

CHAPTER 25

Incentives and Information Security

Ross Anderson, Tyler Moore, Shishir Nagaraja,
and Andy Ozment

Abstract

Many interesting and important new applications of game theory have been discovered over the past 5 years in the context of research into the economics of information security. Many systems fail not ultimately for technical reasons but because incentives are wrong. For example, the people who guard a system often are not the people who suffer the full costs of failure, and as a result they make less effort than would be socially optimal. Some aspects of information security are public goods, like clean air or water; externalities often decide which security products succeed in the marketplace; and some information risks are not insurable because they are correlated in ways that cause insurance markets to fail.

Deeper applications of game-theoretic ideas can be found in the games of incomplete information that occur when critical information, such as about software quality or defender efforts, is hidden from some principals. An interesting application lies in the analysis of distributed system architectures; it took several years of experimentation for designers of peer-to-peer systems to understand incentive issues that we can now analyze reasonably well. Evolutionary game theory has recently allowed us to tie together a number of ideas from network analysis and elsewhere to explain why basing peer-to-peer systems on rings is a bad idea, and why revolutionaries use cells instead. The economics of distributed systems looks like being a very fruitful field of research.

25.1 Introduction

Over the last 6 years, people have realized that security failure is caused at least as often by misaligned incentives as by technical design mistakes. Systems are particularly prone to failure when the person guarding them is not the person who suffers when they fail. The tools and concepts of game theory and microeconomic theory are becoming just as important as the mathematics of cryptography to the security engineer.

In this chapter, we present several live research challenges in the economics of information security, many of which bear on problems in various branches of game theory. We first consider misaligned incentives, and externalities: network insecurity is somewhat like air pollution or traffic congestion, in that people who connect insecure

machines to the net do not bear the full consequences of their actions and so do not make a socially optimal investment in protection. Next we examine the role of asymmetric information and the capacity for hidden action: games where one principal has more knowledge of the game state than her opponent, or games where she can make moves that become known only with a certain probability.

The difficulty in measuring information security risks presents another challenge: these risks cannot be better managed until they can be better measured. Auctions and markets can help in various ways to measure the security of software and thereby reduce the information asymmetry prevalent in the software industry. We also examine the problem of insuring against attacks. The local and global correlations exhibited by different attack types largely determine whether an insurance market in the associated risks is feasible.

The structure of computer networks can also have a great impact on player incentives. One topical example is that the effort devoted to censorship resistance in peer-to-peer systems depends upon whether the application design empowers players to choose which files to share or randomly distributes them. This realization enables us to model solidarity in networks that may come under selective attack.

An even more striking example is how network topology can exacerbate the impact of viruses or susceptibility to targeted attacks. The regular networks, or random networks, commonly used in modeling do not behave the same way as real-world networks, which are better approximated by scale-free models. Scale-free networks turn out to be more robust against random failure but more vulnerable to targeted attack. We finally present a model that uses ideas from evolutionary game theory to explore the interaction between attack and defense strategies, and we provide a framework for evaluating strategies in networks where topology matters.

25.2 Misaligned Incentives

One of the observations that drove initial interest in security economics came from banking. In the United States, banks are generally liable for the costs of card fraud; when a customer disputes a transaction, the bank must either show that she is trying to cheat them or refund her money. In the United Kingdom, the banks had a much easier ride: they could often get away with claiming that the ATM system was “secure,” so a customer who complained must be mistaken or lying. “Lucky bankers,” one might think; yet it turned out that UK banks spent more on security and suffered more fraud. How could this be? Banks appear to have been suffering from a moral-hazard effect: UK bank staff knew that customer complaints would not be taken so seriously, so they became lazy and careless. This situation led to an avalanche of fraud.

Another observation came from the state of the antivirus software market around the year 2000. People were not spending as much money on protecting their computers from infection as would have been ideal. Why not? Well, at that time, a typical virus payload was a service-denial attack against the Web site of a company like Amazon or Microsoft. While a rational consumer might well spend \$20 to prevent a virus from trashing his hard disk, he might not do so just to prevent an attack on Bill.

Legal theorists have long known that liability should be assigned to the party that can best manage the risk. Yet everywhere we look, we see online risks that are poorly

allocated. The result is privacy failures and even protracted regulatory tussles. For example, the United States has seen widespread debate about medical privacy over the last 10 years: from the passage of the Health Insurance Portability and Accountability Act, through the initial regulations made under the Act by President Clinton, the later regulations made by President Bush, and the recent claims that the law fails to protect health privacy while providing a gold-mine for security vendors. The root problem is that medical information systems are purchased by hospital directors and insurance companies, whose interests in account management, cost control and research are not well aligned with the patients' interests in the privacy of their health records.

