

Would a ‘Cyber Warrior’ Protect Us? Exploring Trade-offs Between Attack and Defense of Information Systems

[Position Paper]

Tyler Moore, Allan Friedman and Ariel D. Procaccia
Center for Research on Computation & Society, Harvard University
{tmoore,allan,arielpro}@seas.harvard.edu

ABSTRACT

As information security shifts from the realm of computer science to national security, the priority for safe and secure systems will be balanced against the appeal of using information *in*security as a strategic asset. In “cyber war”, those tasked with defending friendly computer networks are also expected to exploit enemy networks. This paper presents two game-theoretic models of vulnerability discovery and exploitation, where nations must choose between protecting themselves by sharing vulnerability information with vendors or pursuing an offensive advantage while remaining at risk. One game describes a cold war of stockpiling, the other allows for actual attack. In both models, we predict that at least one state will have an incentive to pursue an aggressive cyber war posture, rather than secure its own systems. This finding – that a mutually defensive approach to security is not a stable equilibrium – holds up under a range of assumptions about social risk of cybercrime, technical sophistication, military aggressiveness and the likelihood of vulnerability rediscovery. We conclude with a discussion of the security policy implications of a militarized cyberspace.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*; K.4.1 [Computing Milieux]: Computers and Society—*Public Policy Issues*

General Terms

Economics, Security

1. INTRODUCTION

Computer scientists used to study information security by stating assumptions about the capabilities of an adversary and then building systems to protect against these assumptions. This approach worked well for the design of encryption algorithms and cryptographic protocols. However, it

has not coped as well with the Internet’s rise, and the emergence of a *strategic* adversary capable of adapting to the chosen defenses. Instead, a new perspective was required, which has been met primarily by applying an economic perspective to information security [1]. Attackers and defenders are now understood as being locked in a strategic battle, where the incentives to disrupt and protect systems matter most. An economic approach has been especially helpful for dealing with the rise of the profit-motivated online criminal [6, 15].

In this paper, we argue that the paradigm is shifting once more. The existing literature has treated attack and defense as activities carried out by two mutually exclusive groups, the ‘good’ guys and ‘bad’ guys. In fact, the distinction between attacker and defender is becoming blurred in the context of cyber warfare. As the United States collects responsibility for cybersecurity at a national level under the unified Cyber Command, a single organization assumes responsibility for defending domestic Internet infrastructure and cyber resources, or attacking enemies through offensive operations. In this paper, we present a game-theoretic model that reflects this new paradigm and explores the strategic interactions of actors capable of both attack *and* defense.

1.1 Cyber Command

The strategic use of information technology in the national security context has traditionally been the domain of the National Security Agency (NSA), with an almost legendary capacity for offensive signals intelligence. The establishment of US Cyber Command reflects a compromise between internal forces inside the US national security community, including the desire to avoid duplication of the NSA’s technical capacities, the desire to accommodate new cyber-focused efforts inside the military, particularly the Air Force, and a need to balance legally defined boundaries between the civilian intelligence community and the offensive-focused defense community [4]. The newly created Cyber Command will be placed under the charge of the NSA director, and will coordinate cyber war units inside the armed forces. The goal is to cluster and coordinate US strategic cybersecurity capacity to concentrate efforts in prosecuting national security policy with a united purpose.

Cyber Command, as a single organization, will have to navigate a number of challenging technical and policy hurdles, many of which have been discussed elsewhere [5, 11]. Of particular importance to this paper is the challenge of defending information security systems while still maintaining an offensive readiness. The National Military Strategy for Cyberspace Operations places a strategic priority on “maintaining a robust defense of cyberspace while exploiting ad-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSPW’10, September 21–23, 2010, Concord, Massachusetts, USA.
Copyright 2010 ACM 978-1-4503-0415-3/10/09 ...\$5.00.

versary cyberspace vulnerabilities” [22, p.19]. This paper argues that the nature of cybersecurity imposes a trade-off on those two goals, and that how the trade-offs play out depends on the strategic interactions of the players involved.

1.2 Attack and Defense

The notion of a trade-off between offensive and defensive capacity in the national security context is not new. Intelligence agencies, for example, are responsible for gathering intelligence and providing operational security. If acting on intelligence gained might compromise the source of new information, a rational response might be to accept short run damage to one’s own forces or even civilian infrastructure for the sake of the broader mission. In WWII, for example, the Allies allowed some German attacks to succeed in order to hide their strategic advantage in cryptanalysis and radar technologies [9]. (This trade-off is also found in the geek novel *Cryptonomicon* [23]).

How is this trade-off manifest in the cyber context? Technically, the responsibility for the general security of all non-military public and private information systems falls under the Department of Homeland Security. Yet the Defense Department’s own doctrine stresses that the national security apparatus “must assist in decreasing vulnerabilities to those infrastructures whenever possible through successful partnerships” [22, p. 16]. We saw public evidence of this cooperation when the NSA has lent its expertise to Microsoft during the development of Windows 7 [10] and to Google for protecting the company’s computer networks [16]. This reflects an important component of defense: most common systems are maintained by private vendors or open-source communities, so vulnerabilities discovered by the government will have to be patched by these non-governmental actors.

