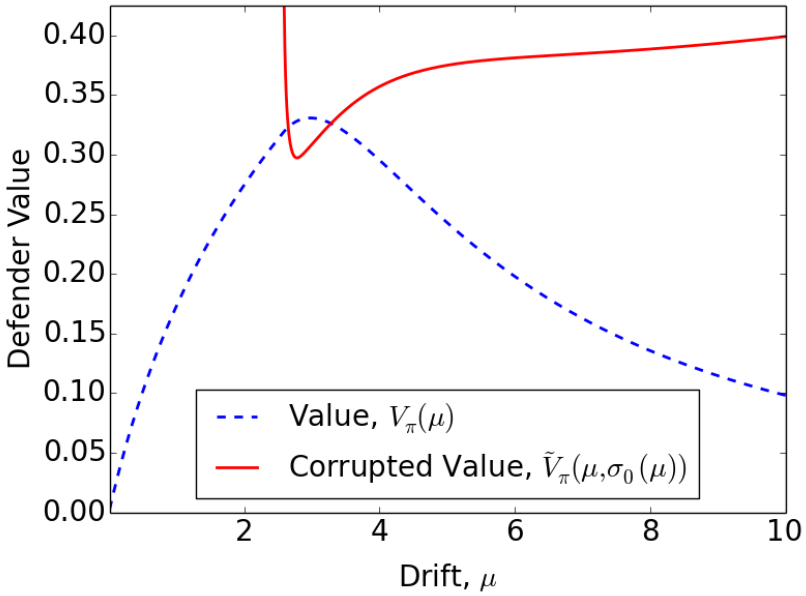

Figure 4.5: Defender's value function $V_\pi(\mu)$ w.r.t. the measure P_π and the defender's corrupted value function $\tilde{V}_\pi(\mu, \tilde{\mu})$ w.r.t. the measure $\tilde{P}_\pi$. Above we show these two functions for varying values of μ with $\tilde{\mu} = \sigma_0(\mu)$, the attacker's best response to the SPRT$(A(\mu), B(\mu))$.

strategic spy in search of information has very different objectives than a strategic bot master stealing computational resources.

We have assumed that the data is generated by exactly one of two possible distributions, and it is the objective of the defender to detect which of the two distributions is generating the data. Alternatively we may assume that the probability distribution changes at some point during the observation process, and the goal of the defender is to detect such a change as quickly as possible. This type of sequential analysis is known as quickest (or change-point) detection [58], and it too has many security applications. Thus it would be of interest to consider the the same type of non-zero-sum, game-theoretic analysis for Bayesian quickest detection as was done here for Bayesian sequential detection.

Simple hypothesis testing might be too limited a test in practice. One response to this is to use composite hypothesis tests. One may test the hypothesis $\theta \in \Theta_0$

versus $\theta \in \Theta_1$ where Θ_0 and Θ_1 are two disjoint subsets of the parameter space Θ. Even with a more robust test such as this, the decision of how to formulate Θ_0 and Θ_1 may depend on the actions of the attacker, thus a similar game-theoretic analysis may lend insight into this problem as well.

Chapter 5

Conclusion

To summarize this dissertation we briefly outline our main contributions to the field of adversarial detection games.

Interdependent detection games and botnet modeling

- To the best of our knowledge we present the first network interdependent detection game. The vast majority of the literature has focussed on security investment as a measure of effort. We contend that the modeling tools of IDS games can be enriched by considering other types of effort. In this vain we propose studying adversarial threat detection in an interdependent environment. While detection is not strictly a type of investment, it can be understood in the larger context of IDS games.

- While strategic adversaries have been modeled in interdependent security games, most of the literature has focussed on zero-sum interactions. That is, the loss incurred by the defenders is the direct gain by the attacker. To our knowledge we are the first to consider the case of non-zero-sum interactions between attacker and defenders in the IDS literature.

- To the best of our knowledge non-cooperative distributed detection has not been considered in the literature. Distributed detection usually refers to interdependent sensors cooperatively trying to detect the presence of a signal. As such distributed detection can be modeled as a cooperative (or coalition) game and is usually approached as a design problem. In our model the sensors are the defenders and they are non-cooperative. They are only interested in detecting the presence of the infection individually. Thus our work does not fit entirely into the distributed detection literature, but is nonetheless related.

Sequential detection games

- We present the first non-zero-sum, continuous-time, sequential detection game. Most sequential games have been restricted to discrete-time, zero-sum games.

- To the best of our knowledge we are the first to use results from the theory of optimal stopping and free-boundary problems [56,64] to construct sequential detection games.

Open Problems

There are several avenues for potential future research. We outline several below.

- We have restricted our attention to Erdős-Rényi random graphs in the epidemic detection games considered here. One criticism of Erdős-Rényi random graphs is that they are not realistic models of some real-world networks. In particular if one is modeling technological networks such as the internet would like to consider other networks that have node degree distributions

with heavier tails. The local mean field models of Lelarge and Bolot [47, 48] are not restricted to these graphs. The only requirement is that the underlying graphs be "locally tree-like". This includes graphs with arbitrary node degree distributions. It would thus be of interest to study the same types of games but on different network topologies to see how the results would differ. One of the main difficulties with this approach is analytical tractability.

- Incorporating security investments directly into our detection game could be accomplished by coupling our LMF model with the original LMF model of Lelarge and Bolot. In Appendix A.9 we generalize the infection dynamics and show how this might help one start to connect the two LMF models.

- Incorporating the intended target of a DDoS attack in the epidemic detection game would help better understand the misaligned incentives of the DDoS problem.

- Nash equilibrium results of the heterogeneous game were shown to be very hard to come by. Further understanding of the effects of defender heterogeneity on the existence of Nash equilibria in a game with a strategic attacker is needed.

- The epidemic detection games considered here are static one shot games. It would be of interest to incorporate dynamics into the model. One straight forward way of doing this is to considered repeated epidemic detection games and allow players to update their strategies after each round. One could study convergence properties of the resulting dynamic game. Alternatively it may interesting to combine the sequential detection games and epidemic detection games considered here.

- The sequential detection game, though dynamic in principle, reduces to a static one-shot game as well. The attacker is limited to choosing a single drift for the stochastic process. It would be of interest to consider allowing the attacker to change his strategy over time. This is a much more difficult problem, but may be approachable utilizing results from optimal stopping and free-boundary problems and other dynamic programming type arguments. Alternatively it may be easier to reformulate the problem as a stochastic game [32].

Appendix A

Appendix

A.1 Proof of Lemma 1

To prove the lemma we fix $T \geq A \geq 0$ and show that for any value of $\lambda q > 0$ the inequality holds. First note that if $T = A$ then $\theta(A, T, \lambda q) = 1$ for all $\lambda q > 0$. Thus we fix $T > A$. We begin by noting that for fixed A and T the implicit function theorem gives us that $h(\lambda q)$ is a differentiable function of the joint parameter λq. Furthermore we can show that $h(\lambda q)$ is monotonically increasing in λq as in [48].From the definition of $h(\lambda q)$ in Prop. 2 it suffices to prove that $\lambda q(F_S(T-A) - h(\lambda q)) < 1$, or equivalently $F_S(T-A) - \frac{1}{\lambda q} < h(\lambda q)$, for all $\lambda q > 0$. The definition in Prop. 2 gives us $F_S(T-A)p \leq h(\lambda q)$. Now, if $0 < \lambda q < \frac{1}{F_S(T-A)(1-p)}$, then

$$\begin{aligned} h(\lambda q) &\geq F_S(T-A)p \\ &= F_S(T-A) - F_S(T-A)(1-p) \\ &> F_S(T-A) - \frac{1}{\lambda q}. \end{aligned}$$

It follows that the claim is true for all $\lambda q \in (0, \frac{1}{F_S(T-A)(1-p)})$. Now suppose there exists a value $y^* \geq \frac{1}{F_S(T-A)(1-p)}$ such that $F_S(T-A) - \frac{1}{y^*} = h(y^*)$. Using the definition of h we have

$$F_S(T-A) - \frac{1}{y^*} = F_S(T-A)[1-(1-p)e^{1-y^* F_S(T-A)}],$$

which gives

$$1 = y^* F_S(T-A)(1-p)e^{1-y^* F_S(T-A)}.$$

It is straightforward to show that for values $0 < \alpha, \beta < 1$ we must have $\alpha\beta x e^{1-\alpha x} < 1$ for all $x > 0$. But this contradicts our result. Hence no such y^* exists. By the continuity of h in λq there are also no values of λq such that $F_S(T-A) - \frac{1}{\lambda q} > h(\lambda q)$. This establishes our result for fixed A and T. Since the choice of A and T was arbitrary this establishes the proposition.

A.2 Proof of Lemma 2

Property 1) is evident from the proof of Prop. 4. To prove property 2) let $0 \leq A \leq A_\infty$. From the monotonicity properties of of $L(A, T_\varnothing)$ in (2.1) and $L_\infty(A)$ in (2.3) we have for all $T_\varnothing \in [A, \infty)$,

$$L(A, T_\varnothing) \leq L_\infty(A) \leq L_\infty(A_\infty).$$

On the other hand using the monotonicity properties of $V(A,T)$ in T and $V_\infty(A)$ in A we have for all $T \in [A, \infty)$,

$$V_\infty(A_\infty) \leq V_\infty(A) \leq V(A,T).$$

By the definition of A_∞ we have $L_\infty(A_\infty) = V_\infty(A_\infty)$. It follows that $L(A, T_\varnothing) \leq V(A, T)$, which corresponds to case 2) in the proof of Prop. 4 impling $\sigma_p(A) = \infty$.

To prove property 3) fix $A > A_\infty$. That $\sigma_p(A)$ is single-valued is established in Prop. 4. For any $(a, t) \in \mathbb{R}^2$ define the function $G(a, t) \triangleq L(a, t) - V(a, t)$. For any point $(a', t') \in \mathbb{R}^2$ satisfying $G(a', t') = 0$ the implicit function theorem gives us the existence of a continuously differentiable function $t(a)$ such that $t(a') = t'$ and $G(a, t(a)) = 0$ for all a in some open neighborhood of a'. Since $\sigma_p(A)$ is the unique value satisfying $G(A, \sigma_p(A)) = 0$ for all $A > A_\infty$, we must have $t(A) = \sigma_p(A)$ for all $A > A_\infty$. Thus $\sigma_p(A)$ is continuously differentiable for all $A > A_\infty$.