25.2.1 Applications of Game Theory

Game theory can provide a means of better understanding the outcome of security decisions made by self-interested individuals. Information security levels often depend on the efforts of many principals, leading to suboptimal security investment whenever decisions are uncoordinated. The level of security investment generally depends on the investor's own costs and benefits, the investment decisions of others, and the manner in which individual investment translates to outcomes. System reliability can depend on the sum of individual efforts, the minimum effort invested, or the maximum effort invested. Programming, for example, might be down to the weakest link (the most careless programmer introducing a fatal vulnerability) while software validation and vulnerability testing might depend on the sum of everyone's efforts. There can also be cases where the security depends on the best effort – the effort of a star cryptanalyst. These different models have interesting effects on whether an appropriate level of defense can be provided and what policy measures are advisable.

A simple model by Varian provides interesting results when players choose their effort levels independently. For the total-effort case, system reliability depends on the agent with the highest benefit-cost ratio, and all other agents free ride. In the weakest-link case, the agent with the lowest benefit-cost ratio dominates, since any additional effort is wasted. Systems become increasingly reliable in the total-effort case as more agents are added, but they become increasingly unreliable in the weakest-link case. What are the implications? One is that software companies should hire more software testers but fewer (more competent) programmers.

Work such as this has inspired other researchers to consider interdependent risk. A recent influential model by Kunreuther and Heal notes that the security of a group often rests on each of its members: an individual taking protective measures creates positive externalities for others that in turn may discourage their own investment. This result has implications far beyond information security. The decision by one apartment owner to install a sprinkler system affects his neighbors' decisions to install their own systems; airlines may decide not to screen luggage transferred from other carriers who are believed to be careful with security; and people thinking of vaccinating their children against a contagious disease may choose to free-ride off the herd immunity instead. In each case, several widely varying Nash equilibrium outcomes are possible, from complete adoption to total refusal, depending on the levels of coordination between independent actors.

25.2.2 Network Effects and Deployment

Game theory is also used to ascertain how network effects impact the level of security investments. In particular, many security technologies face bootstrapping problems. The benefit that these technologies provide to players is dependent upon the number of players that adopt the technology. A bootstrapping problem exists because the cost of the technology is greater than the benefit until a minimum number of players adopt. As a result, each player waits for other players to go first, and the technology is never deployed.

Following the seminal work of Katz and Shapiro, a number of economists have examined the problem of deploying a technology that exhibits network effects. Most of this literature concludes that adoption is a coordination problem. The challenge is to coordinate the different players and to enforce their cooperation. However, the assumptions used in these models do not apply to many security technologies. For example, security technologies that are software-based can often be deployed rapidly, while the economics literature is concerned with coordinating players who must make their decisions far in advance of a slow-moving deployment. Furthermore, security technologies may not provide special benefits to early adopters.

This area is especially topical at the moment. A number of core Internet protocols are considered insecure, such as DNS and routing. More secure protocols exist; the challenge is to bootstrap their adoption. Two examples of security protocols that have already been widely deployed are SSH and IPsec. Both of these protocols overcame the bootstrapping problem because they could provide significant intraorganizational benefits (X session support and VPNs). Adoption was thus driven by organizational needs rather than the benefit that players derived from the global network. The deployment of fax machines also occurred through this mechanism: companies initially bought fax machines to connect their own offices. Limiting the players in a game to the members of some kind of club can also have interesting effects on other aspects of security, as we see below.

25.3 Informational Asymmetries

We now consider two types of informational asymmetries relevant to information security: hidden action, where the difficulty of observing others' actions facilitates certain attacks; and hidden information, where the difficulty of measuring software security has caused vendors to underinvest in quality.

25.3.1 Hidden-Action Attacks

In the theory of asymmetric information, a hidden-action problem arises whenever two parties wish to transact, but one party can take unobservable actions that impact the transaction. The classic example comes from insurance, where the insured party may choose to behave recklessly (which in turn increases the likelihood of a claim) because the insurance company cannot observe her behavior. Crossing to the security domain, this idea generalizes to a class of hidden-action attacks, which are attractive precisely because observation (and therefore punishment) is unlikely.

Computer networks are naturally susceptible to hidden-action attacks. Routers can quietly drop selected packets or falsify responses to routing requests; nodes can redirect network traffic to eavesdrop on conversations; and players in file-sharing systems can hide whether they have chosen to share with others, so some may choose to “free-ride” rather than to help replenish the system. The common element in these examples is that nodes can hide malicious or antisocial behavior from other network elements.

Hidden-action attacks may occur whenever the net utility gain from deviation is greater than the expected penalty enforced when observation is unlikely and less than the expected penalty enforced when observation is likely. (If the expected gain from an attack does not exceed the expected penalty even when actions are likely to remain hidden, then no attack should occur. If the expected gain in attacking exceeds the expected penalty even when observed, then the attack should be launched regardless of whether or not observation is likely.)

In the economics literature, hidden-action problems are dealt with by structuring contracts to induce proper behavior. For example, auto insurers use deductibles to mitigate driver recklessness. By charging customers to file a claim, insurers create an incentive for taking reasonable steps to avoid negative outcomes. The need for observation is eliminated, though not without cost: everyone has to pay, even when the insured did not act recklessly. Mechanism design, as discussed throughout the rest of this book, attempts to create systems that align all of the agents’ incentives so that the agents’ best interest is to operate as intended. A complementary approach is to alter the topology and structure of the interactions to increase observability.