Yet news of the offensive focus of Cyber Command dominates. The NSA boasts of a highly classified ‘cyber-offensive’ capability, such as exploiting vulnerabilities to take over hostile foreign servers controlling botnets [21]. The NSA has trained a cadre of ‘cyber warriors’ for engaging in attacks to “disrupt, deny, degrade, or destroy the information” found in enemy computer systems [7]. Discussions of cyber war inside the defense establishment are laden with “macho rhetoric” [4], and declared policy seeks to “gain and maintain initiative” [22] and a “continued commitment to cyber superiority” [13]. In sum, there is ample evidence that the offensive side of Cyber Command is viewed as least as important as the defensive side. When the two are in conflict, what will be the rational outcome? Below, we present two game-theoretic models that offer insights into the expected imbalance.

2. GAME 1: VULNERABILITY STOCKPILES

Our first model, then, explores this tension between attack and defense in a particular context. What should a military cyber-organization do upon discovery of a previously undisclosed software vulnerability? Is it better to pass the information to the relevant software vendor and improve everyone’s security, or would it be more prudent to keep the vulnerability hidden and develop a zero-day exploit to be saved for an offensive mission against an enemy? Of course, such a decision isn’t taken only by the US, but also by other organizations in different countries. Hence, the best strategy must consider the actions of the other players. In the models

below, we reduce this to a two player game of two adversary states. In reality, of course, multiple countries engage in this dynamic, but we focus on two principle adversaries.

One potential outcome of the militarization of cyber attack is a ‘stockpiling’ of hidden exploits by nation-states to carry out offensive operations, at the expense of the security of civilian computer networks. Another possible outcome, of course, is that choosing to disclose vulnerabilities and improve security is preferred. But under what circumstances might we expect each outcome to prevail? This is the goal of the game-theoretic model presented below.

In the basic game we reduce the problem to a world of two states, as well as a general social risk. This social risk can be seen as the global threat of crime, terrorism, or the general state of insecurity apart from one’s rival. States have the opportunity to discover a particular vulnerability, and must choose whether to stockpile it against the adversary, or to defend their own systems and thereby securing their adversary as well as themselves. The model is slightly complicated by the fact that a state does not know whether it is the first to discover that particular vulnerability, or is rediscovering it after the other state has found it and is already stockpiling it. This game is played for each vulnerability a cyber warrior might discover. (Note that for this game, stockpiling a vulnerability does *not* mean actually launching a cyber-attack.)

This model is built on a few key assumptions. First, both players’ networks rely on the same set of vulnerable applications that can be exploited. There is ample support for this assumption, particularly since both players are likely to use common platforms given the dominance of a few software vendors in most IT systems. Second, we assume that patching one’s own system not only defends against potential attack, but also precludes the defender from using knowledge of this vulnerability to attack in the future. That is, an actor must decide whether to use a vulnerability for defense or offense; it cannot do both. This is reasonable because the duty to patch normally falls on the responsible private vendor, who would release a patch publicly accessible to both sides. Other alternatives exist, such as the reintroduction of export controls for consumer-level software, or behind-the-scenes patching coordinated with the government, but both would arouse suspicion of adversaries. Finally, we assume that a vulnerability has a decent chance of being independently rediscovered by at least two parties. This assumption is common in the economic literature on vulnerability disclosure and patch management [2, 3], and has received empirical support in the context of Windows vulnerabilities [19, Ch. 10].

2.1 Modeling the game

Our game theoretic analysis employs *extensive form games with imperfect information*. For ease of exposition we focus on the specific setting of this paper; the reader is referred to [18] for the general, formal definitions.

We first describe the key *actions*, *payoffs* and *parameters* used in the game, followed by a description of the game itself. For simplicity, the game involves just two players, which can be thought of as nation-states.

Actions and Payoffs.

There are two available actions: S (for “stockpile”) and D (for “defend”).




Figure 1: Vulnerability-stockpiles game. A player discovers a vulnerability with probability p and chooses to either defend or stockpile; stockpiling incurs the risk of the other player rediscovers it and stockpiling in turn.

1. S : The player discovers a software vulnerability but keeps this knowledge secret and stockpiles it into a collection of exploits available for future cyber attacks. The advantage of having a vulnerability that one’s adversary does not is normalized to 1, with a symmetric harm of -1 from the threat of being attacked. The payoff for stockpiling is thus $1 - \delta$ (the meaning of δ is explained below). The cost of being at a disadvantage is $-1 - \delta$.
2. D : The player discovers a software vulnerability and reveals it to the relevant software manufacturer, who immediately fixes the vulnerability. (We realize that this is an unrealistic simplification, as there would in reality be a delay before a patch could be developed and deployed.) The payoff for defense is 0, offering no strategic advantage or risk.

Parameters.

We have selected a few key characteristics whose values may vary, leading to different outcomes.

1. p : This parameter, valued between 0 and 1, is a measure of player one’s technical sophistication to discover a vulnerability. Meanwhile, $(1 - p)$ measures the sophistication of player two. Smaller values of p indicate that player one is less sophisticated compared to player two, while larger values indicate player one is more sophisticated. If the two players are evenly matched, then $p = 0.5$.
2. δ : As discussed above, security threats can come from a more general social risk, such as the threat of cyber-crime. Valued between 0 and 1, δ captures the harm to the general public in each state if no one explicitly chooses to defend their systems. In other words, the negative externalities of insecurity [1] are internalized if δ is positive, but they are ignored completely by the players if $\delta = 0$.

We have represented the steps of the game by a tree in Figure 1. Each internal node is graphically represented as a circle, and is either labeled by a player $i \in \{1, 2\}$, or by c , which stands for “chance”. There are two edges between a

player node and its children, which are labeled by the two available actions, S and D . The edges between a chance node and its children are labeled with probabilities. The leaves of the tree, which are represented as rectangles, contain pairs of numbers: the first is the payoff to player 1 and the second is the payoff to player 2.

