

The College of Wooster

Open Works

Senior Independent Study Theses

2022

Gaming The System: A Game Theoretic Approach To Risk Management And Cybersecurity

Abigail Breitenbucher

The College of Wooster, abreitenbucher22@wooster.edu


Follow this and additional works at: <https://openworks.wooster.edu/independentstudy>

Recommended Citation

Breitenbucher, Abigail, "Gaming The System: A Game Theoretic Approach To Risk Management And Cybersecurity" (2022). *Senior Independent Study Theses*. Paper 9910.

This Senior Independent Study Thesis Exemplar is brought to you by Open Works, a service of The College of Wooster Libraries. It has been accepted for inclusion in Senior Independent Study Theses by an authorized administrator of Open Works. For more information, please contact openworks@wooster.edu.

© Copyright 2022 Abigail Breitenbucher



GAMING THE SYSTEM: A GAME THEORETIC APPROACH TO RISK MANAGEMENT AND CYBERSECURITY

INDEPENDENT STUDY THESIS

Presented in Partial Fulfillment of the Requirements for
the Degree Bachelor of Arts in the
Department of Mathematics at The College of Wooster

by
Abigail Breitenbucher
The College of Wooster
2022

Advised by:

Dr. Colby Long (Mathematics)



THE COLLEGE OF

WOOSTER

© 2022 by Abigail Breitenbucher

ABSTRACT

In this independent study, we detail two fields of game theory, traditional and evolutionary. After establishing foundational concepts in both, we explore risk management and cybersecurity through a game theoretic lens, focusing on some current applications in both fields. The research culminates in an application case of a game theoretic model for the 2017 Equifax data breach. To conclude, we discuss the implications of this model and future work involving evolutionary game theory in cybersecurity/risk management.

ACKNOWLEDGMENTS

Thank you to my parents who have listened to me ramble about Equifax and the different ways we can apply game theory since September! Thank you as well to my advisor, Dr. Long, for helping me navigate the theoretical component of this research and providing support throughout this process.

CONTENTS

Abstract	iii
Acknowledgments	iv
Contents	v
List of Figures	vii
List of Tables	viii
CHAPTER	PAGE
1 Game Theory	6
1.1 Definitions	6
1.2 Static Games	7
1.2.1 Prisoner's Dilemma	8
1.2.2 Dominant Strategies and Nash Equilibrium	10
1.2.2.1 Preparation for Nash's Theorem	12
1.2.2.2 Proof of Nash's Theorem	16
1.2.3 Battle of the Spouses	18
1.2.4 Rock, Paper, Scissors	20
1.2.5 Conclusion	24
1.3 Finite Dynamic Games	24
1.3.1 Dinner Party Game	26
1.3.2 Refining Nash Equilibria	28
1.3.3 Conclusion	32
2 Evolutionary Game Theory	33
2.1 Evolutionary Stability	33
2.1.1 Hawk-Dove Game	36
2.1.2 Rock, Paper, Scissors Revisited	39
2.1.3 Equivalent Concepts to Evolutionary Stability	40
2.2 Evolutionary Stability and Dynamics	44
2.2.1 Prisoner's Dilemma Revisited	44
2.2.2 Learning Rules	49
2.3 Dynamics	50
2.3.1 Dynamic Types	50
2.3.2 Connection between ESSs and Dynamics	52
2.3.3 Dominance and Dynamics	53

2.3.4	Usefulness of Evolutionary Game Theory	54
2.4	Conclusion	55
3	Applications of Game Theory in Risk Management and Cybersecurity	57
3.1	Classifications of Deceptions	57
3.2	Physical Surveillance	62
3.2.1	Background	62
3.2.2	Physical Surveillance Model	65
3.2.2.1	Basic General Model for Physical Surveillance Games	65
3.2.2.2	Model with Uncertainty	67
3.2.3	Implementing the Uncertainty Model	70
3.2.4	Example: Fare Evasion	74
3.2.4.1	Establish Context	74
3.2.4.2	Identify Strategies	76
3.2.4.3	Identify Goals	78
3.2.4.4	Assess Strategies	78
3.2.4.5	Find Optimal Configuration	80
3.2.4.6	Implement Optimal Configuration	82
3.3	Conclusion	82
4	Equifax Data Breach as a Game-Theoretic Situation	84
4.1	Background	84
4.2	Why use game theory?	86
4.3	How do we approach this situation?	88
4.4	Equifax Model	89
4.4.1	Defining Actions	89
4.4.2	Goals and Strategies	90
4.4.3	Assessing Strategies	94
4.4.4	Determining the Optimal Configuration	94
4.4.5	Implementation	96
4.5	What does this tell us?	97
4.6	Directions for the Future	98
	References	103

LIST OF FIGURES

Figure		Page
1.1	Clothes Decision Tree	25
1.2	Dinner Party Game Tree	27
1.3	Dinner Party Game Tree with Three Starting Options	31
2.1	Replicator Dynamical Model of the Prisoner’s Dilemma [10]	48
3.1	Models used in Moving Target Defense, Mixing, Perturbation, and Obfuscation [22]	59
3.2	Models used in Honey-X and Attacker Engagement [22]	60
3.3	Taxonomy of Defensive Deception [22]	61
3.4	Graph of Defender and Attacker Pure Strategies	66
3.5	Example Public Transportation System [11]	75
3.6	Expert Judgment of Strategies for Inspection Intensity	79
3.7	Expert Judgment of Strategies for Detection Intensity	79
3.8	Expert Judgment of Strategies for Cost	79
3.9	Results of HyRiM: Optimal Inspection Strategy [11]	80

LIST OF TABLES

Table		Page
1.1	Payoff Matrix for Prisoner's Dilemma	9
1.2	Payoff Matrix for Battle of the Spouses	18
1.3	Payoff Matrix for Rock, Paper, Scissors	21
1.4	Maximin Example	22
2.1	Payoff Matrix for Population of Incumbents and Mutants	35
2.2	Payoff Matrix for Hawk-Dove Game	37
2.3	Payoff Matrix for Rock, Paper, Scissors	39
2.4	Payoff Matrix for General Prisoner's Dilemma	45
2.5	Payoff Matrix for Rock, Paper, Scissors, Twin	54
3.1	Action Set for Transit Company	77
4.1	Strategy Set for Equifax with N-number of Employees	93

INTRODUCTION

As an increased number of people, businesses, and governments have started to rely on the Internet and technology, their data has become a target for attacks. Attacks differ in both form and scope, but their results can be devastating. In recent years, news about data breaches where attackers gained millions of people's data have become commonplace, so how can businesses protect their data?

There are two primary fields that deal with the question of data safety and protection: risk management and cybersecurity. Both of these fields try to find ways to stay ahead of attackers to protect consumers, employees, and employers. They approach the problem in different ways; with risk management focusing on prevention at the company level and cybersecurity focusing on protecting systems. Since the approaches in these fields differ slightly, they complement each other when implemented by a business.

Risk management encompasses the techniques companies use to identify, analyze, and respond to various risk factors. The ultimate goal of risk management is to control as many factors that could negatively affect the company as possible. In this way, we can think of risk management solutions as responses to a predictive modeling approach that emphasizes proactive action instead of reaction. The structures behind risk management should provide calculations for uncertainty and predict how that would affect the business. For instance, if an analyst at a company noticed that economists were predicting an economic downturn in the next six months, they could utilize risk management structures to mitigate the effect

such a downturn would have on the company. The final products generated by risk management are choices where a company has to decide between accepting or rejecting risks. The acceptance or rejection of risks depends on a company's tolerance level, the amount of loss that a company is willing to experience.

There are three primary ways to respond to predicted risks: avoidance, mitigation, and acceptance. Each of these methods respond to risks differently. In avoidance strategies, companies try to eliminate the root cause of a risk. In mitigation strategies, they attempt to decrease the projected financial value associated with the risk, and in acceptance strategies, companies accept the risk, but they create a contingency plan to mitigate its impact. Having risk management strategies in place is especially important now because of how much information passes between companies and consumers. In general, these techniques can help a business weather various risks, and in the case of data security, they can help prevent or limit the damage done by data breaches. This information has been adapted from [1].

A related field to risk management is cybersecurity. It focuses on protecting systems, as well as confidential and valuable information from digital attacks. Cybersecurity consists of the various methods a corporation can use to prevent or limit the damage done by attackers on their networked systems. Based on the goal of protecting systems, we could see cybersecurity as an application area for risk management strategies. In addition to general risks caused by external attackers, in cybersecurity, a corporation also needs to be able to account for internal attackers as well. Generally, there are a few broad categories of protection strategies in the field: the protection of computer systems and networks, the protection of information, disaster recovery, and end-user education. The first two categories deal with implementing practices that specifically protect computer systems, networks, and other items on which society relies, as well as protecting sensitive data from unauthorized access/theft. Disaster recovery focuses more on how companies

react during an unforeseen event (like a natural disaster or cybersecurity incident) without disrupting daily operations and end-user education focuses on providing consumers information on how to protect themselves and their data from attacks.

There are several different types of cyber threats and risks currently. For instance, malware, a type of program or file that can damage a computer or the data it houses, is a common threat that has evolved to skirt around familiar detection methods like antivirus software. Other common threats are ransomware, phishing schemes, distributed denial-of-service (DDoS) attacks, and advanced persistent threats. Ransomware locks down files, data, or systems with the threat of deletion unless the target company or government pays the attackers a specified sum. Ransomware is like an extension of malware and is transmitted in a similar manner through email attachments or web browsers. Phishing schemes rely on emails or messages that appear to be from legitimate companies asking for personal information on the target. DDoS attacks are coordinated efforts to crash servers, websites, or networks by overloading them with traffic. Different from the rest of these threats is the advanced persistent threat which involves an individual (or a group) that infiltrates a system and remains undetected for an extended stretch of time. These types of threats are more insidious because they do not affect the normal operations of a system, allowing the attackers to gather more information on a company and steal data. This information has been adapted from [3].

While there are many types of attacks, there are also a good number of strategies to prevent or decrease their damage. For example, two-factor authentication, a method which requires end-users to input a password as well as some other form of verification that they are the correct user, are a simple way to help secure data. There are also more high-level methods such as a company's information technology (IT) department assigning particular roles and access privileges to employees. This

method also helps protect the company from attacks by internal agents since access to sensitive information can be restricted.

Even though corporations have been able to find combinations of risk management and cybersecurity strategies that work for them, introducing game theory techniques could be more useful than current methods. Game theory, a mathematical framework that analyzes how players with conflicting interests interact in situations of interdependence, what strategies they choose, and how they assess the outcomes of their chosen strategies, fits quite well into both risk management and cybersecurity [24]. In both fields, we start with a corporation that wants to protect its data from an attacker, so we could generically view this as a two person game where both parties have conflicting goals. The attackers want to steal the company's data while attracting as little notice as possible, whereas the company wants to protect the data and immediately identify any possible threats to that goal. Continuing to think of this in terms of game theory, each player will implement a strategy based on their goals and will need to assess the outcomes of their strategies to determine if they are worth keeping in the future.

There has been some previous work connecting game theory with risk management and cybersecurity, however, it primarily focuses on traditional game theoretic methods to model situations in these fields. This thesis will continue to build on that work while also examining evolutionary game theory as another possible modeling approach. Evolutionary game theory, a field that started because of biology, has traditionally been used to model the strategic element of evolution in nature. However, social scientists have also used it to explore the change in cultural phenomenon over time [10]. This thesis argues that extending current frameworks for understanding game situations in a risk management/cybersecurity context and incorporating some aspects of evolutionary game theory will allow companies to gain a better understanding of security threats and ways to prevent attacks. To show

this, we will first establish important concepts in both traditional and evolutionary game theory which we will then build upon by talking about recent applications of game theory in risk management and cybersecurity. This research will culminate in an application of a risk management framework in a cybersecurity context, that of the 2017 Equifax data breach, and we will discuss the ways in which evolutionary techniques could be introduced to the two fields through this application.

CHAPTER 1

GAME THEORY

Game theory encompasses a large number of situations that can be thought of as games. Because of this, it has become a popular way to think about problems in economics, political science, and many other fields. With risk management and cybersecurity, game theoretic techniques have only started to be applied recently. This chapter will introduce some common terms used in game theory and discuss both static and dynamic games to illustrate the variety of models that exist in traditional game theory.

1.1 DEFINITIONS

The following definitions are standard ones used in games and allow us to extend our understanding of what a game can be to a variety of contexts.

Definition 1.1 PLAYER PAYOFF: *The payoff for a player is a function mapping the set of actions they could take to the real number line, connecting every action in the action set to a numerical value. The goal for players is to maximize individual payoff [6].*

Definition 1.2 RATIONAL PLAYER: *In a game, a player is assumed to be rational, meaning they have preferences, beliefs about the world, and that they will try to optimize their individual payoff. They do this keeping in mind that the other players in a game are also attempting to maximize their payoff [23].*

Definition 1.3 STRATEGY: *Strategies are plans that determine what players choose to do in a given position [23].*

In our application fields, we can view a company as a single, rational player that is attempting to maximize its security against unforeseen threats. The company is then going against a similarly rational attacker whose goal is to maximize the damage (the amount of information they can take) inflicted on a company. The most simple understanding of this game would be in the context of scams or attacks on security systems since there are obviously at least two opposing parties. The same could not be said in the case of a natural disaster where we could not characterize the disaster as a rational player. Thus, we will be limiting the games we look at to ones between an attacker or groups of attackers and corporate entities.

1.2 STATIC GAMES

Static games, also known as simultaneous decision games, involve players choosing an action to take with no knowledge of the decision made by their opponents. These kinds of games could map well to situations in a risk management context since corporations and attackers may not have certain knowledge about how the other party will respond to a threat or new security measure. What makes these games interesting then is how a player determines the best course of action without being able to react to their opponent's strategy. Some examples of these types of games would be the Prisoner's Dilemma, a popular game of discussion in game theory, the Battle of the Spouses, and Rock, Paper, Scissors. In each of these games, the players have several strategies they could choose from, however, they do not know what their opponent will pick, so they need to be able to find a solution which works out the best for them.

To describe static games, we need to specify three items:

1. The set of players in the game, which we index by $i \in \{1, 2, \dots\}$
2. A pure strategy set, S_i , for each player.
3. The payoffs for each player for every possible combination of pure strategies used by all players [8].

Typically when thinking about static games, we create a matrix showing what each player's payoff would be given the choice they make. This visualization allows us to see what strategies are available for each player and it gives us a quick way to determine which strategies would be viable. To illustrate what these matrices look like and how we can use them to solve problems in game theory, we will look at the Prisoner's Dilemma.

1.2.1 PRISONER'S DILEMMA

In the setup of the Prisoner's Dilemma, the police are questioning two suspects in connection to a crime. The suspects are held in different cells and cannot consult each other about possible courses of action at any point in time. In this example, the police do not have enough evidence to arrest the two suspects for the crime they are accusing the suspects of committing, but, they could arrest them on a lesser charge. The police make this offer to the suspects separately:

1. If one suspect confesses that they both participated in the more serious crime:
 - (a) Confessor receives no jail time
 - (b) Other person receives a prison sentence for the serious crime and additional time for obstructing justice.
2. If both confess to the serious crime:
 - (a) They receive the same amount of jail time for it.
3. If neither suspect confesses to the serious crime:
 - (a) Both receive the same amount of jail time for the lesser charge.

In terms of the game's setup, the length of the sentences in this game are flexible since we are focused on finding optimal solutions as opposed to specific sentence lengths. We will give the serious crime a prison sentence of 15 years, the lesser crime one of 5 years, and the sentence for obstructing justice 1 year. The values will be treated as negatives since the suspects will lose their freedom for however many years the situation dictates. Our payoff matrix for this scenario is shown in Table 1.1.

Table 1.1: Payoff Matrix for Prisoner's Dilemma

Decisions	Confess	Stay Silent
Confess	(-15,-15)	(0, -16)
Stay Silent	(-16, 0)	(-5,-5)

We will consider the payoff pairs as coordinates such that the row player's payoff will be the first coordinate and the column player's payoff will be the second. Looking at the table, we can see that the worst outcome for a player would be if their opponent confessed while they stayed silent. Since that is the case, a player would be better off confessing no matter what their opponent does. Thus, both suspects would confess and end up serving 15 years in prison. This result may seem counter-intuitive since there is a clearly better option, that being the one where they both stay silent. If neither confessed, then they would each serve 5 years in prison which is a much lighter sentence in comparison. However, since this is a simultaneous game where neither player can consult with the other, they have no way of knowing if their opponent will confess. That uncertainty is what leads to (-15, -15) being the answer since, in terms of a player's self-interest, their best scenario is getting 0 years in prison and their worst is 15 years, both of which are better than if their opponent were to confess and they stayed silent, receiving 16 years in prison. The solution to the Prisoner's Dilemma is socially inefficient, meaning

that cooperation would have been the more efficient solution, since the sentence for both players would only be 5 years. The competitive aspect of these kinds of games is what makes them interesting since players acting in their own interest may have a worse outcome than if they had cooperated with their partner, yet they do not have the opportunity to cooperate.

1.2.2 DOMINANT STRATEGIES AND NASH EQUILIBRIUM

As the Prisoner's Dilemma illustrates, there are particular strategies that will be more beneficial to a player than others. Since we are attempting to solve these games (find the strategies which result in the best payoff for a player of interest), we can eliminate the weaker strategies and follow the stronger ones. The stronger strategies are known as dominant strategies and they can dominate over others to varying degrees. We look at payoff functions, which are functions of the strategies played by those in a game and show an individual's payoff given that play scenario, to determine the dominance of strategies. We also think of the strategies as "living" in a strategy space, defined by Σ_i , which houses the strategies available to player i .

Definition 1.4 STRICTLY DOMINATED STRATEGY: A strategy for a player, σ_i , is strictly dominated by another strategy σ'_i if

$$\pi_i(\sigma'_i, \sigma_j) > \pi_i(\sigma_i, \sigma_j) \quad \forall \sigma_j \in \Sigma_j$$

where i represents a player, j represents their opponent, and π_i represents the payoff function for player i .

Definition 1.5 WEAKLY DOMINATED STRATEGY: A strategy for a player, σ_i , is weakly dominated by another strategy σ'_i if

$$\pi_i(\sigma'_i, \sigma_j) \geq \pi_i(\sigma_i, \sigma_j) \quad \forall \sigma_j \in \Sigma_j$$

and

$$\exists \sigma'_j \in \Sigma_j \quad \text{subject to} \quad \pi_i(\sigma'_i, \sigma'_j) > \pi_i(\sigma_i, \sigma'_j)$$

where π_i represents the payoff function for player i [8].

These definitions are useful in solving a variety of games since we can find strongly or weakly dominant strategies and eliminate what they dominate from our strategy set. There is a caveat that the solution for a game may depend on the order in which we eliminate strategies.

Another approach is solving games using Nash Equilibria.

Definition 1.6 NASH EQUILIBRIUM: A Nash equilibrium for a two-player game is a pair of strategies (σ_1^*, σ_2^*) such that

$$\pi_1(\sigma_1^*, \sigma_2^*) \geq \pi_1(\sigma_1, \sigma_2^*) \quad \forall \sigma_1 \in \Sigma_1$$

and

$$\pi_2(\sigma_1^*, \sigma_2^*) \geq \pi_2(\sigma_1^*, \sigma_2) \quad \forall \sigma_2 \in \Sigma_2$$

where π_i is the payoff for player i [8].

Essentially, given a strategy one's opponent plays, neither player could do strictly better by adopting another strategy. We can extend this to games with more than two players as well. In our Prisoner's Dilemma example, a Nash equilibrium would be both suspects confessing. Another way we could define Nash equilibrium is to say a strategy is the best response to the fixed strategy played by another player. We can therefore think about Nash equilibria in terms of payoff and strategy.

Definition 1.7 ALTERNATE DEFINITION OF NASH EQUILIBRIUM: A pair of strategies, (σ_1^*, σ_2^*) , is a Nash equilibrium for a two-player game if

$$\sigma_1^* \in \arg \max \pi_1(\sigma_1, \sigma_2^*) \quad \sigma_1 \in \Sigma_1$$

and

$$\sigma_2^* \in \arg \max \pi(\sigma_1^*, \sigma_2) \quad \sigma_2 \in \Sigma_2 \quad [8].$$

To use this definition to solve games, we determine the best responses for each player and then find the pair of strategies that are the best responses to each other.

In addition to equilibrium points being the logical solution for players to gravitate towards since they will each maximize their payoff, Nash equilibria are useful because of Nash's theorem:

Theorem 1.1 (Nash's Theorem).

Every game with a finite strategic form (game with a finite number of players and a finite number of pure strategies for each player) has at least one mixed-strategy Nash equilibrium. [8]

With this theorem, we know we can find a Nash equilibrium in any finite game, meaning we can find at least one solution for a finite game in any case.