One telling example comes from peer-to-peer systems. These exploit network externalities to the fullest by having large member populations with a flat topology: joining one creates the potential for collaboration with every other peer in the system. High turnover is also expected; nodes may join and leave the system rapidly. These properties lower the prospects for repeated interactions, which in turn makes cheating more likely. Inexpensive or even costless identities exacerbate the problem of unrepeated interactions while also making penalties harder to implement. In a network with these properties, nodes are predisposed to hidden action.

One solution is to change the network topology. In most peer-to-peer systems, any node can transact with any other on joining the network. While this flat topology maximizes transaction possibilities, it makes repeated transactions unlikely and observation difficult. An alternative is to adopt a network topology based on clubs of nodes with common interests. Here, nodes first transact with other members of their club to establish legitimacy. Once trust has been established inside the cluster, outside transactions can happen through established channels between groups. Such a topology facilitates self-enforcement by establishing a credible threat of observation to forestall hidden action, and by creating long-lived principals (clubs) against whom sanctions hurt.

Social networks can also be used to create better topologies. When honest players can select their friends as neighbors rather than having their neighbors randomly assigned, they minimize the informational asymmetry present during neighbor interactions. This can raise the cost of entry for an attacker as well as align the incentives between normal players. However, social networks can also lead to inefficient outcomes as players may not be exposed to diverse information and isolated players may be marginalized.

25.3.2 Hidden Information-Measuring Software Security

Another information asymmetry in information security is caused by our inability to effectively measure the security of software. Most commercial software contains design and implementation flaws that could easily have been prevented. Although vendors are capable of creating more secure software, the economics of the software industry provide them with little incentive to do so. In many markets, “ship it Tuesday and get it right by version 3” is perfectly rational behavior. Consumers generally reward vendors for adding features, for being first to market, or for being dominant in an existing market. These motivations clash with the goal of writing more secure software, which requires time-consuming testing and a focus on simplicity. Nonetheless, the problems of software insecurity, viruses, and worms are frequently in the headlines; why does the potential damage to vendor reputations not motivate them to invest in more secure software?

Vendors’ lack of motivation is readily explained: the software market is a “market for lemons.” In a Nobel prize-winning work, economist George Akerlof employed the used car market as a metaphor for a market with asymmetric information. His paper imagines a town in which 50 good used cars (worth \$2,000) are for sale, along with 50 “lemons” (worth \$1,000 each). The sellers know the difference but the buyers do not. What is the market-clearing price? One might initially think \$1,500, but at that price no-one with a good car will offer it for sale; so the market price quickly ends up near \$1,000. Because buyers are unwilling to pay a premium for quality they cannot measure, only low quality used vehicles are available for sale.

The software market suffers from the same asymmetry of information. Vendors may have some intuition about the security of their products, but buyers have no reason to trust them. In some cases, even the vendor might not have a truly accurate picture of its software’s security. As a result, buyers have no reason to pay the premium required to obtain more secure software, and vendors are disinclined to invest in protection.

Three broad research approaches have attempted to provide useful measures of the security of software: statistical, market-based, and insurance-based. The former approach relies on the application of reliability growth models to vulnerabilities and is not be discussed here. The latter two approaches are discussed below.

25.3.3 Market-Based Approaches

One possible way to measure the security of software is to rely on a market: let buyers and sellers establish the actual cost of finding a vulnerability in software or merely estimate the security of software according to their own knowledge. For example, banking standards for PIN-entry terminals specify a minimum cost of various kinds of technical compromise.

In the software business, open markets for reports of previously undiscovered vulnerabilities could provide a security metric. The bid, ask, and most recent sale prices in such a market approximate the labor cost to find a vulnerability. These prices can establish which of two products the market deems to have vulnerabilities that are less expensive to find. Alternatively, a vulnerability market of this type could be designed as an auction.

Several organizations are now actively purchasing vulnerabilities, so an open market or auction actually exists. Unfortunately, these organizations are not publishing their prices. Their business model is to provide the vulnerability information simultaneously to their customers and to the vendor of the affected product (in contrast to the previous practice of waiting until after a patch is released and then making the existence of the vulnerability public). The business models of these organizations are thus not socially optimal: they always have an incentive to leak vulnerability information without proper safeguards.

A market for software security derivatives could also enable security professionals to reach a price consensus on the level of security for a product. Contracts could be issued in pairs: the first pays a fixed value if no vulnerability is found in a program by a specific date, and the second pays the same value if vulnerabilities have been found in that program by that date. If these contracts can be issued as desired and traded via some market, then their trading price indicates the consensus opinion on the security of the program. Software security derivatives could thus conceivably be used to hedge risks by software vendors, players, software company investors, and insurance companies.