The game starts at the root of the tree, which is represented by a double circle, and progresses as follows. If the current node is a chance node, we randomly proceed to one of its children, where the probability of reaching a child is the probability associated with the corresponding edge. Alternatively, if the current node is labeled by a player (1 or 2), that player must choose to stockpile or defend; we follow the edge labeled by S if the former action was taken and the edge labeled by D if the latter action was taken. Finally, when a leaf is reached the game ends and the players receive the payoffs that are specified in this leaf.

Let’s step through the game in Figure 1 first to explain how the tree represents the vulnerability-stockpiling game. The game starts at v_1 . With probability p , player 1 discovers the vulnerability first, moving to node v_2 . From here, player 1 must decide between actions S and D . If player 1 chooses D , then both players receive a payoff of 0 and the game concludes. If, instead, player 1 stockpiles the vulnerability (action S), then the game moves to a second chance node v_3 . With probability p , player 2 does not rediscovers the same vulnerability. Consequently, player 1 has added a vulnerability it alone knows to its stockpile for use in a future cyber-attack, and so derives utility $1 - \delta$, inflicting harm $-1 - \delta$ on player 2. With probability $1 - p$, however, player 2 rediscovers the vulnerability, moving to node v_4 . In this case, player 2 is faced with the same choice player 1 received in v_2 : stockpile the vulnerability (S) or disclose it (D). If player 2 chooses to defend, then both players receive utility 0. Where things get interesting is if player 2 also chooses to stockpile the vulnerability. In this case, the advantage of a stockpiled vulnerability is canceled out by the harms of being threatened. However, there is still a harm in keeping the vulnerabilities hidden – everyone’s computers remain insecure. Criminals can exploit these weaknesses to defraud victims. Consequently, when both players stockpile, they both suffer a loss $-\delta$.

Going back to the root node, suppose that with proba-

bility $1 - p$ player 2 discovers the vulnerability first, not player 1. In this case, the game moves to v_5 , not v_2 , and progresses through a symmetric series of steps to the ones described above, only this time it is player 2 who moves first.

The player nodes are grouped into two *information sets*, one containing the two nodes of player 1 (v_2 and v_7) and the other containing the two nodes of player 2 (v_4 and v_5). In Figure 1, two nodes in the same information set are connected by a dashed line. Conceptually, when it is a player's turn to take an action the player does now know which of the nodes in the player's information set is the current node. The use of information sets is crucial here because even though the game is played sequentially, both players do not know if they are the first one to discover a vulnerability or not. For instance, player 2 only gets to choose her action S or D once – she just doesn't know whether she's at node v_4 or v_5 in the game when the choice is made.

2.2 Finding equilibria

An ordered pair of strategies (x, y) , where $x \in \{S, D\}$ is the strategy of player 1 and $y \in \{S, D\}$ is the strategy of player 2, is called a *strategy profile*. The *utility* of player i for the strategy profile (x, y) , denoted $u_i(x, y)$, is the expected payoff of player i given that player 1 uses the strategy x and player 2 uses the strategy y , where the expectation is taken over the randomness of the chance nodes. To calculate $u_1(S, D)$ for the vulnerability-stockpiles game, we return to the tree in Figure 1. The game starts at v_1 . With probability p we move to v_2 , where player 1 plays S , leading the game to v_3 . Next, with probability p we reach a leaf with a payoff of 1 with respect to player 1. With probability $1 - p$ we reach v_4 , which is labeled by player 2; player 2 then plays D , which leads us to a leaf with a payoff of 0 with respect to player 1. Returning to the root v_1 , with probability $1 - p$ the first move of the game goes right and reaches v_5 . Player 2 then plays D , and the game ends with a utility of 0 with respect to player 1. Hence the expected payoff is

$$u_1(S, D) = p(p \cdot (1 - \delta) + (1 - p) \cdot 0) + (1 - p) \cdot 0 = p^2(1 - \delta).$$

Similarly, we can compute the expected payoffs for all strategy profiles for both players:

$$\begin{aligned} u_1(S, S) &= p^2(1 - \delta) - (1 - p)^2(1 + \delta) - 2p(1 - p)\delta \\ u_1(D, S) &= -(1 - p)^2(1 + \delta) \\ u_1(S, D) &= p^2(1 - \delta) \\ u_1(D, D) &= 0 \\ \\ u_2(S, S) &= (1 - p)^2(1 - \delta) - p^2(1 + \delta) - 2p(1 - p)\delta \\ u_2(D, S) &= (1 - p)^2(1 - \delta) \\ u_2(S, D) &= -p^2(1 + \delta) \\ u_2(D, D) &= 0 \end{aligned}$$

A strategy profile is called a *Nash equilibrium* [17] if, informally, no player can gain by unilaterally deviating. In our setting, this means that neither player can gain from switching to the other strategy. Formally, (x, y) is a Nash equilibrium if and only if $u_1(x, y) \geq u_1(x', y)$, where $x' \in \{S, D\} \setminus \{x\}$, and $u_2(x, y) \geq u_2(x, y')$, where $y' \in \{S, D\} \setminus \{y\}$.