To show $\lim_{A \downarrow A_\infty} \sigma_p(A) = \infty$ it suffices to show that for any $M > 0$ there exists an $\epsilon > 0$ such that $\sigma_p(A) > M$ whenever $0 < A - A_\infty < \epsilon$. From the definition of $\sigma_p(A)$ and the monotonicity and continuity of $L(A, T)$ and $V(A, T)$ in T, it suffices to show that for any $M > 0$ there exists $\tilde{\epsilon} > 0$ such that for any $\epsilon < \tilde{\epsilon}$ there exists $\delta > 0$ such that $L(A_\infty + \epsilon, M) > V(A_\infty + \epsilon, M)$ and $L(A_\infty + \epsilon, M + \delta) < V(A_\infty + \epsilon, M + \delta)$.

Since $V(A, T)$ is decreasing in T we must have $V(A, M) > \lim_{T \to \infty} V(A, T) = V_\infty(A)$ for all finite $A \geq 0$ and $M > A$. To simplify notation we use the following definitions:

$$\rho(A, T) \triangleq (1 - p) e^{-\lambda q h(A, T-k)}, \tag{A.1}$$

$$\rho_\infty \triangleq (1 - p) e^{-\lambda q h_\infty}. \tag{A.2}$$

Notice that since $h(A, T)$ depends on A and T only through the difference $T - A$

we have for any $T > A > 0$ and any constant m the following:

$$\begin{aligned}\frac{V(A, T-m)}{V_\infty(A)} &= \frac{\rho(A, T-m)}{1-\rho(A, T-m)} \frac{1-\rho_\infty}{\rho_\infty} \\ &= \frac{\rho(A+m, T)}{1-\rho(A+m, T)} \frac{1-\rho_\infty}{\rho_\infty} \\ &= \frac{V(A+m, T)}{V_\infty(A+m)}.\end{aligned}$$

Furthermore $V_\infty(A+m) = \frac{\ell(A)}{\ell(A+m)} V_\infty(A)$ finally giving us

$$\frac{\ell(A)}{\ell(A+m)} V(A, T-m) = V(A+m, T).$$

Now fix $M > A_\infty$ and choose some $\epsilon > 0$. Then

$$\begin{aligned}\frac{V(A_\infty+\epsilon, M)}{L(A_\infty+\epsilon, M)} &> \frac{V(A_\infty+\epsilon, M)}{L_\infty(A_\infty+\epsilon)} \\ &= \frac{V(A_\infty+\epsilon, M)}{L_\infty(A_\infty+\epsilon)} \frac{V_\infty(A_\infty+\epsilon)}{V_\infty(A_\infty+\epsilon)} \\ &= \frac{V(A_\infty+\epsilon, M)}{V_\infty(A_\infty+\epsilon)} \frac{\frac{\ell(A_\infty)}{\ell(A_\infty+\epsilon)} V_\infty(A_\infty)}{L_\infty(A_\infty+\epsilon)}\end{aligned}$$

Taking the limit as $\epsilon \downarrow 0$ of the right hand side we have

$$\lim_{\epsilon \downarrow 0} \frac{V(A_\infty+\epsilon, M)}{V_\infty(A_\infty+\epsilon)} \frac{\frac{\ell(A_\infty)}{\ell(A_\infty+\epsilon)} V_\infty(A_\infty)}{L_\infty(A_\infty+\epsilon)} = \frac{V(A_\infty, M)}{V_\infty(A_\infty)} > 1$$

Thus we arrive at

$$\lim_{\epsilon \downarrow 0} \frac{V(A_\infty+\epsilon, M)}{L(A_\infty+\epsilon, M)} > 1.$$

Therefore there exists a $\epsilon' > 0$ such that $V(A_\infty+\epsilon, M) > L(A_\infty+\epsilon, M)$ for all

$\epsilon < \epsilon'$. The definition of A_∞ and the monotonicity properties of $L(A_\infty + \epsilon, T)$ and $V(A_\infty + \epsilon, T)$ in T guarantee the existence of a value $T_\epsilon > M$ such that $V(A_\infty + \epsilon, T_\epsilon) = L(A_\infty + \epsilon, T_\epsilon)$. This implies $\sigma_p(A_\infty + \epsilon) = T_\epsilon > M$.

Now fix $A \geq A_0$. Consider the case where $L(A, T)$ is constant in T. Then $L(A, T) \equiv L_0(A) \equiv L_\infty(A)$. By the monotonicity of $L_\infty(A)$ we have for any $T_\varnothing \geq A$

$$L(A, T_\varnothing) = L_\infty(A) \geq L_\infty(A_0).$$

On the other hand the monotonicity of $V(A, T)$ in T and of $V_0(A)$ in A give us for any $T > A$

$$V(A, T) < V_0(A) \leq V_0(A_0).$$

By definition $L_\infty(A_0) = V_0(A_0)$ which implies $L(A, T_\varnothing) < V(A, T)$ for all $T_\varnothing, T > A$. This corresponds to case 1) with $\tilde{T} = A$ or case 3) in the proof of Prop. 4, which implies $\sigma_p(A) = A$ in either case.

Now consider the case where $L(A, T)$ is strictly monotonically increasing in T. For the sake of contradiction suppose $\lim_{A\to\infty} \sigma_p(A) - A > 0$. For all $T \geq A$ we have the following.

$$\begin{aligned} L_0(A) = 0 \leq V_\infty(A) < V(A, T) \\ L_\infty(A) \geq L_\infty(A_0) = V_0(A_0) \geq V_0(A) \geq V(A, T) \end{aligned}$$

Since $L(A, T_\varnothing)$ is strictly monotonically increasing and $V(A, T)$ is monotonically decreasing there must exist a $\tilde{T}(A)$ such that $L(A, \tilde{T}) = V(A, \tilde{T})$ and by the proof of Prop. 4 we must have $\sigma_p(A) = \tilde{T}(A)$. Similarly there must exist a

$T_0(A)$ such that $L(A, T_0(A)) = V_0(A)$. By the monotonicity of $L(A, T_\varnothing)$ and $V(A, T)$ we must have $A \leq \tilde{T}(A) = \sigma_p(A) \leq T_0(A)$ for all $A \geq A_0$. Taking the limit as $A \to \infty$ we have $\lim_{A\to\infty} V_0(A) = 0$. Since $L(A, T_0(A)) = V_0(A)$ we have $\lim_{A\to\infty} L(A, T_0(A)) = 0$. Since $A \leq T_0(A)$ for all $A > A_0$, we must have $\lim_{A\to\infty} T_0(A) = \infty$. Since $f_S(\cdot)$ is a pdf we must have $\lim_{A\to\infty} f_S(T_0(A)) = 0$. But $\lim_{A\to\infty} \frac{f_S(T_0(A)-A)}{f_S(T_0(A))} = \lim_{A\to\infty} L(A, T_0(A)) = 0$. Notice that this implies $\lim_{A\to\infty} f_S(T_0(A) - A) = 0$, for otherwise $\lim_{A\to\infty} L(A, T_0(A)) = \infty$. By assumption the only possibility is $\lim_{A\to\infty} T_0(A) - A = \infty$. Moreover we must have $f_S(T_0(A)-A)$ converge to 0 faster than $f_S(T_0(A))$. However, since $\lim_{x\to\infty} f_S(x) = 0, f_S(x) > 0$ for $x > 0$ and $f_S(\cdot) \in C^1$ there must exist a value x_0 such that $f_S(x)$ is monotonically decreasing for all $x > x_0$. This implies that for all large enough A we will have $f_S(T_0(A) - A) > f_S(T_0(A))$. It follows that for all large enough A we will have $L(A, T_0(A)) > 1$. But this contradicts the convergence of $L(A, T_0(A)$ to 0, implying our initial assumption was incorrect, thus establishing the final property.

A.3 Proof of Lemma 3

That $\sigma_b(T) < T$ for $T > 0$ is apparent in the proof of Prop. 5. To prove property 2) for any $(A, T) \in \mathbb{R}^2$ define

$$y(A, T) \triangleq \frac{g(A)}{g'(A)} - \frac{F_S(T - A)}{f_S(T - A)}\theta(A, T).$$

For any values $A' < T'$ satisfying $y(A', T') = 0$ the implicit function theorem gives us the existence of a continuously differentiable function $A(T)$ such that $A(T') = A'$ and $y(A(T), T) = 0$ for all T in some open neighborhood of T'. By the strict quasi-concavity of $U(A, T)$ established in Prop. 5 we have $\sigma_b(T)$ as the

unique value satisfying $y(\sigma_b(T), T) = 0$ for any $T > 0$. Therefore $A(T) = \sigma_b(T)$ for all $T > 0$ and $\sigma_b(T)$ is continuously differentiable for all $T > 0$.

Suppose 3) is false, i.e. assume $\limsup_{T\to\infty} \sigma_b(T) < \infty$. Then there exists some value $N > 0$ such that $\sigma_b(T) < N$ for all T. By the optimality of $\sigma_b(T)$ we should have $U(\sigma_b(T), T) \geq U(A, T)$ for all A, T. However,

$$\begin{aligned}\limsup_{T\to\infty} U(\sigma_b(T), T) &= \limsup_{T\to\infty} g(\sigma_b(T))h(\sigma_b(T), T) \\ &< g(N)h_\infty.\end{aligned}$$

But for any $\epsilon > 0$ we have

$$\begin{aligned}\lim_{T\to\infty} U(N+\epsilon, T) &= g(N+\epsilon) \lim_{T\to\infty} h(N+\epsilon, T) \\ &= g(N+\epsilon)h_\infty > g(N)h_\infty.\end{aligned}$$

It follows that there exists some T_0 such that $U(\sigma_b(T_0), T_0) < U(N+\epsilon, T_0)$. This violates the optimality of $\sigma_b(T)$, hence $\limsup_{T\to\infty} \sigma_b(T) = \infty$. Now suppose

$$\limsup_{T\to\infty} (T - \sigma_b(T)) = 0.$$

Then $\limsup_{T\to\infty} \frac{T}{\sigma_b(T)} = 1$ which implies $\limsup_{T\to\infty} \frac{T-1}{\sigma_b(T)} = 1$. Therefore there exists a function $\delta(T)$ such that $\limsup_{T\to\infty} \delta(T) = 0$ and $\frac{T-1}{\sigma_b(T)} = 1 + \delta(T)$.

Recall that $h(A, T)$ only depends on the difference $T - A$ and $\lim_{T\downarrow A} h(A, T) = 0$. It follows that $h(T-1, T) = h(0, 1) > 0$ for all T. Furthermore by our assumption that $\limsup_{T\to\infty}(T - \sigma_b(T)) = 0$ we must have $\limsup_{T\to\infty} h(\sigma_b(T), T) = 0$.

