

Cyber-sword v. cyber-shield: The Dynamics of US Cybersecurity Policy Priorities

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

ABSTRACT

Recent efforts to address cybersecurity risks have focused on leveraging the immense technical capacity of the American intelligence community to protect the nation's information technology infrastructure, and to project power in a new domain. This creates a potential conflict of interest: the joint duties of breaking into foreign systems while securing our own raises questions about competing goals. This paper highlights that tension, and introduces two game-theoretic models of the strategic decisions faced in security vulnerability discovery and disclosure. The country must both protect itself in the new domain and pursue an offensive advantage while still remaining at risk. One game describes a cold war of stockpiling, while 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. The model can also be used to explore the broader national and international cybersecurity context, including explain some observed behaviors, and make predictions about the effects of different policy interventions. Recognizing that securing our infrastructure should be a priority for cyber policy makers, we propose policy recommendations that create the opportunity for more defensive equilibria to take hold.

1. INTRODUCTION

As more attention is paid to cybersecurity, policy researchers must understand the range of policy options and have the tools to evaluate policy proposals and understand how decisions at the forefront of national policy will impact the broader world of information technology. The question of governance in the cyber domain is particularly complex.

The issues span the boundaries of public and private, civilian and defense, virtual and concrete. This paper ties a set of specific management decisions in a new cybersecurity organization, United States Cyber Command, to the broader challenge of securing the IT infrastructure on which our modern economy runs. Cybercommand occupies the central national role in both cyber defense and cyber offense. We argue that the nature of cybersecurity imposes a trade-off on those two goals, and that the trade-offs play out depends on the strategic interactions of the players involved. Game theory modeling can help us understand how the joint cybercommand likely to impact the overall security of the national—and indeed the global—information infrastructure. Moreover, it can be used to study the incentives of the actors involved and identify and evaluate cybersecurity policy options.

It is important to state upfront that this paper is not explicitly about deterrence. The challenges and promises of adapting traditional deterrence theory to cybersecurity have been explored elsewhere [14], [10]. Moreover, there is little consensus on the validity of deterrence games in other settings, and they tend to be driven by strong behavioral assumptions [27]. Nor is this paper a discussion of the legal intricacies and uncertainties that have emerged with the new cyber domain [24]. Instead, this paper seeks to highlight a key tradeoff that must be made in mapping out the nation's cybersecurity policy: how do we balance the ability to project power in the cyber domain with the ability to protect our own information systems.

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, and deterring or attacking enemies through offensive operations. In this paper, we present a game-theoretic model that reflects this new paradigm and explores the policy challenges that arise from this tension. Following a motivating discussion further expanding on how attack and defense may serve conflicting purposes, we present the modeling approach and explain the assumptions on which it is built. In Sections 4 and 5, we present a detailed explanation of each game and derive the equilibria predicting how the United States would behave in response to an adversary certain conditions. We then explore the policy implications of these games, and propose a set of recommendations on policies that support and guide strategic decision making for a more secure internet.

2. MOTIVATION: MANAGEMENT OF ATTACK AND DEFENSE IN INFORMATION SYSTEMS

Information systems touch almost every aspect of modern life. What is meant by securing them? The term 'cybersecurity' can apply at many levels: maintaining national economic advantages and state secrets, allowing agencies, firms and infrastructures to operate normally while safeguarding sensitive information, and the continued operation of individual computer systems.

The management of cybersecurity policy focuses on the higher end of the spectrum, maintaining the status quo and preventing harms to social institutions. At the same time, the actual harms and defenses are initiated at a much lower level. Attacks target specific computer systems and exploit specific pieces of software and hardware. Cybersecurity policy is the process of bridging these two levels.

2.1 Information Security Basics

The field of computer security and information assurance is vast and complicated. For the purposes of this paper, we will focus on a set of attacks that exploit a weakness in software programs that run on computers. All complex software systems will have some unanticipated weaknesses and potential vectors of attack. Better software engineering can reduce the likelihood of vulnerabilities, but will never eliminate the risk completely [12].

A flaw in a software program that can be exploited to some potentially malicious end called a vulnerability. Finding vulnerabilities in new or deployed software takes work, but it is a rapidly growing industry, both inside and outside the government [17]. We treat these as information flows that require management decisions to determine how they will be used. Vulnerabilities are sought after by both "good guys" and "bad guys." What do they do with them?

Defenders seek vulnerabilities to fix them. Commercial software vendors and open source communities refer to closing vulnerabilities as "patching." In this paper, we gloss over the steps involved between identifying a vulnerability and securing a computer system. In the real world, this involves heavy testing to see if the fix breaks other components, distributing the patch and convincing users to install the new patch.

Attackers seek vulnerabilities to exploit the system running the software. A successful attack often be thought of as the ability to perform arbitrary tasks on a computer system without legitimate approval (although there are several other definitions). A mechanism for exploiting an undiscovered vulnerability is referred to as a 'zero day' attack. Just as attackers may have a wide range of motives, the uses of exploited machines are similarly varied. A common definition of security is a combination of confidentiality, integrity and availability. Absent an adequate level of security, an attacker can violate confidentiality by accessing sensitive data, the integrity of the system by feeding false information into the system, or the availability by preventing legitimate use.