25.3.4 Insurance-Based Approaches

Another approach to measuring the security of software is to rely on insurers. The argument for insurance is that cyber-insurance underwriters assign premiums based upon a firm's IT infrastructure and the processes by which it is managed. This assessment results in both detailed best practices and, over the long run, a pool of data by which the insurance company can accurately assign a monetary value to the risks associated with certain practices or software. At the moment, however, the cyber-insurance market is both underdeveloped and underutilized. Why should this be?

One reason is the problem of interdependent risk, which takes at least two forms. Firms are 'physically interdependent' because their IT infrastructure is connected via the Internet to other entities – which implies that the work a firm performs to secure itself may be undermined by failures at other firms. Firms are "logically interdependent" because cyber attacks often exploit a vulnerability in a system used by many firms. For example, viruses or worms may have a global impact upon a specific software platform. This interdependence makes certain cyber-risks unattractive to insurers – particularly those where the risk is globally rather than locally correlated, such as worm and virus attacks, and systemic risks such as Y2K. We note in passing that many writers have called for cyber-risks to be transferred to the responsible software vendors; if this were the law, it is unlikely that Microsoft would be able to buy insurance. So far, vendors have succeeded in dumping almost all risk; but this outcome is also far from being socially optimal.

Because a firm's security depends in part on the efforts of others, firms underinvest in both security technology and in cyber insurance. At the same time, insurance companies must charge a higher premium because the risks against which they are insuring are highly correlated: this higher premium may prevent the vast majority of firms from adequately insuring themselves. As a result, cyber insurance markets may lack the volume and liquidity to become economically efficient.

25.4 The Economics of Censorship Resistance

We have seen that misaligned incentives and information asymmetries are important problems in information security that are amenable to a game theoretic analysis. Another such problem is censorship resistance.

Early peer-to-peer systems were oriented toward censorship resistance rather than music file sharing. They put all content into one pot, with the effect that quite different groups would end up protecting each others' free speech – be they Falun Gong members, critics of scientology, or aficionados of sado-masochistic imagery that is legal in California but banned in Tennessee. The question then arises whether such groups might not be better off with their own peer-to-peer systems. Perhaps they would fight harder to defend their own type of dissident, rather than people involved in struggles in which they had no interest and where they might even be disposed to side with the censor. In the file-sharing context, it might make sense to have a constellation of fan clubs, rather than one huge system – as musicians take widely different views of music sharing, remixing and other activities on the fringes of classical copyright practice.

Such questions are also of topical interest to social theorists and policy people, who wonder whether the growing diversity of modern societies is undermining the social solidarity on which modern welfare states are founded. A related question in guerrilla warfare is when combatants should aggregate or disperse.

We find peer-to-peer systems providing a “single pot,” with widely and randomly distributed functionality, such as Eternity, Freenet, Chord, Pastry, and OceanStore. Other systems, like the popular Gnutella and Kazaa, allow peer nodes to serve content they have downloaded for their personal use, without burdening them with random files. The comparison between these architectures originally focused on purely technical aspects: the cost of search, retrieval, communications, and storage. However, it turns out that incentives matter here too.

25.4.1 Red–Blue Utility Model

Danezis and Anderson introduced the Red–Blue model to analyze the trade-off between diversity and solidarity in distributed systems. We consider a network of N peer nodes. Each node n_i has a preference among two types of resource, say red and blue; one node might prefer to serve 20% red and 80% blue, while another prefers 80% red and 20% blue and the network overall contains 50% red and 50% blue. A censor who attacks the network tries to impose his own preference, perhaps 80% red and 20% blue. This action may meet the approval of some nodes, but usually not most of them.

We assign to each node n_i a preference for red $r_i \in [0, 1]$ and a preference for blue $b_i = 1 - r_i$ (note that $r_i + b_i = 1$). While each node likes having and serving resources, it prefers to have or serve a balance of resources according to its own preference r_i and b_i . So we define the utility function of a node holding T resources out of which R are red resources and B are blue resources (with $T = R + B$) as

$$U_i(R, B) = -T \left(\frac{R}{T - r_i - 1} \right) \left(\frac{R}{T - r_i + 1} \right). \quad (25.1)$$

This is a quadratic function with its maximum at $R = r_i T$, scaled by the overall number of resources T that the node n_i holds. This utility function increases as the total number of resources does, but is also maximal when the balance between red and blue resources matches the preferences of the node ($R = r_i T$ and $B = b_i T$). When nodes choose the file distribution to serve, their utility is naturally maximized.