It follows that $\lim_{T\to\infty} \frac{h(\sigma_b(T),T)}{h(T-1,T)} = 0$. But

$$\begin{aligned}\frac{h(\sigma_b(T),T)}{h(T-1,T)} &= \frac{U(\sigma_b(T),T)}{U(T-1,T)}\frac{g(T-1)}{g(\sigma_b(T))} \\ &= \frac{U(\sigma_b(T),T)}{U(T-1,T)}\frac{g((1+\delta(T))\sigma_b(T))}{g(\sigma_b(T))}.\end{aligned}$$

The concavity of $g(\cdot)$ guarantees us that $\limsup_{T\to\infty} \frac{g((1+\delta(T))\sigma_b(T))}{g(\sigma_b(T))} = 1$. Therefore we must have

$$\limsup_{T\to\infty} \frac{U(\sigma_b(T),T)}{U(T-1,T)} = 0.$$

It follows that there exists a T_0 such that $\frac{U(\sigma_b(T_0),T_0)}{U(T_0-1,T_0)} < 1$ which violates the optimality of $\sigma_b(T)$ and establishes property 3). Property 4) follows from Properties 1)-3).

A.4 Proof of Lemma 5

Note that the result is trivial for $\phi = 0$. Thus we assume $\phi > 0$. Define the functions

$$\begin{aligned}f_\infty(\phi) &= 2(1-p)\phi e^{-\phi h_\infty(\phi)}, \\ y(\phi) &= 2(1-p)\phi e^{-\phi+\frac{1}{2}}.\end{aligned}$$

To establish our result it suffices to prove that

$$f_\infty(\phi) \overset{<}{\underset{>}{=}} 1 \iff y(\phi) \overset{<}{\underset{>}{=}} 1. \tag{A.3}$$

By the implicit function theorem both $f_\infty(\phi)$ and $y(\phi)$ are differentiable in ϕ. It is then straight forward to show that both $f_\infty(\phi)$ and $y(\phi)$ are strictly

quasi-concave with unique global maxima at $\phi = 1$. Furthermore

$$\begin{aligned} f_\infty(\phi) < y(\phi) &\iff 2(1-p)\phi e^{-\phi h_\infty(\phi)} < 2(1-p)\phi e^{-\phi+\frac{1}{2}} \\ &\iff -\phi h_\infty(\phi) < -\phi + \frac{1}{2} \\ &\iff h_\infty(\phi) > 1 - \frac{1}{2\phi} \\ &\iff f_\infty(\phi) < 1. \end{aligned}$$

By similar reasoning we arrive at

$$f_\infty(\phi) \overset{<}{\geq} y(\phi) \iff f_\infty(\phi) \overset{<}{\geq} 1. \tag{A.4}$$

First consider the case $p \geq 1 - \frac{1}{2}e^{\frac{1}{2}}$. It follows that for all $\phi > 0$

$$y(\phi) \leq \max_{\phi>0} y(\phi) = y(1) = 2(1-p)e^{-\frac{1}{2}} \leq 1,$$

with $y(\phi) = 1$ if and only if $p = 1 - \frac{1}{2}e^{\frac{1}{2}}$ and $\phi = 1$. Suppose there exists a $\phi_0 > 0$ such that $f_\infty(\phi_0) > 1$. By Lemma 4

$$\lim_{\phi\to 0} f_\infty(\phi) = \lim_{\phi\to 0} 2(1-p)\phi e^{-\phi h_\infty(\phi)} = 0.$$

By the continuity of $f_\infty(\phi)$ there must exist a value $\phi_1 \in (0, \phi_0)$ such that $f_\infty(\phi_1) = 1$. By (A.4) this implies $y(\phi_1) = 1$. On the other hand, by Lemma 4

$$\lim_{\phi\to\infty} f_\infty(\phi) = \lim_{\phi\to\infty} 2(1-p)\phi e^{-\phi h_\infty(\phi)} = 0.$$

Again by the continuity of $f_\infty(\phi)$ there must exist a value $\phi_2 > \phi_0 > \phi_1$ such that $f_\infty(\phi_2) = 1$. By (A.4) this implies $y(\phi_2) = 1$. But by our assumption

$y(\phi) = 1$ if and only if $p = 1 - \frac{1}{2}e^{\frac{1}{2}}$ and $\phi = 1$. This contradicts the result that $y(\phi_1) = y(\phi_2) = 1$ and $\phi_2 > \phi_1$. Thus no such ϕ_0 exists and we must have $f_\infty(\phi_0) \leq 1$. It is not hard to see that if $p > 1 - \frac{1}{2}e^{\frac{1}{2}}$ then both $y(\phi)$ and $f_\infty(\phi)$ are strictly less than one and (A.3) holds. Similarly if $p = 1 - \frac{1}{2}e^{\frac{1}{2}}$ then $y(\phi)$ and $f_\infty(\phi)$ are strictly less than one if and only if $\phi \neq 1$ and $y(1) = f_\infty(1) = 1$. Again (A.3) holds.

Now consider the case $p < 1 - \frac{1}{2}e^{\frac{1}{2}}$. In this case there exist two values ϕ_1 and ϕ_2 with $0 < \phi_1 < 1 < \phi_2$ that are solutions to $y(\phi) = 1$. Furthermore by the continuity and strict quasi-concavity of $y(\phi)$ we must have

$$\begin{aligned} y(\phi) > 1 &\iff \phi \in (\phi_1, \phi_2), \\ y(\phi) < 1 &\iff \phi \notin [\phi_1, \phi_2], \\ y(\phi) = 1 &\iff \phi \in \{\phi_1, \phi_2\}. \end{aligned}$$

For $\phi \notin (\phi_1, \phi_2)$ the same contradiction arguments used in the case $p \leq 1 - \frac{1}{2}e^{\frac{1}{2}}$ can be used to establish condition (A.3). Thus we need only establish the result for $\phi \in (\phi_1, \phi_2)$.

Let $\phi \in (\phi_1, \phi_2)$. Then it must be that $y(\phi) > 1$. Recall that both $y(\phi)$ and $f_\infty(\phi)$ take their maximum values at $\phi = 1$. We claim that for $p < 1 - \frac{1}{2}e^{\frac{1}{2}}$ we have $f_\infty(1) > y(1)$. Let $\tilde{p} = 1 - \frac{1}{2}e^{\frac{1}{2}}$. By the implicit function theorem $h_\infty(1)$ and $y(1)$ are differentiable functions in p. Furthermore that $h_\infty(1)$ is monotonically increasing in p and $\lim_{p\to 0} h_\infty(1) = 0$. By the continuity and monotonicity of $h_\infty(1)$ in p there exists some value $p_0 \in [0, 1]$ such that $h_\infty(1) = \frac{1}{2}$ when $p = p_0$, $h_\infty(1) < \frac{1}{2}$ when $p < p_0$ and $h_\infty(1) > \frac{1}{2}$ when $p > p_0$. It is straight forward to also

show that

$$f_\infty(1) \overset{>}{\underset{<}{=}} y(1) \iff h_\infty \overset{<}{\underset{>}{=}} \frac{1}{2}.$$

Suppose $p_0 > \tilde{p}$. Then for $p = p_0$ we would have $f_\infty(1) = y(1)$. By (A.4) this implies $f_\infty(1) = y(1) = 1$, but this contradicts the fact that $y(\phi) < 1$ for all $p > \tilde{p}$. Suppose on the other hand that $p_0 < \tilde{p}$. Then at $p = p_0$ we would have $f_\infty(1) = y(1)$. By (A.4) this implies $f_\infty(1) = y(1) = 1$, but this contradicts the fact that $y(\phi) > 1$ for all $\phi \in (\phi_1, \phi_2)$ when $p < \tilde{p}$. It follows that $p_0 = \tilde{p}$ and $f_\infty(1) > y(1)$ for all $p < \tilde{p}$.

Now suppose there exists a value $\phi_0 \in (\phi_1, 1)$ such that $f_\infty(\phi_0) < y(\phi_0)$. Since $f_\infty(\phi) < y(\phi)$ for $\phi < \phi_1$ the continuity of $f_\infty(\phi)$ and $y(\phi)$ imply the existence of a point $\phi_3 > \phi_0$ such that $f_\infty(\phi_3) = y(\phi_3)$. Again by (A.4) this implies $f_\infty(\phi_3) = y(\phi_3) = 1$. This contradicts the fact that $y(\phi) > 1$ for all $\phi \in (\phi_1, \phi_2)$. A similar arguments shows there is no such point in $(1, \phi_2)$. It follows that $f_\infty(\phi) > y(\phi)$ for all $\phi \in (\phi_1, \phi_2)$. Since $y(\phi) > 1$ for all $\phi \in (\phi_1, \phi_2)$ we have $f_\infty(\phi) > 1$. Thus $y(\phi) > 1 \implies f_\infty(\phi) > 1$. The other direction is trivial since $y(\phi) > 1$ for all $\phi \in (\phi_1, \phi_2)$.

A.5 Proof of Lemma 6

It follows from Lemma 5 that

$$1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} > 0 \iff 1 - 2\phi e^{-\phi + \frac{1}{2}} > 0.$$

Thus if $p \geq 1 - \frac{1}{2}e^{\frac{1}{2}}$ then $1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} + (1-p)e^{-\phi h_\infty(\phi)} > 0$ and we are done. Now let $p < 1 - \frac{1}{2}e^{\frac{1}{2}}$. Again let ϕ_1 and ϕ_2 be solutions to $2\phi e^{-\phi + \frac{1}{2}} = 1$ with

$0 < \phi_1 < 1 < \phi_2$. By Lemma 5 if $\phi \notin (\phi_1, \phi_2)$ then $1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} > 0$ again giving us the result. Now let $\phi \in (\phi_1, \phi_2)$. Then $1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} < 0$.

