CHAPTER 10

Game Theory Case Study: Ransomware

Introduction

Game theory is an excellent tool for analyzing complex, competitive situations. Cybersecurity is concerned with just such situations involving attackers, defenders, and others like regulating entities. Within game theory, "a particular game is defined when the choices open to the players in each situation, the situations defining the end of play, and the payoffs associated with each play-terminating situation have been specified."[1] This does not apply to a general cybersecurity situation. We use the term "infinite game" to describe something like a game such that: players may start or stop playing the game at any time; the end of play may never be defined; and the rules governing play may change at any time without all (or any) players knowing about the changes. While infinite games do not meet the strict definition of a game, both games and infinite games can be analyzed with the rigor and tools of game theory.

To analyze ransomware attacks as part of the infinite cybersecurity game, we describe and analyze several games (in the strict game theory sense) and discuss how these games relate to the infinite game.

[1] Rapoport, Anatol. Two-person Game Theory. Courier Corporation, 1999.

J. G. Oakley et al., *Theoretical Cybersecurity*, https://doi.org/10.1007/978-1-4842-8300-4_10

Our focus is on enabling the reader to understand the value of applying game theory techniques to real cybersecurity problems. Some of the techniques discussed in the previous chapter are applied here to situations from a defender's perspective, and some from an attacker's perspective.

In all cases, the applications are intended to help the reader make better security-focused decisions. This does not mean the applications will all be in the cybersecurity domain. For example, understanding applications of game theory to negotiations is an important tool for minimizing the payment a defender will be required to give an attacker. The expected value of such payments, together with the likelihood of ending up in such a negotiation and many other considerations, can be used to accurately reason over potential ransomware attack outcomes. Of course, the level of accuracy depends on the data available to estimate certain unknown quantities. Here, a government agency, such as the FBI, can be instrumental in helping estimate these quantities in order to perform a proper analysis.

A simplified version of the global problem of ransomware is modeled as involving the steps outlined as follows. While no single ransomware attack will involve all the steps that follow, each step is relevant to a possible ransomware attack. Also, some of these steps may occur simultaneously or in very quick succession. For example, payload activation and making a ransom demand may happen together. Note that while there are many effects that could be considered during the payload activation step, we narrow our attention to data encryption only.

- **Attack capability development.** The attacker develops, purchases, or otherwise obtains the ability to implement some attacker steps of a ransomware attack. For example, this may include the ability to detect whether a defender has initiated a cybersecurity mitigation measure after the payload has been activated.

- **Defense capability development.** The defender implements security measures to mitigate the threat of a ransomware attack and enables some defender steps of a ransomware attack. For example, this may include the ability to recover data from a backup.
- **Target selection.** The attacker determines whether to attack the defender and, if so, which defender information systems to target.
- **Payload deployment.** This involves initial access to an information system, network discovery, vulnerability exploitation, and many other activities normally associated with a cybersecurity attack of an information system. The end result of this is a deployed payload and the ability to activate the payload.
- **Payload activation.** A portion of the information is disabled via data encryption.
- **Ransom demand.** The ransom demand is presented to the target.
- **Cybersecurity mitigation.** The defender recovers or replaces the data without receiving the encryption key from the attacker.
- **Retaliation.** If a mitigation attempt is detected, the attacker renders further damage and permanently withholds the key.
- **Ransom response.** The defender decides whether to pay the ransom demand, enter into negotiations with the attacker, or neither.
- **Payoff negotiation.** The defender and attacker negotiate terms of a ransom payoff.

- **Payoff.** The terms of the ransom payoff are enacted or bluffed. This may include providing the encryption key, paying the negotiated payoff amount, or both.
- **Recovery.** The defender takes measures to ensure the payoff deactivation is complete, the data is usable and uncorrupted, the payload itself is wiped from the information system, and the information system vulnerabilities that allowed payload deployment are patched or otherwise mitigated.

It is beyond the current scope to give a complete analysis that thoroughly considers each step. Instead, we examine only certain steps in isolation or in combination with each other. Our aim is to help the reader gain an intuition for the kind of considerations involved in a game theory analysis and to enable the reader to understand the value of applying game theory techniques to real cybersecurity problems.

In addition to considering steps in combination, it is valuable to consider attacks in combination. One ransomware attack involves one attacker entity (person, state, group, etc.) and one defender entity. Many ransomware attacks involving the same attacker may take place in concert with each other. For a complete analysis, it is important to consider many potential attacker and defender types. Likewise, it is important to consider many attacks that may occur simultaneously or in sequence rather than just a single attack in isolation.