Distributed hash tables and architectures such as Eternity, by contrast, scatter the red and blue resources randomly across all nodes n_i . If the system has a total of $\mathcal{R}$ red resources and $\mathcal{B}$ blue resources, we can define a systemwide distribution of resources (r_s, b_s) so that each node in the system holds on average:

$$r_s = \frac{\mathcal{R}}{\mathcal{R} + \mathcal{B}} \quad b_s = \frac{\mathcal{B}}{\mathcal{R} + \mathcal{B}}. \quad (25.2)$$

Each node n_i has on average a utility equal to $U(r_s T, b_s T)$. The utility each node attains in the random case is always less than or equal to the utility a node has under the discretionary model:

$$U_i(r_i T, b_i T) \geq U_i(r_s T, b_s T). \quad (25.3)$$

$U_i(r_i T, b_i T) = U_i(r_s T, b_s T)$ when $r_s = r_i$ and $b_s = b_i$ – in other words, when the system's distribution of resources aligns with a particular node's preferences. However, this cannot hold true for all nodes unless they share the same preferences. Moreover, it is in every node's self-interest to try to tip the balance of $\mathcal{R}$ and $\mathcal{B}$ toward its own preferences. With a utility function slightly more biased toward serving, the network could be flooded with red or blue files, depending on the dominant preference.

25.4.2 Comparing Censorship Resistance

We model censorship as an external entity's attempt to impose a particular distribution of files r_c, b_c on a set of nodes. The censor's effect is not fixed; rather, it depends on the amount of resistance the affected nodes offer.

Assume a node that is not receiving attention from the censor can store up to T resources. A node under censorship can choose to store fewer resources ($T - t$) and invest an effort level t to resist censorship. We define the probability that a node successfully fights censorship (and reestablish its previous distribution of resources) as $P(t)$. With probability $1 - P(t)$, the censor prevails and imposes the distribution r_c, b_c .

We first consider the discretionary case, in which nodes select the content they serve. Knowing the nodes' preferences r_i, b_i , the censor's distribution r_c, b_c , the total resource bound T , and the probability $P(t)$ that it defeats the censor, we can calculate the optimal amount of resources a node invests in resisting censorship. The expected utility of a node under censorship is the probability of success, times the utility in that case, plus the probability of failure times the utility in that case:

$$U = P(t)U_i(r_i(T - t), b_i(T - t)) + (1 - P(t))U_i(r_c(T - t), b_c(T - t)). \quad (25.4)$$

Our utility functions U_i are unimodal and smooth, so if the functions $P(t)$ are sufficiently well-behaved, there is a single optimal investment in resistance t in $[0, T]$ by setting $\frac{dU}{dt} = 0$.

We begin with the simplest example, namely where the probability $P(t)$ of resisting censorship is linear in the defensive effort t . Assume that if a node invests all its resources in fighting, it definitely prevails but has nothing left to serve any files. At the other extreme, if it spends nothing on lawyers (or whatever the relevant mode of combat) then the censor prevails for sure. Therefore we define $P(t)$ as

$$P(t) = \frac{1}{T}t. \quad (25.5)$$

By maximizing (25.4) with $P(t)$ defined as in (25.5), we find that the optimal defense budget t_d :

$$t_d = \frac{T}{2} \frac{2U_i(r_c, b_c) - U_i(r_i, b_i)}{U_i(r_c, b_c) - U_i(r_i, b_i)}. \quad (25.6)$$

The node diverts t_d resources from serving files to fighting censorship. We also assume, for now, that the cost of the attack for the censor is equal to the node's defense budget t .

We now turn to the case of Eternity or DHTs where resources are scattered randomly around the network, where each node is expected to hold a mixture of files r_s, b_s . As in the previous example, the utility of a node under censorship depends on its defense budget t , the censor's choice of r_c, b_c , and the system's distribution of files r_s, b_s :

$$U = P(t)U_i(r_s(T - t), b_s(T - t)) + (1 - P(t))U_i(r_c(T - t), b_c(T - t)). \quad (25.7)$$

A similar approach is followed as above to derive the optimal defense budget t for each node:

$$t_s = \frac{T}{2} \frac{2U_i(r_c, b_c) - U_i(r_s, b_s)}{U_i(r_c, b_c) - U_i(r_s, b_s)}. \quad (25.8)$$

However, not all nodes are motivated to resist the censor! Some may find that $U_i(r_s T, b_s T) \leq U_i(r_c T, b_c T)$, which means that their utility under censorship increases. This is not an improbable situation: in a network where half the resources are red and half are blue ($r_s = 0.5, b_s = 0.5$) a censor that shifts the balance to $r_c = 0$ benefits the blue-loving nodes, and if they are free to set their own defense budgets then they select $t = 0$.

Who fights censorship harder? The aggregate defense budget, and thus the cost of censorship, is greater in the discretionary model than in the random one, except in the case in which all nodes have the same preferences (in which case equality holds).