To prove this theorem, we first need to define some key terms.

1.2.2.1 PREPARATION FOR NASH'S THEOREM

Definition 1.8 PURE STRATEGY: *A pure strategy is a strategy without randomization. In other words, a player chooses one of the strategy options.* [7]

Looking back at the Prisoner's Dilemma, an example of a pure strategy would be choosing to confess. The player is in complete control of that decision and would not randomly choose to confess.

Definition 1.9 MIXED STRATEGY: *A mixed strategy, s_i , for a player, i , will be a collection of non-negative numbers which sum to 1 and is in one to one correspondence with its pure strategies. We represent mixed strategies as:*

$$s_i = \sum_{\alpha} c_{i\alpha} p_{i\alpha} \quad c_{i\alpha} \geq 0$$

with

$$\sum_{\alpha=1}^k c_{i\alpha} = 1,$$

where $p_{i\alpha}$ represents the i^{th} player's α^{th} strategy. In contrast to pure strategies, the player randomly chooses the strategy they play [18].

Mixed strategies are essentially weighted pure strategies, which we could also think of as linear combinations. We are taking some probability of using a pure strategy, $c_{i\alpha}$, and adding those weighted values together. These mixed strategies then can be treated as points in a k -dimensional simplex, which takes the form $C = \{c_{i1}p_{i1} + c_{i2}p_{i2} + \dots + c_{ik}p_{ik} \mid \sum_{\alpha=1}^k c_{i\alpha} = 1, c_{i\alpha} \geq 0\}$, with vertices that are the pure strategies. This simplex is a convex subset, a set that contains all of the line segments connecting any pair of points, of a real vector space, which allows us to use linear combinations for the mixed strategies. We will denote an n -tuple of mixed strategies for n players as s , which will be regarded as a point in vector space, the product space of the vector spaces containing mixed strategies. The set of all of those n -tuples is the product of the simplices representing the mixed strategies.

We will use the notation $\pi_i(s; r_i)$ to denote the payoff for the i^{th} player using a mixed strategy, r_i , from the mixed strategy set s .

Definition 1.10 EQUILIBRIUM POINT: An n -tuple $s = (s_1, s_2, \dots, s_n)$ of mixed strategies is an equilibrium point if and only if for every player i

$$\pi_i(s) = \max_{\forall r_i' \in S} [\pi_i(s; r_i')]$$

where the r_i 's represent mixed strategies.

In other words, an equilibrium point is an n -tuple s such that a player's mixed strategy maximizes their payoff if the strategies of the other players are fixed.

We say a mixed strategy uses a pure strategy if $s_i = \sum_{\beta} c_{i\beta} p_{i\beta}$ and $c_{i\alpha} > 0$. If $s = (s_1, s_2, \dots, s_n)$ and s_i uses $p_{i\alpha}$, then s uses $p_{i\alpha}$.

There are some conditions that a point s needs to meet to be considered an equilibrium point:

1. Because of the linearity of $\pi_i(s_1, s_2, \dots, s_n)$ in s_i ,

$$\max_{\forall r'_i s} [\pi_i(s; r_i)] = \max_{\alpha} [\pi_i(s; p_{i\alpha})].$$

Essentially, the maximum payoff over all of the mixed strategies available to a player equals the maximum payoff of the pure strategies.

2. Defining $\pi_{i\alpha}(s) = \pi_i(s; p_{i\alpha})$, we then can say

$$\pi_i(s) = \max_{\alpha} \pi_i(s; p_{i\alpha}).$$

Equivalently, the payoff of a mixed strategy for the i^{th} player equals the maximum payoff of using a pure strategy.

3. For the previous condition to hold, we must have $c_{i\alpha} = 0$ whenever $\pi_{i\alpha}(s) < \max_{\beta} \pi_{i\beta}(s)$. This means if the payoff for the i^{th} player switching to a pure strategy from a mixed one is less than the maximum payoff of a pure strategy, then the probability $c_{i\alpha} = 0$. Thus, s does not use $p_{i\alpha}$ unless it is an optimal pure strategy for player i . Therefore if $p_{i\alpha}$ is used in s , then $\pi_{i\alpha}(s) = \max_{\beta} \pi_{i\beta}(s)$.

To better illustrate what this definition entails, we will examine our known Nash equilibrium point from the Prisoner's Dilemma. We determined that the equilibrium would result in a payoff of -15 for both prisoners. We can think of "confess" and "stay silent" as pure strategies, which occur with probability 0 or 1, the players could utilize in the game. Let us assume that we do not know for sure that both players confessing is an equilibrium point.

Claim: $s = (1, 0, 1, 0)$ is a Nash equilibrium point for the Prisoner's Dilemma.

Proof. From Definition 1.10, we know an equilibrium point is where the payoff for the i^{th} player is maximized when the others are fixed. Therefore, assuming $(1, 0, 1, 0)$ is our point, the equilibrium will occur when:

$$\pi_1(s) = \max_p [\pi_1(p, (1-p), 1, 0)]$$

$$\pi_2(s) = \max_q [\pi_2(1, 0, q, (1 - q))]$$

where p and q represent mixed strategies. From our payoff matrix (Table 1.1), we can find the payoffs for each of these players:

$$\pi_1(p, (1 - p), 1, 0) = -15p - 16(1 - p) = -15p - 16 + 16p = p - 16$$

$$\pi_2(1, 0, q, (1 - q)) = -15q - 16(1 - q) = -15q - 16 + 16q = q - 16.$$

To find these payoffs, fix the opposing player's strategy so that they confess, and then look across the corresponding row/column in the payoff matrix to determine the coefficients on the strategies. We then want to solve

$$\pi_1(s) = \max_{p \in [0,1]} [p - 16]$$

$$\pi_2(s) = \max_{q \in [0,1]} [q - 16].$$

We know p and q are in the set $[0, 1]$ because mixed strategies must sum to one from the definition. To maximize our functions, we will want a higher p/q value since both of our functions are decreasing. This means that when $p = q = 1$, both equations will be maximized and we will have our maximum payoff. Thus,

$$\pi_1(s) = -15 \quad \pi_2(s) = -15,$$

which proves that $(1, 0, 1, 0)$ is a Nash equilibrium point for the Prisoner's Dilemma. \square

This also shows why condition one is true; since the pure strategies are the "edges" of the mixed strategy space, those are the places where the payoff functions would logically be maximized. Therefore, the maximum payoff achieved at some

other mixed strategy will be equivalent to the maximum payoff achieved with a pure strategy.

The last item we need is the following theorem.

Theorem 1.2 (Brouwer's Fixed Point Theorem).

For some compact, convex set S , and a continuous map $f : S \rightarrow S$, there is a point x_0 such that $f(x_0) = x_0$ [8].

1.2.2.2 PROOF OF NASH'S THEOREM

We can now prove Nash's theorem using the definitions we have established above.

Proof. We want to prove that every finite game has an equilibrium point. We will let s be an n -tuple of mixed strategies, $\pi_i(s)$ be the payoff for a player i using a mixed strategy, and $\pi_{i\alpha}(s)$ be the payoff for the i^{th} player if they change from a mixed strategy to the α^{th} pure strategy, $p_{i\alpha}$, while the other players continue to use their mixed strategies from s . To gauge whether it is in the player's best interest to switch to playing a pure strategy, we want to see the difference between the i^{th} player's payoff with a pure strategy and their payoff with a mixed strategy. We will therefore define a set of continuous functions of s :

$$\phi_{i\alpha}(s) = \max(0, \pi_{i\alpha}(s) - \pi_i(s)). \quad (1.1)$$

For each $s_i \in s$ we will create a modification s'_i :

$$s'_i = \frac{s_i + \sum_{\alpha} \phi_{i\alpha}(s) p_{i\alpha}}{1 + \sum_{\alpha} \phi_{i\alpha}(s)}, \quad (1.2)$$

and call s' the n -tuple $(s'_1, s'_2, \dots, s'_n)$. By creating this modification, we have created a set that is a linear transformation of our original set of n -tuples, s . We now want to show that the fixed points of the mapping $T : s \rightarrow s'$ is the equilibrium point.

Fixed Point to Equilibrium Point: We will first consider the case where we have any fixed point (n -tuple) s . If s is fixed under T , the proportion of $p_{i\alpha}$ used in s_i must not be decreased by T . This is because the maximum probability of using a strategy is 1, which means that when we add our modification based on the set of $\phi_{i\alpha}$, we must also divide by the sum of $\phi_{i\alpha}$ to normalize the elements of s' . Based on Equation 1.2, we can see for all pure strategies, β , $\phi_{i\beta}(s)$ has to be zero to prevent the denominator from being greater than one which would decrease the proportion by T . Therefore, if s is fixed under T , for any i and β , $\phi_{i\beta}(s) = 0$. Thus, no player can improve their payoff by moving to a pure strategy. This matches the criteria for being an equilibrium point under the definition.

Equilibrium Point to Fixed Point: Let s be an equilibrium point. Then from the definition, all ϕ 's equal 0, making s a fixed point under T .

Now that we have shown a strategy is an equilibrium point if and only if it is a fixed point under T , we will show there exists a fixed point of the map T . The space of n -tuples is the $n - 1$ dimensional probability simplex, which is closed, compact, and convex. Thus, the Brouwer Fixed Point Theorem requires T to have at least one fixed point s since we are mapping from the product space of the vectors containing the mixed strategies to itself. This means this fixed point must be an equilibrium point [18]. □

While Nash equilibria are useful for solving a variety of games, they may not always be the best method for finding a solution. This is because games can have multiple Nash equilibrium points and therefore have multiple solutions. A typical example of a game with multiple Nash equilibria is the Battle of the Spouses.

1.2.3 BATTLE OF THE SPOUSES

Battle of the Spouses is a coordination game, a game with more than one Nash equilibrium point [23], and is named after a, usually, minor conflict between a couple. For our example, spouse 1 would like to watch a mystery show in the evening while spouse 2 would like to watch a science fiction one. Each of their payoffs will be made up of two components: the payoff for watching television together, which will be a payoff of 2 for each of them and 0 otherwise, and a payoff of 1 if they watch their preferred program. If the couple does not watch a player of interest's preferred show, then that player gets a payoff of 0. The total payoff of a player will then be the addition of the payoffs from watching a program together and the payoff from watching their preferred program. The payoff table for this situation is below:

Table 1.2: Payoff Matrix for Battle of the Spouses

Decisions	Mystery	Science Fiction
Mystery	(3,2)	(1, 1)
Science Fiction	(0, 0)	(2,3)

While the upper right cell of the payoff matrix could match the lower left, we have a payoff of 1 for each player since they each got to watch their preferred program, they just did not watch it together. From the matrix, we can see that there are two Nash equilibria: when both choose to watch the mystery program and when both watch the science fiction one. In addition to these pure strategies, we can also find the mixed strategies for this game.

We will let the mixed strategy of spouse 1 be that they choose the mystery show with a probability of p and the science fiction show with a probability of $1 - p$. Spouse 2's mixed strategy will have them chose the mystery show with a probability

of q and the science fiction one with a probability of $1 - q$. Given spouse 2's mixed strategy, the payoff for spouse 1 will be:

$$\pi_1(\text{Mystery}, (q, 1 - q)) = 3q + 1(1 - q) = 2q + 1$$

$$\pi_1(\text{Science Fiction}, (q, 1 - q)) = 0(q) + 2(1 - q) = 2 - 2q.$$

So the equilibrium point would be where those two payoffs equal each other:

$$2q + 1 = 2 - 2q$$

$$4q = 1$$

$$q = \frac{1}{4}.$$

Similarly, the payoff for spouse 2 given spouse 1's mixed strategy will be:

$$\pi_2((p, 1 - p), \text{Mystery}) = 2p + 0(1 - p) = 2p$$

$$\pi_2((p, 1 - p), \text{Science Fiction}) = 1(p) + 3(1 - p) = 3 - 2p.$$

$$2p = 3 - 2p$$

$$4p = 3$$

$$p = \frac{3}{4}.$$

Thus the mixed strategy Nash equilibrium will be $(\frac{3}{4}, \frac{1}{4})$ for the first spouse and it will be $(\frac{1}{4}, \frac{3}{4})$ for the second. However, we still have the same issue of multiple Nash equilibria. In this type of situation, players can not easily determine a strategy to use since they would need to be able to decide between the Nash equilibria. Attempts to

solve the problem with multiple equilibrium points have used cultural assumptions, redefinitions of Nash equilibria to eliminate some points, and evolution.

For the Battle of the Spouses, we could assume one partner would be more accommodating to the other as a cultural assumption, but we would still have the problem of choosing that as our assumption. How do we pick which cultural assumption to use? Redefining Nash equilibria has had similar issues where the attempts to do so have not been able to eliminate all but one of the Nash equilibria. The evolutionary concept approaches this problem from a different direction, assuming there is a population of players who pair up to play each other in the game. As more players play against each other, the proportion of players using any given strategy will change over time depending on how successful the strategy is [8]. That is where evolution comes in since it is akin to the idea of the “survival of the fittest” with “the fittest” referring to the most successful strategies. This approach is not without its flaws and will be discussed further in Chapter 2.

1.2.4 ROCK, PAPER, SCISSORS

Along with our previous examples, we can also have a case where a player’s choice results in them gaining exactly what their opponent loses. These types of games are called zero-sum games.

Definition 1.11 ZERO-SUM GAME: *A game in which the payoff of the players (for any strategy) adds up to 0 [8].*

As an example, suppose we have two players playing a game of Rock, Paper, Scissors where the winner wins \$15 dollars from the loser. The payoff matrix of this game will be Table 1.3.

In this case, we can see that the best outcome for one player would be one in

Table 1.3: Payoff Matrix for Rock, Paper, Scissors

Decisions	Rock	Paper	Scissors
Rock	(0,0)	(-15, 15)	(15, -15)
Paper	(15, -15)	(0,0)	(-15, 15)
Scissors	(-15, 15)	(15, -15)	(0,0)

which they ended with 15 and their opponent ended with -15. Thus, the sum of the payoffs equals 0 and Rock, Paper, Scissors is a zero-sum game.

Looking at Table 1.3, the solution is not as clear as the previous examples we have discussed; a player either loses everything, gains everything, or ties, so trying to determine what their opponent will do is difficult. There is not a pure strategy solution which would result in both players adding to their payoff. We can approach the solution of this game via two methods: finding the game's Nash equilibrium or finding the maximin solution. The maximin solution has two components; in the first, the player must find their minimum payoff in each strategy and in the second, they determine which action to take to make sure that small payoff value is the highest available. We can actually use the Nash equilibrium as a way to build the maximin solution.

Let the payoffs of our Rock, Paper, Scissors game be defined as $\pi(\sigma_1, \sigma_2) = \pi_1(\sigma_1, \sigma_2)$ and $-\pi(\sigma_1, \sigma_2) = \pi_2(\sigma_1, \sigma_2)$. Recall the conditions for a Nash equilibrium defined in Definition 1.6,

$$\pi_1(\sigma_1^*, \sigma_2^*) \geq \pi_1(\sigma_1, \sigma_2^*) \quad \forall \sigma_1 \in \Sigma_1$$

$$\pi_2(\sigma_1^*, \sigma_2^*) \geq \pi_2(\sigma_1^*, \sigma_2) \quad \forall \sigma_2 \in \Sigma_2.$$

These conditions can be rewritten in the following form:

$$\pi(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1 \in \Sigma_1} \pi(\sigma_1, \sigma_2^*)$$

$$\pi(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2 \in \Sigma_2} \pi(\sigma_1^*, \sigma_2).$$

By rewriting the conditions in this format, we can more easily see that in order to win the game, a player should play a best response to their opponent's action.

For our maximin solution, we need to keep in mind that the player of interest must find their highest, minimum-valued payoff and take actions to ensure that the strategy pair which results in that gets played. With that established and the rewritten Nash equilibrium conditions, we can combine our redefined Nash equilibrium into one item:

$$\begin{aligned} \pi(\sigma_1^*, \sigma_2^*) &= \max_{\sigma_1 \in \Sigma_1} \pi(\sigma_1, \sigma_2^*) \\ &= \max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} \pi(\sigma_1, \sigma_2). \end{aligned}$$

The maximin solution can then be thought of as a saddle point where a player of interest determines the worst possible outcome and then chooses the action which makes that outcome the highest possible. To illustrate this, we will examine a non-zero sum game so we can focus on what this process looks like with pure strategies. Let us consider a payoff matrix: Table 1.4. We will let player 1, the row

Table 1.4: Maximin Example

Decisions	Strategy 1	Strategy 2
Strategy 1	(3,1)	(2, -20)
Strategy 2	(5, 0)	(-10,1)

player, be our player of interest. Looking at the payoffs, they could get 3 or 2 using strategy 1 or they could get 5 or -10 playing strategy 2. The worst outcome for player 1 would be if they played strategy 2 and their opponent played strategy 2 since that would result in a payoff of -10. We would therefore not consider strategy 2 as a viable option. This means player 1 will play strategy 1 since the worst payoff

they could receive would be 2. By playing this strategy, that will force player 2 to play strategy 1 as well since, if they do not, their payoff would be -20. Thus, the maximin solution for this game is (3,1).

If no saddle point exists and the game includes some randomness (a player can choose what percentage of the time they will play a pure strategy), then the maximin strategy relies on probability, similar to what we saw with mixed strategies. The reliance on probability comes from the lack of pure strategy equilibria. In our game of Rock, Paper, Scissors, we can see that there is no pure strategy equilibria. A losing player could always change their strategy and win back the money they had lost. Thus, players need to play a mixed strategy to avoid having their strategy exploited by their opponent.

We will let the probability that a player chooses rock be p , paper be q , and scissors be $1 - p - q$. Player 1's payoff for each decision will be:

Rock: $p(0) + q(-15) + (1 - p - q)(15) = -15p - 30q + 15$

Paper: $p(15) + q(0) + (1 - p - q)(-15) = 30p + 15q - 15$

Scissors: $p(-15) + q(15) + (1 - p - q)(0) = -15p + 15q.$

Player 2's payoffs will be the exact same as player 1's. Since we are looking for an equilibrium point, we will set each of these payoffs equal to each other

$$-15p - 30q + 15 = 30p + 15q - 15 = -15p + 15q.$$

After some algebraic manipulation, we find that $q = \frac{1}{3}, \frac{2}{3} - p$ and $p = \frac{1}{3}$. While q has two values, since $p = \frac{1}{3}$, both values will be $\frac{1}{3}$, thus the mixed strategy will be $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

1.2.5 CONCLUSION


We can think of many situations in the same framework as static games, allowing us to use the concepts we have discussed in this section. We discussed a few different types of games, those that are more straight-forward like the Prisoner's Dilemma where there is one solution, others like the Battle of the Spouses which have multiple solutions, and Rock, Paper, Scissors, which rely on mixed strategies to find solutions. The most important item we discussed though, was the Nash equilibrium. Nash equilibria are the foundation of solutions for all types of games housed under game theory. When we discuss how to apply game theoretic models in risk management and cybersecurity, we will be looking for strategies which result in an equilibrium so knowing the different variations that we have seen in this section will give us the context we need to understand how different scenarios might play out. We also discussed Nash's theorem which will allow us to continue finding solutions in every game with a finite strategic form, such as in finite dynamic games.

1.3 FINITE DYNAMIC GAMES

In contrast to static games where players make simultaneous decisions, in dynamic games players are able to make decisions at various times knowing at least some of the earlier decisions. These games are then useful for modeling situations where players can be reactionary and determine the best courses of actions based on the knowledge they possess. Dynamic games lend themselves more easily to cybersecurity and risk management than static games since attackers or companies would be implementing solutions to problems knowing the past decisions they have made. Since this is a new form of game, we also represent the payoffs of players differently. Instead of a matrix, we will primarily use trees since that will give us the freedom to make decisions at different points in time. Trees consist of

decision nodes, points where a player makes a decision between some number of options, and will end in terminal nodes where no more decisions can take place. To find solutions for these trees, we can use a process called backward induction where we essentially look at the end solutions with the worst payoffs, eliminate them from our consideration, and continue to move back up the tree following the same method. A basic example of a tree is presented below.

Figure 1.1: Clothes Decision Tree



We can still create payoff matrices for these games, but the use of a tree highlights what information a player has at a point in the game. They therefore emphasize the difference between static and dynamic games better than a "traditional" payoff matrix would.

We will formally define what information a player knows at a given point in time by using information sets.

Definition 1.12 INFORMATION SET: *An information set is a set of decision nodes in a game tree such that only the player concerned is making a decision and the player does not know which node has been reached, they only know it is one of the ones in the set [5].*

Essentially, information sets allow us to understand what options are available to a player at a decision node. Once the player reaches the next decision node in the game tree, they may not know what action their opponent took before them so

any of their options are valid. This somewhat ties back to the idea of simultaneous decision making from static games since a player has to rationalize the decisions their opponent would make to reach a conclusion about what their own decision should be.