Suppose $\phi \in (\phi_1, 1]$. It follows that $(1-p)e^{-\phi h_\infty(\phi)} - (1-p)\phi e^{-\phi h_\infty(\phi)} \geq 0$. At the same time $1 - (1-p)\phi e^{-\phi h_\infty(\phi)} > 0$. Combining these inequalities we arrive at $1 - 2(1-p)\phi e^{-\phi h_\infty(\phi)} + (1-p)e^{-\phi h_\infty(\phi)} > 0$. Now suppose $\phi \in (1, \phi_2)$. Define the function

$$u(\phi) \triangleq (2\phi - 1)(1-p)e^{-\phi h_\infty(\phi)},$$

which is clearly differentiable in ϕ. It suffices to show that $u(\phi) \leq 1$ for $\phi \in (1, \phi_2)$. Notice that $u(1) = (1-p)e^{-h_\infty(1)} < 1$. Suppose there exists a value ϕ_0 such that $u(\phi_0) > 1$. By continuity there must exist a value $\phi_c \in (1, \phi_0)$ such that $u(\phi_c) = 1$, or equivalently

$$(2\phi_c - 1)(1-p)e^{-\phi_c h_\infty(\phi_c)} = 1,$$

from which it follows that

$$h_\infty(\phi_c) = 1 - \frac{1}{2\phi_c - 1}.$$

Plugging this into the definition of $h_\infty(\phi)$ gives

$$1 = (2\phi_c - 1)(1-p)e^{-\phi_c + \frac{\phi_c}{2\phi_c - 1}}.$$

Define the function $v(\phi) \triangleq (2\phi - 1)(1-p)e^{-\phi + \frac{\phi}{2\phi-1}}$. Differentiating we obtain

$$\frac{\partial v}{\partial \phi} = -4\frac{(1-p)e^{-\phi + \frac{\phi}{2\phi-1}}}{2\phi - 1}(\phi - 1)^2.$$

Notice that for $\phi > 1$ we have $\frac{\partial v}{\partial \phi} < 0$. Thus $\max_{\phi \in (1,\phi_2)} v(\phi) = v(1) = (1-p)e^{-1+\frac{1}{2-1}} = 1-p < 1$. But this contradicts our assumption that $u(\phi_c) = 1$. Thus there is no such ϕ_0 such that $u(\phi_0) > 1$ and we conclude that $u(\phi) \leq 1$. This establishes our result.

A.6 Proof of Lemma 7

Define the function

$$g(A,T) \triangleq 1 - 2F_S(T-A)\lambda q \rho(A,T) + \rho(A,T).$$

First note that $g(A,A) = 1 + \rho(A,A) = p > 0$. Furthermore

$$\frac{\partial g}{\partial T} = -\lambda q f_S(T-A)\frac{\rho}{\theta}\left(2(1-\lambda q F_S(T-A)) + (1-\rho)\right).$$

Notice that $\frac{\partial g}{\partial T} = 0$ if and only if $2(1-\lambda q F_S(T-A)) + (1-\rho) = 0$, or equivalently if and only if $\rho(3-\rho) = 2\lambda q F_S(T-A)\rho$. Suppose $\frac{\partial g}{\partial T} = 0$ at some point T_0. Then

$$\begin{aligned} g(A,T_0) &= 1 - 2F_S(T_0-A)\lambda q \rho(A,T_0) + \rho(A,T_0) \\ &= (1-\rho(A,T_0))^2 \geq 0. \end{aligned}$$

It follows that if $g(T) < 0$ for some value T' then $g(T) < 0$ for all $T > T'$. Otherwise by the mean value theorem there would exist a value T_c such that $g(T_c) < 0$ and $\frac{\partial g}{\partial T}\Big|_{T=T_c} = 0$ which is a contradiction. Specifically if there exists a T' such that $g(T') < 0$ then it must be that $\lim_{T\to\infty} g(T) < 0$. We will show that this is not possible.

Recall from (2.18) that h_∞ satisfies $h_\infty = 1 - (1-p)e^{-\lambda q h_\infty}$ and $\rho_{\text{inf}} = (1-$

$p)e^{-\lambda q h_\infty}$. Defining $g_\infty \equiv \lim_{T\to\infty} g(T)$ we have

$$\begin{aligned} g_\infty &= \lim_{T\to\infty} [1 - 2F_S(T-A)\lambda q\rho(A,T) + \rho(A,T)] \\ &= 1 - 2\lambda q\rho_\infty + \rho_\infty \end{aligned}$$

It suffices to show that $g_\infty \geq 0$. But this is exactly the result in Lemma 6, giving us our result.

A.7 Proof of Lemma 8

Property 1) follows from the fact that $\frac{\partial \bar{C}_c}{\partial T} < 0$ for $T < A$. To prove the remaining properties we use the function $M(A,T)$ in (2.34). Recall that $M(A,T)$ is monotonically increasing and differentiable in T and $\text{sign}\,(M) = \text{sign}\left(\frac{\partial \bar{C}_c}{\partial T}\right)$ for all $T > A$.

Let $A > A^c_\infty$. That $\sigma_c(A)$ is single-valued follows from the strict quasi-concavity of $\bar{C}_c(A,T)$ in T. For any $(A',T') \in \mathbb{R}^2$ satisfying $M(A',T') = 0$ the implicit function theorem gives us the existence of a continuously differentiable function $T(A)$ such that $T(A') = T'$ and $M(A,T(A)) = 0$ for all A in some open neighborhood of A'. Since $\sigma_p(A)$ is the unique value satisfying $M(A,\sigma_p(A)) = 0$ for all $A > A^c_\infty$, we must have $T(A) = \sigma_p(A)$ for all $A > A^c_\infty$. Thus $\sigma_p(A)$ is continuously differentiable for all $A > A^c_\infty$.

The strict monotonicity of $M(A,T)$ in T and the definition of A^c_∞ imply that for any $\epsilon > 0$ we must have $M_\infty(A^c_\infty + \epsilon) > 0$. Furthermore there must exist a value $T^* > A^c_\infty + \epsilon$ such that $M(A^c_\infty + \epsilon, T^*) = 0$. This is precisely the definition of $\sigma_c(A)$, i.e. $T^* = \sigma_c(A^c_\infty + \epsilon)$. We wish to show that $\lim_{\epsilon\downarrow 0} \sigma_c(A^c_\infty + \epsilon) = \infty$.

To do so we will show that for any $N > 0$ there exists an $\epsilon > 0$ such that $M(A^c_\infty + \epsilon, N) < 0$. Notice that the strict monotonicity of $M(A,T)$ in T and the

definition of A^c_∞ then imply $\sigma_c(A^c_\infty+\epsilon) > N$, giving us the desired result. Suppose there exists some $N_0 < \infty$ such that for all $\epsilon > 0$ we have $M(A^c_\infty + \epsilon, N_0) \geq 0$. By the continuity of $M(A,T)$ in A we must then have $\lim_{\epsilon\downarrow 0} M(A^c_\infty + \epsilon, N_0) \geq 0$. But the definition of A^c_∞ implies $\lim_{\epsilon\downarrow 0} M_\infty(A^c_\infty + \epsilon) = 0$. This violates the strict monotonicity of $M(A,T)$ in T, thus no such N_0 exists and property 3) is proved.

To prove property 4) we begin with the special case that $\frac{\partial L}{\partial T} \equiv 0$. In this case $M_0(A) > -\infty$ for all A. Moreover the monotonicity of $M_0(A)$ in A implies that $M_0(A) \geq 0$ for all $A > A^c_0$. Then for any $A > A^c_0$ the strict monotonicity of $M(A,T)$ in T then implies $M(A,T) > 0$ for all $T > A$. It follows that $\frac{\partial C}{\partial T} > 0$ for all $T > A$, hence $\sigma_c(A) = A$.

Now consider the case $\frac{\partial L}{\partial T} > 0$. In this case $M_0(A) = -\infty$ for all A. For any $A > A^c_\infty$ we have $M_\infty(A) > 0$ and $\sigma_c(A)$ is the unique value satisfying $M(A, \sigma_c(A)) = 0$. Furthermore by the continuity of $\sigma_c(A)$ we have

$$\lim_{A\to\infty} M(A, \sigma(A)) = 0.$$

We also have $\sigma_c(A) > A$ so $\lim_{A\to\infty} \sigma_c(A) = \infty$. By assumption $\ell(A) \to \infty$, thus we have the following:

$$\lim_{A\to\infty} M(A,\sigma_c(A))\rho(A,\sigma_c(A)) = \lim_{A\to\infty} \left[1 - \frac{c}{\ell(A)} \frac{\theta(A,\sigma_c(A))}{L(A,\sigma_c(A))} \frac{\rho(A,\sigma_c(A))}{1-\rho(A,\sigma_c(A))}\right].$$

Since $0 < \rho(A,T) < 1$ for all A, T we must have $\lim_{A\to\infty} M(A,\sigma_c(A))\rho(A,\sigma_c(A)) = 0$, from which it follows that

$$\lim_{A\to\infty} \frac{c}{\ell(A)} \frac{\theta(A,\sigma_c(A))}{L(A,\sigma_c(A))} \frac{\rho(A,\sigma_c(A))}{1-\rho(A,\sigma_c(A))} = 1.$$

Furthermore, we have the bounds

$$0 < \theta(A, \sigma_c(A)) \frac{\rho(A, \sigma_c(A))}{1 - \rho(A, \sigma_c(A))} < \infty,$$

while $\lim_{A\to\infty} \frac{c}{\ell(A)} = 0$. It follows that $\lim_{A\to\infty} L(A, \sigma_c(A)) = 0$. At this point the same considerations apply as in the decentralized case and the property 4) follows.

A.8 Extension of Equilibrium Results to $G(n, \lambda/n)$

A.8.1 Convergence Results for the Centralized Botnet Game

The preceding analysis is applicable to the limiting object of a sequence of random rooted Poisson Branching Process $T_n(\lambda) \to \mathcal{T}(\lambda)$. In this section we show that Nash equilibria on $\mathcal{T}(\lambda)$ are also Nash equilibria in the same game played on the limiting graph of a sequence of Erdős-Rényi random graphs $G(n, \lambda/n)$, which we denote by $G_\infty(\lambda)$. The proof relies on the objective method [3] and follows the proof in [48].