A (pure) *strategy* in an extensive form game with imperfect information is a function from information sets to ac-




Figure 2: Different equilibria are possible for different values of p and δ in the vulnerability-stockpiles game.

tions.¹ For our purposes, this means that a player's strategy is either S or D , and the same action would be played when either of the player's nodes is reached. In this paper we do not consider *mixed strategies*, where players are allowed to randomize over pure strategies.²

In order for strategy (S, S) to be a Nash equilibrium, player one must prefer not to deviate to (D, S) , while at the same time player two must prefer not to deviate to (S, D) . Consequently, it must hold that $u_1(S, S) \geq u_1(D, S)$ and $u_2(S, S) \geq u_2(S, D)$. These inequalities reduce to

$$\begin{aligned} u_1(S, S) \geq u_1(D, S) &\Leftrightarrow p^2(1 - \delta) - 2p(1 - p)\delta \geq 0 \\ u_2(S, S) \geq u_2(S, D) &\Leftrightarrow (1 - p)^2(1 - \delta) - 2p(1 - p)\delta \geq 0 \end{aligned}$$

The following inequalities must hold for (D, S) to be a Nash equilibrium:

$$\begin{aligned} u_1(D, S) \geq u_1(S, S) &\Leftrightarrow p^2(1 - \delta) - 2p(1 - p)\delta \leq 0 \\ u_2(D, S) \geq u_2(D, D) &\Leftrightarrow \text{True } \forall p \in [0, 1], \delta \in [0, 1] \end{aligned}$$

Similarly, the following inequalities must hold for (S, D) to be a Nash equilibrium:

$$\begin{aligned} u_1(S, D) \geq u_1(D, D) &\Leftrightarrow \text{True } \forall p \in [0, 1], \delta \in [0, 1] \\ u_2(S, D) \geq u_2(S, S) &\Leftrightarrow (1 - p)^2(1 - \delta) - 2p(1 - p)\delta \leq 0 \end{aligned}$$

¹In particular all the nodes in an information set must have identical available actions.

²Note that in some settings the analysis of extensive form games of imperfect information is quite subtle, and calls for significantly more refined equilibrium concepts (e.g., perfect Bayesian equilibrium or sequential equilibrium). However, our setting is rather straightforward and it seems that the generally coarser concept of Nash equilibrium captures the strategic aspects of our games perfectly.

Finally, the strategy profile (D, D) is a Nash equilibrium only when $\delta = 1$.

Unsurprisingly, which equilibrium outcome will happen depends on the values assigned to p and δ . Figure 2 plots the range of equilibria that can occur for different values of p and δ , based on the inequalities just described. For middling values of p and small values of δ , the equilibrium strategy is for both players to stockpile. If p becomes too large or small then it makes sense for one of the players to defend. Whenever $\delta \geq 1/3$ and p is close to 0.5 then both (D, S) and (S, D) are equilibria simultaneously. We discuss the implications of the different equilibrium outcomes in greater detail in the following section.

2.3 Discussion

Without any social cost, both actors will pursue an aggressive strategy of always stockpiling, regardless of one’s technical advantage. This is because neither has a strong incentive to defend: the worst case is that both end up with large stockpiles pointing at each other without any explicit cost. Even with a low degree of technical sophistication, there is always a positive probability that the other state will not discover the vulnerability, leading to a pure advantage.

Increased social costs impose an externality. Note that for any equilibria under substantial social cost ($\delta > \frac{1}{3}$), some one will elect to share the vulnerability information with the vendor, making the world safer. Who ends up bearing the cost of this externality? It will be borne by the less technically sophisticated nation. As the technical advantage grows, the amount of social harm that player 1 can absorb increases before being tempted to defend. Thus, the likelihood of anyone sharing their vulnerability information with the world is lowest when both actors have similar technical capacities. As the imbalance grows, the weaker party fears the social cost more.

Note that there is no equilibria for the mutually secure world (D, D) (apart from the special case where $\delta = 1$). If a state knows that the other side will defend, it is always in its interest to attempt to stockpile a new vulnerability, as long as there is some chance of discovering one. Even when failure to fix one’s own systems is very costly (large δ), stockpiling can be the safest decision. Suppose a player expects its adversary to defend due to large expected costs. In these circumstances the best reaction is to stockpile, yielding a payoff of $1 - \delta$ or 0 rather than the costly $-\delta$. But why would one party then commit to defending? They would defend if the other actor is likely to *not defend*, which could lead to even higher social costs. Hence, the best we can hope for is one actor to defend, leaving some singly-discovered vulnerabilities unsecured.

3. GAME 2: CYBER HAWK

The first game examined the trade-off between stockpiling vulnerabilities for later use in offensive operations and protecting society by fixing vulnerabilities. This model explicitly focuses on the costs and benefits of being aware of potential exploits, without considering the outcomes of an actual conflict. What happens when there is a chance that some one might choose to attack? Cyber conflict holds many risks, such as the likelihood of escalation, but it can also bring benefits to the aggressor. If there is reason to believe that escalation is unlikely, cyber conflict allows for strategic

engagement of an adversary with less risk to military forces. Cyber conflict can also aid traditional mission goals, ranging from obtaining an advantage in intelligence gathering or espionage to crippling an enemy in advance of – or even in lieu of – conventional attack.

Strategic decision of the cyber commander, then, must reflect beliefs of an enemy’s likelihood to actually exploit a given vulnerability, as well as an understanding of one’s own plans and objectives in using it. In the second game, we include this aggression component to explore the strategic implications. The game still revolves around the same core decision of whether to defend or not. However, instead of stockpiling, one might expect an adversary to actually weaponize the vulnerability to attack. This leads to an added level of uncertainty: not only is the cyber commander uncertain about whether the adversary has discovered the vulnerability, but he is uncertain about whether the enemy will attack before the commander does. This parameter can be seen as a willingness to attack: if the other player is more likely to, this should alter calculations. It can also be understood as a time component: who will be the first to launch an attack after discovering a commonly known vulnerability. We call this game *cyber hawk* because it captures the interplay between proficiency in identifying vulnerabilities and the aggressiveness of players in launching attacks.