Information sets allow us to categorize games as perfect information or imperfect information games.

Definition 1.13 PERFECT INFORMATION GAME: *A perfect information game is one in which each information set, of which there will be multiple since they will be the set of decisions available to a player when they go to take their turn, contains one node.*

Definition 1.14 IMPERFECT INFORMATION GAME: *An imperfect information game is one in which each information set, of which there will be multiple since they will be the set of decisions available to a player when they go to take their turn, contains multiple nodes [5].*

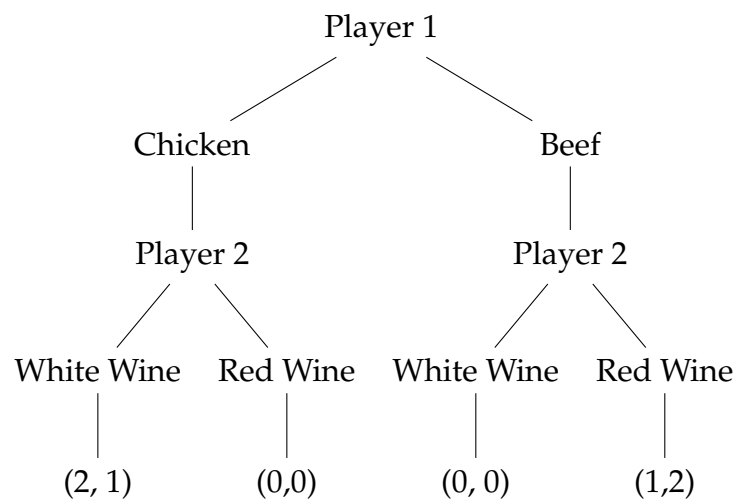
In perfect information games, a player knows all of the decisions made at every decision node and can therefore make an educated choice on how to proceed. In imperfect information games, there is a measure of uncertainty since a player may not know which decision their opponent made. This uncertainty allows us to use mixed strategies as we previously did with static games. We can also use behavioral strategies where the opportunity for randomization occurs at each information set.

1.3.1 DINNER PARTY GAME

As an example of a finite dynamic game, we will consider the Dinner Party game. In this situation, two players are hosting a dinner party; one player chooses the protein for the meal and the other responds to that decision by choosing an appropriate wine. We will let player 1 pick the protein, either chicken or beef, and player 2 pick the wine, white or red. In this case, player 1 prefers chicken and player 2 prefers beef. We will therefore define the payoffs as 2 if the pair is correct (chicken with

white wine and beef with red) and it is the player's preferred protein, 1 if the pairing is correct but it is not the player's preference, and 0 if it is the wrong pairing. What makes this situation dynamic is that player 2 needs to respond to player 1's actions in order to create a successful pairing between the wine and protein. The game tree for this situation is in Figure 1.2.

Figure 1.2: Dinner Party Game Tree



We will treat this game as though it is one with perfect information; our information sets will be the actions available to player 1, choosing either chicken or beef, and our information set for player 2 will be choosing either white wine or red wine. In this case we only have two information sets since this is a small game. In larger games, we would have more information sets since they would correspond to the actions available to a player at each point in our game tree. By treating this game as though it has perfect information, player 2 will know what player 1 bought before making their decision. Using backward induction, if player 1 buys chicken, then player 2 would choose white wine since that will get them a payoff of 1 as opposed to 0. If player 1 chose beef, then player 2 would choose red wine since that would give them a payoff of 2 as opposed to 0. Considering player 1's perspective, the best choice they could make would be to buy chicken since they

would receive a payoff of 2 since it would be in their opponent's best interest to choose white wine. This solution is a Nash equilibrium. However, similarly to the Battle of the Spouses, there are multiple Nash equilibria points. One occurs when player 1 chooses chicken and player 2 chooses white wine and the other occurs when player 1 chooses beef and player 2 chooses red wine. In both cases, the total payoff of the game is the same.

1.3.2 REFINING NASH EQUILIBRIA

As we discussed previously, one of the challenges of Nash equilibrium solutions is the possibility that there will be more than one in a game. To find one solution for a game, we then need to find ways to refine the concept of Nash equilibria or use cultural assumptions/other methods to eliminate solutions. In finite dynamic games, one of the ways mathematicians have tried to refine this concept has been through subgames.

Definition 1.15 SUBGAMES: *Subgames are parts (called sub-trees) of a game tree that satisfy three conditions.*

1. *A subgame begins at a decision node for any player.*
2. *The player knows all of the decisions that have been made previously.*
3. *The sub-tree contains every decision node that occurs after the initial node and no others [5].*

We can then define a new type of Nash equilibria called subgame perfect Nash equilibria.

Definition 1.16 SUBGAME PERFECT NASH EQUILIBRIA: *Subgame perfect Nash equilibria are Nash equilibria in which the specified behavior in every subgame is a Nash equilibrium point for the subgame [5].*

In our previous Dinner Party game, there are two subgames: the parts of the tree where the second player makes a decision and the whole game. At each of the decision nodes for the second player, there would only be one decision made (beef or chicken) up to that point. Each of the following decision nodes will follow the initial decision node of player 1 with no overlap since player 1 can not buy both beef and chicken. Thus, the sub-tree of player 2's decision meets the definition of a subgame. The whole game also meets the criteria of the definition for a subgame since it begins with player 1's decision, no decisions were made previous to player 1's move, and the game contains all of the nodes that stem from player 1's decision.

Using the same example, we can find the subgame perfect Nash equilibria. In the subgame where the first player chooses beef, the Nash equilibrium for the subgame would occur when the second player chooses red wine. Similarly, when the first player chooses chicken, the Nash equilibrium for the subgame would occur if the second player chooses white wine. This example shows that, if we utilize backward induction, any Nash equilibrium in a finite dynamic game will be a subgame perfect one. This then leads us to Theorem 1.3.2.

Theorem 1.3 (Existence of Subgame Perfect Equilibria).

Every finite dynamic game has a subgame perfect Nash equilibrium [5].

To prove this, we will use Theorem 1.1 and Definition 1.15.

Proof. We would like to show that every finite dynamic game has a subgame perfect Nash equilibrium. From Theorem 1.1, we know that every finite game has at least one mixed-strategy Nash equilibria. Thus, any finite dynamic game will have at least one mixed-strategy Nash equilibrium point.

We now want to show that that mixed-strategy equilibrium point is a subgame perfect equilibrium. To do this, we first need to show that a finite dynamic game will be a subgame of itself. This will allow us to move from the overall game's Nash equilibria to the subgame perfect one. Since we are looking at finite dynamic games,

we know that we can represent the decisions available to a player as a game tree. Each game tree must have some starting decision since the player who goes first has to make some kind of choice that another player responds to. We have therefore met the first condition of Definition 1.15. Since a player will respond to the decision of the first, they will know what decision has already been made. They will therefore know all of the previous decisions made in the game, even with higher numbers of players. This is because every player will respond to what their opponents have decided and will therefore need to know how decisions have impacted the game play. Thus the second condition of the subgame definition has been met. Lastly, since every decision will stem from the first one made, the sub-tree will contain the decision nodes that follow the first decision and no extraneous ones.

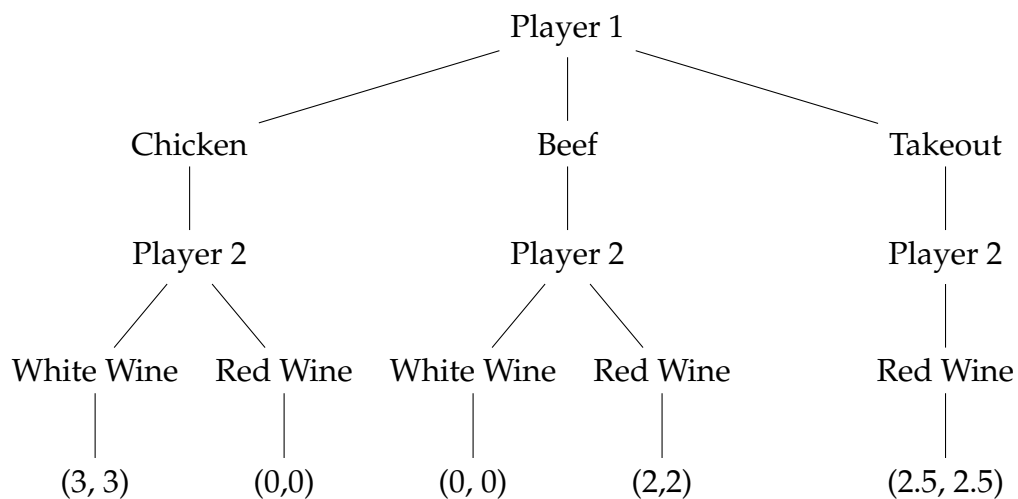
Thus, we know that a game is a subgame of itself and that there will be at least one mixed-strategy equilibrium point. From Definition 1.16, we know that a Nash equilibrium point in which the specified behavior in every subgame is also an equilibrium for the subgame is subgame perfect. Thus, the mixed-strategy Nash equilibrium is an equilibrium point for every subgame and meets Definition 1.16. Thus, every finite dynamic game will have a subgame perfect Nash equilibrium. \square

Subgame perfection has been one of the methods suggested for refining the Nash equilibria. As we discussed earlier in Section 1.2.3, mathematicians have attempted to narrow or tweak the definition of a Nash equilibrium point in an effort to reduce the number of them present in a game. While subgame perfection is a way to possibly limit the number of Nash equilibria present in a finite dynamic game, there is still the problem of determining how many Nash equilibria can reasonably be in a game and that the extra conditions introduced by subgame perfection may still leave multiple equilibria points. The first problem of having a reasonable amount of equilibrium points is hard to get rid of since the models produced for different situations will lead to variation in how many points we could expect.

In addition to subgame perfection, the use of forward induction has been introduced as a way to limit the number of equilibrium points. As the name suggests, this is another way to approach solving game trees where instead of moving backward through the tree to reach a desired payoff, one would move forward toward the specific payoff. We flip our understanding of rationality in forward induction by assuming past decisions were rational while in backward induction, we assume future decisions would be rational.

For instance, in a new Dinner Party game, represented by Figure 1.3, our first player has three options: chicken, beef, or ordering takeout. We will let the payoff of player 1 choosing beef and player 2 choosing red wine be (2,2) and if player 1 decides to buy chicken and player 2 pairs it with white wine, the payoff will be (3,3). The payoff for ordering takeout will be (2.5, 2.5) and will only have a red wine pairing.

Figure 1.3: Dinner Party Game Tree with Three Starting Options



In this instance, we have three Nash equilibria: (Chicken, White Wine), (Beef, Red Wine), and (Takeout, Red Wine). These are all Nash equilibria since player 2 can not do better than a payoff of 3, 2, or 2.5 after the decision made by player 1. Using forward induction, we can eliminate two of these equilibrium points. If

player 1 acts rationally, they will never choose to buy beef since they can at most get a payoff of 2, which is worse in comparison to chicken or takeout. Similarly, after eliminating beef the best player 1 could do is either 3 or 2.5 so they would logically choose chicken since that maximizes their payoff. This in turn means player 2 would choose white wine since that will maximize their payoff, leaving us with only one Nash equilibrium.

Forward induction seems like a viable solution to the problem of multiple Nash equilibria, however, the rationality assumption introduces a problem. By assuming rationality, we may assume that the players would reach a part of the game tree that they logically would not reach if they acted in a way prescribed by the equilibrium. They would therefore be exhibiting irrational behavior and we could not definitively establish that rationality would reappear in the game. So while forward induction may help us limit equilibrium points in particular instances, they are not the “ultimate” way to refine Nash equilibria.

1.3.3 CONCLUSION

Like static games, finite dynamic games allow us to model a variety of situations in different contexts. However, finite dynamic games allow us to account for decisions made by players at different points in time, which is more realistic to the situations that occur in risk management and cybersecurity. The idea of perfect and imperfect information also brings reality into the theoretical framework of game theory since we can account for gaps in different players' information sets which would explain seemingly odd decisions. The way we solve finite dynamic games also introduces some ways in which we can limit the amount of Nash equilibria present in a scenario. While those methods have their flaws, they can be useful in specific situations and are a good tool to keep in mind when considering different game trees.

CHAPTER 2

EVOLUTIONARY GAME THEORY

An offshoot of traditional game theory, evolutionary game theory is a field that specifically focuses on the strategic aspect of evolution. It has primarily been used in biologic and social scientific contexts, however, we could extend this to fit cybersecurity and risk management. Evolutionary game theory originally developed as an application of game theoretic techniques, much like the ones discussed last chapter, in the context of populations. R. A. Fisher was the first to use what would later become evolutionary game theory when he tried to determine why there was approximate equality in the sex ratios of mammals [10]. His work could be understood from a game theory perspective, but, he did not explicitly discuss it through that lens. As the field gained more traction with the works of R. C. Lewinton and Maynard Smith, it was seen as a possible avenue for alleviating the problems we have encountered with Nash equilibria.

2.1 EVOLUTIONARY STABILITY

From our work in Section 1.2.3, we know that the existence of Nash equilibria does not guarantee unique equilibrium points. One of the ways to remedy the issues caused by multiple equilibria, namely the difficulty for players to choose the “best”

decision, was to strengthen the definition. This stronger understanding of strategies came from evolutionary game theory and was focused on evolutionary stability.

Definition 2.1 EVOLUTIONARY STABLE STRATEGIES (ESS): *A strategy σ is an evolutionary stable strategy if and only if for all other strategies $\mu \neq \sigma$ either $\pi(\sigma|\sigma) > \pi(\mu|\sigma)$ or $\pi(\sigma|\sigma) = \pi(\mu|\sigma)$ and $\pi(\sigma|\mu) > \pi(\mu|\mu)$ [10].*

In this definition, the notation $\pi(\text{strategy 1}|\text{strategy 2})$ denotes the payoff obtained for a player playing strategy 1 against an opponent playing strategy 2. This is similar to how we denoted payoffs in Chapter one, we have just changed the notation slightly to get at the “playing against” concept more clearly. We could alternatively define ESSs as

Definition 2.2 ALTERNATE DEFINITION OF ESS: *A strategy σ is an evolutionary stable strategy if and only if for all other strategies*

1. $\pi(\sigma|\sigma) \geq \pi(\mu|\sigma)$
2. *If $\pi(\sigma|\sigma) = \pi(\mu|\sigma)$, then $\pi(\sigma|\mu) > \pi(\mu|\mu)$ [10].*

Essentially, evolutionary stable strategies are ones in which a player’s payoff is better than if they had played some alternate strategy, given any strategy played by their opponent. While this may seem similar to the definitions of Nash equilibria we previously discussed (Definition 1.6), those definitions were not strong enough to capture the concept of evolutionary stability. Nash equilibria will not allow for any players to do better than the payoff given by the equilibrium, however, there could be a point where a player could divert from the equilibrium strategy and receive the same payoff.

As an example, we will assume that there exists some population where all individuals follow the same strategy (called the incumbent population). We will call this strategy σ_1 and we will assign a payoff of 1. We will introduce a mutant to

this population who plays strategy σ_2 with payoff of 1 against those playing σ_1 and 2 against those playing σ_2 . The payoff matrix is shown in Table 2.1.

Table 2.1: Payoff Matrix for Population of Incumbents and Mutants

Strategy	σ_1	σ_2
σ_1	(1,1)	(1, 1)
σ_2	(1,1)	(2, 2)

From the matrix, we can identify two pure strategy Nash equilibria, one corresponding to (σ_1, σ_1) and another to (σ_2, σ_2) . Since we have two equilibria, it is hard to determine which strategies would be the best for our players to play. Going back to the alternate definition of an ESS, Definition 2.2, we can check to see which of our equilibria is evolutionarily stable. For (σ_1, σ_1) , we know that both players are using the same strategy, so the payoff must be greater than or equal to the payoff obtained from player 1 playing σ_2 against σ_1 . We are therefore comparing a payoff of (1,1) to a payoff of (1,1), meaning that we will now check the second condition of the definition since the payoffs are equivalent. We now need the payoff of player 1 playing σ_1 against σ_2 to be greater than the payoff obtained from both players playing σ_2 . Our payoffs will be (1,1) for $(\sigma_1|\sigma_2)$ and (2,2) for $(\sigma_2|\sigma_2)$. The first Nash equilibrium, (σ_1, σ_1) , is not evolutionarily stable.

Following a similar process for our second equilibria, (σ_2, σ_2) , we can see the payoff of playing that strategy will be greater than the payoff received from playing $(\sigma_1|\sigma_2)$, meaning that the second equilibrium point meets the criteria for evolutionary stability. We could use that fact to eliminate (σ_1, σ_1) from our solutions since we have a solution that is a stable equilibrium point. In terms of what this means in context, it essentially tells us that bringing in mutants to the incumbent population would cause the number of individuals using the incumbent strategy to disappear. This will in turn generate more mutants playing a mutant strategy, which will cause

the mutant group to dominate the incumbent group. This idea will reappear when we discuss dynamics later in the chapter.

We then could try to restrict the definition of Nash equilibria to avoid situations like the previous example, for instance by focusing on strict Nash equilibria. That is a situation where deviation from the equilibrium leads to a worse outcome for the player. However, that approach is generally too strong to capture evolutionary stability, as was shown by Smith and Price in “The Logic of Animal Conflict” [10]. They focused on the Hawk-Dove game, which is a game where individuals compete to gain some resource with a fixed value.

2.1.1 HAWK-DOVE GAME

We will define the hawk player as one who initiates aggressive behavior and does not stop until they are injured or their opponent backs down. The dove player will retreat if their opponent exhibits aggressive behavior. In addition to these player definitions, we will also define some assumptions for the game:

1. When both players initiate aggressive behavior, there will eventually be a conflict resulting in an injury. The injury will have an equally likely probability of happening to either player.
2. The conflict will reduce an injured player's fitness by some constant value referred to as C .
3. When a hawk player meets a dove, the dove will retreat and the hawk will acquire the resource.
4. When two dove players come in contact, they will share the resource equally.
5. The fixed value of the resource, V , is assumed to be less than C .

From these assumptions and the player types, we can generate a payoff matrix. Since we have two player types and this will be a two-person game, we know our potential pairs will be hawk-hawk, hawk-dove, dove-hawk, and dove-dove.

In the hawk-hawk case, both players will initiate aggressive action, which, from assumption one, we know will result in injury, so each player's payoff will be $\frac{V-C}{2}$. Each hawk player has an equally likely reduction in fitness (C) so even if they are able to retrieve the resource, we need to account for that complication in the payoff. Each player also has an equally likely chance to retrieve the resource which is why we divide by 2.

In the case of a dove-hawk or hawk-dove pairing, we know from the definition of a dove and assumption three that the dove will automatically retreat. Thus, the dove's payoff will always be 0 in that case and the hawk's will be V . From assumption four, we know in a dove-dove pair that the resource will be shared equally so the payoff for both players will be $\frac{V}{2}$.

Based on this, we can generate payoff matrix Table 2.2.

Table 2.2: Payoff Matrix for Hawk-Dove Game

Strategy	Hawk	Dove
Hawk	$(\frac{V-C}{2}, \frac{V-C}{2})$	$(V, 0)$
Dove	$(0, V)$	$(\frac{V}{2}, \frac{V}{2})$

The Hawk-Dove game does not have any pure strategy Nash equilibria. A hawk player will always do better against a dove and a dove will always do better cooperating with another dove. However, if a player were to play a dove strategy, the risk of their opponent playing a hawk strategy would be too great to take. A hawk player would not want to play another hawk player since they would be losing a good amount of the resource to potential injury so there is not a pure strategy equilibrium here. However, from Nash's Theorem 1.1, we know that there will be a mixed strategy Nash equilibrium. Following a similar algebraic procedure to what we did in Section 1.2.3, we can find the probabilities of playing hawk (probability labeled p) and dove (probability labeled $1 - p$). To be a Nash equilibrium, like we saw in Section 1.2.3, the following should be equal:

$$\begin{aligned}
p\left(\frac{V-C}{2}\right) + V(1-p) &= p(0) + (1-p)\left(\frac{V}{2}\right) \\
\frac{pV - pC - 2Vp + pV}{2} &= \frac{V}{2} - V \\
-pC &= -V \\
p &= \frac{V}{C}
\end{aligned}$$

Thus, the mixed strategy equilibrium will occur when hawk is played with probability $\frac{V}{C}$ and dove is played with probability $1 - \frac{V}{C}$. We will denote this strategy using σ .

While we know that the above strategy results in an equilibrium, to understand why it is also an ESS, we will use the fundamental theorem of mixed strategy equilibria:

Theorem 2.1 (Fundamental Theorem of Mixed Strategy Equilibria).

Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a mixed strategy profile for an n -player game. Let σ_{-i} represent the mixed strategies played by all other players in a game for any player $i = 1, \dots, n$. A strategy σ is a Nash equilibrium if and only if for any player, $i = 1, \dots, n$ with pure strategy set S_i ,

1. If $s, s' \in S_i$ occur with positive probability in σ , then the payoffs to s and s' , when played against σ_{-i} , are equal.
2. If s occurs with positive probability in σ_i and s' occurs with zero probability in σ_i , then the payoff to s' is less than or equal to the payoff to s when played against σ_{-i} [15].

This theorem gives us further insight into how to find Nash equilibria, which we know will also help us find ESSs since an ESS will be an equilibrium point. A proof of the theorem can be found in [15]. From the first condition, we know that strategies in the pure strategy set played with positive probability for a particular player will be equal to the payoff obtained from the pure strategies played by an opponent. For the Hawk-Dove game, that means $\pi(\text{Hawk}|\sigma) = \pi(\text{Dove}|\sigma) = \pi(\sigma|\sigma)$.

Following from this, we can say that any other mixed strategy μ will have the same payoff as σ : $\pi(\mu|\sigma) = \pi(\sigma|\sigma)$. This means the equilibrium is not strict, but the payoff from using this incumbent strategy will be higher when played against a mutant strategy than the payoff from the mutant strategy played against itself [10]. From this example, we can see the benefit of an ESS since it will be stronger than strategies that result in Nash equilibria, but it will not be as restrictive as a strict equilibrium.

2.1.2 ROCK, PAPER, SCISSORS REVISITED

Referring back to Definitions 2.1 and 2.2, we know that every ESS will result in a Nash equilibrium point and Definition 2.1 specifically tells us that a strict Nash equilibrium is evolutionary stable. This is because the ESS will result in a better payoff than other strategies given what one's opponent does, which is the central idea of a Nash equilibrium. The extra criterion to be evolutionarily stable then strengthens the Nash equilibrium because ESSs may not exist in every game. We know from Theorem 1.1 that we will always have at least one Nash equilibrium resulting from a mixed strategy so if the equilibrium is also an ESS, that could help us determine the best solution in a situation where multiple exist since an ESS is a stronger concept.

An example of a game without an ESS is Rock, Paper, Scissors, which we have already discussed in Section 1.2.4. We found that there was no pure strategy equilibrium for the game, but, there was a mixed strategy where rock, paper, and scissors were each played a third of the time. Our payoffs were Table 2.3.

Table 2.3: Payoff Matrix for Rock, Paper, Scissors

Decisions	Rock	Paper	Scissors
Rock	(0,0)	(-15, 15)	(15, -15)
Paper	(15, -15)	(0,0)	(-15, 15)
Scissors	(-15, 15)	(15, -15)	(0,0)

The payoff we would get from playing the Nash equilibrium strategy was 0 for rock, paper, and scissors. From Definition 2.1, we know to have an ESS, $\pi(\sigma|\sigma) > \pi(\mu|\sigma)$ or if $\pi(\sigma|\sigma) = \pi(\mu|\sigma)$, $\pi(\sigma|\mu) > \pi(\mu|\mu)$. We therefore are looking at comparisons between the payoff of our mixed strategy and the payoff of the pure strategies.

We will define $\pi(\sigma|\sigma)$ as the payoff of the mixed strategy. We know this payoff is 0 since that is what we found in Section 1.2.4. To meet the definition of an ESS, we want to check the equality of payoffs with different strategy pairs. We will choose rock as our example so $\pi(Rock|\sigma)$ will be the quantity we are interested in. If a player plays rock, their payoff will be $0 + -15 + 15$ (row player) or $0 + 15 + -15$ (column player) both of these equal 0 so the payoff is $\pi(Rock|\sigma) = 0$. Since $\pi(\sigma|\sigma) = \pi(Rock|\sigma)$, we will need $\pi(\sigma|Rock) > \pi(Rock|Rock)$ to meet the definition. The payoff of the first quantity will still be 0 and from the payoff matrix, we can see that the payoff from both playing rock will also be 0. Since the mixed strategy does not meet the requirements to be an ESS, Rock, Paper, Scissors does not have an evolutionarily stable strategy.

2.1.3 EQUIVALENT CONCEPTS TO EVOLUTIONARY STABILITY

Two other concepts that are equivalent to evolutionary stability are the idea of an invasion barrier and local superiority. The first concept relates to the mutants and incumbents we explored earlier in the section. We define σ and μ to be two strategies and we will let ε represent the probability of playing σ and $1 - \varepsilon$ represent the probability of playing μ . In this case, we are looking solely at mixed strategies so ε can only take on values between 0 and 1 exclusively. We can then define a mixed strategy as $\varepsilon\mu + (1 - \varepsilon)\sigma$ where the player plays μ with a probability ε and σ with a probability $(1 - \varepsilon)$. In addition to this interpretation, we could also view that strategy definition as a strategy used by a randomly selected individual in a

population where the majority of players play the $(1 - \varepsilon)$ strategy and the minority play the ε strategy [10]. This view of the strategy is beneficial since it allows us to look at variants and how the incumbent population handles their introduction. This is where the idea of an invasion barrier comes in; it is essentially the threshold that the mutant population would need to reach in order to get a higher expected fitness (a more biological concept).

Definition 2.3 UNIFORM INVASION BARRIER: *A strategy σ will have a uniform invasion barrier if there exists some $\bar{\varepsilon} > 0$ such that for all $\mu \neq \sigma$ and $\varepsilon \in (0, \bar{\varepsilon})$,*

$$\pi(\sigma|\varepsilon\mu + (1 - \varepsilon)\sigma) > \pi(\mu|\varepsilon\mu + (1 - \varepsilon)\sigma) \quad [10].$$

This definition essentially states that when a large population all follows the same strategy, σ , except for a small group who follow a single other strategy, μ , the strategy followed by the majority will have a strictly higher expected fitness level than the strategy followed by the minority. In a more grounded context, we would expect that if a large population followed the same strategy and a small group started following a different one that the majority strategy would dominate the minority one in the population. This logically makes sense when thinking in biologic terms since a mutant population would need a significant amount of individuals passing on a mutated gene to dominate the original, non-mutated group. The strategy σ would then be evolutionarily stable with a sufficiently large population and a sufficiently small group of mutants.

The second concept deals more with the possibility of “drifting” away from a Nash equilibrium; it hones in on the problem where a player could move away from the equilibrium strategy and receive the same payoff with a different strategy.

Definition 2.4 LOCAL SUPERIORITY: *A strategy σ is considered to be locally superior if*

there is a neighborhood, U , around σ such that for all strategies $\mu \in U$, where $\mu \neq \sigma$, it is the case that $\pi(\sigma|\mu) > \pi(\mu|\mu)$ [10].

This definition deals with the problem of drifting by defining σ to be the strategy that results in the highest payoff when played against another strategy in the neighborhood around it. In this case, we are using the term neighborhood to define a group of strategies.

These two concepts and evolutionary stability are then equivalent, and we will prove that an ESS will have a uniform invasion barrier. Proofs of the other directions can be found in [16] and [19].

Theorem 2.2 (Equivalence of Evolutionary Stability).

The following statements are equivalent:

1. σ is an evolutionary stable strategy.
2. σ has a uniform invasion barrier.
3. σ is locally superior.

Proof. We want to prove that an evolutionary stable strategy, σ , will have a uniform invasion barrier. We will therefore assume that σ is an ESS. From Definition 2.3, we know for σ to have a uniform invasion barrier, there needs to exist some $\bar{\varepsilon} > 0$ such that for all $\mu \neq \sigma$ and $\varepsilon \in (0, \bar{\varepsilon})$, $\pi(\sigma|\varepsilon\mu + (1 - \varepsilon)\sigma) > \pi(\mu|\varepsilon\mu + (1 - \varepsilon)\sigma)$. We can think of this statement as:

$$\pi(\sigma|\varepsilon\mu + (1 - \varepsilon)\sigma) > \pi(\mu|\varepsilon\mu + (1 - \varepsilon)\sigma) \quad (2.1)$$

$$\Leftrightarrow \pi(\sigma|\varepsilon\mu + (1 - \varepsilon)\sigma) - \pi(\mu|\varepsilon\mu + (1 - \varepsilon)\sigma) > 0 \quad (2.2)$$

$$\Leftrightarrow \varepsilon(\pi(\sigma|\mu)) + (1 - \varepsilon)(\pi(\sigma|\sigma)) - (\varepsilon(\pi(\mu|\mu)) + (1 - \varepsilon)(\pi(\mu|\sigma))) > 0. \quad (2.3)$$

We can factor out ε and $(1 - \varepsilon)$ in Equation 2.3 because those values are simply telling us how often the strategy is played; the payoff will not be affected by how often a strategy gets played, it depends on what strategies are chosen for play. We can then continue manipulating Equation 2.3:

$$\Leftrightarrow \varepsilon(\pi(\sigma|\mu) - \pi(\mu|\mu)) + (1 - \varepsilon)(\pi(\sigma|\sigma) - \pi(\mu|\sigma)) > 0. \quad (2.4)$$

In a game, we will know the payoffs for each strategy pair, hence the only unknown quantity in this equation will be ε . We can therefore define a function of ε , $f(\varepsilon)$ as

$$f(\varepsilon) = \varepsilon(\pi(\sigma|\mu) - \pi(\mu|\mu)) + (1 - \varepsilon)(\pi(\sigma|\sigma) - \pi(\mu|\sigma)). \quad (2.5)$$

So there exists a uniform invasion barrier if and only if there exists an $\bar{\varepsilon}$ such that for all $\varepsilon \in (0, \bar{\varepsilon})$, $f(\varepsilon) > 0$. Because we are dealing with a mixed strategy, we know ε and $(1 - \varepsilon)$ represent the probability of playing σ or μ . We will therefore be limited between 0 and 1 so we will find $f(0)$ and $f(1)$:

$$\begin{aligned} f(0) &= 0 * (\pi(\sigma|\mu) - \pi(\mu|\mu)) + (1 - 0)(\pi(\sigma|\sigma) - \pi(\mu|\sigma)) \\ &= \pi(\sigma|\sigma) - \pi(\mu|\sigma) \end{aligned} \quad (2.6)$$

$$\begin{aligned} f(1) &= 1(\pi(\sigma|\mu) - \pi(\mu|\mu)) + (1 - 1)(\pi(\sigma|\sigma) - \pi(\mu|\sigma)) \\ &= \pi(\sigma|\mu) - \pi(\mu|\mu). \end{aligned} \quad (2.7)$$

Since we assumed σ was an ESS, we know that either $f(0) > 0$, or both $f(0) = 0$, and $f(1) > 0$.

This means we will have three cases to consider; one where $f(0) > 0$ and $f(1) > 0$, one where $f(0) = 0$ and $f(1) > 0$, and one where $f(0) > 0$ and $f(1) < 0$. In the first two cases $f(\varepsilon)$ will be greater than 0 for all $\varepsilon \in [0, 1]$. If $f(0) > 0$ and $f(1) < 0$, then

define $a = f(0)$ and $b = f(1)$. We will therefore solve $f(\varepsilon) = 0$:

$$\varepsilon(a) + (1 - \varepsilon)(b) = 0$$

$$\varepsilon(a) + b - \varepsilon(b) = 0$$

$$\varepsilon(a - b) = -b$$

$$\varepsilon = \frac{-b}{a - b}.$$

Thus in the case where $f(0) > 0, f(1) < 0, f(\varepsilon) > 0$ for any $\varepsilon \in (0, \frac{-b}{a-b})$. Thus, if a strategy σ is an ESS, it will have a uniform invasion barrier. \square

2.2 EVOLUTIONARY STABILITY AND DYNAMICS

In addition to the static type of evolutionary stability we have discussed in this chapter, we can also take a dynamic approach. This type of thinking is similar to what we encountered in traditional game theory where we were able to model many situations as static games, but, dynamic game models gave us more flexibility to understand player reactions. In this approach to stability, we begin to think about how individuals evolve over time by using a game model.

2.2.1 PRISONER'S DILEMMA REVISITED

For an example of how dynamics play a role in evolutionary stability, we will reexamine the Prisoner's Dilemma. In this case, we will focus on how these types of games function generally so our assigned payoffs will be given a letter "value" as opposed to a payoff defined in years. For a refresher on the Prisoner's Dilemma, refer to Section 1.2.1. In this type of game, our players will have the option to cooperate or defect, each with assigned payoffs following this order: $W > X > Y > Z$ and we will

also define $\frac{W+Z}{2} < X$ since that ensures, in the context of an infinitely repeated game, that a group playing cooperate-defect and defect-cooperate strategies alternatively do not have an advantage. Our payoffs will look as follows in Table 2.4.

Table 2.4: Payoff Matrix for General Prisoner's Dilemma

Decisions	Cooperate	Defect
Cooperate	(X,X)	(Z, W)
Defect	(W, Z)	(Y,Y)

We will introduce some assumptions about the players to better track the situation in our population. Our first assumption is that our population of players is large. We will also assume that the probability for an individual to interact with another playing cooperate or defect is equal to the proportion of the population playing that strategy. The proportion assumption is especially helpful since it means we can represent the state of the population by looking at those proportions [10]. We will let $prop_C$ and $prop_D$ represent the proportion of cooperators and defectors respectively. Since we are looking at how well each of these strategies do over time, we also want to consider their “fitness” or how well they survive. We will represent the expected values for cooperate and defect as E_C and E_D . The average fitness of the population as a whole will then be represented by \bar{E} .

While the quantities defined above can vary over time, for simplicity we will assume the values remain the same. Based on our assumptions, we would expect that the fitness of each group would depend on the proportion in the population and their success represented by the payoff received from each strategy. We can then define E_C , E_D , and \bar{E} as:

$$E_C = F_0 + \pi(\text{Cooperate}|\text{Cooperate}) * prop_C + \pi(\text{Cooperate}|\text{Defect}) * prop_D$$

$$E_D = F_0 + \pi(\text{Defect}|\text{Cooperate}) * prop_C + \pi(\text{Defect}|\text{Defect}) * prop_D$$

$$\bar{E} = prop_C * E_C + prop_D * E_D.$$

In these equations, F_0 represents the starting fitness level of an individual before any interactions have taken place [10]. As the population continues to play the game, it logically follows that the previous generation's results will affect the future outcomes. The strategies that do well would be passed on and played more frequently than weaker strategies. We will then assume that the proportion of cooperators and defectors in a new generation will be related to the previous so we will represent that as:

$$prop'_C = \frac{prop_C * E_C}{\bar{E}}$$

$$prop'_D = \frac{prop_D * E_D}{\bar{E}}$$

We are defining the new proportions this way based on the assumption that the fitness of cooperators in the population will be lower than the average fitness of the population. The idea that cooperating will have a lower fitness comes from the way the payoffs are defined; while the payoff of both parties cooperating is the second best payoff, the worst comes from a situation where one defects and one cooperates. The disparity between defect and cooperate would then result in a difference in the fitness levels of the defector and cooperator groups since it would be “riskier” to cooperate, so we would expect the fitness of that group to be lower than the average fitness.

This difference in fitness level suggests that there would be a group switching from the cooperate strategy to the defect strategy since it would be more advantageous. The rate at which that group defects would be proportional to how much worse the cooperate group does than the overall population [10]. That relationship would be proportional because there would be some proportion of cooperators switching strategies based on a comparison between fitness levels; if the fitness level of the cooperate group remained the same, there would be no incentive to

switch strategies. But, if the defector group is able to skew the average fitness level of the population, we would expect individuals to switch at a proportional rate to the comparative fitness levels. Since $E_C < \bar{E}$, we know $\frac{E_C}{\bar{E}} < 1$, which implies $prop'_C < prop_C$.

What we are interested in then is how different the proportions of cooperators are between generations. This difference will tell us how the population has changed strategically since we can see what portion of the population switched from a cooperate strategy to a defect strategy. We therefore want to rewrite our equations for $prop'_C$ and $prop'_D$ to represent the difference:

$$\begin{aligned} prop'_C &= \frac{prop_C * E_C}{\bar{E}} \\ prop'_C - prop_C &= \frac{prop_C * E_C}{\bar{E}} - prop_C \\ prop'_C - prop_C &= \frac{prop_C * E_C}{\bar{E}} - \frac{prop_C * \bar{E}}{\bar{E}} \\ prop'_C - prop_C &= \frac{(prop_C)(E_C - \bar{E})}{\bar{E}}. \end{aligned}$$

We will have the same end equation for the defector group; we will just replace our Cs with Ds:

$$prop'_D - prop_D = \frac{(prop_D)(E_D - \bar{E})}{\bar{E}}.$$


Since we are looking at the difference in proportion from one generation to the next, we are essentially looking at the rate of change of $prop_C$ and $prop_D$. We could therefore approximate the difference between $prop'_C - prop_C$ and $prop'_D - prop_D$ using the derivative with respect to time if the change is sufficiently small. In this case, we can use this approximation since the population proportions should have continuous growth or decay; we would not expect the proportions of individuals using the different strategies to dramatically change from generation to generation.

We will now write our equations as:

$$\frac{dprop_C}{dt} = \frac{(prop_C)(E_C - \bar{E})}{\bar{E}} \quad \text{and} \quad \frac{dprop_D}{dt} = \frac{(prop_D)(E_D - \bar{E})}{\bar{E}}.$$

The differential equations presented above are known as replicator dynamics [26]. Since $E_C < \bar{E}$, we know that all future populations will have fewer cooperators than the previous. We could think of the replicator dynamics as a spectrum following a population from a state of total cooperation to total defection, with the in-between showing a mix of defectors and cooperators based on the effect of the dynamics. This idea is shown in Figure 2.1.

Figure 2.1: Replicator Dynamical Model of the Prisoner's Dilemma [10]



Because we expect to lose cooperators every generation, we will move from a state of full cooperation in the population to one of increasingly mixed cooperation/defection, until the population all defects. In a real world scenario, both of the endpoints are fairly extreme, however, in a theoretical sense this movement shows us different levels of stability at equilibrium points. If we consider full cooperation to be an equilibrium point, the replicator dynamics show us that it is an unstable one since any defection will move us further away from the equilibrium. At the other end of the spectrum, full defection in the population would be a stable equilibrium point since there is no incentive for members of the population to revert to cooperation. If a group were to try and revert, the dynamics would just push the population back to full defection.

2.2.2 LEARNING RULES

Replicator dynamics were the first kind of dynamics used in evolutionary game theory, but there are other dynamics which have become a part of the mathematical framework for understanding these types of situations. All of these dynamics have some basic building block known as a learning rule. These learning rules are essentially the individual level decisions that characterize the different dynamic forms. They can be represented as functions of the current population state (what strategies are currently in practice in the population) and the payoffs associated with each strategy given the state of the population. These functions then map to a matrix of the conditional switch rates associated with each strategy; essentially, what happens that causes an individual to switch strategies based on the state of the population and the expected payoffs [10]. These learning rules can then be used to determine the overall population dynamics, which follow a similar methodology for how we calculated the replicator dynamics in the Prisoner's Dilemma.

In the Prisoner's Dilemma example, we compared the difference between generations and realized that we were basically finding the rate of change for the proportions. In general, we can do the same kind of approximation since, at the population level, we are interested in what proportion of the population is switching from some strategy to another. Our population dynamics can then be modeled by the equation:

$$\frac{dp_i}{dt} = (\text{Rate at which players start to use } S_i) - (\text{Rate at which players stop using } S_i),$$

where p_i represents the proportion using the i^{th} strategy and S_i represents a strategy [10]. We can then substitute learning rules into this framework, solve the system of equations given by doing so, and get the population dynamics.

2.3 DYNAMICS

Dynamics are the other key concept in evolutionary game theory because they give us an understanding of how populations change over time. While we discussed replicator dynamics in the previous section, there are other dynamic frameworks also in use. We can also generalize our dynamic definitions to better encompass a wide range of scenarios.

2.3.1 DYNAMIC TYPES

In general, the learning rule for the replicator dynamic model consists of a player randomly selecting another individual from the population and then comparing their payoff earned previously to what they earned by picking the individual they did. If there was a higher payoff with the individual selected, the player would adopt the strategy used by them [10]. With this definition of the dynamic, we could expand this model to cybersecurity problems since an attacker would logically want to compare payoffs earned from different strategies. We could also model problems as if the attacker compared how they did against different companies' security systems. For instance, if an attacker were to send out a phishing email to a company A, but they were only able to get a few people to click through the email (due to individual skepticism, spam filters, etc.), they would want to learn from that experience. They would know more about their opponent's strategy and would know that if they adopted a strategy adapted from the company's security measures, they would most likely have a higher payoff in the next attempt.