Our exploration starts with the final step involving a simple attacker versus defender game arrived at via a number of simplifying assumptions. This approach provides the analyst with results that assist with understanding more realistic behaviors in a more complex scenario. Throughout the chapter, we introduce additional complexity as we include additional steps in reverse order. When analyzing earlier steps, this approach allows us to treat later steps as a subgame that can be summarized with expected values of outcomes, as presented in the previous chapter.

Payoff and Recovery

The first step we consider is the payoff. In this step, the attacker and defender both take actions. The payoff negotiation could include any variety of potential actions for either side. For now, we assume the agreement is two-fold. First, if the defender provides a payment of the agreed amount, then the attacker will provide an encryption key that will allow the defender to recover their data. Second, if the defender does not provide the payment of the agreed amount by the agreed time, the attacker – as punishment for not upholding the agreement – will permanently withhold the encryption key. The potential actions and outcomes for this game are shown in Figure 10-1. The outcomes are described in terms of attacker and defender payoffs.

Looking at this game in isolation, there is no incentive for the attacker to give the key after the defender provides the payment. That is, the payment to the attacker is the same whether the key is provided or not.

However, if the key is not provided, then the attacker can immediately return to the ransom demand step to try to get an additional payment. For the attacker who is not concerned about reputation, there is no downside to this.

Regardless of whether the key is provided, the recovery step is critical to preventing the attacker from reactivating the payload – which may survive a data wipe – at a later time.

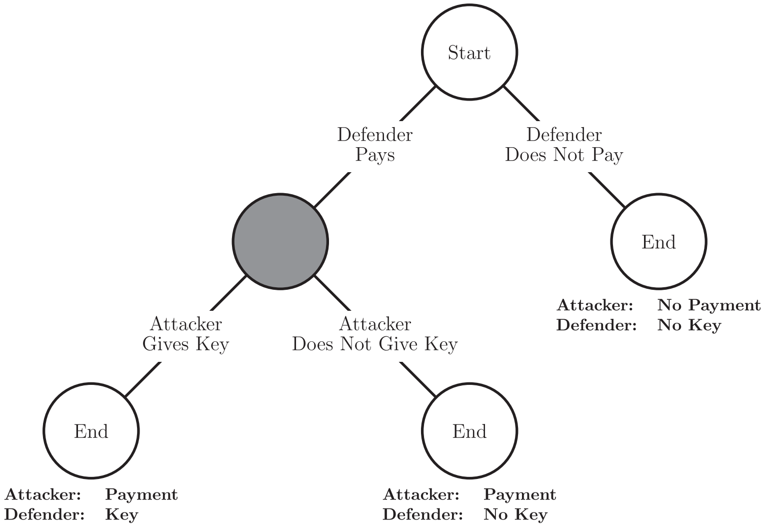

Figure 10-1. *Payoff game*

Reputation

Let's take a closer look at the role of reputation, beginning with the game presented in Section Payoff and Recovery. Before deciding whether to pay, the defender has some belief about how the attacker would respond were a payment provided. (Even if the defender does not actually have a belief about this, this can be modelled mathematically as having an equal belief in each possible outcome.) If the belief that the attacker will not provide the key is strong enough, then this belief justifies the defender withholding payment. In this case, justification for the belief itself ultimately determines how justified the defender is in withholding payment.

For this reason, the attacker may choose to care about reputation. For reputation to matter to the attacker, there must be a way for subsequent defenders targeted by the attacker to know or at least form beliefs about the outcome of the current payoff. This may happen via government-provided statistics on ransomware payoffs in general. It may happen via statistics associated with this specific attacker. In this latter case, the attacker would have to establish an identity and make that identity known at some step prior to the payoff step.

This identity, or rather the marker of the identity, now presents an opportunity to other attackers. Let's say the attacker so far is attacker A, who wants to have a reputation of honoring the terms of the payoff negotiation. Suppose there is an attacker who is motivated to withhold the key. For example, consider attacker B, who wants to take advantage of A's reputation. If B can use the same marker of identity used by A, then B can masquerade as A. Then, when it comes to a payoff structure like Figure 10-1, B can reap the benefits of A's reputation. That is, despite B having no intention of providing the key, the defender may still pay because A has a strong reputation of honoring the terms agreed to in the previous step.