3.1 Modeling the game

We model this game in the same fashion as the stockpiling game described in Section 2. We begin by describing the key *actions*, *payoffs* and *parameters* used in the game, followed by a description of the game itself. Again, the game involves just two players, which can be thought of as nation-states.

Actions and Payoffs.

There are two available actions: A (for “attack”) and D (for “defend”).

1. A : The player discovers a software vulnerability but keeps this knowledge secret and converts it into an exploit for use in a future cyber attack. As will be explained below, using the attack action does not necessarily mean that the player launches a successful attack. An attack is successful when the player discovers the vulnerability and uses it before the other player does. The payoff for being the first to attack using the vulnerability is normalized to 1. The cost of being attacked is -1.
2. D : the player discovers a software vulnerability and reveals it to the relevant software manufacturer, who immediately fixes the vulnerability. The payoff for defense is 0. (The defend strategy is the same as for game 1.)

Parameters.

We have selected a few key characteristics whose values may vary, leading to different outcomes.

1. p : This parameter, valued between 0 and 1, is a measure of the technical sophistication of player one in discovering vulnerabilities (p has the same meaning as for game 1).
2. q : This parameter captures the relative likelihood that a player will choose to attack after discovering a vul-




Figure 3: Cyber-hawk game. Similar to the vulnerability-stockpiles game, but when both players have the vulnerability, the winner of a zero-sum game is determined by the aggressiveness parameter q .

nerability. Valued between 0 and 1, q indicates how fast player one will act, and $(1 - q)$ indicates how fast player two will act. Smaller values of q indicate that player one is more restrained in launching attacks, while larger values indicate player one is ‘trigger-happy’. If the two players are evenly matched, then $q = 0.5$.

We have chosen to omit the social cost variable δ included in the first game. We instead assume that the attack exclusively harms the losing player since the vulnerability is ultimately exploited.

We can step through the tree in Figure 3 to explain how the cyber-hawk game proceeds. In fact, the structure is the game closely resembles the stockpiling game, except that ‘attack’ branch in nodes v_4 and v_8 now lead to a chance node rather than a payout of $-\delta$.

The game starts at v_1 . With probability p , player 1 discovers the vulnerability first, moving to node v_2 . From here, player 1 must decide between actions A and D . If player 1 chooses D , then both players receive a payoff of 0 and the game concludes. If, instead, player 1 chooses to weaponize the vulnerability for an attack (action A), then the game moves to a second chance node v_3 . As above, this involves player 1 keeping the vulnerability a secret. With probability p , player 2 does not rediscover the same vulnerability. Consequently, player 1 alone knows the vulnerability and uses it in a cyber attack at some point in the future, deriving utility 1 and inflicting harm -1 on player 2. Since player 2 will not discover the vulnerability, the relative aggressiveness q is not an issue.

With probability $1 - p$, however, player 2 rediscovers the vulnerability, moving to node v_4 . In this case, player 2 is faced with the same choice player 1 received in v_2 : keep the vulnerability secret for launching a cyber-attack (A) or disclose it (D). If player 2 chooses to defend, then both players receive utility 0.

If player 2 also chooses to attack with the same vulnerability, then it’s a race to see which player launches an attack based on the hidden vulnerability first. This is captured by the chance node at v_5 and the parameter q . With probability q , player 1 launches the attack first, gaining utility 1 while player 2 suffers a loss of utility -1 . Alternatively, player 2 will launch the first attack with probability $1 - q$ and the fortunes will be reversed.

3.2 Finding equilibria

From the tree in Figure 3, we can derive the expected utility of player 1 for different strategy profiles:

$$\begin{aligned} u_1(A, A) &= p^2 - (1 - p)^2 + 2p(1 - p)q - 2p(1 - p)(1 - q) \\ u_1(D, A) &= -(1 - p)^2 \\ u_1(A, D) &= p^2 \\ u_1(D, D) &= 0 \end{aligned}$$

Because this is a zero-sum game, the expected utility of player 2 is the same as for player 1 except the signs are reversed, i.e., $u_2(x, y) = -u_1(x, y)$ for every $x, y \in \{A, D\}$.

In order for strategy (A, A) to be a Nash equilibrium, player one must prefer not to deviate to (D, A) , while at the same time player two must prefer not to deviate to (A, D) . Consequently, the following must hold: $u_1(A, A) \geq u_1(D, A)$ and $-u_1(A, A) = u_2(A, A) \geq u_2(A, D) = -u_1(A, D)$. These inequalities reduce to

$$\begin{aligned} u_1(A, A) \geq u_1(D, A) &\Leftrightarrow p^2 + 2p(1 - p)(2q - 1) \geq 0 \\ u_2(A, A) \geq u_2(A, D) &\Leftrightarrow (1 - p)^2 + 2p(1 - p)(1 - 2q) \geq 0 \end{aligned}$$

As in game 1, both $u_1(A, D) \geq u_1(D, D)$ and $u_2(D, A) \geq u_2(D, D)$ for any $p, q \in [0, 1]$. Therefore, the strategy profile (D, A) is a Nash equilibrium when

$$u_1(A, A) \leq u_1(D, A) \Leftrightarrow p^2 + 2p(1 - p)(2q - 1) \leq 0.$$

Similarly, the strategy profile (A, D) is a Nash equilibrium when

$$u_2(A, A) \leq u_2(A, D) \Leftrightarrow (1 - p)^2 + 2p(1 - p)(1 - 2q) \leq 0.$$

Finally, the strategy profile (D, D) can never be a Nash equilibrium, as if $p > 0$ then $u_1(A, D) > u_1(D, D)$, and if $p = 0$ then $u_2(D, A) > u_2(D, D)$.