Notice that for a given A and T a defender's cost and the botmaster's utility are random variables. Fixing $A \in \mathbb{R}^+$ and $T \in \mathbb{R}^+$ let $C_i^{(n)}(A, T)$ be the random cost of defender i, $(i = 1, 2, ..., n)$ and $U_b^{(n)}(A, T)$ the random utility of the bot master on $G(n, \lambda/n)$. Let $X_i^{(n)}(A, T)$ be the indicator random variable for a false alarm and $Y_i^{(n)}(A, T)$ be the indicator random variable for a missed detection for defender i on $G(n, \lambda/n)$. Furthermore let $W_i^{(n)}(A, T)$ be the indicator random variable for infection of defender i on $G(n, \lambda/n)$ and let $D_i^{(n)}(A, T)$ be the indicator random variable for a detection event by defender i on $G(n, \lambda/n)$. If defender i and defender j are neighbors in $G(n, \lambda/n)$ then we write $i \sim j$. We will suppress

the A, T dependence notation from here on. With the above notation we have the following relations.

$$W_i^{(n)} = 1 - (1 - \chi_i^{(n)}) \prod_{i \sim j} (1 - B_{ki}^{(n)} Y_i^{(n)}) \tag{A.5}$$

$$D_i^{(n)} = \mathbb{1}_{\{T < W_i^{(n)} + S_i^{(n)} A\}} \tag{A.6}$$

$$X_i^{(n)} = (1 - W_i^{(n)}) D_i^{(n)} \tag{A.7}$$

$$Y_i^{(n)} = W_i^{(n)} (1 - D_i^{(n)}) \tag{A.8}$$

Let

$$\begin{aligned}
C_i^{(n)} &= c X_i^{(n)} + \ell Y_i^{(n)}, \\
C^{(n)} &= \frac{1}{n} \sum_{i=1}^{n} C_i^{(n)} = c \frac{1}{n} \sum_{i=1}^{n} X_i^{(n)} + \ell \frac{1}{n} \sum_{i=1}^{n} Y_i^{(n)} \\
U_b^{(n)} &= A \frac{1}{n} \sum_{i=1}^{n} Y_i^{(n)}.
\end{aligned}$$

The expected cost and utilities are then

$$\begin{aligned}
E[C_i^{(n)}] &= c E[X_i^{(n)}] + \ell E[Y_i^{(n)}], \\
E[C^{(n)}] &= c \frac{1}{n} \sum_{i=1}^{n} E[X_i^{(n)}] + \ell \frac{1}{n} \sum_{i=1}^{n} E[Y_i^{(n)}] \\
E[U_b^{(n)}] &= A \frac{1}{n} \sum_{i=1}^{n} E[Y_i^{(n)}].
\end{aligned}$$

Because the underlying graph $G(n, \lambda/n)$ is random the labeling of nodes is interchangeable and by exchangeability we have for all $i \neq j$

$$\begin{aligned}
E[X_i^{(n)}] &= E[X_j^{(n)}], \\
E[Y_i^{(n)}] &= E[Y_j^{(n)}].
\end{aligned}$$

In particular the root node of $G(n, \lambda/n)$, say node $i = 0$ is chosen uniformly at random, thus we have for all $i = 0, 1, 2, 3..., n-1$

$$E[C_i^{(n)}] = cE[X_0^{(n)}] + \ell E[Y_0^{(n)}], \tag{A.9}$$

$$E[C^{(n)}] = c\frac{1}{n}\sum_{i=1}^{n} E[X_0^{(n)}] + \ell\frac{1}{n}\sum_{i=1}^{n} E[Y_0^{(n)}] = cE[X_0^{(n)}] + \ell E[Y_0^{(n)}] \tag{A.10}$$

$$E[U_b^{(n)}] = A\frac{1}{n}\sum_{i=1}^{n} E[Y_0^{(n)}] = AE[Y_0^{(n)}]. \tag{A.11}$$

Proposition 18. *For any $(A, T) \in \mathbb{R}^+ \times \mathbb{R}^+$ if the processes $\left\{X_i^{(n)}(A, T)\right\}_{i=0}^{n-1}$ and $\left\{Y_i^{(n)}(A, T)\right\}_{i=0}^{n-1}$ satisfy (A.5) - (A.8) on $G(n, \lambda/n)$, then*

$$\lim_{n\to\infty} E\left[X_i^{(n)}(A, T)\right] = [1 - F_S(T)](1-p)e^{-\lambda q h(A,T)}$$
$$\lim_{n\to\infty} E\left[Y_i^{(n)}(A, T)\right] = h(A, T).$$

Proof. For $d > 0$ let $N_d(1, G(n, \lambda/n))$ be a neighborhood of radius d about the root node $i = 1$ of $G(n, \lambda/n)$. For fixed d we have $G(n, \lambda/n) \xrightarrow{\mathcal{D}} T(\lambda, d)$ as $n \to \infty$. By the Skorohod Representation Theorem we can consider the two random graphs to be defined on the same probability space and with probability one, there is a finite random variable N such that $N_d(0, G(n, \lambda/n)) = T(\lambda, d)$ for all $n \geq N$. Fix $d > 0$ and denote the leaves of $T(\lambda, d)$ by $\partial T(\lambda, d)$. We now construct two depth-d recursive tree processes, $L_i^{(d)}$ and $U_i^{(d)}$. For $i \in \partial T(\lambda, d)$ let

$$L_i^{(d)} = \chi_i \mathbb{1}(T \leq S_i + \chi_i A)$$
$$U_i^{(d)} = 1.$$

For any recursive tree process (RTP) R_i defined for each $i \in \mathcal{T}(\lambda)$ define the

functionals $W(\cdot)$ and $D(\cdot)$ as follows.

$$W\left(R_i\right) = 1 - \left(1 - \chi_i\right) \prod_{j \to i} \left(1 - B_{ji} R_j\right)$$
$$D\left(R_i\right) = \mathbb{1}_{\left\{T < S_i + W\left(R_i\right) A\right\}}$$

Thus the functional $W(\cdot)$ and $D(\cdot)$ are actually functionals of all children of the argument R_i. For all $i \notin \partial T(\lambda, d)$ we define

$$L_i^{(d)} = W\left(L_i^{(n)}\right) D\left(L_i^{(n)}\right)$$
$$U_i^{(d)} = W\left(U_i^{(n)}\right) D\left(U_i^{(n)}\right).$$

For $n > N$ we can consider $N_d(0, G(n, \lambda/n)) = \mathcal{T}(\lambda)$. We can then define the corresponding RTP $\left\{\tilde{Y}_i^{(n)}(A,T)\right\}_{i=0}^{n-1}$ for $n > N$ by

$$\tilde{Y}_i^{(n)}(A,T) = \begin{cases} Y_i^{(n)}(A,T) \text{ if } i \in \partial T(\lambda, d) \\ W\left(\tilde{Y}_i^{(n)}(A,T)\right) D\left(\tilde{Y}_i^{(n)}(A,T)\right) \text{ o.w.} \end{cases}$$

Observe that for $n > N$ we have $E[Y_0^{(n)}] = E[\tilde{Y}_0^{(n)}]$. This is not necessarily true for $i \neq 0$, but we are only concerned about the root here.

First observe that for all $n \geq N$ and for all $i \in \partial T(\lambda, d)$ we have $L_i^{(d)} \leq \tilde{Y}_i^{(n)} \leq U_i^{(d)}$. We will show that in fact $L_i^{(d)} \leq \tilde{Y}_i^{(n)} \leq U_i^{(d)}$ holds for all i of equal depth in the tree, in particular

$$L_\varnothing^{(d)} \leq \tilde{Y}_0^{(n)} \leq U_\varnothing^{(d)}. \tag{A.12}$$

We prove (A.12) by showing that the functionals $W(\cdot)$ and $D(\cdot)$ are monotonic, i.e. for any indicator random variables Q_i, R_i defined for each $i \in \mathcal{T}(\lambda)$, if $Q_j \leq$

R_j for each j such that $j \to i$, then $W(Q_i) \leq W(R_i)$ and $D(Q_i) \leq D(R_i)$. To prove this we consider the different cases. First note that if $\chi_i = 1$ then $W(Q_i) = W(R_i) = 1$. Suppose $\chi_i = 0$. If $\prod_{j\to i}(1 - B_{ji}Q_j) = \prod_{j\to i}(1 - B_{ji}R_j)$ then $W(Q_i) = W(R_i) = 1$. Suppose $\prod_{j\to i}(1 - B_{ji}Q_j) \neq \prod_{j\to i}(1 - B_{ji}R_j)$. Then there are two possibilities. Either

$$0 = \prod_{j\to i}(1 - B_{ji}Q_j) < \prod_{j\to i}(1 - B_{ji}R_j) = 1 \tag{A.13}$$

or

$$1 = \prod_{j\to i}(1 - B_{ji}Q_j) > \prod_{j\to i}(1 - B_{ji}R_j) = 0. \tag{A.14}$$

Suppose (A.13) is true. Then $B_{ji}R_j = 0$ for all j such that $(i, j) \in E$ while at the same time $B_{ji}Q_j = 1$ for some j such that $(i, j) \in E$. Let j^* be such that $B_{j^*i}Q_{j^*} = 1$. Then we must have $B_{j^*i} = Q_{j^*} = 1$. But then $R_{j^*} = 0$ giving us $R_{j^*} < Q_{j^*}$. This contradicts our assumption that $Q_j \leq R_j$. It follows that (A.14) must hold, which implies $0 = W(Q_j) < W(R_j) = 1$. This exhausts all possibilities.

The proof for the monotonicity of the functional $D(\cdot)$ follows directly from the monotonicity of $W(\cdot)$. Specifically if $W(Q_i) \leq W(R_i)$ then we need only consider the two cases. If $W(Q_i) = W(R_i)$ then $D(Q_i) = D(R_i)$. If $W(Q_i) \neq W(R_i)$ then we have $W(Q_i) = 0$ and $W(R_i) = 1$, in which case $D(Q_i) = \mathbb{1}_{\{T_i<S_i\}}$ and $D(R_i) = \mathbb{1}_{\{T_i<S_i+A\}}$. If $T_i < S_i$ then $D(Q_i) = D(R_i) = 1$. If $S_i \leq T_i < S_i + A$ then $0 = D(Q_i) < D(R_i) = 1$. Finally if $T_i \geq S_i + A$ then $D(Q_i) = D(R_i) = 0$. Hence $D(Q_i) \leq D(R_i)$.

By the monotonicity of both $W(\cdot)$ and $D(\cdot)$ we have the monotonicity of $W(\cdot)D(\cdot)$. Thus for all i at depth $d-1$ from the root we must have $L_i^{(d)} \leq \tilde{Y}_i^{(n)} \leq$

$U_i^{(d)}$. If $d = 1$ we then have (A.12) trivially. By induction on d we obtain the result for any finite d.

We now have

$$E[L_\varnothing^{(d)}] \leq E[\tilde{Y}_0^{(n)}] = E[Y_0^{(n)}] \leq E[U_\varnothing^{(d)}]. \tag{A.15}$$

In order to finish the proof we show that $L_\varnothing^{(d)}$ and $U_\varnothing^{(d)}$ both converge in distribution to Bernoulli random variables with parameter $h(A,T)$ as $d \to \infty$.