Is it in B's interest to withhold the key in this case? That depends on B's goals, which may be incompatible with A's goals. If B wants to cause chaos or has a grudge against the defender, then it may be worth it to B to withhold the key and allow A's reputation to change because of it. If B wants to make as much money as possible over the course of many attacks, it may be worth it to B to provide the key in such situations to keep A's reputation strong. There very well could be situations in which there is a high cost of building a reputation, but maintaining one carries no cost and has other benefits. This would explain why B would masquerade as A rather than just build a reputation independently.

Now suppose B is a nation state and wants to inflict as much damage on another nation as possible. Here, B's strategy is to target defenders in the opposing nation and masquerade as A to take payments without providing keys to defenders. Masquerading allows for a high chance of receiving payments without incurring the cost of building a reputation. Receiving payments and withholding keys damages the opposing nation.

For this analysis, we simplify by holding fixed the cost of getting to the payoff step while masquerading as attacker A and the structure of the payoff game (fixed to that of Figure 10-1). To make calculations quite easy, we also simplify our model of A's reputation and our model of the defender. A's reputation is a number between zero and one calculated as a ratio. The top of the ratio is the total number of times A or B has provided the key to a defender. The bottom of the ratio is the total number of times A or B has been paid by a defender. The defender will pay if the reputation of A is above a threshold and will not pay otherwise. The primary variable is whether B will provide a key after the defender has provided payment.

Suppose A has been paid by a large number of defenders and has stopped attacking altogether. Before B's first attack, A's reputation is one. Once B starts attacking, A's reputation drops more or less quickly depending on how often B withholds the key from a defender. Figure 10-2 shows three possibilities. The solid line shows the reputation decay if B withholds the key for every defender. This quickly drops below the threshold for the defender providing a payment. The dotted line shows the decay if B withholds the key for every other defender. This drops below the threshold more slowly. If B provides the key for a percentage of defenders that matches the threshold, A's reputation never drops below the threshold. This represents nearly the greatest damage B can do to the opposing nation in this simplified scenario.

There is a much richer depth of reputation analysis possible even for this question of how often to provide or withhold the key from a defender; more so when including all the steps of a ransomware attack. Indeed, the analysis presented is embarrassingly oversimplified. Proper analysis requires more advanced techniques and the relaxation of simplifying assumptions. While this deeper analysis is beyond the scope of this work, the fact that this analysis could be done is immanently relevant. Considering reputation over sequences of attacks is more realistic and more complex. The results of a proper analysis offer a great advantage to whichever players (attacker, defender, or regulator) obtain them. Given the expertise (i.e., expense) required to do this, it is likely only justifiable for very large corporations, very large criminal organizations, or nation-states (which includes attackers, defenders, and regulators) to conduct.

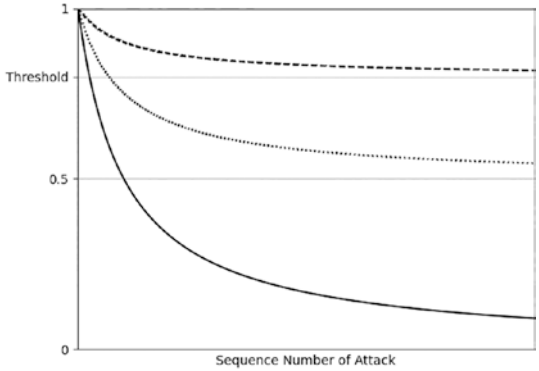

Figure 10-2. *Reputation Decay*

Payoff Negotiation

A proper treatment of game theory concepts for negotiation is well beyond the scope of this work. Whole textbooks have been written on just this topic. Instead, we consider how to incorporate an important finding of the payoff step analysis into the strategy for payoff negotiation. Here, the players' strategies consist of three elements: the structure of the payoff step, the terms the attacker agrees to enact, and the terms the defender agrees to enact.

The structure of the game of Figure 10-1 that puts the defender at a particular disadvantage is the sequential nature of the payoffs. The defender makes a payment or not. Even after a payment is made, the attacker is free to choose to provide the key or change a prior decision about providing the key.

Taking this into account, the payoff negotiation should have a different expected outcome for payoff structures that permit payoff terms to be fulfilled sequentially versus structures that employ one of the many protocols for simultaneous implementation of terms.