While the replicator dynamics are useful for understanding how populations behave, there are some potential issues with using them to model different scenarios. In the learning rule for replicator dynamics, we assume that imitation will be a reliable way to model the population in future states. This assumption can be

problematic because we might be seeing payoffs from specific strategies that have more to do with the population than strategic merit. For instance, if we were looking at a population of companies with one using a specific security strategy that was not common in the group of their peers, they might be better at protecting their data solely because their system is not as common. The security strategy itself may not have the best strategic value in terms of the possible amount of data protected, but the company would not change its strategy unless there was a viable threat to the system. From a strategic merit standpoint, the company would want to choose the security system that best protects their data, but, in practice, this could be disregarded because there is not another incentive to spark a change. Another limitation is that the strategies need to be present in the population to start with; if it is not in the population at the start, there is no way it can be imitated in future generations.

In addition to replicator dynamics, there are also the Brown-Nash-von Neumann (BNN) dynamics. The BNN dynamics address some of the possible issues present in the learning rule for the replicator dynamics. To remedy the problems discussed previously, we could consider the learning rule where the rate at which players switch from a strategy S_i is only dependent on the expected payoff of S_i exceeding the average payoff of the population at the point in time [10]. This learning rule then leads to the BNN dynamics when plugged into the modeling equation we discussed in Section 2.2.2.

Besides addressing some of the problems with our assumption in creating the learning rule for the replicator dynamics, the BNN dynamics also give our players a higher degree of rationality. This rationality comes from a similar place to what we saw in finite dynamic games; our players in this scenario would know the entire set of possible strategies and the associated payoffs so they would be able to make the best-informed decision. This could include bringing a strategy that is not currently

in the population into play [10]. Being able to bring in new strategies is a big benefit of the BNN dynamics since we can not do the same with the replicator dynamics.

The Smith dynamics then build off of the learning rule used to create the BNN dynamics. The learning rule to create the Smith dynamics compares the expected payoff of a player's current strategy in the present population state with the expected payoff of the other strategies. However, only the alternative strategies that have a higher expected payoff have a nonzero probability of getting adopted by the player [10]. Taking this learning rule and plugging it in to the framework described in Section 2.2.2, we get the Smith dynamics.

2.3.2 CONNECTION BETWEEN ESSs AND DYNAMICS

A big difference in evolutionary stability and dynamics is what these concepts entail; evolutionary stability applies to strategies, whether they be pure or mixed, while dynamics model populations where individuals employ pure strategies. To bridge this gap, we can interpret the probabilities that appear in an evolutionarily stable strategy as population frequencies, which then gives us an evolutionarily stable state. This understanding of the probabilities expands our idea of stability to the population level and we can then start to examine under what conditions an evolutionary dynamic will converge to an evolutionarily stable state [10]. Like what we discussed in Section 2.1.2, dynamics may not converge to an ESS. A clear example of this would be the replicator dynamics because they have a restriction on the strategies in the population. If an evolutionarily stable state needed specific pure strategies to be present in the population but they were not present to begin with, there is no way to reach the ESS. Interestingly, this can lead to scenarios where strictly dominated strategies (see Section 1.2.2) can continue to be used in the population.

While those scenarios can occur, they are fairly rare since if all strategies are

present to begin with, even at a small level, the replicator dynamics will cause the strictly dominated strategies to disappear. Placing this in the context of the Prisoner's Dilemma, if there was even one person defecting, the population would move further away from full cooperation as time goes on. It is important to note that even if a weakly dominated strategy appears in a Nash equilibrium, it can never be an ESS (refer to Section 1.2.2). Weibull proves this in [27].

2.3.3 DOMINANCE AND DYNAMICS

While a weakly dominated strategy can not be an ESS, the replicator dynamics can still keep the strategy in play rather than eliminating it. This may seem counter-intuitive if we think about what this means when compared to traditional game theory. In that field, we look at dominated strategies as something to eliminate from consideration, whereas in the case of evolutionary dynamics, we do not have the same idea. The elimination of a strategy, in the case of replicator dynamics, depends on the mix of strategies present in the population. Examining BNN and Smith dynamics, we also find that they do not necessarily guarantee the elimination of strictly dominated strategies [10]. As an example, we will look at a modified version of Rock, Paper, Scissors.

In this form of the game, we add a new strategy which is a twin to the paper strategy. The only difference in payoffs between the two is that the payoff from playing the twin strategy will always be the payoff of paper minus some small $\varepsilon > 0$. This scenario gives us a familiar game with a strictly dominated strategy so we can better understand dominance and evolutionary dynamics. We will assign payoffs of $-1, 1$, and 0 for the base strategies of rock, paper, and scissors. This results in the payoff matrix Table 2.5.

In this game, we would want to consider the Smith dynamics since they compare the expected payoff of the current strategy in the present state to the expected payoff

Table 2.5: Payoff Matrix for Rock, Paper, Scissors, Twin

Strategies	Rock	Scissors	Paper	Twin
Rock	(0,0)	(1, -1)	(-1, 1)	(-1, $1 - \varepsilon$)
Scissors	(-1,1)	(0, 0)	(1, -1)	(1, $-1 - \varepsilon$)
Paper	(1,-1)	(-1, 1)	(0, 0)	(0, $-\varepsilon$)
Twin	($1 - \varepsilon$, -1)	($-1 - \varepsilon$, 1)	($-\varepsilon$, 0)	($-\varepsilon$, $-\varepsilon$)

of the other strategies. In theory, each of these strategies is a viable option for a player to start with, so in an initial state, an individual could choose to play the twin strategy. As we iterate the game, players could still choose to start with twin and then switch to a strategy with a better payoff [10]. This means that the twin strategy would remain present in the population even though it is strictly dominated by other strategies (in this case it is dominated by paper). A player logically would not want to switch to twin since it is dominated, but the new entrants into the game could keep the strategy in play.

2.3.4 USEFULNESS OF EVOLUTIONARY GAME THEORY

So far in this chapter, we have looked at the static approach in evolutionary game theory (evolutionary stability) and the dynamic approach (evolutionary dynamics). In both cases, the theory built upon existing concepts in game theory, like the Nash equilibrium or dominated strategies, but in doing so, they were able to give us stronger tools with which to examine games. We have discussed evolutionary stability as being a possible way to remedy the problem of multiple equilibrium points. However, much like the refinements discussed in Section 1.3.2, there are still complications with evolutionary stability. Depending on our approach, we may have competing understandings of stability so instead of choosing from our pool of equilibrium points, we would just shift the problem to choosing between refinements. There is also a slight issue in using evolutionary stability as our sole

justification for preferring one equilibria over another for a player in a game since there are differences in what the concept means in the static and dynamic settings [10]. While there are these issues with using evolutionary stability as a way to find preferred Nash equilibria, we should note that it is still a good tool to narrow down solutions, particularly in the static case, since it strengthens the definition of Nash equilibrium.

Another benefit of evolutionary game theory is the lighter emphasis placed on rational agents. In traditional game theory, we assume high levels of rationality that do not really match what we would see in actual game play. Thinking back to the Battle of the Spouses (Section 1.2.3), one of problems with finding a solution was how cultural assumptions may affect players. Previous notions and cultural influences would logically affect how an individual approaches a game scenario so the evolutionary approach of learning rules which then inform population-level behavior more closely matches that. This less-restrictive rationality assumption also allows for a more dynamic understanding of game scenarios. For instance, in our discussion of finite dynamic games (Section 1.3) we looked at game trees which were built on the assumption that a player would know the possible paths a game could take. That would require a lot of upfront knowledge and would require players to follow strictly defined paths. Evolutionary approaches do not have this same concept since it is more “organic;” players are able to move fluidly from approach to approach.

2.4 CONCLUSION

In this chapter, we examined two of the core concepts in evolutionary game theory, evolutionary stability and evolutionary dynamics. These concepts build off of concepts we discussed in our chapter on traditional game theory, strengthening

definitions and broadening our understanding of what situations can be understood through a game theory approach. In terms of real-world uses, evolutionary game theory techniques will allow us to model scenarios in risk management and cybersecurity more flexibly and could provide stronger reasoning for solutions. We can also gain insights on populations (corporate fields or types of attackers) based on micro-level analyses, learning rules, between individual companies and attacks on their data. Evolutionary game theory is a new avenue to explore which will provide a good framework for modeling problems in those fields because of the evolutionary aspect of security systems. Security measures need to be one step ahead of the technology and techniques available to attackers so they need to be able to learn from competitor strategies and attackers, which evolutionary game theory techniques would allow them to do.

CHAPTER 3

APPLICATIONS OF GAME THEORY IN RISK MANAGEMENT AND CYBERSECURITY

Previous work in cybersecurity and risk management has looked for ways to apply game theory techniques and models to defensive security measures. These measures span from physical security systems to network protections. This chapter will examine some of these examples in-depth to better illustrate the possible connections between the two fields and traditional game theory, allowing us to see how these techniques might be refined through an evolutionary approach. The introduction of game theory in these areas is still relatively new so this chapter will provide an insight into how approaches in real-world scenarios have grown in the past few years.

3.1 CLASSIFICATIONS OF DECEPTIONS

To start, we will discuss how game theory techniques have appeared in cybersecurity/risk management work that is focused on deception. We will draw from Pawlick and Zhu's taxonomy of defensive deception to understand how scenarios get classified [22]. This taxonomy gives us a clear starting point for which games to consider as feasible models for different situations and will allow us to more strictly define nebulous terms like *deception*.

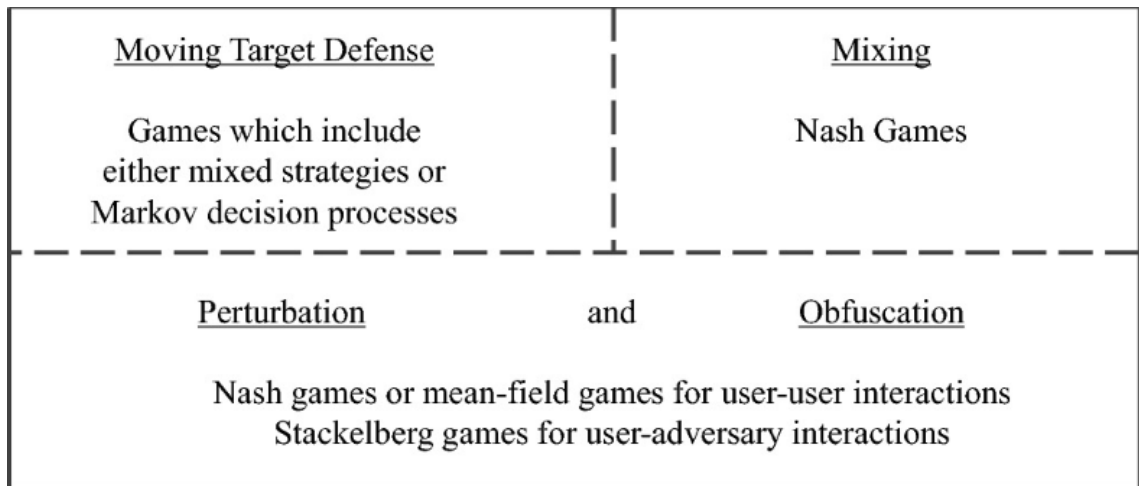
In the case of cybersecurity, defensive deception, a type of deception that seeks to protect information from attackers, primarily appears in three forms in the current literature: obfuscation, moving target defense, and honeypot. Obfuscation approaches seek to waste the resources and effort of attackers by directing them to decoy targets or revealing useless information alongside real information [22]. Essentially, the defending group attempts to confuse the attackers, and even if they are not distracted by the decoy targets, the false information mixed into the real information will make it harder for the attackers to find and sell the actual data. Moving target defense utilizes randomization and network, asset, and defense tool reconfiguration to limit the effectiveness of attackers [22]. This process makes it harder for attackers to maintain their understanding of a security system because it will be constantly changing through the processes of randomization and reconfiguration. The last type is the honeypot which consists of drawing attackers toward specific systems by disguising them as valuable assets. Because of its broad way of catching attackers, different types of honeypots will be represented by the term honey-x [22]. Honey-x defenses are a good way to gauge what attackers are interested in attacking and where attacks may be coming from since some honey-x programs like honeyfiles can protect data from internal attacks. If one notices that someone has fallen for a honeyfile, they would be aware of an insider threat which is helpful when trying to determine next step protective measures.

In addition to defensive deception, Pawlick and Zhu also found similarities in deception techniques used in cybersecurity. They highlighted perturbation, which adds noise to data to make it difficult to pull out useful information. Noise disguises trends and makes data appear more random than it may actually be. The usage of mixing in networks or zones is another technique that allows defenders to hide the linkability between different objects and individuals. An example they highlight was for location privacy where identifying pseudonyms in car networks would

switch between vehicles multiple times, disguising the location of a particular vehicle at different points [22]. These six different methods of protecting data then become our way to classify protective strategies.

Examining scenarios from previous literature in terms of these six classifications then makes it easier to see any modeling trends present. In addition to the types of games we have discussed, other models present in these articles were Stackelberg games and Bayesian-Nash games. In a Stackelberg game, a defender would make the first move and an attacker would choose the best response to that move. This is similar to how finite dynamic games functioned in Section 1.3. In Bayesian-Nash games, we are looking at incomplete information scenarios where we bridge the gap created by the lack of knowledge with probability distributions. At the beginning of Bayesian games, players are assigned characteristics and probability distributions get modeled to those characteristics. The outcome of a game is then calculated using the Bayesian probability. The following Figures, 3.1 and 3.2, show which models frequently appeared in each of the six classifications.

Figure 3.1: Models used in Moving Target Defense, Mixing, Perturbation, and Obfuscation [22]



While Pawlick and Zhu found some trends in the models used in each of these defense strategies, the nebulous nature of how deception is defined does not make

Figure 3.2: Models used in Honey-X and Attacker Engagement [22]

<u>Honey-X</u>	<u>Attacker Engagement</u>
Emphasis on attacker belief: Signaling games	Multiple-period games
Emphasis on defender defense allocation: Bayesian Nash games	Interaction between games and MDPs
	One-sided stochastic games

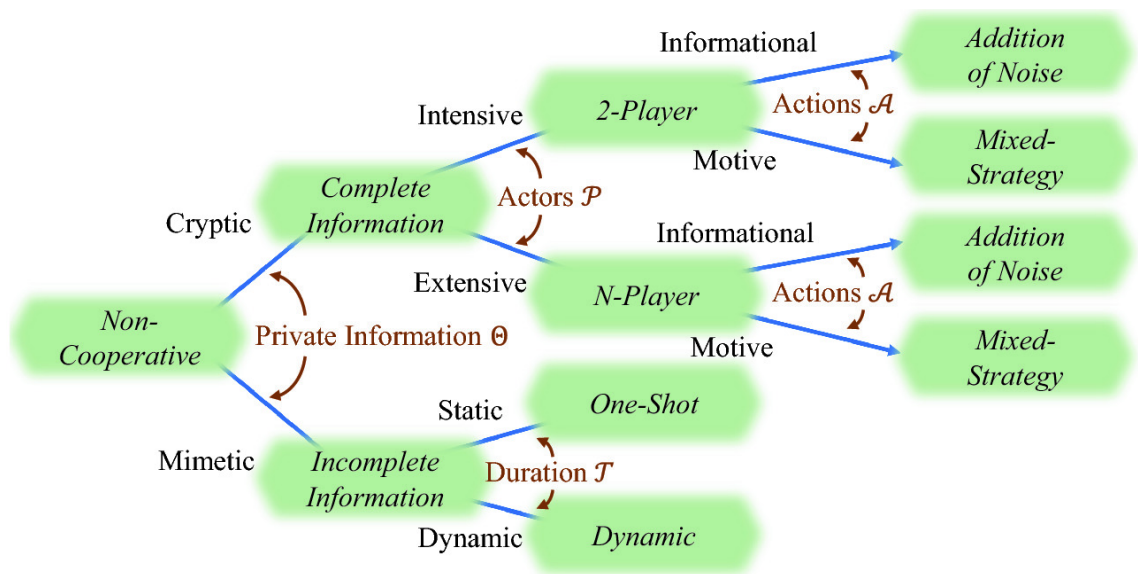
the modeling techniques listed the definitive methods for each category. Different models could be used for each classification; these are just the trends they noticed in their examination of the previous literature. Another important thing to note is that they did not see a hierarchical structure between the deception types. We therefore should not give precedence to certain types of deception over others; they should all be seen on the same level.

The differences between each of these types of deceptions will be based on game theory components such as private information, players, actions, and the time-horizon. In this case, private information covers both the intentional attempts to cause an opposing player to believe/acquire false information or preventing them from acquiring true information. Pawlick and Zhu borrow terminology from biology to distinguish between these two concepts with *crypsis* referring to the first and *mimeis* referring to the second [22]. We can also differentiate between types of players by whether they participate in a deception and to what extent. Cryptic deception can be divided into intensive and extensive forms, with players in the intensive form being modified to hide from an opponent and in the extensive form, multiple users are used to hide information. This difference is somewhat small but the goals are the important differentiating piece; in an intensive deception game, our players want to hide themselves to achieve a goal while in an extensive deception game, the players want to hide information. Along with the players themselves, we also need to consider the actions available to them. In the case

of cybersecurity, deceptions can either use information or use motion. Deception that uses information manipulates data released about players' properties while deception that uses motion modifies those properties over time [22]. The last category we need to consider is the duration of a game. Mimetic games can be either static or dynamic with static games having only one interaction and dynamic ones having multiple. Most defensive game theoretic models currently use static models, including honey-x games, and there is currently a push to research more viable dynamic models to use in cybersecurity settings.

The taxonomy structure as a whole appears in Figure 3.3.

Figure 3.3: Taxonomy of Defensive Deception [22]



This taxonomy then gives us a foundation for understanding what types of deception-based cybersecurity games currently exist in the field. This then can help inform our approach to different cybersecurity scenarios since it allows us to more easily define the game's form and results. Pawlick and Zhu work through a few types of deception games in their work *Game Theory for Cyber Deception*, exploring new types of game forms to model the complexity of cybersecurity scenarios involving deception. The field combining deceptive cybersecurity and game theory

is still new and developing, which is why we wanted to discuss this taxonomy to show what steps are currently being taken to understand these scenarios.

For brevity and relevance to our application case in Chapter 4, we will be working through a risk management example as opposed to one of the examples from Pawlick and Zhu. Interested readers can find in-depth explanations of the game theoretic applications being pursued in obfuscation, honey-x, and attacker engagement games in *Game Theory for Cyber Deception* [20].

3.2 PHYSICAL SURVEILLANCE

3.2.1 BACKGROUND

In addition to deception, another area of interest in cybersecurity and risk management is surveillance. In this case, we will be looking at a physical surveillance system as opposed to internet-related risks. There is a wide variety of surveillance systems currently in use to protect information, including camera-based or guard-based surveillance. A camera-based system can be limited in terms of protection because the defender will need a large number of cameras to thoroughly document an area and the cameras can be disabled in many ways (turned off, damaged, etc.). To circumvent these limitations, a defender should remain aware of potential threats, both internally and externally, and could use game theory techniques to strategically place cameras. Using game theory to determine camera position could help better protect the defender since they will know which areas would be of interest to an attacker and, if they have limited resources, they will maximize the benefits of their security system by not overextending it. For a surveillance system that relies on security guards, there is a similar need to understand where to place the guards to maximize protection without overcommitting resources.

We can think of physical surveillance systems as distribution-valued games

that model interactions between two or more players (attackers and defenders of varying kinds) that each have a finite strategy set. This means that the payoffs associated with the game will come from some probability distribution, hence it will be distribution-valued. The reason these games will be based on distributions is because of the inherent randomness involved in a physical surveillance scenario. A defender could deploy some number of security guards or have cameras watching over a specific area, but there is still a chance that both measures could miss an attacker if they understand how the defending company has set up their security system. This chance of missing an attacker is seen as a hypothetical, additional player that causes randomness in a real player's outcome, with the real players being the defender and the attacker. A distribution-valued game will take in a random outcome distribution as the payoff itself to avoid losing any information, which essentially means that we will compute the behavior that shapes the payoff distribution the best [11]. In this scenario, since we want the defenders to prevail, the equilibrium strategy will be the one which provides the defenders with optimal surveillance policies and the best strategic allocation of resources.

These kinds of games will take place over a large environment (the area being surveilled) that has some number of security guards watching for violations. Because of the limitations in resources available, including the guards, a defending group needs to be able to prioritize sensitive information by doing frequent patrols or having some security measure in place that provides almost constant protection. Less important information will then be less surveilled, and attackers could learn the surveillance patterns used on the prioritized data. We could think of this as a more realistic version of Capture the Flag.

The basic game of Capture the Flag consists of two teams, who have opposing goals. One team (referred to as Team 1) is charged with protecting some item that the opposing team (referred to as Team 2) attempts to steal. Team 2 has a safe area

where they can put the stolen item and there may be some “safe haven” areas where they can not be tagged while retreating to their safe zone. Capture the Flag is an interesting game because it lies at the corner of game theory and graph theory since strategies will also rely on the position of players [17]. For instance, a member of Team 1 would need to consider which player from Team 2 would be the best to attempt to tag, which would mean the Team 1 player would need to determine something about that player’s position and if it is a rational move to target them.