Define $h_i^{(d)} = P\left(L_i^{(d)} = 1\right)$. For $d = 1$ we have

$$\begin{aligned}
L_\varnothing^{(1)} &= 1 - \max\left\{1 - W\left(L_\varnothing^{(1)}\right), \mathbb{1}\{T_\varnothing \leq S_\varnothing + W\left(L_\varnothing^{(1)}\right) A\}\right\}, \\
W\left(L_\varnothing^{(1)}\right) &= 1 - (1-\chi_i) \prod_{j \to i} \left(1 - B_{ji} L_j^{(1)}\right).
\end{aligned}$$

By definition for $j \in \partial T(\lambda, 1)$

$$\begin{aligned}
h_j^{(1)} &= P\left(L_j^{(1)} = 1\right) \\
&= P\left(\chi_j = 1, T_j > S_j + \chi_j A\right) \\
&= F_S(T_j - A)p.
\end{aligned}$$

It is then possible to show by a similar derivation as we did to get $h(A,T)$ that

$$\begin{aligned}
h_\varnothing^{(1)} &= F_S(T_\varnothing - A)[1 - (1-p)e^{-\lambda q h_1^{(1)}}], \\
&= F_S(T_\varnothing - A)[1 - (1-p)e^{-\lambda q F_S(T_\varnothing - A)p}].
\end{aligned}$$

Define the function $g(x, A, T) = F_S(T - A)[1 - (1-p)e^{-\lambda q x}]$. The above gives $h_\varnothing^{(1)} = g(h_1^{(1)}, A, T)$. By induction on d it is straight forward to show that $h_\varnothing^{(d+1)} = g(h_1^{(d+1)}, A, T) = g^d(F_S(T-A)p, A, T)$ where superscript d represents composition

in x. Thus as $d \to \infty$ repeated composition of the function $g(\cdot, A, T)$ will converge to the unique fixed point solution $h(A, T)$. The proof for $U_{\varnothing}^{(d)}$ is analogous.

With the above we have $\lim_{d\to\infty} E[L_{\varnothing}^{(d)}] = h(A, T)$ and $\lim_{d\to\infty} E[U_{\varnothing}^{(d)}] = h(A, T)$. Then in the limit as $d \to \infty$ we must have $E[Y_0^{(n)}] = h(A, T)$ for $n \geq N$. With these results a similar argument shows that in the limit as $d \to \infty$ for $n \geq N$ we must have

$$E[X_0^{(n)}] = [1 - F_S(T_{\varnothing})](1 - p)e^{-\lambda q h(A,T)}.$$

□

Applying the above proposition to (A.10) and (A.11) we get the following corollary.

Corollary 2. *For any $(A, T) \in \mathbb{R}^+ \times \mathbb{R}^+$ and corresponding processes $\{X_i^{(n)}(A, T)\}_{i=0}^{n-1}$ and $\{Y_i^{(n)}(A, T)\}_{i=0}^{n-1}$ satisfying (A.5) - (A.8) on $G(n, \lambda/n)$ we have*

$$\lim_{n\to\infty} E[C^{(n)}(A, T)] = c[1 - F_S(T)](1 - p)e^{-\lambda q h(A,T)} + \ell h(A, T) = C(A, T),$$
$$\lim_{n\to\infty} E[U_b^{(n)}(A, T)] = Ah(A, T) = U(A, T).$$

Given the above proposition and corollary we have the following.

Proposition 19. *Any pure, symmetric Nash equilibrium (A^*, T^*) in the centralized botnet game on $\mathcal{T}(\lambda)$ is a pure, symmetric Nash equilibrium in the centralized botnet game on $G^{\infty}(\lambda)$.*

A.8.2 Convergence Results for the Decentralized Botnet Game

Now consider the case for the decentralized game. We still work in the same probability space but our strategy space is now $\mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$.

As before for a given $(A, T, T_\varnothing) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$ a root defender's cost and the botmaster's utility are random variables. Let $C_0^{(n)}(A, T, T_\varnothing)$ be the random cost of a root defender and $U_b^{(n)}(A, T)$ the random utility of the bot master on $G(n, \lambda/n)$. As before let $X_i^{(n)}(A, T)$ be the indicator random variable for a false alarm and $Y_i^{(n)}(A, T)$ be the indicator random variable for a missed detection for defender $i > 0$ on $G(n, \lambda/n)$ and denote by $X_0^{(n)}(A, T, T_\varnothing)$ and $Y_0^{(n)}(A, T, T_\varnothing)$ the indicator random variables for false alarm and missed detection, respectively, for a root defender. The defining relations analogous to (A.5)-(A.8) are as follows.

$$W_i^{(n)} = 1 - (1 - \chi_i^{(n)}) \prod_{i \sim j} (1 - B_{ki}^{(n)} Y_i^{(n)}) \tag{A.16}$$

$$D_i^{(n)} = \begin{cases} \mathbb{1}_{\{T_\varnothing < W_0^{(n)} + S_i^{(n)} A\}} & \text{if } i = 0 \\ \mathbb{1}_{\{T < W_i^{(n)} + S_i^{(n)} A\}} & \text{if } i > 0 \end{cases} \tag{A.17}$$

$$X_i^{(n)} = (1 - W_i^{(n)}) D_i^{(n)} \tag{A.18}$$

$$Y_i^{(n)} = W_i^{(n)} (1 - D_i^{(n)}) \tag{A.19}$$

The random cost to the root defender and the random utility to the bot master are then

$$\begin{aligned} C_0^{(n)}(A, T, T_\varnothing) &= c X_0^{(n)}(A, T, T_\varnothing) + \ell Y_0^{(n)}(A, T, T_\varnothing), \\ U_b^{(n)}(A, T, T_\varnothing) &= A \frac{1}{n} \sum_{i=0}^{n-1} Y_i^{(n)} \\ &= A \frac{1}{n} Y_0^{(n)}(A, T, T_\varnothing) + A \frac{1}{n} \sum_{i=1}^{n-1} Y_i^{(n)}(A, T). \end{aligned}$$

The expected cost and utilities become

$$E[C_0^{(n)}(A,T,T_\varnothing)] = cE[X_0^{(n)}(A,T,T_\varnothing)] + \ell E[Y_0^{(n)}(A,T,T_\varnothing)], \tag{A.20}$$

$$E[U_b^{(n)}(A,T,T_\varnothing)] = A\frac{1}{n}E[Y_0^{(n)}(A,T,T_\varnothing)] + A\frac{1}{n}\sum_{i=1}^{n-1} E[Y_i^{(n)}(A,T)]. \tag{A.21}$$

Since a root node is chosen uniformly at random we have by exchangeability for all $i,j \neq 0$

$$E[X_i^{(n)}] = E[X_j^{(n)}],$$
$$E[Y_i^{(n)}] = E[Y_j^{(n)}].$$

Thus we can write $E[U_b^{(n)}(A,T,T_\varnothing)] = A\frac{1}{n}E[Y_0^{(n)}(A,T,T_\varnothing)] + A\frac{n-1}{n}E[Y_1^{(n)}(A,T)]$. Then $\lim_{n\to\infty} E[U_b^{(n)}(A,T,T_\varnothing)] = \lim_{n\to\infty} AE[Y_1^{(n)}(A,T)]$ provided this limit exists. Thus we can consider the limiting expected utility of the bot master as a function of A and T only. In addition if we can show that $\lim_{n\to\infty} E[C_0^{(n)}(A,T,T_\varnothing)] = C_\varnothing(A,T,T_\varnothing)$, then by our previous equilibrium results there will exist an optimal population strategy $T^*(A)$. In this case all defenders will play the same strategy, i.e. $T = T_\varnothing = T^*(A)$ and by exchangeability we will have for all $i \neq j$

$$E[X_i^{(n)}] = E[X_j^{(n)}],$$
$$E[Y_i^{(n)}] = E[Y_j^{(n)}].$$

In particular we have

$$E[U_b^{(n)}(A,T)] = A\frac{1}{n}\sum_{i=0}^{n-1} E[Y_0^{(n)}] = AE[Y_0^{(n)}].$$

Thus it suffices to prove the that

$$\lim_{n\to\infty} E[X_0^{(n)}(A,T,T_\varnothing)] = E[X_\varnothing(A,T,T_\varnothing)],$$
$$\lim_{n\to\infty} E[Y_0^{(n)}(A,T,T_\varnothing)] = E[Y_\varnothing(A,T,T_\varnothing)].$$

The proof of this convergence is exactly as in the centralized case. Thus we state the corresponding propositions for the decentralized game without proof.

Proposition 20. *For any* $(A,T,T_\varnothing) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$ *if the processes* $\left\{X_i^{(n)}\right\}_{i=0}^{n-1}$ *and* $\left\{Y_i^{(n)}\right\}_{i=0}^{n-1}$ *satisfy (A.16) - (A.19) on* $G(n,\lambda/n)$*, then*

$$\lim_{n\to\infty} E\left[X_0^{(n)}\right] = [1-F_S(T_\varnothing)](1-p)e^{-\lambda q h(A,T)},$$
$$\lim_{n\to\infty} E\left[Y_0^{(n)}\right] = F_S(T_\varnothing - A)[1-(1-p)e^{-\lambda q h(A,T)}].$$

Applying the above proposition to (A.20) and (A.21) we get the following corollary.

Corollary 3. *For any* $(A,T,T_\varnothing) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$ *and corresponding processes* $\{X_i^{(n)}(A,T)\}_{i=0}^{n-1}$ *and* $\{Y_i^{(n)}(A,T)\}_{i=0}^{n-1}$ *satisfying (A.16) - (A.19) on* $G(n,\lambda/n)$ *we have*

$$\begin{aligned}\lim_{n\to\infty} E[C^{(n)}(A,T)] &= c[1-F_S(T)](1-p)e^{-\lambda q h(A,T)} \\ &\qquad + \ell F_S(T_\varnothing - A)[1-(1-p)e^{-\lambda q h(A,T)}] \\ &= C_\varnothing(A,T,T_\varnothing), \\ \lim_{n\to\infty} E[U_b^{(n)}(A,T)] &= Ah(A,T) \\ &= U(A,T).\end{aligned}$$

Given the above proposition and corollary we have the following.

Proposition 21. *Any pure, symmetric Nash equilibrium* (A^*, T^*) *in the decentralized botnet game on* $\mathcal{T}(\lambda)$ *is a pure, symmetric Nash equilibrium in the decentralized botnet game on* $G^\infty(\lambda)$.