In considering ransomware steps that occur before negotiation, we can treat negotiation and payoff as a subgame that can be reduced to expected values over possible outcomes. The significant outcome variations to consider will include the attacker providing the encryption key and the attacker not providing the key. For example, in the response game of Figure 10-3, the left, gray node is the beginning of this subgame. The child nodes of this are the two significant outcome variations just described. One represents the average overall negotiations and payoffs that include the attacker providing the encryption key, the other that do not include providing the key. The payments associated with these child nodes are expected values based on a full analysis of the subgame.

Note well that to calculate the expected values for a summarized game, the likelihood of each outcome must be calculated (or at least estimated). This is best done by considering reputation and beliefs of the players.

Ransom Response, Mitigation, and Retaliation

Once the ransom demand has been made, it is up to the defender to decide whether to engage the demand, try to mitigate the damage, or just to accept the loss of data. As mentioned before, we simplify the sub-games by representing only the expected value of outcomes. Here, the important outcomes cover cases in which the key is provided or not, different payment amounts are provided, different amounts of effort are spent, and whether the attacker successfully inflicts additional damage.

The effort required by a player refers to the effort required to negotiate or mitigate. To make the example manageable, we assume the effort required is similar for all negotiations, is similar but different for all mitigation outcomes, and is zero when not negotiating or mitigating. In the language of the previous chapter, effort is a kind of cost. There may be other kinds of cost to consider, too.

The additional damage an attacker may inflict is relevant only when the defender attempts to mitigate the attack instead of negotiate. This damage may be as simple as corrupting or deleting the encrypted data. This is relevant if the defender does not pair the mitigation attempt with a successful recovery step or if the defender tries to recoup value from the encrypted data (like waiting until quantum computers can crack the encryption). It is not relevant if the defender recovers the data via another means (like a backup) and successfully implements a recovery step.

The salient outcome possibilities are shown in Figure 10-3.

To reason over these outcomes, the players must have a way to rank the outcomes. One way is to convert the outcome types into common units of payoff for each player. However, it is not necessary to have a formula that produces a single number to be used for the ranking. Such a formula that has certain properties is called a von Neumann-Morgenstern utility function and is useful for a great many game-theoretic techniques involving rationality, risk aversion, expected utility, and more. In the language of the previous chapter, utility is equivalent to payoff.

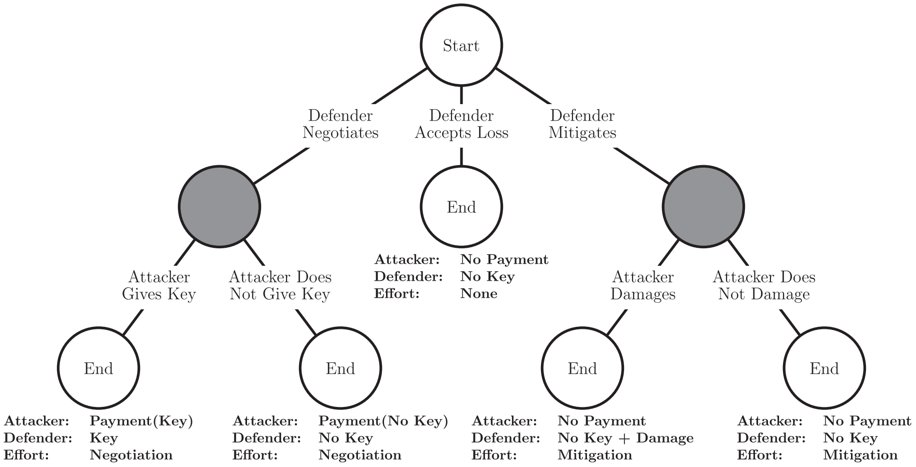

Figure 10-3. *Response game*

Having such a function available is common in this situation. This is because cybersecurity decisions relevant to ransomware attacks are often phrased in terms of monetary cost. This means the cost of having or not having the encryption key and of expending a negotiation- or mitigation-level of effort are explicitly converted into expected dollar amounts. Even when schedule is an important cost to consider, the impact of schedule can often be converted into an equivalent monetary cost. A formula that produces a monetary cost for each outcome can be treated as a von Neumann-Morgenstern utility function.

In the absence of a von Neumann-Morgenstern utility function, simply having a ranking of the outcomes provides the foundation to apply quite powerful game theory tools that enable excellent insights into optimal choices and expected behavior. For a complete analysis based on rankings alone, each player must know or estimate the outcome rankings of all players. There is an important variation on this in which explicit rankings are not known or estimated, but the belief about different rankings is known. This also provides the necessary foundation for a robust game-theoretic analysis.