Of course, there are many reasons why having the capacity to violate these principles of security would be good for the nation. Our adversaries have secrets we wish to know, and systems that could be used to threaten us, and vice versa. Cybersecurity management requires institutions capable of balancing the demands of attack and defense.

2.2 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 in 2010 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 mission boundaries between the civilian intelligence community and the offensive-focused defense community [5]. 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, some of which have been discussed elsewhere [26, 13]. Of particular importance to this paper is the challenge of protecting information 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 adversary cyberspace vulnerabilities" [25, 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.

2.3 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 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 [7].

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" [25, p. 16].

On the defense side, the NSA is involved in a number of projects to protect American information infrastructures, including the recently announced Perfect Citizen program [1].

In the context of commercial software and systems, the NSA has lent its expertise to Microsoft during the development of Windows 7 [9] and to Google for protecting the company's computer networks [18]. 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. This will come into play as a key part of the model below.

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 [22]. 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 [6]. Discussions of cyber war inside the defense establishment are laden with "macho rhetoric" [5], and declared policy seeks to "gain and maintain initiative" [25] and a "continued commitment to cyber superiority" [16]. 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.

3. MODELING CYBERSECURITY POLICY TRADEOFFS

The models presented in this paper explore the tension between attack and defense in a particular context. Specifically, what should a cybersecurity organization do upon discovery of a previously unknown software vulnerability? We argue that a civilian or uniformed manager faces two conflicting options: to use the knowledge as a weapon in a cyber arsenal, or to treat the knowledge as an opportunity to secure our own systems. The choice is to behave aggressively or defensively.

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? We explore the best strategy in a game context, where the optimal American policy will also reflect the strategy of others. As a two-player game, we can represent relationships with an arbitrary adversary. How the model can capture the attributes of the adversary and the relationship is discussed further below. A multiparty game is outside the scope of this paper.

The model's power derives from three critical, empirically-grounded assumptions. First, both players' networks rely on the common vulnerable software that can be exploited. That is, our adversary uses the same tools that we do, and therefore is can be harmed through the same weaknesses that we can. There is ample support for this assumption, particularly since both players are likely to use common platforms given the concentration in many sectors of the software industry. Around 90 % of the world's internet users use Microsoft Windows [8]; over 75% of all websites are run using one of two server software [20].

Second, we assume that patching, or fixing, one's own sys-

tem 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. Alternatively, a close study of a computer system can detect when it has been altered, opening the possibility of discovering how, and reverse engineering the defense.

Third, 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 [3, 4], and has received empirical support in the context of Windows vulnerabilities [21, Ch. 10].

A game follows this basic model, capturing the organizations' responses to the discovery of a vulnerability. Either the US or its adversary initially discovers a vulnerability. That country must decide whether to play defensively or offensively. A defensive action will lead to a patch, precluding malicious use of a system through that vector. However, an aggressive state cannot preclude its rival from also having a chance of independently rediscovering the same vulnerability. The rediscoverer must then make the same choice between offense and defense. The two games below walk through two types of offensive actions, built on either the capacity to attack or actually taking advantage of the vulnerability. The game is played sequentially but, as will be further explained below, a country may not have a clear idea whether it was the first or second to discover a vulnerability. The game is not iterative or cumulative; it is meant to capture the policy approach of a cybersecurity organization dealing with a single vulnerability.

4. GAME 1: VULNERABILITY STOCKPILES

One potential outcome of a searching for vulnerabilities is the 'stockpiling' of hidden exploits by nation-states to carry out future 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. Recall that these are mutually exclusive, since fixing ones own system can reveal knowledge of the vulnerability. Under what circumstances might we expect a patching initiative to win out over stockpiling the vulnerability? We use game theory to predict the optimal behavior in response to an adversary nation's cyber posture.