Which equilibrium outcome will happen depends on the values assigned to p and q . Figure 4 plots the range of equilibria that can happen for different values of p and q . The equilibrium strategy is for both players to attack whenever p and q are both middling or when one is large and the other is small. If p and q are both small, then (D, A) is in equilibrium, while if p and q are both large then (A, D) is in equilibrium.




Figure 4: Different equilibria are possible for differing values of p and q in the cyber-hawk game.

3.3 Discussion

For two perfectly paired adversaries, where neither has a clear technical advantage or a greater proclivity to attack, both choose to attack. This is because the opportunity of uniquely discovering the vulnerability trumps the risk of being attacked. Similarly, when the technical advantage is high enough, the risk of being less aggressive is dominated by the likelihood of having the jump on one's opponent in being the only actor to have that weapon.

This yields an important finding: there is no level of technical advantage that will dissuade an adversary that knows it is more likely to use a weapon from attacking. In terms of the model parameters, player 1's equilibrium strategy is always to attack whenever $q > 0.5$ for all values of p .

It is only when the chance of discovery and the likelihood of using the weapon are both small enough that a player will select to act defensively. The dangers of being without the first-mover advantage (small q) and being surprised by an attack (small p) can lead a state to pursue a defensive strategy.

However, even when a state holds as large a technical advantage as 2:1, if it is sufficiently unlikely to actually use a cyber weapon, it will pursue a defensive strategy. This is reflected in the shaded area of (D, A) in Figure 4. In this case, reticence to attack creates an opening for a technologically-weaker player to attack when the other defends.

Why is the peaceful, security-friendly world of common defense never an equilibrium? As long as a state knows that the other player will always select a defensive posturing, then it can interpret the fact that it is at a decision node as evidence of being the first to discover the vulnerability. Thus, there is no harm in planning to attack: the worst outcome is that the other actor will make the world safe for everyone.




Figure 5: Varying p , δ and r – the probability that a vulnerability will be rediscovered – for the vulnerability-stockpiles game. The graph shows the range of values of r where the equilibrium strategy is (S, S) .

How do these findings compare to current events in cyber international relations? Russia and the US have been at loggerheads over the best way to combat online crime internationally. Russia has pushed for the US to adopt a treaty banning the use of offensive cyber operations, which the US has firmly resisted [12]. Why might the Russians push for such a ban? The US likely has an edge over the Russians in terms of technical sophistication for cyber attacks ($\frac{1}{2} \leq p \leq \frac{2}{3}$). Compliance with the ban would be difficult to verify, and the US suspects that Russia might continue to develop offensive capabilities despite agreeing otherwise. If the US held up their end of the bargain and restricted the conditions under which it would launch attacks but the Russians did not respect the pact, then q would be small. Therefore, we could easily end up in an equilibrium where the US chooses to defend and the Russians attack (strategy profile (D, A)).

4. THE IMPACT OF VARYING VULNERABILITY REDISCOVERY RATES

In our initial modeling of the game, we used probability p as a single parameter for discovery and rediscovery. Yet the value of this parameter is critical. Many economic models of software patching assume that the main incentive for fixing vulnerabilities is the threat of independent rediscovery [2, 3], but the frequency with which vulnerabilities are rediscovered has been hotly debated in the literature. Rescorla argued that vulnerability rediscovery is very unlikely given the vast number of bugs in software [20]. Others, though, have demonstrated rediscovery to be quite possible. Ozment, for instance, studied Microsoft vulnerability reports from 2002–




Figure 6: Varying p , q and r – the probability that a vulnerability will be rediscovered – for the cyberhawk game. The graph shows the range of values of r where the equilibrium strategy is (A, A) .

2004 and found that Microsoft vulnerabilities were rediscovered and re-reported around 8% of the time before the vulnerability was publicly disclosed [19, Ch. 10].

Consequently, we can choose to model the rediscovery rate r as an exogenous variable. The likelihood of rediscovery now depends only on r and not on player 1’s technical advantage p . The tree in Figure 1 changes as follows. In the top chance branch at node v_3 , r replaces the probability $1-p$ and $1-r$ replaces p in the bottom branch. In the top chance branch at node v_6 , r replaces the probability p and $1-r$ replaces $1-p$ in the bottom branch. A similar substitution takes place at nodes v_3 and v_7 in Figure 3’s tree.

Figure 5 plots for a range of values of r the values of p and δ where the strategy profile (S, S) is an equilibrium for the vulnerability-stockpiles game. Similar to Figure 2, both players are more likely to stockpile when they are evenly matched. However, what we learn from the introduction of an exogenous rediscovery parameter is that players prefer to stockpile when rediscovery is unlikely. Players are more likely to ignore substantial social costs and keep stockpiling because the value of a uniquely acquired vulnerability is too great. As vulnerability rediscovery becomes more likely (bigger values of r), states tolerate only smaller social costs before choosing to defend. This is because rediscovery ‘cheapens’ the value of discovered vulnerability.