For the maximum value of the defense budget t to be positive in the interval $[0, T]$, the following condition must be true:

$$0 < \frac{T}{2} \frac{2U_i(r_c, b_c) - U_i(r_s, b_s)}{U_i(r_c, b_c) - U_i(r_s, b_s)}. \quad (25.9)$$

In other words,

$$2U_i(r_c, b_c) < U_i(r_c, b_c). \quad (25.10)$$

When this is not true, a node maximizes its utility by not fighting at all and choosing $t = 0$. Given these observations, it follows that

$$\forall i \in \mathcal{S}, t_{d_i} \geq t_{s_i} \Rightarrow \sum_{i \in \mathcal{S}} t_{d_i} \geq \sum_{i \in \mathcal{S}} t_{s_i}. \quad (25.11)$$

Whatever the attacker's strategy, it is at least as costly or more so, to attack a network's architecture via the discretionary rather than the random model. Equality holds when for each node, $t_d = t_s$, which in turn means that $r_i = r_s$. This is the case of homogeneous preferences. In all other cases, the cost to censor a set of nodes is maximized when resources are distributed according to their preferences rather than randomly.

25.5 Complex Networks and Topology

The final area of information security that we discuss is the topology of complex networks. Computer networks from the Internet to decentralized peer-to-peer networks are systems of great complexity that emerge from ad hoc interactions of many entities on the basis of simple ground rules that are minimally restrictive. The emergent complexity, coupled with heterogeneity on every relevant scale, is similar to networks found "in the wild" – from the social networks made up from interactions between people to metabolic pathways in living organisms. Recently a discipline of network analysis has emerged at the boundary between sociology and condensed-matter physics. It takes ideas added from other disciplines like graph theory, which provides tools and concepts for modeling and investigating such networks. Our interest here is the interaction of network science with information security; as we shall see, we can build an interesting bridge to evolutionary game theory.

Network topology can strongly influence conflict dynamics. Often an attacker tries to disconnect a network or increase its diameter by destroying nodes or edges, while the defender counters using various resilience mechanisms. Examples include a music industry body attempting to close down a peer-to-peer file-sharing network; a police organization trying to decapitate a terrorist organization; and a totalitarian government conducting surveillance on political activists. Police forces have been curious for some years about whether network science might be of practical use in covert conflicts – whether to insurgents or to counterinsurgency forces.

Different topologies have different robustness properties with respect to various attacks. Albert, Jeong, and Barabási famously showed that certain real-world networks with scale-free degree distributions are more robust to random attacks than targeted attacks. This is because scale-free networks – like many real-world networks – get much of their connectivity from a minority of nodes that have a high vertex order. This resilience makes them highly robust against random upsets; but remove the 'kingpin' nodes, and connectivity collapses.

This is the static case – for example, when a police force becomes aware of a criminal or terrorist network, and sets out to disrupt it by finding and arresting its key personnel. The result of Albert et al. models this well. But what about the dynamic case – where at each round the attacker can remove a certain number of nodes, but the defenders can

recruit other nodes to replace them? How do attack and defense interact: what is the interplay of tactics and strategy?

We built a simulation in which a network game is played with a number of rounds. Each round consists of attack followed by node replenishment and adaptation. The attacker can remove a proportion of nodes; his choice of nodes is his strategy. The defenders' strategy lies in the adaptation phase; the way they rewire their network after each round of attack and replenishment. This rewiring must be done using only local knowledge.

An attack strategy is more efficient, for a given defense strategy, if an attacker using it requires a smaller budget to disrupt the network. Similarly, a defense strategy is more efficient if, for a given attack strategy, it compels the attacker to expend a higher budget to achieve network disruption.

We started off by considering the static attacker of Albert et al., whereby high vertex order nodes are removed, and a defense strategy of either random replenishment, forming rings, or forming cliques. In the ring strategy, defenders replace high-order nodes with rings – as in P2P systems such as Chord. In the clique strategy, high-order nodes are replaced with cliques – clusters of nodes all connected to each other.

The results of the initial three simulations are given in Figure 25.1.

Random replenishment (line with circles) in Figure 25.1 provides a calibration baseline. As seen above, it is ineffective: within three rounds the size of the largest connected component has fallen by a half, from 400 nodes to well under 200. The line with crosses shows that rings give only a surprisingly short-term defense benefit. They postpone network collapse from about two rounds to about a dozen rounds. Thereafter, the network is almost completely disconnected.

Cliques (indicated by the caret symbol), on the other hand, work well. A few vertices are disconnected at each attack round, but the network itself remains robustly connected. This may provide some insight into why, although rings have seemed




Figure 25.1. Vertex order decapitation attack in rings, cliques, and with no adaptation.




Figure 25.2. Rings and cliques defense under vertex order and centrality attacks.

attractive to theoreticians, those real revolutionary movements that have left some trace in the history books have used a cell structure instead.

We then proceeded through several rounds of attack evolution. As cliques are a good defense against the simple vertex-order attack, we looked for a good way to attack cliques. The best performer we found is an attack based on centrality. We used Brandes' algorithm to select the highest-centrality nodes for destruction at each round. As before, our calibration baseline is random replenishment.

Figure 25.2 shows that the same holds for rings (the squares and crosses): the network collapses completely after about a dozen rounds. Centrality attacks are more effective against cliques; they significantly reduce the size of the largest connected component.