Our situation with a physical surveillance system will be a simplified version of Capture the Flag, but we still have to account for potentially missing an attacker. To do that, we will assume there is some likelihood of missing the attacker every round, meaning there is some chance that the attacker can cause some amount of damage in a specific area. These areas where damage would occur would logically be the ones with the least amount of surveillance. To mitigate that, the goal of the defenders would be to avoid damage suffered from intrusions by managing surveillance activities. We will therefore quantify how well the surveillance systems are working in terms of damage prevention [11].

It is important to keep in mind that quantifying this kind of damage is difficult. Realistically, we can not assign a definitive payoff to preventing an attacker from stealing information since we can not assign a value to the information itself. Likewise, we can not assign a payoff for an attacker successfully stealing something because we still may not be able to quantify the value of the stolen item. This would then make it seem like we can not accurately capture the stakes in a physical surveillance model, but, if we aggregate data from multiple sources, we can still get a detailed picture of what we need to achieve the defender’s goal of optimizing their surveillance.

3.2.2 PHYSICAL SURVEILLANCE MODEL

The approach to this kind of game will fall under empirical game theory, but, the game will be over a function space as opposed to the real numbers [11]. In empirical game theory, modeling is assisted by the use of real world data or outcomes from highly trusted simulations [2]. Using this function space allows for an integration of the uncertainty present in the situation, the use of distribution-valued payoffs in game theory, and allows us to optimize for different goals [11]. The model used will also not make assumptions for the attacker besides the different ways they could attack. Instead of attempting to determine which attack strategies would be more appealing, we will examine a “worst case” scenario since we want our defenders to be able to counter even the worst outcome. This will essentially become a zero-sum game (see Section 1.2.4) since in the worst case, the defender will have a maximum amount of damage and the attacker will have the opportunity to cause that damage. We will examine the model that Ali Alshawish and others provided for this scenario, but we will also introduce the general framework they provided for modeling physical surveillance systems.


3.2.2.1 BASIC GENERAL MODEL FOR PHYSICAL SURVEILLANCE GAMES

As mentioned previously with Capture the Flag, in a physical surveillance system we need to be aware of surveillance and player positions. We therefore will think of the environment as a finite, undirected graph, $G = (V, E)$, where V is the set of nodes documenting physical areas and E documents the set of edges connecting the physical areas. This graph essentially defines a space over which the game can be played, telling us which items may be targeted (V) and how an attacker may approach (E). Without loss of generality, we can assume that edges are without surveillance. We can make this assumption because we can model paths under surveillance as another node in the middle of an edge. Essentially, if something is

under surveillance we will consider it to be an area; for instance, if a location A and another location B were connected by a pathway that was under surveillance, we would treat that pathway as an area C with an edge sequence $A - C - B$ [11]. In practice, an attacker will attempt to walk through our graph to reach a targeted area, with the possibilities being that they avoid detection or get removed from the graph (found by piece of surveillance and removed from the area).

More formally, we will define a single pure strategy in the standard model as being a circle in the graph G so that the strategy space for the surveillance person is a set of circles, $\{C_1, \dots, C_n\}$, that spans G . We will then define the attacker's action set to be a set of paths $\{P_1, \dots, P_m\}$, without loss of generality, which end at a specific valuable target node $v_0 \in V$. This is represented in Figure 3.4.

Figure 3.4: Graph of Defender and Attacker Pure Strategies



The payoff of our game then corresponds to if the attacker gets detected. Our game will become a matrix with stochastic payoffs. The payoff matrix for this game will be stochastic in the sense that a Bernoulli probability distribution describes the matrix.

A Bernoulli distribution is discrete and offers two options: a “success,” denoted as 1, and a “failure,” denoted as 0. Those options happen with some assigned probability such that if the success happens with probability p , the failure will

happen with probability $q = 1 - p$. A basic example of a Bernoulli situation would be one where someone was flipping a coin one time and called heads; heads would appear with probability p , and tails would appear with probability $q = 1 - p$. In the coin toss scenario, we could have a $p = q = .5$ probability if the coin is weighted fairly, otherwise we would have $p \neq q$.

In the case of our payoff matrix, $A = (A_{ij})_{(i,j=1)}^{(n,m)}$, where i, j, n , and m represent the range of strategies for the defender (i to n) and attacker (j to m) respectively, an entry $A_{ij} \sim \text{Bernoulli}(p_{ij})$ with:

$$A_{ij} := \begin{cases} 0 & \text{if the intruder is missed;} \\ 1 & \text{if the intruder is caught.} \end{cases} \quad (3.1)$$

where p_{ij} represents how likely a detection of the path P_j is along the circle C_i . In this case, to determine the Nash equilibria we will need to change the matrix into a real-valued one from its original random variable format. One of the ways to do that would be to look at the expected value of each entry (the mean) which will result in a matrix $B = (p_{ij})_{i,j=1}^{n,m} = (E[A_{ij}])_{i,j=1}^{n,m}$ [11]. Matrix B can then be treated with the minimax approach and optimization, but, since this is a basic general model as opposed to the one for our situation of interest, we will stop the process here.

3.2.2.2 MODEL WITH UNCERTAINTY

Building off of the basic ideas established by the previous model, we will now bring in the uncertainty present in real-world surveillance situations. In an actual surveillance model, an assigned numerical payoff like we have been using throughout this research will not be able to capture the nuances of the situation; our understanding of payoff should come from simulations, surveys, and/or expert opinion. The combination of these sources of information will give us the best

understanding of how a surveillance system functions and if the one currently in place is the best for what the defenders need.

While the Bernoulli distribution approach works in a basic model, it is too simplistic to capture the nuance of what uncertainty brings to a surveillance game. We only had two options in the basic model, so we could not account for partial successes or partial failures, both of which would definitely be of interest to a defender. We therefore want a categorical distribution that will allow us to account for more than two scenarios.

Our model will then look at T_1, \dots, T_{Max} different types of areas that are tagged with their security demands (this allows for degrees of surveillance as opposed to totally surveilled or no surveillance). Our pure strategies will then be a set of frequencies $f = (f_{T_1}, \dots, f_{T_{Max}})$ that represent how many times a security guard performs a check in the different areas [11]. The overall strategy space will then be the collection of all of the reasonable and doable frequency tuples. Essentially, the space will consist of the frequencies a defender could reasonably implement based on resources available in the different areas. The attacker strategy space will then be composed of the various paths to the set of targeted security zones $\{Z_1, \dots, Z_m\}$ where the attacker wants to inflict damage.

We will then want to account for the factors affecting how we define payoff in this scenario so we will want to incorporate data generated through some means (like simulation or expert opinion). We will refer to this data as dat_{ij} and it represents the effectiveness of a defense strategy i against an attack strategy j [11]. We will then construct our payoff matrix, $A = (A_{ij})_{(i,j=1)}^{(n,m)}$, using the probability distributions as the payoff as opposed to a single number. Using the distribution allows us to maintain a degree of uncertainty, but, we will still be able to solve our game since we will know things about the shape and spread of the distribution.

Kernel density estimation is a way to estimate the probability density function

(pdf) of a distribution. The pdf provides the relative likelihood that the value of a random variable would be close to a point in the sample space. This function can then be integrated to get the cumulative distribution function (cdf) which gives us the probability of a random variable taking on a value less than or equal to the point at which it is evaluated. The kernel density estimation technique will then allow us to make estimates F_{ij} , which will be inferences about the likelihood given our data. The random payoffs, A_{ij} , will then follow $A_{ij} \sim F_{ij}(dat_{ij})$.

As opposed to the minimax and optimization approach suggested in the basic model, for a situation with uncertainty, we want to use a more flexible approach. Doing so will allow us to maintain the randomness encompassed by the probability distributions, which would be lost in the other approaches since we would need some representative number. We used the Bernoulli distribution in our basic model because it is a special case where we can convert the 0-1 random values into their averages (expected values) for a game-theoretic treatment. As mentioned, the Bernoulli distribution is too simplistic for these models with uncertainty so we will utilize a different probability distribution in this scenario. Using a different distribution will allow us to capture more options for game play.

To do actual analysis of these models with uncertainty, one can use the statistical computing platform R. The HyRiM package allows someone trying to manage risk to model these types of scenarios with zero-sum games and is able to directly take in expert opinion or other empirical data [11]. The package itself contains functions that calculate the cdf, various loss distributions, and compute multi-goal security strategies. These functions then allow someone to use all of the available data and rely on the theory and algorithms present in the package.

3.2.3 IMPLEMENTING THE UNCERTAINTY MODEL

The next step in this example is to examine the decision-making framework that applies the model detailed in the previous section. This places the model in the context of making risk management decisions and details an approach one can take when trying to model and solve physical surveillance games. Alshawish and others present a six-step framework that applies this game theoretic approach to physical surveillance to find an optimal solution. As a reminder, the optimal solution in physical surveillance games is the minimization of risk, which encompasses effectively monitoring high profile targets while still protecting other information that is not as important. If there are holes in a surveillance system, whether it is due to a lack of security guards or camera blindspots, that will increase the defender's risk.

The first piece of this framework is determining where exposure to risk can occur in a system. A risk manager needs to consider the physical boundaries they are working within along with who is playing the game and the resources available. Alshawish calls this "context establishment" since the risk assessor needs to understand the context of the situation to provide the best guidance. For instance, if one were to operate without knowing the players, they could not accurately predict strategies, which would ultimately result in some damage. The information used to establish the context can come from a variety of sources; as examples, one could use vulnerability assessments, ethnographic surveys, and/or business process analysis to get an understanding of potential risks [11].

The next two steps build logically off of that context; the first is identifying available strategies and the second is identifying goals. In the identification of strategies step, the defender will examine possible surveillance layouts, configurations, and operational patterns. In addition to identifying that, they will also find a way to parameterize the layouts [11]. The ultimate goal of this second step is to get the

action sets of the different players involved in the game. The third step then looks at the goals of the game, which can take many forms. As we touched on earlier, finding the probability of an attacker being detected or escaping could be a goal since that would give the defenders a good idea about how at-risk their information is. Along with that, one could also want to determine costs of different surveillance layouts, consider potential privacy breaches for end users, willingness to support a surveillance system, or legal regulation that concerns employees.

Cost is a logical thing to consider when setting goals for a surveillance model since resources may be limited, so being able to find a cost-effective but high performing solution would be the best outcome. In terms of privacy breaches, with the categorization of information into sensitive and non-sensitive data, a defender needs to consider how their security system will be viewed by their users. If an individual experiences multiple data leaks because their information was deemed “non-sensitive” by the defender, they would most likely lose the individual’s trust and business. Another thing to consider is how comfortable people are with a surveillance system since discomfort could lead to people, not necessarily attackers, seeking ways to bypass the system, which if they succeed, means there is a hole in the surveillance system. The last potential goal mentioned with legal regulation is primarily focused on making sure the defender complies with laws around data and privacy to avoid legal challenges.

The next step would be to determine the effectiveness of different strategies. This assessment may differ from scenario to scenario based on the specific goals of the defender. For instance, if a defender wanted to ascertain the level of comfort individuals had with their surveillance system, they would not be able to simulate that data. One can not reliably predict human emotions so to test the effectiveness of a surveillance strategy with that goal in mind, they would need to consider some type of real-world data. That data could come from surveys of end users

or from expert opinions and past statistical data. Because of the need to look at these harder to quantify goals, this step will quantify damage in terms of categorical or continuous probability distributions [11]. Those distributions must include all available information subject to:

1. All assessments need to be made on the same scale. This is needed to make the multicriteria optimization aspect of the model work. Numeric indicators are then changed to fit into a common categorical scale.
2. The data source should be reliable in the context of the intended risk assessment [11].

After determining the effectiveness of a strategy, we want to identify the best surveillance configuration. We will find equilibrium points in order to do this, and in this case, they will be Pareto-Nash equilibria. *Pareto efficiency* refers to a game where it is impossible to reallocate resources to give one player a better outcome without making another player's outcome worse. This concept focuses on making the best use of resources, but that does not mean that the payoff will be equivalent for both players in a two-person game. To illustrate what this means, we will refer back to the Prisoner's Dilemma (Section 1.2.1). In that scenario, our Nash equilibrium payoff was (-15, -15). We also had payoffs of (-16,0), (0,-16), and (-5,-5). Compared to our equilibrium payoff, the (-5,-5) payoff, which corresponded to both prisoners not confessing, would be Pareto efficient since we could not change actions from that point without decreasing one of the player's payoffs. The Prisoner's Dilemma is a Pareto inefficient game because the Nash equilibrium we get is not Pareto efficient.

In the case of a surveillance game, the need for the equilibria to be Pareto-Nash makes sense since the defender will attempt to get the maximum payoff possible given the attacker's actions (maximum information protection), but, if the attacker were to switch strategies away from the equilibrium in an attempt to increase payoff, there would be a decrease in the defender's payoff (attacker would be able to take

some information). The outcome of assessing the effectiveness of the different strategies available should be used to build the distribution-valued payoff matrix of the applied model.

Determining the Nash equilibria will follow a similar process to what we discussed in chapter one, but there will be some qualitative differences. Since we are trying to model a situation with multiple goals, we want to find the Pareto-Nash equilibria which will require us to make some adjustments to how we approach solving games. The primary change is scalarizing or compressing the game into a single-goal game and finding the Nash equilibria of the modified game. The scalarization will be a weighted sum of the goal payoffs with the weight reflecting the priority of each goal [11]. We will limit the weights to be strictly positive since in a surveillance game, we would not consider goals that would be useless or negatively impact the defender. By reducing the game into a single goal form, still taking into consideration the other goals, we are able to see the importance of each goal which is great practically. Our introduction of variables for theoretical reasons translates into a meaningful representation of what is important to the defender so the model shows the priorities of the defender throughout the process.

We will then finish off our framework for these types of problems by implementing the optimal surveillance configuration. In this phase, the defender would consider what has been shown as an effective strategy for optimizing their physical surveillance system and decide on if/how they will implement it. This implementation could take the form of physically building the suggested configuration or it could be simulated and compared to a current system. If the company were to simulate, they could see the effectiveness of a model before expending significant amounts of resources to implement it. The form of the simulation could center around times and frequencies since that is ultimately what we are looking at in surveillance games. We want to predict when an attacker will attempt to infiltrate a

defender and how often they will do so. The simulation could then consider some time unit, an equilibrium which has an action happening with some probability, and consider the pauses between actions. This kind of simulation would give an idea of when an action would take place.

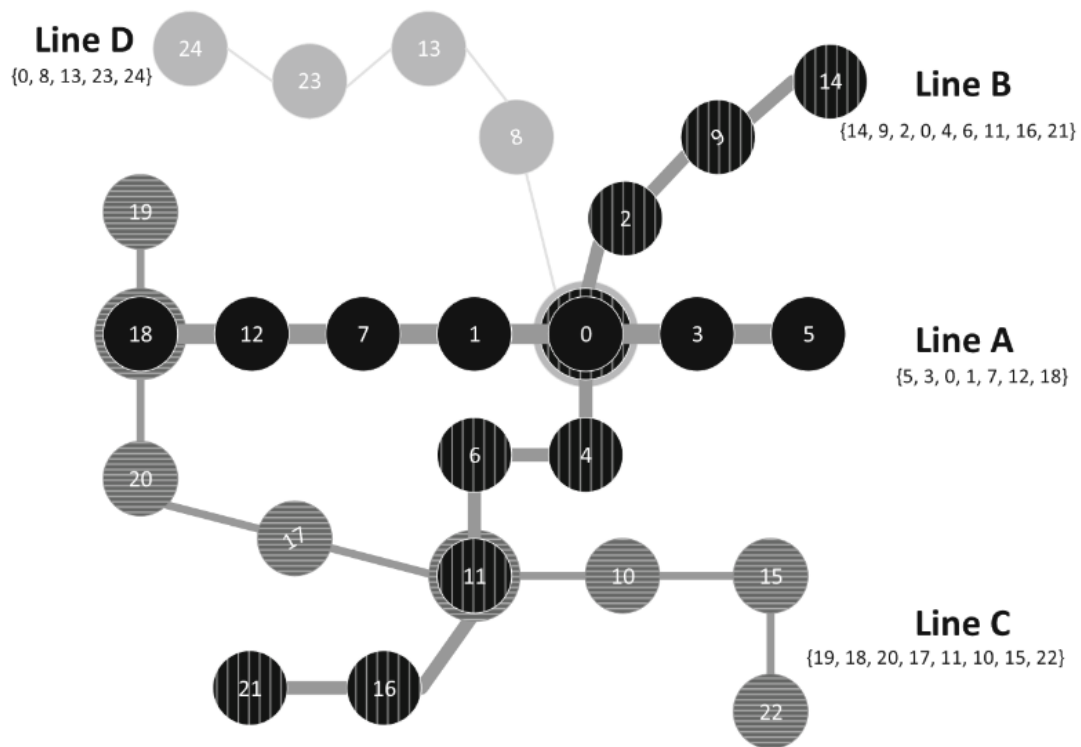
3.2.4 EXAMPLE: FARE EVASION

To see this framework and model in a real-world setting, we will look at fare evasion in public transit systems. In a public transit situation, the defending group who runs the system needs to know when to schedule fare checks to reduce the chance of commuters evading fares. In this case, we are trying to defend against the end users since the people who take public transit may also attempt to evade paying the fare. We could possibly view this as an “internal” attacker scenario since there is a small group within our overall user population that is going against the defending transit group. The goal of this game would be to create an inspection schedule that minimizes the risk of fare evasion but also minimizes costs.

3.2.4.1 ESTABLISH CONTEXT

Going back to the first step of our framework, we want to establish the context of this situation. A few things to consider in this would be how big our transit system is, how often particular lines are used, and what items inspectors check to make sure fares are paid. We will examine the transit system described by Alshawish which is a public transportation network consisting of 24 stops served by four lines, *A*, *B*, *C*, and *D*. The figure representing the system is shown in Figure 3.5 and the lines will exhibit statistically different utilization rates based on the thickness of the connecting lines [11]. This means that the thicker the line is, the higher its average passenger volume will be.

In addition to the transportation system itself, we will also want an understanding

Figure 3.5: Example Public Transportation System [11]

of how it works from a business perspective. Most public transport requires some kind of ticket. For instance, one generally needs to get a ticket/card and use that to swipe into turnstiles to access subways. This is where a person pays the fare for a ride or evades it so an understanding of how this process works is crucial for knowing where risk will come into play. Transportation companies will use inspections on “randomly” selected trips to find fare evaders, and the number of the inspectors who carry those checks out depends on how big the transportation network is and the passenger volume [11]. With larger networks and number of passengers, there would logically be more chances for fare evasion so a company would need to hire more inspectors to mitigate the problem. If there are more inspectors/inspections on high volume lines, we would expect more passengers to purchase tickets to avoid getting caught for fare evasion. However, the problem with more inspectors/inspections is that it will incur higher costs for

the transportation company. We will need to keep this context in mind and weight the benefits of deterring fare evasion through increased inspections and the cost of those inspections.

3.2.4.2 IDENTIFY STRATEGIES

With the context established, we want to move into the second step of the framework and identify potential strategies. There are a variety of strategy parameters that could exist for this scenario, but we will focus on a few of them for brevity. The first parameter would be how an inspector chooses which line to inspect. Referring back to Figure 3.5, we can see there is a clear difference in line usage so an inspector would need to determine whether monitoring a popular line, like line A, would be better than monitoring a less used line like C. There would be two strategies as a result of the inspector's line selection; an inspector could play the strategy where each line is inspected an equal amount or they could check the lines with higher passenger volumes more frequently. As we mentioned earlier, the risk of fare evasion will be proportional to the passenger volume. For the preferential strategy, we would therefore want to consider statistical data to help define a probability distribution over the transportation grid [11]. For the strategy where each line has an equal chance of inspection, we would not use the same process because the inspector will have no line preference.

There are also a few other parameters we can consider, including inspection frequency, inspection duration, the total number of inspectors, and collaboration between inspectors. We could consider the type of clothing, movement during checks, and the number of possible switches between lines in an inspection schedule. For this example, we will only focus on a few of these parameters including the total number of inspectors, collaboration between inspectors, and line selection. When we discuss collaboration in this context, we mean that inspectors will either

work individually or carry out their inspections as part of a group. We will limit our number of inspectors to be less than four, and we will define the set $\{2, 4\}$ as the number of inspectors having a daily inspection duty of 8 hours. We will denote this parameter in the strategy label $xPERS$ where x will be the number of inspectors [11]. Our collaboration parameter will represent the possibility of the inspector carrying out an inspection individually (referred to as I) and the possibility of them working as a group (referred to as T).