One insight from this outcome-ranking consideration is that the defender can only be expected to negotiate if the effort required to negotiate is not too high and the chance of a mitigation attempt being successful is not too high. To be more precise, it is the defender's belief about this effort and chance that matter most. Carefully considering beliefs and reputation come into play in calculating or estimating these.

As the subgames of the response game were summarized using expected values, the response game itself is a subgame of a larger game. Based on the goal of a larger analysis, this can be summarized a number of ways. In one sense, the game cannot be simplified as the outcomes each have a unique combination of key/no key/damage and level of effort. However, if the focus of analysis is on the value of mitigation, it would be useful to summarize the game into two outcomes given as expected cost when the defender mitigates or does not mitigate. This kind of analysis would help justify a potential defender spending money on cybersecurity mitigation equipment, software, and personnel prior to a ransomware attack.

Activation and Demand

We assume activation and demand steps occur near-simultaneously, so they can be considered together as one step. To simplify further, consider a situation in which the presentation of a demand to even one defender leads to a widespread update to antivirus signatures that would render further attempts to deploy payloads unlikely to succeed.

In this idealized (and admittedly unrealistic) case, an attacker can choose to activate payloads at different times. This may help space out the load on resources required for negotiation and payoff steps. However, the more time passes between deployment and activation for a given defender, the more likely that defender will discover and remove the payload prior to activation. As the number of deployed payloads increases, the elapsed time between the first and last activated payload increases, resulting in an increase in the chance of a defender discovering and removing the last payload to be activated. This diminishing return limits the maximum expected value of the number of payments an attacker will receive.

In a more realistic case, an attacker is not limited to a single or even a fixed number of deployment techniques. There are also techniques that are not thwarted by software of any kind – for example, exploiting humans to compromise the network. Still, the lesson of the idealized case applies. Once a payload is deployed, it is in the attacker's interest to activate it before too much time passes.

Notice the analysis of this step is about optimizing expected payment for the attacker. In this sense, it is an optimization analysis, not a game-theoretic analysis. It does depend on game theory analysis, though. This is because the determination of expected value of payment relies on the game theory analysis of steps that come after deployment and activation and demand steps. This mix of analysis types is common in evaluating complex games.

Deployment

Deployment involves initial access to an information system, network discovery, vulnerability exploitation, and many other activities normally associated with a cybersecurity attack of an information system. As with negotiation, a proper treatment of game theory concepts for Deployment is beyond the scope of this work, as the research in this area goes back decades.

Of particular interest and usefulness are the graph-based methods. A graph is a collection of nodes with edges between the nodes. A protected network can be modelled as a graph with each asset (router, domain controller, host, etc.) represented as a node and steps in an attack path from one node to another represented as directed edges. These nodes and edges can be imbued with amplifying information, such as installed software, known exploits, firewall rules, etc.. Such a graph can be used to perform a static analysis of the overall level of security of a network, perform a dynamic assessment of the risk of a specific set of assets being compromised given real-time alerts, and provide a great deal of insight of many other kinds.

For the current work, the details of such analyses are less important than knowing they can be carried out in a structured, meaningful manner. We can model this step as having two outcomes, the attacker is successful or not successful in deploying the payload. Both of these outcomes can be associated with a probability, defender cost, and attacker cost.

On the defender side, these probabilities are primarily based on the specific network being compromised (it's topology, controls, policies, vulnerabilities, etc.), the manner in which the network is protected (as determined by the defender's capability development step), and the skill of the people monitoring and protecting the network. On the attacker side, these probabilities are primarily based on the specific network being compromised, the manner in which the network is attacked (as determined by the attacker's capability development step), and the skill of the people monitoring and attacking the network.

Selection

At this step, there are no defenders, only potential defenders. In this sense, while the ransomware attack has begun, the ransomware game (in the game theory sense) has not begun. That is, the infinite ransomware game (itself a sub-infinite-game of the infinite cybersecurity game) has begun, but the strict conditions for a game have not yet been met.

A number of factors could be considered here. How much financial liquidity does a potential defender have? How good is their network security? How well-trained are the people using or protecting the network?

Which software or configuration vulnerabilities are unmitigated? How likely are they to hire a professional negotiator or turn to the FBI for help? While all of these considerations, and many more, are relevant, they can be difficult or impossible to assess prior to the deployment step.