A.9 Generalized Infection Dynamics

One assumption in our LMF model is that if a defender detects and removes the bot infection then he cannot infect any of his neighbors. Clearly this is not going to always be the case. For example, if the virus is particularly fast then it is likely it will infect the entire network before any defender detects it. We can generalize our model to include an entire range of propagation scenarios of which our original model is simply an extreme case.

A.9.1 Detection

In our model each defender is an intrusion detection system. Taking this into account we can put our variables in the traditional detection framework. As we have it now

$$Y_i : \text{False Negative},$$

$$X_i : \text{False Positive}.$$

We introduce a new indicator random variable, D_i, indicating a True Positive, i.e. $D_i = 1$ iff $W_i = 1$ and $Z_i \geq T_i$. Otherwise $D_i = 0$. For completeness we introduce $R_i = (1 - Y_i)(1 - X_i)(1 - D_i)$ as the indicator random variable for True Negative. Recall B_{ki} is an indicator random variable for the event that a propagating virus or worm on host k successfully infects the neighboring defender i. Because our model is not dynamic we must deal with the issue of the ordering of events that take

place. Consider the case where defender i is infected and successfully detects the infection, subsequently removing it, i.e. the case $D_i = 1$. It is realistic to assume that in this case defender i may still transmit the virus before he is able to detect and remove it. We introduce the indicator random variable η_{ki} to indicate that defender k was not able to remove the infection before transmitting it to defender i. Since B_{ki} indicates there was a successful attempt to transmit the virus we can think of η_{ki} as indicating *when* the attempt took place. We have $P(B_{ki} = 1) = q$ and $P(\eta_{ki} = 1) = \hat{q}$. The new equations of our model are

$$\begin{aligned}
W_i &= 1 - (1 - \chi_i) \prod_{k \to i} (1 - B_{ki} Y_k)(1 - B_{ki} \eta_{ki} D_k), \\
D_i &= 1 - \max \{1 - W_i, \mathbb{1}(T_i > S_i + W_i A), \} \\
Y_i &= 1 - \max \{1 - W_i, \mathbb{1}(T_i \leq S_i + W_i A)\}.
\end{aligned}$$

The new Random Distributional Equations are

$$\begin{aligned}
W &\stackrel{d}{=} 1 - (1 - \chi) \prod_{k=1}^{N} (1 - B_k Y_k)(1 - B_k \eta_k D_k), \\
D &\stackrel{d}{=} 1 - \max \{1 - W, \mathbb{1}(T > S + WA), \} \\
Y &\stackrel{d}{=} 1 - \max \{1 - W, \mathbb{1}(T \leq S + WA)\}.
\end{aligned}$$

Again let $P(Y = 1) = h$ and also $P(W = 1) = \gamma$. Notice that

$$P(D = 1) = 1 - P(Y = 1) - P(W = 0) = \gamma - h.$$

Proposition 22. *For fixed $A, T, p, q, \hat{q}$ the random variables Y and W are Bernoulli*

with parameters h and γ, respectively, that satisfy

$$\gamma = 1 - (1-p)e^{-\lambda q\gamma[\hat{q}+F_S(T-A)(1-\hat{q})]},$$
$$h = F_S(T-A)\gamma.$$

Proof. As before we can derive the equation for the value h and obtain

$$h = F_S(T-A)P(W=1) = F_S(T-A)\gamma.$$

In a similar manner we can also derive an equation for γ.

$$\begin{aligned}
1-\gamma &= P(W=0) \\
&= P\left((1-\chi)\prod_{k=1}^{N}(1-B_kY)(1-B_k\eta_kD)=1\right) \\
&= P(\chi=0)\sum_{n=0}^{\infty}\left[P(B_kY=0,B_k\eta_kD=0)\right]^n P(N=n) \\
&= (1-p)\sum_{n=0}^{\infty}\left[P(B_kY=0,B_k\eta_kD=0)\right]^n \frac{e^{-\lambda}\lambda^n}{n!}
\end{aligned}$$

We obtain $P(B_kY=0,B_k\eta_kD=0)$ using the fact that

$$\begin{aligned}
P(B_kY=0,B_k\eta_kD=0) = {} & P(B_kY=0,B_k\eta_kD=0|B_k=1)P(B_k=1) \\
& + P(B_kY=0,B_k\eta_kD=0|B_k=0)P(B_k=0).
\end{aligned}$$

We then have

$$\begin{aligned} P(B_k Y = 0, B_k \eta_k D = 0) &= P(Y = 0, \eta_k D = 0 | B_k = 1) P(B_k = 1) + P(B_k = 0) \\ &= P(Y = 0, \eta_k D = 0 | B_k = 1) q + 1 - q. \end{aligned}$$

Similarly

$$\begin{aligned} P(Y = 0, \eta_k D = 0 | B_k = 1) = P(Y = 0, \eta_k D = 0 | \eta_k = 1, B_k = 1) P(\eta_k = 1) \\ + P(Y = 0, \eta_k D = 0 | \eta_k = 0, B_k = 1) P(\eta_k = 0). \end{aligned}$$

Our assumption of independence leads us to

$$\begin{aligned} P(Y = 0, \eta_k D = 0 | B_k = 1) &= P(Y = 0, \eta_k D = 0) \\ &= P(Y = 0, D = 0) \hat{q} + (1 - h)(1 - \hat{q}). \end{aligned}$$

Notice that $P(Y = 0, D = 0) = P(W = 0) = 1 - \gamma$. Combining these results we obtain

$$\begin{aligned} 1 - \gamma &= (1 - p) \sum_{n=0}^{\infty} \left[1 - q(1 - ((1 - \gamma)\hat{q} + (1 - h)(1 - \hat{q}))) \right]^n \frac{e^{-\lambda} \lambda^n}{n!} \\ &= (1 - p) e^{[\lambda - q\lambda(1 - (1-\gamma)\hat{q} - (1-h)(1-\hat{q}))]} e^{-\lambda} \\ &= (1 - p) e^{-\lambda q [1 - (1-\gamma)\hat{q} - (1-h)(1-\hat{q})]} \\ &= (1 - p) e^{-\lambda q [h + \hat{q}(\gamma - h)]}. \end{aligned}$$

Plugging in $h = F_S(T - A)\gamma$ gives us the result. □

A.9.2 Relating our LMF to the Lelarge LMF

Notice if $\hat{q} = 0$ then we arrive at

$$\gamma = 1 - (1-p)e^{-\lambda q h},$$
$$h = F_S(T-A)\gamma,$$

which is equivalent to our LMF. On the other hand $\hat{q} = 1$ leads to

$$\gamma = 1 - (1-p)e^{-\lambda q \gamma},$$
$$h = F_S(T-A)\gamma.$$

This is in some sense equivalent to the Lelarge-Bolot LMF model without investment and our detection model played after the epidemic process is complete. Thus we see that the probability that the infection reaches a typical defender, γ, does not depend on T or A. The only dependence on the strategies is in h, the probability of a missed detection. If all players are homogeneous in their cost functions this case is equivalent to the 2 player game with a probability γ of being infected. This would provide an easy way to connect our two models. Including the Lelarge-Bolot investment model in the case $\hat{q} = 0$ would be a first step in coupling the two distinct LMG models. In this case the models are not truly coupled but are are independent of one another. Selecting $\hat{q} > 0$ would truly couple them. This problem would be much harder, but the case $\hat{q} = 1$ may be tractable.

A.9.3 Open Problem

We have the existence of pure Nash equilibria for the case $\hat{q} = 0$ but we would like to establish similar results for $0 \leq \hat{q} < 1$. For $\hat{q} = 0$ the situation is equivalent

to the two-player game. For $0 < \hat{q} < 1$ the situation is more complicated. Due to the coupling of γ and h a different approach is needed to obtain such a result.

Bibliography

[1] Daron Acemoglu, Azarakhsh Malekian, and Asu Ozdaglar. Network security and contagion. *SIGMETRICS Perform. Eval. Rev.*, 42(3):38–38, December 2014.

[2] David Aldous and Antar Bandyopadhyay. A survey of max-type recursive distributional equations. *Annals of Applied Probability*, 15:1047–1110, 2005.

[3] David Aldous and J. Michael Steele. The objective method: Probabilistic combinatorial optimization and local weak convergence. *Probability on Discrete Structures*, 110:1–72, 2004.

[4] Tansu Alpcan and T. Basar. A game theoretic approach to decision and analysis in network intrusion detection. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 3, pages 2595–2600 Vol.3, Dec 2003.

[5] Tansu Alpcan and T. Basar. A game theoretic analysis of intrusion detection in access control systems. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 2, pages 1568–1573 Vol.2, Dec 2004.

[6] Tansu Alpcan and Tamer Basar. An intrusion detection game with limited observations. *12th Int. Symp. on Dynamic Games and Applications, Sophia Antipolis, France*, 2006.

[7] Tansu Alpcan and Tamer Basar. *Network Security: A Decision and Game Theoretic Approach*. Cambridge University Press, 1 edition, 2011.

[8] Saurabh Amin, Galina A. Schwartz, and S. Shankar Sastry. Security interdependencies for network control systems with identical agents. In *Decision and Game Theory for Security*, November 2010.

[9] Saurabh Amin, Galina A. Schwartz, and Hamidou Tembine. Incentives and security in electricity distribution networks. In *Decision and Game Theory for Security*, November 2012.

[10] R. Anderson. Why information security is hard - an economic perspective. In *Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual*, pages 358–365, Dec 2001.

[11] Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799):610–613, 2006.

[12] Roy M. Anderson, Robert M. May, and Robert McCredie. *Infectious diseases of humans : dynamics and control*. Oxford science publications. Oxford University Press, Oxford, New York, 1991. Includes indexes.

[13] K. J. Arrow, D. Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3/4):pp. 213–244, 1949.

[14] Y. Bachrach, M. Draief, and S. Goyal. Contagion and observability in security domains. In *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*, pages 1364–1371, Oct 2013.

[15] Ning Bao, O.Patrick Kreidl, and John Musacchio. A network security classification game. In Rahul Jain and Rajgopal Kannan, editors, *Game Theory for Networks*, volume 75 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 265–280. Springer Berlin Heidelberg, 2012.

[16] Tamer Basar. On the Relative Leadership Property of Stackelberg Strategies. *Journal of Optimization Theory and Applications*, 11(6):655–661, 1973.

[17] Tamer Basar. A general theory for Stackelberg games with partial state information. *Large Scale Systems*, 3(1):47–56, 1982.