For our last parameter, we will define the probabilities of selecting a specific line as $P(A), P(B), P(C)$, and $P(D)$ where A, B, C, D refer to the lines in Figure 3.5. Since we know inspectors could inspect all lines with the same frequency or they could give preference to more traveled lines, we will define UR to represent the scenario with equal inspections and CRW to refer to the preferential line method. Since UR corresponds to equal inspections, we know that $P(A) = P(B) = P(C) = P(D) = 25\%$ in that case. For the CRW method, we will define the probabilities as: $P(A) = 45\%, P(B) = 30\%, P(C) = 17\%$, and $P(D) = 8\%$ [11]. We can now construct our action set for the defending transit company (Table 3.1):

Table 3.1: Action Set for Transit Company

Strategy Number	Strategy Label
Strategy 1	2PERS-T-UR
Strategy 2	2PERS-I-UR
Strategy 3	2PERS-T-CRW
Strategy 4	2PERS-I-CRW
Strategy 5	4PERS-T-UR
Strategy 6	4PERS-I-UR
Strategy 7	4PERS-T-CRW
Strategy 8	4PERS-I-CRW

This gives us the possible strategies the defending transit company can use to manage the risk of fare evasion. Our next step is to determine the strategies for the “attackers,” also referred to as fare evaders. For simplicity, an evader’s strategy will

be to choose a single line, *A*, *B*, *C*, or *D*, where they will not pay fares [11]. While these strategies are simple, they also make sense from the perspective of a fare evader. If one were to switch lines multiple times, there would be higher risk of getting caught since one would need to track inspections across all of the lines rather than just one. Since an attacker will choose a singular line, they will have four available strategies that we call *line A*, *line B*, *line C*, and *line D*.

3.2.4.3 IDENTIFY GOALS

Now that we have the possible actions for both players, we want to explicitly identify the goals of the game. Remembering back to the identification step in Section 3.2.3, these goals will set us up for the assessment of our strategies. From the transit company's side, we have multiple goals to consider, including maximizing the spot-checking missions on specific lines, minimizing how much those checks cost, and maximizing the number of penalty fares claimed from fare evaders [11]. On the attacker's side, they would like to avoid paying a fare without detection so they want to maximize the number of fares they can evade. Since this is a risk management game, we want to focus on achieving the transit company's goals as opposed to the attacker's goal.

3.2.4.4 ASSESS STRATEGIES

In the assessment stage, we have multiple approaches we could take to determine which strategy will allow us to achieve the goals we have set up. The method chosen was to rely on evaluations done by experts which gave a ranking on a scale from very low to very high for each of the defender's goals. In the case of maximizing spot-checking missions on specific lines (inspection intensity) and maximizing the number of penalty fares claimed (detection intensity), we would want a strategy to be ranked "very high" on whichever line we were examining. For cost, we will

want that to be ranked as “very low.” The rankings from 15 experts are shown in Figures 3.6, 3.7, and 3.8. Each of the abbreviations corresponds to a rank so “VL” = very low, “L” = low, “LM” = low to medium, “M” = medium, “MH” = medium to high, “H” = high, and “VH” = very high.

Figure 3.6: Expert Judgment of Strategies for Inspection Intensity

		Line A					Line B					Line C				Line D				
		expert ID																		
		1	2	6	7	10	3	8	9	12	14	2	5	10	11	12	13	14	15	
Strategies	2PERS-T-UR	VL	L	L	VL	L	M	L	VL	L	L	L	LM	VL	VL	L	M	L	M	
	2PERS-I-UR	M	LM	LM	L	L	M	LM	L	L	L	L	M	L	L	L	M	LM	M	
	2PERS-T-CRW	M	LM	M	LM	LM	LM	LM	L	L	L	VL	L	VL	VL	VL	VL	L	VL	
	2PERS-I-CRW	M	LM	M	M	LM	LM	LM	L	LM	M	L	LM	VL	VL	VL	L	L	VL	
	4PERS-T-UR	VL	L	LM	VL	L	LM	LM	VL	L	L	LM	LM	L	VL	VL	LM	L	M	
	4PERS-I-UR	H	H	LM	M	MH	M	LM	LM	LM	LM	M	LM	LM	M	MH	H	MH	M	
	4PERS-T-CRW	VH	H	H	M	M	M	LM	H	M	LM	LM	LM	L	VL	L	L	L	LM	
4PERS-I-CRW	VH	H	VH	MH	MH	MH	VH	VH	MH	M	LM	L	L	L	LM	L	L	M		

Figure 3.7: Expert Judgment of Strategies for Detection Intensity

		Line A					Line B					Line C					Line D			
		expert ID																		
		1	2	6	7	10	3	8	9	12	14	2	5	10	11	12	13	14	15	
Strategies	2PERS-T-UR	VL	L	L	VL	L	M	L	VL	L	L	L	LM	VL	VL	L	M	L	M	
	2PERS-I-UR	L	LM	LM	L	LM	M	LM	LM	LM	LM	L	LM	L	VL	L	M	L	M	
	2PERS-T-CRW	L	L	VL	L	L	M	VL	M	L	L	L	M	L	VL	M	M	M	LM	
	2PERS-I-CRW	M	M	H	M	LM	MH	LM	M	L	VL	L	VL	LM	L	L	L	VL		
	4PERS-T-UR	M	M	L	M	LM	M	LM	L	H	M	LM	L	M	M	M	L	L	L	
	4PERS-I-UR	LM	M	LM	M	L	M	VL	M	M	M	LM	M	M	LM	M	MH	M	M	
	4PERS-T-CRW	H	MH	MH	VH	L	LM	M	LM	M	LM	VH	H	MH	LM	L	M	M	MH	
	4PERS-I-CRW	M	H	H	MH	LM	MH	H	H	H	H	H	H	MH	M	L	M	M	M	


Figure 3.8: Expert Judgment of Strategies for Cost

		Line A					Line B					Line C				Line D				
		expert ID																		
		1	2	6	7	10	3	8	9	12	14	2	5	10	11	12	13	14	15	
Strategies	2PERS-T-UR	L	LM	L	VL	LM	M	L	LM	LM	L	VL	LM	L	VL	VL	LM	L	M	
	2PERS-I-UR	L	LM	LM	L	LM	M	LM	LM	LM	L	LM	L	VL	L	M	L	M		
	2PERS-T-CRW	L	LM	L	VL	LM	M	LM	LM	LM	L	VL	LM	VL	VL	M	L	M		
	2PERS-I-CRW	L	LM	LM	L	LM	M	LM	LM	LM	LM	LM	LM	L	VL	L	LM	M		
	4PERS-T-UR	H	MH	M	MH	M	M	LM	MH	M	M	L	M	MH	H	LM	M	LM		
	4PERS-I-UR	H	MH	M	MH	MH	M	LM	MH	M	MH	L	M	H	H	LM	M	L		
	4PERS-T-CRW	H	H	MH	MH	H	H	M	MH	M	MH	LM	L	VL	VL	VL	L	M		
	4PERS-I-CRW	VH	H	H	VH	MH	MH	M	MH	H	H	LM	LM	L	VL	L	VL	VL		

3.2.4.5 FIND OPTIMAL CONFIGURATION

Our next step in the framework is to actually play the game and determine the optimal configuration. In this context, an optimal configuration will be the best combination of inspectors, how they conduct inspections, and where they conduct those inspections. We will determine this using the HyRiM package in R on our distribution-valued payoffs which come from our collected data. The first step that we take is to use our data to define the distribution-valued payoffs required by the framework we have been following. In this case, the assessments provided by the experts in Figures 3.6, 3.7, and 3.8 will be the data we use. We then need to establish the priorities of the goals, so we need to consider what is most important to the transit company. After both of those steps, we can use the HyRiM package to find the Pareto-Nash equilibrium. Using this process, the Pareto-Nash equilibrium will be a nontrivial, mixed strategy with a probability distribution over the inspection strategies [11]. The results from this process are in Figure 3.9.

Figure 3.9: Results of HyRiM: Optimal Inspection Strategy [11]



From the figure, we can see that the four person, independent, uniform strategy has a high positive probability of .699. This is the highest probability we see in our mixed strategy, meaning it would be played the most compared to the others. In addition to that strategy, we would also expect the transit company to play the four person, team, high frequency line strategy with a probability of .105 and the four person, independent, high frequency line strategy with a probability of .196. Essentially, the probability distribution tells us that the transit company should always have four inspectors, since each of the strategies that make up the mixed strategy have four people, and they would most likely benefit the most from having the inspectors conduct inspections separately. While the optimal configuration has two components where inspectors would examine high volume lines, the usage probability is still low compared to the uniform strategy so it would most likely not be worth it to focus on the high volume lines.

We also know from Figure 3.9 that the transit company should not use any strategy with only two inspectors since the optimal configuration plays all of those strategies with a probability of 0. That essentially means that from our position as risk analysts that there is no benefit to using any of those strategies to curb fare evasion. Interestingly, we also have a 0 probability for the four person, team, uniform strategy. This, especially in conjunction with the other four person team strategy, tells us that having the inspectors work in teams is not as effective generally compared to when they work individually. Thinking about the scenario logically, this also makes sense because we would expect inspectors to find more fare evaders if they are not all in the same area. There will only be so many fare evaders in one location, so spreading the inspectors out would open up the possibility of inspecting multiple lines, thereby increasing how many fare evaders they could catch in the system as a whole. Overall, the optimal configuration described by the mixed strategy shows us the most important factors for the company to consider, how the

inspectors carry out their inspections and what lines they focus on, and confirm that a higher number of inspectors is better for catching fare evaders.

3.2.4.6 IMPLEMENT OPTIMAL CONFIGURATION

The last step in our framework is to implement the optimal configuration described in the previous section. Based on the mixed strategy, we would recommend that they operate using four inspectors that work independently and inspect each line with the same probability around 70% of the time. The rest of the time, we would recommend that they continue using four inspectors that work independently, primarily on the high volume lines like *A* and *B* (refer back to Section 3.2.4.2 for the probabilities of checking specific lines) around 20% of the time. The last piece would then be to use a four person group that worked in a team on the high volume lines around 10% of the time. The transit company could choose to implement this strategy in a variety of ways over time to end up with these percentages. For instance, it could randomly assign one of the effective strategies to the inspectors so the probabilities of selecting each of them corresponds to the probabilities in the equilibrium [11]. The implementation step affords a lot of freedom to the company since they know from the optimal configuration provided by the previous section which factors are the most important (always want four inspectors) and they can see the effects of changing how the inspectors work. With that knowledge, they are in a much better position to counter fare evasion than they would have been had they not gamified their problem.

3.3 CONCLUSION

Within the last 20 years, more people have turned to game theory techniques as a potential way for finding optimal strategies to solve risk management and

cybersecurity problems. These examples span from physical infrastructure and people to digital tracking and data protection. The two application areas we explored in this chapter are some of the more recent applications, coming from within the last four years. Since this way of approaching risk management and cybersecurity problems is still developing, we can continue trying to find new ways to incorporate different aspects of traditional game theory, and potentially finding ways to apply evolutionary game theory outside of biologic and social science contexts.

CHAPTER 4

EQUIFAX DATA BREACH AS A GAME-THEORETIC SITUATION

In this chapter, we will examine the 2017 Equifax data breach and model it in terms of game theory. This chapter will build off of the previous three and will give insight into how we can approach problems involving data breaches from this perspective. When we discuss data breaches, we will be referring to the exposure of sensitive, confidential, or protected data that is accessed in some way by someone who is not authorized to do so. This definition covers a wide variety of situations, including the copying, transmitting, viewing, or use of that data by unauthorized individuals. The goal of this chapter is to show how to extend a risk management framework to a cybersecurity setting.

4.1 BACKGROUND

We first want to establish the context surrounding the data breach and discuss the aftermath. This will allow us to fully understand Equifax's security policies, the lapses that led up to the breach, and how that ultimately impacted millions of individuals.

Equifax is one of the credit reporting agencies that determine an individual's financial health in the United States. They also have sizable operations internationally and are focused on using data analytics to help their consumers make

decisions about their financial situations [4]. At the time of the breach in March of 2017, Equifax was using a popular backend software for web applications called Apache Struts [25]. Apache Struts is an open-source framework for creating Java web applications and it is a popular software for companies like Equifax to use. A cybersecurity researcher, Nike Zheng, discovered a flaw in Apache Struts that could be used to steal data from any company using the software. Apache published this finding along with a patch on March 6, which then led to hackers attempting to find systems vulnerable to the attack [25]. The hackers found Equifax and were able to breach its security system through the flaw.

After gaining access to Equifax's data, the first group who breached the system brought in another group to gather the data. Gathering the data, which consisted of Social Security numbers, birth dates, addresses and more, for at least 143 million people in the United States, did not start in earnest until May of 2017 [25]. The attackers remained undetected until July of 2017, when Equifax became aware of an additional vulnerability in their security system that had remained undetected.

To monitor internal network traffic, Equifax had tools to decrypt, analyze, and re-encrypt that traffic, which would help the company discover data exfiltration. The attackers had been encrypting the stolen data they were moving since that makes it harder for security systems to find. Equifax had not renewed the public-key certificate they needed to re-encrypt the traffic internally, which allowed the attackers to remain undetected. This meant that encrypted traffic was not being reviewed, which is a major flaw in security. In addition to these flaws, Equifax was also working with a security consulting company called Mandiant, which warned Equifax about unpatched and misconfigured systems [14]. This shows that there were multiple points where Equifax could have acted to prevent the leak, but they decided not to.

As mentioned, at least 143 million people had their personal data leaked, however,

that data was not used to steal identities which was a major concern for those affected. The data did not surface, which has led to the theory that the data breach was a state-sponsored job [25]. That means that internal actors are not thought to be a part of the attack. We will not focus on the suspected groups or Equifax's response to the breach, but more information can be found in [14] and [25].

4.2 WHY USE GAME THEORY?

Now that we have the context and an understanding of the data breach, we want to discuss why game theory may be an appropriate way to model the attack. In this situation, we want to understand what actions were available to Equifax to prevent this kind of data breach and we want to know how an attacker could counter those actions. This is fairly similar to what we discussed in the previous chapter where we wanted to look from a preventative perspective at specific threats to security. We could then view this situation between Equifax, our "defender," and our attackers, the group of hackers, as a strategy-based game where each group pulls from their available actions to achieve their best outcome. For Equifax, that outcome would be the maximum amount of data protection and for the attackers, that would be the maximum amount of data extracted.

Like our previous examples, this game would be a two-person game between Equifax and the hackers, which would be a catch-all group comprised of the infiltration group and the extraction group. In addition to the goals mentioned above, it would be logical to assume that there would be other goals involving cost and detection as well. We know that we can have a multi-goal game theoretic situation from our section on physical surveillance (Section 3.2) so we could extend some of the techniques shown in that section to this example. Perhaps the best reason to approach this situation from a game theoretic perspective is the ability to

see what different strategies Equifax could have used to prevent this attack, which could in turn inform security policies for the future.

Using an understanding of previous strategies to inform future ones also links back to evolutionary game theory techniques. In a larger context, we could view the Equifax breach as an instance of game play where they lost using an “incumbent” strategy. Other members of the population (other credit companies or companies that use similar software) could then look at what strategy Equifax used to inform how they approach their own security. In this way, the companies would act like the populations we discussed throughout chapter two, particularly in terms of the replicator dynamics which we explored through Section 2.2.1. With the Equifax breach, we could assume that other companies used a similar security system at this point in 2017, primarily meaning they used similar software, so we could view them as using the same strategy. After witnessing this massive breach, they would see the deficiencies in that similar strategy, which could cause them to “defect” (switch to a different security system/strategy). As more breaches happen, we would expect more companies to learn from the mistakes of their peers and change their security systems, so we would be moving from a state of “cooperation” where many companies share similar security measures to one of “defection” where there are a variety of strategies in play.

Using a game theoretic model then would provide an optimal strategy profile which would give a company the most important areas to focus on in their security system. By understanding what the optimal strategy profile is, the company could focus on security measures that have a greater impact on the successful protection of data.

4.3 HOW DO WE APPROACH THIS SITUATION?

We want to consider this problem using the methodological approach from the public transit example detailed in Section 3.2.3. This framework will allow us to clearly define what actions, strategies, and goals exist in this scenario and it will give us a basis for understanding how to assess the different strategies, build a model, and posit how the data breach's outcome could have differed based on the optimal strategy found.

Our model will be hypothetical as opposed to fully worked out based on data. This is because we do not have publicly available data from experts weighing in on the various strategies we will explore. While we do not have this data, the congressional report from the Oversight and Government Reform committee provides some further information about what Manidant found when examining Equifax, including the suggestions they made to improve security systems. There were eleven suggestions in total, primarily focused on enhancing vulnerability scanning and deploying different detection and monitoring technologies [9]. These findings will help inform what aspects of Equifax's security system should be focused on in the model, which will in turn give us an idea of what goals are the most important. However, since the reports and documentation for the situation come from after the breach, we do not have expert opinion on to what degree different preemptive actions (strategies) should have been taken to prevent the breach. The holes found in Equifax's security will therefore act as the basis for the model, but we will not be providing hard, numeric results.

4.4 EQUIFAX MODEL

4.4.1 DEFINING ACTIONS

We have already established the context of the problem in Section 4.1, which has given us insight into the actions available to both Equifax and the attackers; mainly, Equifax's actions will revolve around maintaining/updating existing security measures while the attackers' actions will center on finding holes in those defenses. Holes will appear any time a security measure is not updated, properly detected in diagnostics as being outdated, or expert opinion is not considered. This idea captures the lapse in not updating Apache, the failed diagnostics that did not identify vulnerabilities, the failure to renew the public key, and the decision to not act on Manidant's suggestions. We will parameterize Equifax's action set as the set, A , containing actions X , Y , and Z with X corresponding to automatic software updates, Y corresponding to running diagnostic tests, and Z corresponding to automatic renewals of essential technology.

The action with the highest usage probability would be X since software typically updates regularly since new issues may appear or some functionality may need to be implemented to meet the needs of the users. We would expect Y to have the second highest usage probability since a company with large amounts of protected data would logically check to make sure their systems are operating without issue on some schedule. If the company does not have a diagnostic schedule, we would still expect diagnostic checks to happen with some regularity since people in the company may have issues or notice problems in a system, which would lead to an IT specialist checking it. This leaves Z as our least utilized action, primarily because renewals would most likely happen on a yearly basis, which would be much less frequent than what we expect for the other two strategies.

Our attackers will then have different actions centered on infiltrating systems,

not necessarily targeting Equifax, but generally. An attacker group will have some degree of information that will inform their actions; they will know about some exploit/flow/weakness in a particular system and will then act upon that information to find companies using software they can infiltrate. Their action then is to attack a company with software they can infiltrate. Our attacker's strategy will be to choose a vulnerability, such as outdated software, infrequent diagnostics (able to escape detection if attacker's actions affect functionality), or companies that rely on public security in their systems, and use that to steal data.

4.4.2 GOALS AND STRATEGIES

With an idea of what actions are available to both Equifax and the attackers, we will now want to consider their goals. Equifax would want to prevent data breaches entirely at the lowest technology cost and with the lowest number of people. One of the goals then would be to proactively prevent flaws from appearing in their security system since that would be the way to stop data breaches from occurring. In terms of cost, as a company, they would want the best technology for minimizing the cost of upkeep while not sacrificing the protective capabilities. If they achieved this goal, the company would be able to focus money on another aspect of the company's strategic plan and/or increase their profit. Similarly, Equifax would want the minimum number of people required to successfully operate their security system since that would contribute to minimizing costs. As of February 10, 2022, there were 17 jobs open in Equifax's Security, Technology, Governance, and Compliance department [12]. This gives us an insight into the size of their department maintaining security, which we could use to inform our own estimates of how many people worked in the department, with the necessary caveat that these numbers are only estimations and most likely are different than what was available in 2017. Again, since we are

describing a modeling procedure as opposed to working it through with numbers, we will just keep this mind as an additional aspect affecting strategies in the game.

The attackers will have fewer goals to consider than Equifax. Their goal is to extract data from a company with a vulnerability they know how to exploit. We could think of this goal as an amalgamation of two; with the first being to find a company with a vulnerability, and the second being to extract its high-value data. Since we are modeling the Equifax data breach, we will only assume that the goal is to extract their data since we already have named a target. If we wanted to complicate the model further, we could also consider the goal of avoiding detection while taking the data. Again, since this is an illustrative model, we will keep it fairly simple and consider the one goal of maximizing the amount of valuable data stolen.

The goals of this situation are a little more complicated than what we saw in Section 3.2. We will need to balance the three goals for the defender and one goal for the attacker, like we saw in that section, but the goals themselves are harder to quantify. How do we quantify the proactive measures? What is the difference between having a higher number of diagnostics vs software updates in terms of a game theoretic value we can use? To grapple with this quantification problem, we will assign usage probabilities to each of the different defender strategies, X , Y , and Z , based on our assumption of how often Equifax would run them. We had stated previously that software updates would occur the most frequently, followed by diagnostics, and then renewals. To approach the problem posed by the second question, we could also consider alternative usage probabilities where the strategy based on diagnostics has a higher usage probability than the others, with software updates and renewals being used with increasingly lower probabilities.