The attacker's capabilities may dictate the deployment method. In turn, this may dictate or restrict the selection criteria. For example, an attacker may choose to compromise a security information and event management (SIEM) product. In this case, the potential defenders available to be attacked are limited to those who utilize the compromised SIEM.

To conduct a high-quality analysis of the ransomware game as a whole, it is necessary to collect and incorporate statistics on the methods used to deploy payloads. This can help the analyst better understand the selection step. The goal of analysis here is to understand the conditions that make it more or less likely to be selected for an attack.

Capability Development

Attacker capability development involves developing, buying, training, or otherwise obtaining the skills and tools necessary to perform each of the steps of a ransomware attack. Defender capability development involves obtaining the skills and tools necessary to defend against a ransomware attack. We examine how the different steps of a ransomware attack inform the decisions made during capability development. While we focus on the defender's decisions, we will consider some attacker decisions along the way.

Deployment

This is where the bulk of the budget is spent for a cybersecurity operations center. Selecting a SIEM, an antivirus solution, personnel, training, policies, security controls, and much more comes into play. Even scratching the surface of the mass of considerations for this is beyond the scope of this work. In aggregate, these decisions impact the likelihood that a randomly selected defender will succumb to the exploits of the attacker during payload deployment.

Activation, Demand, Mitigation, Retaliation, Recovery

At this point in the attack, the network is compromised. The defender has failed to prevent the data from being encrypted against the defender's will. Whether or not the demand has been made, the defender can cut the overall attack short and minimize the impact of the attack by detecting and mitigating the encryption.

The mitigation may involve restoring data from a backup, using a separate fail-over network with duplicate data, or a variety of other solutions – all of which can be cost-justified if the threat of ransomware attacks becomes too great. Regardless of the mitigation, there is the threat of retaliation. We assume the attacker can at least delete the data that is encrypted. Undoubtedly, there is more the attacker can delete and corrupt. If all that can be affected by the attacker is covered by the mitigation plan, then retaliation is an empty threat. Implicit in this condition is that the mitigation plan includes a full recovery plan. Otherwise, the attacker may be able to reactivate or redeploy with minimal effort and a higher demanded payment.

Response, Negotiation, Payoff

If mitigation is not an option, the response options are to accept the loss or negotiate. At this point, the most notable aspect of capability development that may help is training in negotiation tactics, not a cybersecurity discipline.

Attacker Types

Before getting to the target selection step, we consider four types of attackers. First, there is the low-capability attacker. They can only take advantage of defenders who have not provided a minimum foundation of protection.

Next, there is the high-capability attacker. This could range from the cyber arm of an organized crime operation to a very capable individual. This type of attacker can develop robust capabilities and use them to compromise a target network. These attackers are a real threat to any organization.

The third type is the optimized attacker. This is a high-capability attacker who designs campaigns to optimize the use of resources. This is distinct from the high-capability capability attacker who may attack opportunistically without regard for the efficient use of resources. In game theory terms, the optimized attacker is a rational player – this is a technical term defined mathematically. Contrasting this, the (non-optimized) high-capability attacker is either non-rational or has bounded rationality.

Finally, we have the APT attacker. This can be a high-capability attacker or an optimized attacker. The distinguishing feature is the APT attacker has a very high budget of time and money, whereas the other attacker types are comparatively quite limited.

In terms of the 1-9-90 principle in Chapter 3, the APT attacker is the 1, the high-capability and optimized attackers are the 9, and the low-capability attacker is the 90.

Target Selection

Here, we ask the question, "how can the likelihood of attack be minimized?" We consider one aspect of this with a simplified game. The purpose of this game is to demonstrate the weakness of only considering aggregate probabilities in analyzing potential cybersecurity outcomes.

To narrow our focus, we look at options for selecting software to protect a network. Intentionally leaving the type of software vague, suppose there are three options. Each option has different levels of cost, popularity, and provided security, as shown in Table 10-1.

Suppose a new defender is making a cybersecurity plan and needs to select one of these options. We assume the popularity of each option is sufficiently high that the defender's choice will have a negligible effect on popularity. We also assume the cost of each option is within the budget of the defender. Next, we consider the security of each option.