Then, knowing that centrality attacks are powerful, we tried a number of other possible defenses. The most promising at present appears to be a compound defense based on cliques and delegation.

The idea behind delegation is simple. A node that is becoming too well-connected selects one of its neighbors as a "deputy" and transfers some of its links to it. This reflects normal human behavior even in peacetime: busy leaders pass new recruits on to colleagues. In wartime, and with an enemy that might resort to vertex-order attacks, the incentive to delegate is even greater. Thus a terrorist leader who gets an offer from a wealthy businessman to finance an attack might simply introduce him to a young militant who wants to carry one out. The leader need now maintain communications with at most one of the two.

The delegation defense on its own, however, is rather like the ring defense. Network fragmentation is postponed (about 14 rounds with the parameters used here) though not ultimately averted. However, when we form a network and run the delegation strategy for some rounds before attacks start, then run a clique defense as well from




Figure 25.3. Component size: clique, immunization by delegation, and combined clique and delegation defenses against centrality attack.

the initiation of hostilities, this compound strategy works rather better than ordinary cliques. Figure 25.3 shows the simulation results.

Delegation results in shorter path lengths under attack: it postpones and slows down the growth of path length that otherwise results from hub elimination. As a result, equilibrium is achieved later, and with a larger minimum connected component.

Finally, we note that clique formation and delegation do not make the attacks in the earlier rounds of attack evolution any easier. Specifically, the effectiveness of a vertex-order attack depends on the skewness of the distribution of vertex order. Both delegation and clique formation lead to lesser skewness, and this is partly why they are an effective defense against a vertex-order attack in the first place. Hence these defensive manoeuvres will not make the earlier attacks any more effective than in the case where no defense actions are taken.

25.6 Conclusion

Information security has seen a number of interesting applications of game theory over the last 5 years. These have largely taken place in the context of a research program on the economics of security, which has built many cross-disciplinary links and has produced many useful (and indeed delightful) insights from unexpected places.

We have discussed how many information security failures are caused by incentive failures, where the people who guard a system are not the people who suffer when it fails; and how externalities make many security problems somewhat reminiscent of environmental pollution. Some aspects of information security are public goods, like clean air and water. Externalities also play a key role in determining which security products succeed in the market, and which fail.

Games with incomplete information also play an important role: where either information or action is hidden, things can go wrong in interesting ways. Markets, and auctions, can sometimes be used as information-processing mechanisms to tackle the resulting problems; we discussed software dependability and the problems of cyber-insurance.

Finally we looked at effects on distributed system architectures. The designers of early peer-to-peer systems adopted a flat architecture, which promoted free-riding and made attacks easy; later, more successful, systems used a discretionary architecture that mitigated these problems. We now know how to analyze cooperation in heterogeneous distributed systems, and the tools have wider implications for understanding human societies.

The second aspect of architecture is topology. Albert, Jeong, and Barabasi showed that scale-free networks are more robust than random networks against random failure, but more vulnerable to targeted attack; by extending their analysis from the static to the dynamic case, we have shown why revolutionaries organize in cells – and why building peer-to-peer systems based on rings was a bad idea. At the conceptual level, we have provided a framework for analyzing such problems systematically, and started to build a bridge between network analysis and evolutionary game theory.

25.7 Notes

Anderson (2001) was the first security researcher to identify the importance of incentives and economics. In earlier work he described misaligned incentives with respect to ATM security (Anderson, 1994) and the Eternity Service, the first peer-to-peer system designed to offer censorship resistance (Anderson, 1996). With George Danezis, he considered the role of economics on censorship resistance (Danezis and Anderson, 2005).

Hal Varian was the first economist to pay attention to information security. He noted that users lacked sufficient incentive to protect themselves from viruses because much of the resulting harm was suffered by others (Varian, 2000). He also created a game-theoretic model to describe the impact of independent security decisions: whether system defense depended on the best effort of the defenders, on their worst effort, or on the sum of their efforts (Varian, 2004).

Howard Kunreuther and Geoffrey Heal extended the result to the case where the security of group rests upon the efforts of interdependent members (Kunreuther and Heal, 2003). Katz and Shapiro (1985) famously noted how network externalities affected the adoption of technology. Akerlof (1970) won a Nobel prize for his articulation of the effect of asymmetric information on markets.

Schechter (2002) was the first to propose vulnerability markets. Ozment (2004) argued that those markets could be better designed as auctions. In joint work, they have proposed statistical measures of software security based upon software engineering approaches (Ozment and Schechter, 2006a). They have also analyzed the bootstrapping problems faced by those who would deploy security technologies (Ozment and Schechter, 2006b).

Banking standards for PIN-entry terminals assume a cost-based analysis of vulnerability (PIN management requirements, 2004). Karthik Kannan and Rahul Telang

have analyzed the social utility of the organizations currently purchasing software vulnerabilities and found it to be less than ideal (Kannan and Telang, 2004). Rainer Böhme (2006) has argued that software derivatives are a better tool than markets or auctions for the measurement of software security. With Gaurav Kataria, he analysed how the interdependence of cyber-risks could cause insurance market failure (Böhme and Kataria, 2006). Hulusi Ogut, Nirup Menon, and Srinivasan Raghunathan showed that the interdependence of cyber-risk results in firms underinvesting in both security and insurance (Ogut et al., 2005).