With the actions and usage defined, we can then start thinking strategically. In this scenario, an employee at Equifax would have two choices: update software more frequently than running diagnostics, or run diagnostics more frequently than

checking for software updates. Each employee would then be able to employ each of these strategies with some preference in mind weighted towards software or diagnostics. This means we would want to consider statistical data to define a probability distribution for each strategy. Another thing we could consider would be whether employees were assigned to specific tasks, so they could work in a team on updating software/running diagnostics/renewing licenses, or if the department does not assign tasks and employees have to notice outdated software and lapses in security. If there were assigned tasks, we would expect there to be less opportunity for holes in security to appear since a group would be working on maintaining various aspects of the system and could provide checks on each other. If the employees do not have a recurring, assigned task involving maintenance, we would expect more holes to appear. To build our strategy set, we will use some of the same notation as in the transit example with *XPERS* representing the number of employees working on maintaining systems, *I* representing independent work (the strategies involving one person will only be independent), *T* representing working in a team, and we will introduce *DIAG* for the situation where diagnostic checks have a higher usage probability than the other actions, and *SOFT* to represent the strategy where updating software has the highest usage probability. Equifax's strategy set would then be Table 4.1.

Returning to the attackers' perspective, since we defined their goal as stealing data through vulnerabilities, their strategy set would consist of three strategies: *attackX*, *attackY*, and *attackZ* where each attack corresponds to attempting to find a vulnerability in each of the actions Equifax could take. For instance, the attackers could choose to hunt for un-updated software, which would be the *attackX* strategy since it corresponds to the available action of updating software on Equifax's side. In this scenario, the optimal attack strategy was a combination of the three strategies; the attackers were first able to gain entry using strategy *attackX* and they evaded

Table 4.1: Strategy Set for Equifax with N-number of Employees

Strategy Number	Strategy Label
Strategy 1	1PERS-I-SOFT
Strategy 2	1PERS-I-DIAG
Strategy 3	2PERS-T-SOFT
Strategy 4	2PERS-I-SOFT
Strategy 5	2PERS-T-DIAG
Strategy 6	2PERS-I-DIAG
Strategy 7	3PERS-T-SOFT
Strategy 8	3PERS-I-SOFT
Strategy 9	3PERS-T-DIAG
Strategy 10	3PERS-I-DIAG
...	...
Strategy M-3	NPERS-I-SOFT
Strategy M-2	NPERS-T-SOFT
Strategy M-1	NPERS-I-DIAG
Strategy M	NPERS-T-DIAG

detection because of faulty diagnostics (*attackY*) and Equifax's failure to renew important keys (*attackZ*). The strategy combinations for this scenario will therefore be more complicated than what we saw in the transit example since we would also need to consider the optimal, mixed attack strategy. The combination of these strategies would then lead to the attackers achieving their goal of stealing as much data as possible.

Since we are modeling an event that occurred and has publicized documents/actions taken, we could also expand the scope of our model to take more goals into account. For brevity, we will describe the possible directions this could be taken, but we will disregard them to focus on the goals identified previously for the rest of our discussion. If we were to consider the possible perpetrators of the attack or Equifax's response, we could also take goals about specific data and the purpose it would serve the attackers or goals about regaining public trust into account. These goals would need to take the congressional report mentioned previously ([9]) as the primary document for understanding the attackers' potential motivations and the

different approaches Equifax took to remedy the situation publicly. The framework provided for physical surveillance then works as a nice way to understand and build our model for this situation, even though it goes beyond what it was originally used for, since it allows us to account for multiple goals.

4.4.3 ASSESSING STRATEGIES

With the strategy profiles described, we want to determine their effectiveness in achieving Equifax's goals. We will focus on Equifax because they are our defending company in this scenario and we want to understand how they could have prevented the loss of 143 million individuals' data. To that end, we will not build our game model to achieve the attackers' goal, but it is understood that they would attempt to undermine Equifax's actions to steal the data. As mentioned previously, we do not have access to data on the different methods Equifax could have taken to mitigate the damage caused. In the transit example, we have expert rankings comparing each of the defense strategies against the others for each of the stated goals. While we do not have this data here, the goals are similar enough that we would expect somewhat similar rankings for the strategies purposed previously. For instance, to proactively prevent data breaches, we would want to maximize the software updates implemented or diagnostics run, meaning that we would want a strategy involving those actions to have a "very high" ranking. For minimizing cost, we would want a strategy to have a "very low" ranking; similarly for the number of people needed to maintain the security system.

4.4.4 DETERMINING THE OPTIMAL CONFIGURATION

From our assessment, we expect to see an optimal configuration similar to the transit example; that is, we would expect there to be some optimal number of people

working individually or in a group with each of the preferential strategies. Based on the congressional report and the interviews housed within it, we would expect there to be more strategies where teams lead to higher degrees of success since a lot of the security problems stemmed from a lack of communication between departments, particularly the IT and Security departments. At the time of the breach, the Security department was housed under the legal office, which led to gaps in compliance and implementation of security policies [9]. This is why we would expect strategies with a team aspect to appear in an optimal configuration since there would be accountability between team members to maintain the security systems.

Since we do not have defined data for this scenario, we can not provide a concrete number of people we would expect to be working on maintaining security systems. However, based on what we saw in the transit example and the sheer scope of the amount of data Equifax had, we would assume that a higher number of employees working to update required software and determining where there may be problems through diagnostics would be the best strategies picked out to be a part of the optimal configuration. Since we have a total of N employees, we would expect that there would be a trend such that the number of players working together to maintain the security system having success would rise with each employee added up until an optimal number, where adding more people would cause a decrease in the level of success. This may seem counterintuitive since adding more people to a monitoring team should help eliminate holes, but if we had too many people, we would be wasting resources on additional, unnecessary monitoring which could cause us to miss flaws in other aspects of Equifax's security system.

In terms of whether the strategies with an emphasis on software updates or diagnostics would appear more frequently, we would need hard data. Based on what occurred in this situation, we would assume that strategies with software being updated more frequently would appear with higher probabilities since a software

vulnerability was what allowed the attackers to enter Equifax's systems. This is supported by Manidant's suggestions in the congressional report where their first point was to "[e]nhance vulnerability scanning and patch management procedures" [9]. We would expect diagnostic-focused strategies to appear occasionally in the optimal strategy profile, but with lower probabilities than the software-focused ones. We would therefore argue, based on our understanding of the Equifax data breach and the congressional report which further contextualizes it, that the optimal configuration will consist of strategies with a mid-high number of employees working in a monitoring team focused on updating software appearing with high probabilities, followed by strategies with the same number of people in a team focusing on diagnostics. The optimal configuration for Equifax would, on a macro-level, focus on bridging the communication gaps that existed within the company which led to the vulnerabilities in their security system.

4.4.5 IMPLEMENTATION

To proactively prevent a data breach, we argue that Equifax could have implemented a strategy based on a game theoretic approach to cybersecurity and risk management. This strategy would require Equifax to maintain the appropriate number of people needed to monitor their software and key renewals and emphasize cross-department/team work. By doing this, that would minimize Equifax's cost in the long run since any data breaches will have a much more negative financial impact on the company than the cost of employees and maintaining the software. In terms of what aspects of the security system to focus on, we would suggest emphasizing software maintenance for periods of time, possibly alternative quarters, and then focus on diagnostics. By shifting focus throughout the year, the company would benefit from keeping their software up-to-date (even after switching to an emphasis on diagnostics since the team would still monitor the software) and they

could catch problems earlier. Alternating between the strategies would also make it harder for attackers to exploit vulnerabilities since they would be caught by the team before attackers could catch them.

4.5 WHAT DOES THIS TELL US?

The goal of modeling the Equifax breach using a game theoretic approach was to determine what strategies the company could have taken to prevent the breach from occurring. To that end, our model gives us insight into what options were available to Equifax, the possible methods they could have used to maintain their security systems, and it highlights the importance of communication between various aspects of a business dealing with large quantities of personal data. Based on our understanding of the breach, the breakdowns in the security system primarily stemmed from a lack of maintenance and monitoring with regards to software and key renewals. This type of model would allow us to understand how important each part of the maintenance cycle is since it would show the combination of the best preventative steps, focusing on updates, diagnostics, and renewals.

In addition to what the model reveals about the data breach, it can also act as a foundation for future approaches to security. While Manidant provided a recommendation list of remedial security measures to Equifax, doing this kind of model in the context of Equifax's cybersecurity would allow them to better understand which suggestions should take precedence. Supporting that is the mixed strategy approach which gives Equifax a clear insight into how they should vary their approach to security to continuously evade attackers. This kind of thinking harkens back to randomization which has been a common theme in applications of game theory in cybersecurity and risk management as well as in our examples of mixed strategies. Models utilizing game theory in these fields would

further support suggestions from experts while allowing companies to explore their own goals since the models have a flexibility to them. The broadness of game theoretic concepts allows for a fluid definition of games which in turn helps describe these complex security situations. Overall, our game theoretic model for the Equifax breach provides further insight into how Equifax could have prevented the breach since it highlights important security factors and it emphasizes the need for teamwork/communication to maintain the security system.

4.6 DIRECTIONS FOR THE FUTURE

Our model that we constructed in this section is just one way to approach modeling this situation from a game theoretic perspective. As mentioned in the chapter on evolutionary game theory, that field could be useful in this context since it allows for more flexibility than what traditional game theory has and it gives us some stronger definitions for concepts like the Nash equilibrium. We would suggest examining internal attacker models through an evolutionary lens, particularly making use of replicator dynamics, since we think that could allow researchers and security experts to capture more information about how an attack continues.

For instance, if one employee “goes rogue,” we could think of that as defecting from the incumbent strategy that maintains security protections. If that employee is able to show the benefits of defecting (which would most likely be an amount of money equal to the value of the data to outside malicious attackers), either in private conversations or convincing coworkers to join (the game situation), they could cause other employees to defect. While we would not expect the defectors to shift the population of all employees to a completely defecting state, this way of thinking could be helpful for understanding how attacks can scale up.

In addition to dynamics, we also argue that bringing in the concept of evolutionary stability would be a good way to confirm the strength of an optimal configuration. Evolutionary stability strengthens the definition of a Nash equilibrium, so if an optimal configuration were to also be an ESS, that could lend further support as to why a company should implement it.

Furthermore, we argue that both evolutionary stability and dynamics fit nicely with currently used game types like Stackelberg games (see paragraph four in Section 3.1). In those games, there is a kind of call and response between the players where the defender implements a measure that the attackers then have to respond to. In a population, if a mutant appears, the group has to respond in some fashion. Does the mutant die out or continue to spread through the population? There is the same type of call and response in evolutionary game theory so we believe that it would work in a complementary way to currently used game models in the field.

Since the scholarship working on how to implement game theory into cybersecurity and risk management is still very much developing, with both [11] and [21] coming out within the last four years as of 2022, we believe there are many future paths to take. The ones we argue would be the most interesting and useful for security professionals would be evolutionary game theory techniques since they extend concepts of standard game theory and are already used in physical and social science settings. Behavior modeling is already occurring using evolutionary techniques and risk management/cybersecurity also focus on the behavior of companies and attackers, so this extension should work as a viable option for future exploration.

CONCLUSION

Our goal for this independent study was to explore two different game theory fields, traditional and evolutionary, to get an in-depth understanding of how concepts from those fields have appeared in cybersecurity and risk management, and apply what we have learned in the context of the Equifax 2017 data breach. We started with traditional game theory, focusing on two of the major game types, static and dynamic. We were able to define important concepts such as Nash equilibria and dominance, which gave us the tools we needed to solve different game types. We focused on common types such as the Prisoner's Dilemma and Rock, Paper, Scissors and then extended the concepts covered in static games to finite dynamic games. Basic finite dynamic games provided us with an understanding of how one player's actions could determine their opponent's, something which became important in our understanding of game theory applications in risk management and cybersecurity.

Evolutionary game theory then gave us new tools to consider when trying to model various scenarios from a game theoretic perspective. We focused on stability and dynamics which allowed us to strengthen the definition of a Nash equilibrium and gave us more insight into populations. This emphasis on change over time, learning from past generations' success or failure, and ability to expand individual-level results to the population-level made evolutionary game theory an interesting field to consider implementing in risk management/cybersecurity.

With an established understanding of game theory, we then moved into applications in risk management and cybersecurity. We focused on two applications; one that is newly developing, defensive deception, and another which gave us a framework we could extend, physical surveillance. These applications show the variety of game theoretic techniques currently being brought into both cybersecurity and risk management spaces, showing how professionals can manipulate different situations into a game format. The physical surveillance example was particularly useful since the framework used to model the game was also general enough that we were able to extend its usage to a cybersecurity setting.

We finished our work by examining the 2017 Equifax data breach using the framework from the physical surveillance example as our jumping off point. This allowed us to delve deeper into the relationships between various security measures and their failures in this instance, highlighting what aspects of security Equifax could have focused on to prevent the breach. Through creating this model, we also saw different ways game theory could be further integrated into cybersecurity and risk management, particularly with regards to evolutionary game theory. Many of the problems for Equifax came from a lack of monitoring on the part of the employees, which led us to question how insider attacks could function in a game model.

In terms of the future, we argued that internal attacks could be examined through replicator dynamics, which would give a better insight into how internal attacks spread and function. We also argued that incorporating evolutionary stability into risk management/cybersecurity settings would further strengthen situational understanding and provide support for recommendations.

Overall, this research explored the primary tools and popular game models associated with two types of game theory. From that exploration, we were able to take those concepts outside of economics or mathematics and show that there are

other fields where these types of models could flourish. This further expands game theory's usefulness and provides additional methods of understanding complex scenarios in risk management and cybersecurity.

REFERENCES

1. Risk Management. <https://corporatefinanceinstitute.com/resources/knowledge/strategy/risk-management/>. 2
2. Strategic Reasoning Group | Empirical Game-Theoretic Analysis. 65
3. What is Cybersecurity? <https://www.ibm.com/topics/cybersecurity>. 3
4. Who We Are - United States - Evo Prod. <https://www.equifax.com/about-equifax/who-we-are/>. 85
5. Finite Dynamic Games. In James N. Webb, editor, *Game Theory: Decisions, Interaction and Evolution*, Springer Undergraduate Mathematics Series, pages 89–105. Springer, London, 2007. ISBN 978-1-84628-636-0. doi: 10.1007/978-1-84628-636-0_5. 25, 26, 28, 29
6. Simple Decision Models. In James N. Webb, editor, *Game Theory: Decisions, Interaction and Evolution*, Springer Undergraduate Mathematics Series, pages 3–22. Springer, London, 2007. ISBN 978-1-84628-636-0. doi: 10.1007/978-1-84628-636-0_1. 6
7. Simple Decision Processes. In James N. Webb, editor, *Game Theory: Decisions, Interaction and Evolution*, Springer Undergraduate Mathematics Series, pages 23–35. Springer, London, 2007. ISBN 978-1-84628-636-0. doi: 10.1007/978-1-84628-636-0_2. 12
8. Static Games. In James N. Webb, editor, *Game Theory: Decisions, Interaction and Evolution*, Springer Undergraduate Mathematics Series, pages 61–87. Springer, London, 2007. ISBN 978-1-84628-636-0. doi: 10.1007/978-1-84628-636-0_4. 8, 11, 12, 16, 20
9. Committee Releases Report Revealing New Information on Equifax Data Breach. <https://republicans-oversight.house.gov/report/committee-releases-report-revealing-new-information-on-equifax-data-breach/>, December 2018. 88, 93, 95, 96
10. J. McKenzie Alexander. Evolutionary Game Theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2021 edition, 2021. vii, 4, 33, 34, 36, 39, 41, 42, 45, 46, 48, 49, 50, 51, 52, 53, 54, 55

11. Ali Alshawish, Mohamed Amine Abid, Hermann de Meer, Stefan Schauer, Sandra König, Antonios Gouglidis, and David Hutchison. G-DPS: A Game-Theoretical Decision-Making Framework for Physical Surveillance Games. In Stefan Rass and Stefan Schauer, editors, *Game Theory for Security and Risk Management: From Theory to Practice*, Static & Dynamic Game Theory: Foundations & Applications, pages 129–156. Springer International Publishing, Cham, 2018. ISBN 978-3-319-75268-6. doi: 10.1007/978-3-319-75268-6_6. [vii](#), [63](#), [64](#), [65](#), [66](#), [67](#), [68](#), [69](#), [70](#), [72](#), [73](#), [74](#), [75](#), [76](#), [77](#), [78](#), [80](#), [82](#), [99](#)
12. Equifax Careers. Security, Technology Governance & Compliance. <https://careers.equifax.com/en/teams/security-technology-governance-compliance/>. [90](#)
13. Essam EL-Seidy. On the Behavior of Strategies in Hawk-Dove Game. *Journal of Game Theory*, 5(1):9–15, /26/2016. ISSN 2325-0054.
14. Josh Fruhlinger. Equifax data breach FAQ: What happened, who was affected, what was the impact? <https://www.csoonline.com/article/3444488/equifax-data-breach-faq-what-happened-who-was-affected-what-was-the-impact.html>, February 2020. [85](#), [86](#)
15. Herbert Gintis. Game Theory: Basic Concepts. In *Game Theory Evolving*, A Problem-Centered Introduction to Modeling Strategic Interaction - Second Edition, pages 32–51. Princeton University Press, rev - revised, 2 edition, 2009. ISBN 978-0-691-14050-6. doi: 10.2307/j.ctvc4gm4gjh.6. [38](#)
16. J. Hofbauer, P. Schuster, and K. Sigmund. A note on evolutionary stable strategies and game dynamics. *Journal of Theoretical Biology*, 81(3):609–612, December 1979. ISSN 0022-5193. doi: 10.1016/0022-5193(79)90058-4. [42](#)
17. Athanasios Kehagias and Georgios Konstantinidis. Cops and Robbers, Game Theory and Zermelo’s Early Results. *arXiv:1407.1647 [cs]*, July 2014. [64](#)
18. John Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951. ISSN 0003-486X. doi: 10.2307/1969529. [13](#), [17](#)
19. Günther Palm. Evolutionary stable strategies and game dynamics for n-person games. *Journal of Mathematical Biology*, 19(3):329–334, July 1984. ISSN 1432-1416. doi: 10.1007/BF00277103. [42](#)
20. Jeffrey Pawlick and Quanyan Zhu. *Game Theory for Cyber Deception: From Theory to Applications*. Static & Dynamic Game Theory: Foundations and Applications. Birkhäuser, Cham, Switzerland, 2021. ISBN 978-3-030-66065-9. [62](#)
21. Jeffrey Pawlick and Quanyan Zhu. Obfuscation. In Jeffrey Pawlick and Quanyan Zhu, editors, *Game Theory for Cyber Deception: From Theory to Applications*, Static & Dynamic Game Theory: Foundations & Applications, pages 49–58. Springer International Publishing, Cham, 2021. ISBN 978-3-030-66065-9. doi: 10.1007/978-3-030-66065-9_5. [99](#)

22. Jeffrey Pawlick and Quanyan Zhu. A Taxonomy of Defensive Deception. In Jeffrey Pawlick and Quanyan Zhu, editors, *Game Theory for Cyber Deception: From Theory to Applications*, Static & Dynamic Game Theory: Foundations & Applications, pages 37–48. Springer International Publishing, Cham, 2021. ISBN 978-3-030-66065-9. doi: 10.1007/978-3-030-66065-9_4. [vii](#), [57](#), [58](#), [59](#), [60](#), [61](#)
23. Erich Prisner. *Game Theory Through Examples*. American Mathematical Society, Washington, UNITED STATES, 2014. ISBN 978-1-61444-115-1. [6](#), [7](#), [18](#)
24. Lisa Rajbhandari and Einar Snekkenes. Utilizing Game Theory for Security Risk Assessment. In Stefan Rass and Stefan Schauer, editors, *Game Theory for Security and Risk Management: From Theory to Practice*, Static & Dynamic Game Theory: Foundations & Applications, pages 3–19. Springer International Publishing, Cham, 2018. ISBN 978-3-319-75268-6. doi: 10.1007/978-3-319-75268-6_1. [4](#)
25. Michael Riley, Jordan Robertson, and Anita Sharpe. The Inside Story of Equifax’s Massive Data Breach. *Bloomberg.com*, September 2017. [85](#), [86](#)
26. Peter D. Taylor and Leo B. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1):145–156, July 1978. ISSN 0025-5564. doi: 10.1016/0025-5564(78)90077-9. [48](#)
27. Jorgen Weibull. Dominated Strategies. In *Evolutionary Game Theory*. The MIT Press, Cambridge, MA, 1995. ISBN 978-0-262-73121-8. [53](#)