A complementary lesson can be learned by examining the role of rediscovery in game 2. Figure 6 plots for a range of values of r the values of p and q where the strategy profile (A, A) is an equilibrium. When the likelihood of rediscovery is small (say $r = 0.1$), both players choose to attack except for small values of p (< 0.1) and q (< 0.5) or large values of p (> 0.9) and q (> 0.5). In other words, one player needs to

be far inferior technically but reticent to attack in order for defense to become the prevailing strategy. Otherwise, both players choose to attack.

As the likelihood of rediscovery increases, the range of values for p and q where attacking prevails grows smaller. Why? When rediscovery is more likely, the players are less likely to uniquely discover a weakness that can be exploited, and thus the adversary may preemptively attack. This makes defending against a discovered vulnerability more attractive.

So what are reasonable values for the rediscovery rate r ? In the case of software vulnerabilities, the best available empirical data suggests a rediscovery rate closer to 0.1 [19, Ch. 10]. So if the players are collecting zero-day exploits of Windows, we might expect aggressive stockpiling. But there are also situations where r might be closer to 1. For instance, it is widely known that the process control systems that control critical infrastructures such as chemical refineries and the power grid are insecure [8, 14]. Mapped to our game, the likelihood that such weaknesses will be discovered by both parties is high. It is interesting, then, that our model predicts that the players are less likely to take an aggressively offensive position in these circumstances. Such a prediction is consistent with reality, in the sense that few (if any) reported attacks targeting the process control systems controlling critical infrastructures have been realized.

5. POLICY IMPLICATIONS

The games presented in this paper seek to model state response to situations where nations compete for dominance in cyberspace while attempting to balance this with defense of their own systems. In the context of national security, there is evidence that defense gets the short end of the stick. Yet as security researchers, we are interested in the sharing of vulnerability information to improve the security of systems for all users. The above results may be pessimistic with respect to increased information security, but these models can also inform the policy process. How might a state actor move towards a world where sharing vulnerability data is more common?

In the first game, insecurity is driven by a lack of consensus on the value of shared costs. If, as we believe, there are substantial social costs to insecure systems, then the true value of δ may be higher than what is perceived by cyber commanders. A simple approach to reconcile this difference might be a lobbying or public awareness campaign to bring these costs into their calculations.

Perversely, this model also predicts an increase in global government attention to cyber defenses if it increases δ itself by raising the general social threat. Thus, in a game between the US and Russia, Russian tacit support of cyber crime might yield greater levels of vulnerability patching.

Another side effect of the militarization of cyber space is an increased level of secrecy. Is this secrecy good for the security of the internet? At first glance, it might be bad. If we assume that high secrecy leads each actor to estimate median values for their opponent, we move towards the zones where all sides stockpile or attack. From this perspective, secrecy is bad.

Less information, however, can also be strategic, and the effects of over- or underestimating are important to explore. Recent headlines have been dominated by Chinese cyber attacks, breathless with implications of hordes of attackers probing our defenses, and a willingness to exploit any

vulnerability found. The effect of overstating these threats might drive the US to adjust its position towards a more defensive posture than it otherwise should have. This is particularly true for technical sophistication; the analysis of game 2 predicts that once expectation of an attack is past a certain point, no one will risk sharing defensive information. Underestimating the abilities and aggressiveness of an adversary has the opposite effect. If Russia underplays its technical sophistication, and unilaterally commits to cyber nonaggression [12], then the US risks adopting an overly aggressive position, stockpiling and attacking when it should actually defend. Future work should formally consider the cases of information asymmetry.

The strong emphasis on dominance of the cyber domain may actually have positive social value. In both games presented, the highest likelihood of an aggressive strategy of stockpiling or attacking results when both parties are close to evenly matched. However, when both parties believe that one has a technical advantage—superiority—the less dominant party is more likely to adopt a defensive posture. While this *cyber Pax Romana* does have all the connotations of a potential hegemony that accompany military superiority in any domain, it might suggest a period of relative stability as well. Hence, another perspective might be a form of deterrence through strength.

Finally, we note that the existence of Nash equilibria that favor cyber offense is not an endorsement of an aggressive cyber war strategy. Cyber peace is a desirable goal, and one that we believe policymakers should be explicitly working towards. That our model predicts aggressive outcomes for a range of parameters should be interpreted as a sign of how difficult stopping the proliferation of cyber war is likely to be. Instead, we hope that our findings will inspire further investigation into policies which promote cyber de-escalation.

6. LIMITATIONS

It is important to recognize that the models presented in this paper do not capture the entire domain of cyber strategy. We have examined one small part, looking at the decisions faced by a joint cyber command unit on the discovery of a vulnerability, what is quite possibly the first move in a multistage game. We have not considered the challenges of response, escalation or uncertainty, to name just a few. Attacks are always successful in this model, and never pose risk to the attacker through system interdependency. Furthermore, as with any game-theoretic model, strategies are ultimately determined by the payoffs for each outcome; if actor payoffs do not reflect real world outcomes, these games are less useful in guiding policy. Finally, we only model two-party relationships. There is reason to believe that adding another player (such as the dynamics between the US, China and Russia) might distort the outcomes.