Here, it is useful to recall the chapter on Infinite Cybersecurity (Chapter 8). The key lesson to apply here is that the defender should not try to "win" the cybersecurity game. This applies equally to ransomware attacks and other types of attacks. Rather, the defender should try to keep playing as long as possible. To this end, the software options can help the defender capitalize on the access to and information of their own attack surface, systems, and organization.

Table 10-1. *Attributes of software options for network protection*

Option	Cost	Popularity	Security
A	High	Medium	High
B	Medium	High	Highest
C	Low	Low	Medium

We assume that no software choice will make a defender protected against APT threats. This is a cost of playing the infinite game that must be accepted. Software option C provides for the minimum foundation needed to protect against low-capability attackers – when coupled with capable and well-trained professionals (a must that is often overlooked even for the highest-budget operations centers). Software options A and B provide good protection against low- and high-capability attacker types when coupled with capable and well-trained professionals. The difference is that option B has a greater record of success. Keep in mind that this record is in an aggregate measure that does not take into account the frequency of each attacker type. When thinking only about this record of success versus cost, the defender has a clear reason to select option B, as it dominates option A in cost and security.

The most interesting case to consider is the optimized attacker. This attacker wants to optimize their gain per unit of effort. How popularity and security compare with each other numerically determines the optimal choice for the optimized attacker. In this sense, the security value for the software options can be thought of as the level of effort required by the attacker to circumvent or otherwise thwart the software. If the attacker can find or develop an exploit for either option A or option B, then they can pursue attacks against a large part of the population of users for the affected software.

Suppose the effort required to compromise option B is twice that to compromise option A and the number of users of option B is ten-fold greater than the number of users of option A. In this case, the optimized attacker is justified in targeting users of option B if the expected value of payoff is the same for targeting a user of option A and targeting a user of option B. Assuming these expected values are the same, compromising option B gives a 5-fold yield vs. compromising option A (twice the effort for ten times the payoff). Strictly speaking, the attacker's beliefs about popularity, level of effort, and expected payoffs are more important than the actual values.

From the defender's perspective, this insight about the optimized attacker may change the risk assessment for each option. As we have already pointed out, the defender can calculate or estimate the equivalent cost of compromise weighted by the likelihood of being targeted. Now, the defender has to reason over the chance that selecting a more popular software option increases the likelihood of being targeted. This may well outweigh the cost savings for selecting option B over option A.

It is critical in these situations to ensure the basic foundation of security is not compromised on the basis of a popularity argument. That is, if option A's high level of security is just not high enough, then the defender is justified in selecting option B even if this comes with an increased likelihood of being targeted by optimized attackers.

Before leaving this example, we draw a comparison to another decision about how a cybersecurity budget could be spent. In this example, the main decision was about a trade-off between cost, security, and likelihood of being targeted. This same kind of analysis can be used to decide between using high-cost software (like a SIEM) with median-capability staff versus using the same software with high-capability staff. The increased cost in training or salaries can be offset by a lower expected value of the cost of compromise. This lower compromise cost comes from lowering the probability of being compromised.

Summary

We introduced one variation of a ransomware attack and analyzed it using game-theoretic and optimization methods. The presented techniques include a small sample of game theory tools for a skilled analyst to apply. The overall analysis was broken down into smaller analysis steps that were combined together into higher-level analyses incrementally. This required us to consider steps of a ransomware attack in reverse order, where the smallest games to analyze involve the final steps of an attack.

To get the best results, a multitude of aspects must be properly considered. One of the most subtle and difficult aspects to handle rightly is reputation. However, this is especially important for regulators to consider so that the regulations they impose will have a chance of yielding the effects intended.

Another subtle and difficult aspect to consider is the type of attackers. The proper incorporation of this concept requires understanding the methods, reasons, beliefs, and rationality of the attacker types. In addition, the number of each type is important to take into account. However, the number of each type may be difficult to determine. Thus, this kind of information may only be available via nation-state-level entities.

Two areas of research that we were not able to present in depth are the game theory of negotiation and attack graphs. These require great knowledge and skill to apply correctly. However, they offer excellent value when properly applied to the analysis of ransomware attacks or of cybersecurity in general.

A major goal of this work is to enable readers to make disciplined choices about how to use available resources. For the cybersecurity professional, this includes the time and budget within the given scope. For the executive, this includes taking into account the likelihood and severity of outcomes as a function of how much the cybersecurity budget is increased or how restrictions are put on the use of this budget. In either case, reasoning over costs, outcomes, and risk can be significantly aided using the robust tools of game theory.