[18] Tamer Basar and Geert Jan Olsder. Mixed Stackelberg Strategies in continuous-kernel games. *IEEE Trans. Automat. Control*, AC-25(2), 1980.

[19] Tamer Basar and Geert Jan Olsder. Team-optimal closed-loop Stackelberg strategies in hierarchical control problems. *Automatica*, 16(4), 1980.

[20] Tamer Basar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, 2 edition, 1995.

[21] Alain Bensoussan, Murat Kantarcioglu, and SingRu (Celine) Hoe. A game-theoretical approach for finding optimal strategies in a botnet defense model. In *Decision and Game Theory for Security*, November 2010.

[22] James Orvis Berger. *Statistical decision theory and bayesian analysis : with 23 illustrations*. Springer series in statistics. Springer, New York, Berlin, Heidelberg, 1985.

[23] Sourabh Bhattacharya and Tamer Basar. Graph-theoretic approach for connectivity maintenance in mobile networks in the presence of a jammer. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 3560–3565. IEEE, 2010.

[24] B. Bollobas. *Random Graphs*. Cambridge University Press, 2001.

[25] Zheng Bu, Pedro Bueno, Rahul Kashyap, and Adam Wosotowsky. The new era of botnets. Technical report, McAfee, 2010.

[26] A.A. Cardenas, S. Amin, G. Schwartz, R. Dong, and S. Sastry. A game theory model for electricity theft detection and privacy-aware control in ami systems. In *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, pages 1830–1837, Oct 2012.

[27] H. Chan, M. Ceyko, and L. E. Ortiz. Interdependent Defense Games: Modeling Interdependent Security under Deliberate Attacks. *ArXiv e-prints*, October 2012.

[28] David Dagon, Cliff Zou, and Wenke Lee. Modeling botnet propagation using time zones. In *In Proceedings of the 13 th Network and Distributed System Security Symposium NDSS*, 2006.

[29] Lemonia Dritsoula, Patrick Loiseau, and John Musacchio. Computing the nash equilibria of intruder classification games. In *Decision and Game Theory for Security*, November 2012.

[30] Richard Durrett. *Probability: Theory and Examples*. Cambridge University Press, 2010.

[31] Paul W. Ewald. The evolution of virulence: A unifying link between parasitology and ecology. *The Journal of Parasitology*, 81(5):pp. 659–669, 1995.

[32] Jerzy Filar and Koos Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag New York, Inc., New York, NY, USA, 1996.

[33] Neal Fultz and Jens Grossklags. Blue versus red: Towards a model of distributed security attacks. In *Financial Cryptography and Data Security*, February 2009.

[34] Nicola Gatti, Mattia Monga, and Sabrina Sicari. A localization game in wireless sensor networks. In *Decision and Game Theory for Security*, November 2010.

[35] Geoffrey Heal and Howard Kunreuther. Interdependent security: A general model. Working Paper 10706, National Bureau of Economic Research, August 2004.

[36] Herbert W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.

[37] Peter Honeyman and Galina Schwartz. Interdependece of reliability and security. In *WEIS*, Jan 2007.

[38] Libin Jiang, Venkat Anantharam, and Jean Walrand. Efficiency of selfish investments in network security. In *NetEcon '08*, August 2008.

[39] Benjamin Johnson, Jens Grossklags, Nicolas Christin, and John Chuang. Uncertainty in interdependent security games. In *Decision and Game Theory for Security*, November 2010.

[40] Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, 1985.

[41] Jeffrey O. Kephart and Steve R. White. Directed-graph epidemiological models of computer viruses. In *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, pages 343–359. IEEE, 1991.

[42] Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, March 2003.

[43] Howard Kunreuther and Geoffrey Heal. IDS models of airline security. *Journal of Conflict Resolution*, 49(2):201–217, 2005.

[44] R.J. La. Interdependent security with strategic agents and cascades of infection. *Networking, IEEE/ACM Transactions on*, PP(99):1–14, 2015.

[45] Aron Laszka, Mark Felegyhazi, and Levente Buttyan. A survey of interdependent information security games. *ACM Comput. Surv.*, 47(2):23:1–23:38, August 2014.

[46] Anthony Lavigna, Armand M. Makowski, and John S. Baras. A continuous-time distributed version of wald's sequential hypothesis testing problem. In A. Bensoussan and J.L. Lions, editors, *Analysis and Optimization of Systems*, volume 83 of *Lecture Notes in Control and Information Sciences*, pages 533–543. Springer Berlin Heidelberg, 1986.

[47] Marc Lelarge. Economics of malware: Epidemic risks model, network externalities and incentives. In *WEIS '09*, June 2009.

[48] Marc Lelarge and Jean Bolot. Network externalities and the deployment of security features and protocols in the internet. In *SIGMETRICS '08*, June 2008.

[49] Zhen Li, Qi Liao, and Aaron Striegel. Botnet economics: Uncertainty matters. In *Managing Information Risk and the Economics of Security*, pages 245–267. Springer US, 2009.

[50] Yi Luo, Ferenc Szidarovszky, Youssif Al-Nashif, and Salim Hariri. A fictitious play-based response strategy for multistage intrusion defense systems. In *Security and Communication Networks*, 2011.

[51] Rajiv T. Maheswaran and Tamer Basar. Social welfare of selfish agents: motivating efficiency for divisible resources. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 2, pages 1550–1555. IEEE, 2004.

[52] Patrick Maille, Peter Reichl, and Bruno Tuffin. Interplay between security providers, consumers, and attackers: A weighted congestion game approach. In *Decision and Game Theory for Security*, November 2011.

[53] Y. Namestnikov. The economics of botnets. Technical report, Kaspersky, November 2011.

[54] J. Omic, A. Orda, and P. Van Mieghem. Protecting against network infections: A game theoretic perspective. In *INFOCOM 2009, IEEE*, pages 1485–1493, April 2009.

[55] Ranjan Pal and Pan Hui. Modeling internet security investments: Tackling topological information uncertainty. In *Decision and Game Theory for Security*, November 2011.

[56] Goran Peskir and Albert Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, 2000.

[57] Viet Pham and Carlos Cid. Are we compromised? Modelling security assessment games. In Jens Grossklags and Jean Walrand, editors, *Decision and Game Theory for Security*, volume 7638 of *Lecture Notes in Computer Science*, pages 234–247. Springer Berlin Heidelberg, 2012.

[58] H. Vincent Poor and Olympia Hadjiliadis. *Quickest Detection*. Cambridge University Press, 2008. Cambridge Books Online.

[59] S. Radosavac and J.S. Baras. Application of sequential detection schemes for obtaining performance bounds of greedy users in the IEEE 802.11 MAC. *Communications Magazine, IEEE*, 46(2):148–154, February 2008.

[60] Svetlana Radosavac, George Moustakides, John S. Baras, and Iordanis Koutsopoulos. An analytic framework for modeling and detecting access layer misbehavior in wireless networks. *ACM Trans. Inf. Syst. Secur.*, 11(4):19:1–19:28, July 2008.

[61] Kurt R Rohloff and Tamer Basar. Stochastic behavior of random constant scanning worms. In *Computer Communications and Networks, 2005. ICCCN 2005. Proceedings. 14th International Conference on*, pages 339–344. IEEE, 2005.

[62] Vicente Segura and Javier Lahuerta. Modeling the economic incentives of DDoS attacks: femtocell case study. In *WEIS'09*, June 2009.

[63] A. N. Shiryaev. On optimum methods in quickest detection problems. *Theory of Probability & Its Applications*, 8(1):22–46, 1963.

[64] A.N. Shiryayev. *Optimal Stopping Rules*. Springer-Verlag, 1978.

[65] David Siegmund. *Sequential analysis : tests and confidence intervals*. Springer series in statistics. Springer-Verlag, New York, 1985.

[66] B. Soper and J. Musacchio. A botnet detection game. In *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*, pages 294–303, Sept 2014.

[67] B. Soper and J. Musacchio. A heterogeneous botnet detection game. In *NetGCoop*, Oct 2014.

[68] Gilbert N. Sorebo and Michael C. Echols. *Smart Grid Security: An End-to-End View of Security in the New Electrical Grid*. CRC Press, 2012.

[69] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to own the internet in your spare time. In Dan Boneh, editor, *USENIX Security Symposium*, pages 149–167. USENIX, 2002.

[70] B. Stone-Gross, M. Cova, B. Gilbert, R. Kemmerer, C. Kruegel, and G. Vigna. Analysis of a Botnet Takeover. *IEEE Security and Privacy Magazine*, 9(1):64–72, January 2011.

[71] Brett Stone-Gross, Marco Cova, Bob Gilbert, Lorenzo Cavallaro, Martin Szydlowski, Christopher Kruegel, Giovanni Vigna, and Richard Kemmerer. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *Proceedings of the Computer and Communications Security Conference (CCS)*, Chicago, IL, November 2009.

[72] Demosthenis Teneketzis and Yu-Chi Ho. The decentralized Wald problem. *Information and Computation*, 73(1):23 – 44, 1987.

[73] G. Theodorakopoulos, J.-Y. Le Boudec, and J.S. Baras. Selfish response to epidemic propagation. In *American Control Conference (ACC), 2011*, pages 4069–4074, June 2011.

[74] Marten van Dijk, Ari Juels, Alina Oprea, and Ronald L. Rivest. Flipit: The game of "stealthy takeover". Cryptology ePrint Archive, Report 2012/103, 2012. `http://eprint.iacr.org/`.

[75] Hal Varian. Managing online security risks. *The New York Times*, June 2000.

[76] Hal R. Varian. System reliability and free riding. In *Economics of Information Security*, pages 1–15. Kluwer Academic Publishers, 2001.

[77] Nevena Vratonjic, Mohammad Hossein Manshaei, Maxim Raya, and Jean-Pierre Hubaux. ISPs and ad networks against botnet ad fraud. In *Decision and Game Theory for Security*, November 2010.

[78] A. Wald. Sequential tests of statistical hypotheses. *Ann. Math. Statist.*, 16(2):117–186, 06 1945.

[79] A. Wald and J. Wolfowitz. Bayes solutions of sequential decision problems. *Ann. Math. Statist.*, 21(1):82–99, 03 1950.

[80] Abraham Wald. *Sequential Analysis*. Wiley, 1947.

[81] Quanyan Zhu, A. Clark, R. Poovendran, and T. Basar. Deployment and exploitation of deceptive honeybots in social networks. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 212–219, Dec 2013.