One particular component of cyber conflict that has received much attention is the problem of attribution, or how to identify the attacker in order to appropriately respond. Attribution is important when modeling conflict as a multistage game, particularly for deterrence. Our model stops short of explicitly looking at the decision to attack or how to respond to an attack, and only reflects deterrence in that a country may be driven to behave defensively and secure their own systems, rather than build attack capacity. The games in this paper imply that an attack can be advantageous, and being attacked can be disadvantageous. None of

this is altered by the challenges of attribution. Moreover, the game framework could be expanded with more parameters to reflect other dimensions of cyber attacks, including probability of attribution, uncertainty of success or asymmetric harms from attacks.

Some of the findings that depend on a technically unsophisticated actor choosing defense may not have a great deal of impact from a policy standpoint, since the consequent low probability of discovery and rediscovery would minimize frequency that an actor would actually have the opportunity to play D . That is, it does not matter if a state would always share vulnerability information when they are unlikely to have that information to share.

Finally, we acknowledge that this paper does have an explicit US focus, prompted by the recent attention to Cyber Command. The dilemma of cyber attack and defense is not limited to superpowers: 20–30 nations have established offensive cyber units [4]. Expanding the model to include a more accurate picture of the international political environment could offer further contributions. First, modeling other approaches to cyber security governance could yield new theoretical findings towards a more secure environment. Second, the dynamics of the game could shift dramatically if we include multiple parties investing in attack and defense capacities against each other to reflect balances of power or perceived threats. No doubt this further analysis would also include the challenges of attribution discussed above. Approaching the problem from another state’s perspective, or bringing in a multilateral approach may yield further findings.

7. CONCLUSION

The militarization of cyberspace represents a substantial change in our understanding of information security. We have made a first attempt at exploring the dynamics of information security when the attacker is also defender. We presented two game-theoretic models of vulnerability discovery and exploitation. These games capture a trade-off where nations must choose between protecting themselves or pursuing an offensive advantage while remaining at risk. One key finding is that strategic interaction may very well lead to a proliferation of offensive behavior, even if defensive behavior is preferred. The presence of aggressive equilibria is sobering: nations may naturally be tempted to pursue cyber attack, which reinforces the important role policymakers have in promoting cyber defense. Using these models we can better understand the incentives facing states juggling the sometimes conflicting goals of cyber attack and defense, and how best to shape policy that promotes better security investment.

Acknowledgments

Discussions with Melissa Hathaway, Roger Hurwitz and Bruce Schneier inspired this article. Comments from the anonymous reviewers, the participants at NSPW, and our shepherd Brian Snow have helped improve the paper. The authors are also grateful to the participants of the ECIR Cyber Policy Lunch at the Harvard Kennedy School.

8. REFERENCES

- [1] R. Anderson and T. Moore. The economics of

- information security. *Science*, 314(5799):610–613, October 2006.
- [2] A. Arora, R. Telang, and H. Xu. Optimal policy for software vulnerability disclosure. *Management Science*, 54(4):642–656, 2008.
- [3] H. Cavusoglu, H. Cavusoglu, and S. Raghunathan. Emerging issues in responsible vulnerability disclosure. In *4th Workshop on the Economics of Information Security*, 2005.
- [4] R.A. Clarke and R.K. Knake. *Cyber War*. HarperCollins, 2010.
- [5] National Research Council. *Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities*. National Academies Press, 2009.
- [6] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *CCS '07: Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 375–388, New York, NY, USA, 2007. ACM.
- [7] J. Goldsmith. Can we stop the global cyber arms race? *The Washington Post*, February 2010. <http://www.washingtonpost.com/wp-dyn/content/article/2010/01/31/AR2010013101834.html>.
- [8] V.M. Ijure, M. Vinay, S.A. Laughter, and R.D. Williams. Security issues in SCADA networks. *Computers & Security*, (7):498–506, 2006.
- [9] R.V. Jones. *The Wizard War : British Scientific Intelligence*. Coward, McCann & Geoghegan, 1978.
- [10] G. Keizer. NSA helped with Windows 7 development. *Computerworld*, November 2009. http://www.computerworld.com/s/article/9141105/NSA_helped_with_Windows_7_development.
- [11] J.A. Lewis. *Securing Cyberspace for the 44th Presidency*. Center for Strategic and International Studies, 2008.
- [12] J. Markoff and A.E. Kramer. U.S. and Russia differ on a treaty for cyberspace. *The New York Times*, June 2009. <http://www.nytimes.com/2009/06/28/world/28cyber.html>.
- [13] N.A. Schwartz M.B. Donley. 2010 United States Air Force posture statement, 2010.
- [14] A. Miller. Trends in process control systems security. *IEEE Security & Privacy*, (5):57–60, October 2006.
- [15] T. Moore, R. Clayton, and R. Anderson. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, 2009.
- [16] E. Nakashima. Google to enlist NSA to help it ward off cyberattacks. *The Washington Post*, February 2010. <http://www.washingtonpost.com/wp-dyn/content/article/2010/02/03/AR2010020304057.html>.
- [17] J. F. Nash. Equilibrium points in N -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36:48–49, 1950.
- [18] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [19] A. Ozment. *Vulnerability Discovery & Software Security*. PhD thesis, University of Cambridge, 2007.
- [20] E. Rescorla. Is finding security holes a good idea? *IEEE Security & Privacy*, 3(1):14–19, 2005.
- [21] D.E. Sanger, J. Markoff, and T. Shanker. U.S. steps up effort on digital defenses. *The New York Times*, April 2009. http://www.nytimes.com/2009/04/28/us/28cyber.html?_r=2.
- [22] Chairman of the Joint Chiefs of Staff. The national military strategy for cyberspace operations, 2006.
- [23] N. Stephenson. *Cryptonomicon*. Avon, 1999.