Crespo and Garcia-Molina (2002) argue for network topologies based on clubs of nodes with common interests. Moore (2005) has noted the security import of hidden-action attacks. Sparrow (1990) surveyed possible applications of social network theory to law enforcement in 1990; a more recent survey is by Ballester, Calvó-Armengol and Zenou (2004). For the debate on whether the diversity of modern societies is undermining the social solidarity on which welfare systems are based, see Goodhart (2004).

Albert, Jeong and Barabási (2000) showed that scale-free network topology being good for robustness against random failure but bad for security against targeted attack. Finally, Nagaraja and Anderson (2006) extended this from the static to the dynamic case.

Bibliography

- Reka Albert, Hawoong Jeong, and Albert I. Barabási. Error and attack tolerance of complex networks. *Nature*, 406(1):387–482, 2000.
- George A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- Ross J. Anderson. Why cryptosystems fail. *Communications of the ACM*, 37(11):32–40, 1994.
- Ross Anderson. The eternity service. In *First International Conference on the Theory and Applications of Cryptology, PRAGOCRYPT '96*, 1996.
- Ross Anderson. Why information security is hard – an economic perspective. In *17th Annual Computer Security Applications Conference*, December 2001. New Orleans, LA.
- Rainer Böhme. A comparison of market approaches to software vulnerability disclosure. In *Proceedings of ETRICS*, pp. 298–311. Springer Verlag, March 2006. LNCS 2995.
- Rainer Böhme and Gaurav Kataria. Models and measures for correlation in cyber-insurance. In *Proceedings of the Fifth Workshop on the Economics of Information Security*, June 2006. Cambridge, UK.
- A. Calvó-Armengol, C. Ballester, and Y. Zenou. Who’s who in crime networks – wanted the key player. In *IUI Working Paper Series 617*, 2004. The Research Institute of Industrial Economics.
- Arturo Crespo and Hector Garcia-Molina. Semantic overlay networks for p2p systems. Technical report, Stanford University, 2002.
- George Danezis and Ross J. Anderson. The economics of resisting censorship. *IEEE Security & Privacy*, 3(1):45–50, 2005.
- David Goodhart. Too diverse? *Prospect*, February 2004. <http://www.guardian.co.uk/race/story/0,11374,1154684,00.html>.
- Howard Kunreuther and Geoffrey Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, March–May 2003.
- Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *The American Economic Review*, 75(3):424–440, June 1985.

BIBLIOGRAPHY

647

- Karthik Kannan and Rahul Telang. Economic analysis of market for software vulnerabilities. In *Proceedings of the Third Workshop on the Economics of Information Security*, May 2004. Minneapolis, MN.
- Tyler Moore. Countering hidden-action attacks on networked systems. In *Proceedings of the Fourth Workshop on the Economics of Information Security*, June 2005.
- Shishir Nagaraja and Ross Anderson. The topology of covert conflict. In *Proceedings of the Fifth Workshop on Economics of Information Security*, June 2006. Cambridge, United Kingdom.
- Hulisi Ogut, Nirup Menon, and Srinivasan Raghunathan. Cyber insurance and IT security investment: Impact of interdependent risk. In *Proceedings of the Fourth Workshop on the Economics of Information Security*, June 2005. Cambridge, MA.
- Andy Ozment and Stuart Schechter. Milk or wine: Does software security improve with age? In *15th Usenix Security Symposium*, July 2006. Vancouver, BC, Canada.
- Andy Ozment and Stuart E. Schechter. Bootstrapping the adoption of internet security protocols. In *Proceedings of the Fifth Workshop on the Economics of Information Security*, June 2006.
- Andy Ozment. Bug auctions: Vulnerability markets reconsidered. In *Proceedings of the Third Workshop on the Economics of Information Security*, May 2004. Minneapolis, MN.
- Stuart E. Schechter. How to buy better testing. In George I. Davida, Yair Frankel, and Owen Rees, editors, *InfraSec*, volume 2437 of *Lecture Notes in Computer Science*, pp. 73–87. Springer, 2002.
- Malcolm Sparrow. The application of network analysis to criminal intelligence: An assessment of the prospects. *Social Networks*, 13:253–274, 1990.
- Hal Varian. Managing online security risks. *The New York Times*, June 2000. Available at: <http://www.nytimes.com/library/financial/columns/060100econ-scene.html>.
- Hal Varian. System reliability and free riding. In L. Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pp. 1–15. Kluwer Academic Publishers, 2004.
- PIN management requirements: PIN entry device security requirements manual, 2004. Available at: http://partnernetnetwork.visa.com/dv/pin/pdf/Visa_ATM_Security_Requirements.pdf.