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

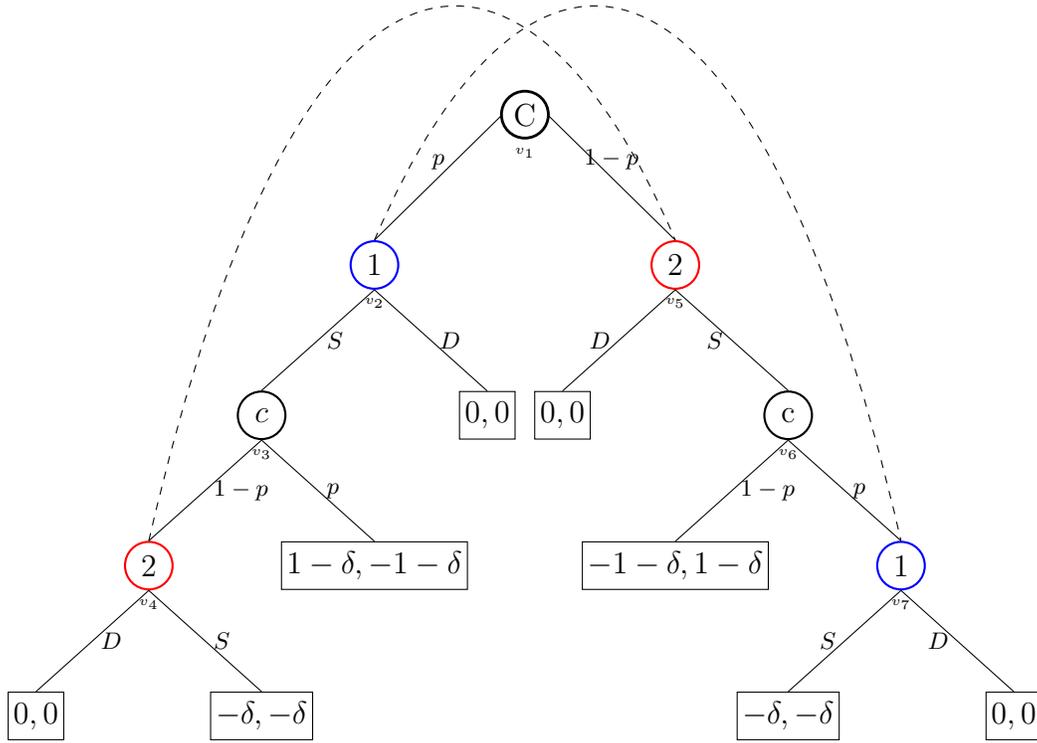


Figure 1: Vulnerability-stockpiles game. A player discovers a vulnerability with probability p and chooses to either defend or stockpile; stockpiling incurs the risk of the other player rediscovers it and stockpiling in turn. The harms of not being patched is δ

vulnerability a cyber warrior might discover. (Note that for this game, stockpiling a vulnerability does *not* mean actually launching a cyber-attack.)

4.1 Modeling the game

We first describe the key *actions*, *payoffs* and *parameters* used in the game, followed by a description of the game itself. Following the convention of game theory, each country is referred to as a player. Player 1 will be the United States in all games..

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

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 : Relative Technical Sophistication. 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 : Social risk from an unsecured machine. As discussed above, security threats can come from a more general social risk, such as the threat of cybercrime. 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 [2] 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 rediscover 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 probability $1 - p$ 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.

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 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. 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.

4.2 Finding equilibria

A strategy is a player’s decision to either stockpile (S) or defend (D). When a strategy is selected, it will be played for all decision opportunities that player has. (In this paper we do not consider *mixed strategies*, where players are allowed to randomize over pure strategies.)¹

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 $u_1(S, D)$, where player 1 is playing aggressively and player 2 is playing deensively, 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* [19] if, informally, no player can gain by unilaterally deviating. In our

¹Note 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.

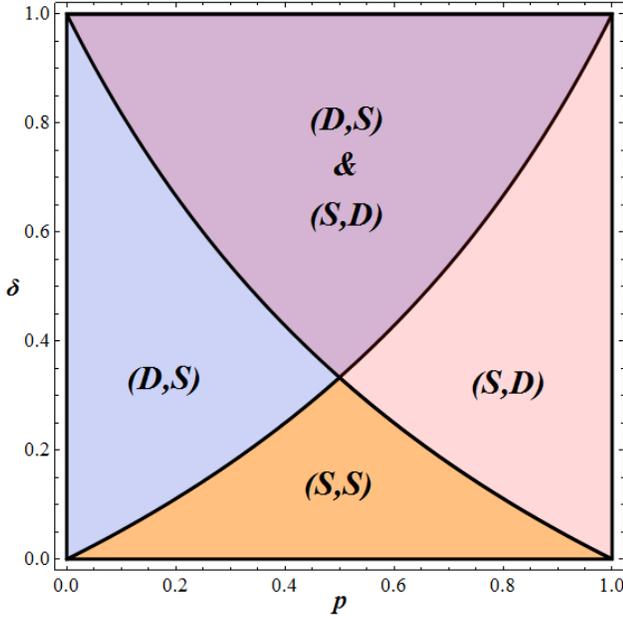


Figure 2: Different equilibria are possible for different values of p and δ in the vulnerability-stockpiles game.

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 $y' \in \{S, D\} \setminus \{y\}$.

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:

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

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

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. This plot traces out the Nash equilibria for different values of p and δ .

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.

4.3 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.

5. 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, rang-

ing 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 about cyber power, 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, q , 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.

5.1 Modeling the game

We model this game in the same fashion as the stockpiling game described in Section 4. We begin by describing the key *actions*, *payoffs* and *parameters* used in the game, followed by a description of the game itself. With the exception of the added dimension of the attack, this game is quite similar to that presented in the previous section.

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 : Relative Technical Sophistication. Valued between 0 and 1, this is a measure of the relative advantage of

player one in discovering vulnerabilities over player 2. (p has the same meaning as for game 1).

2. q : Aggression. This captures the relative likelihood that a player will choose to attack after discovering a vulnerability. 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 $q = 0.5$.

To maintain a small parameter space, 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.

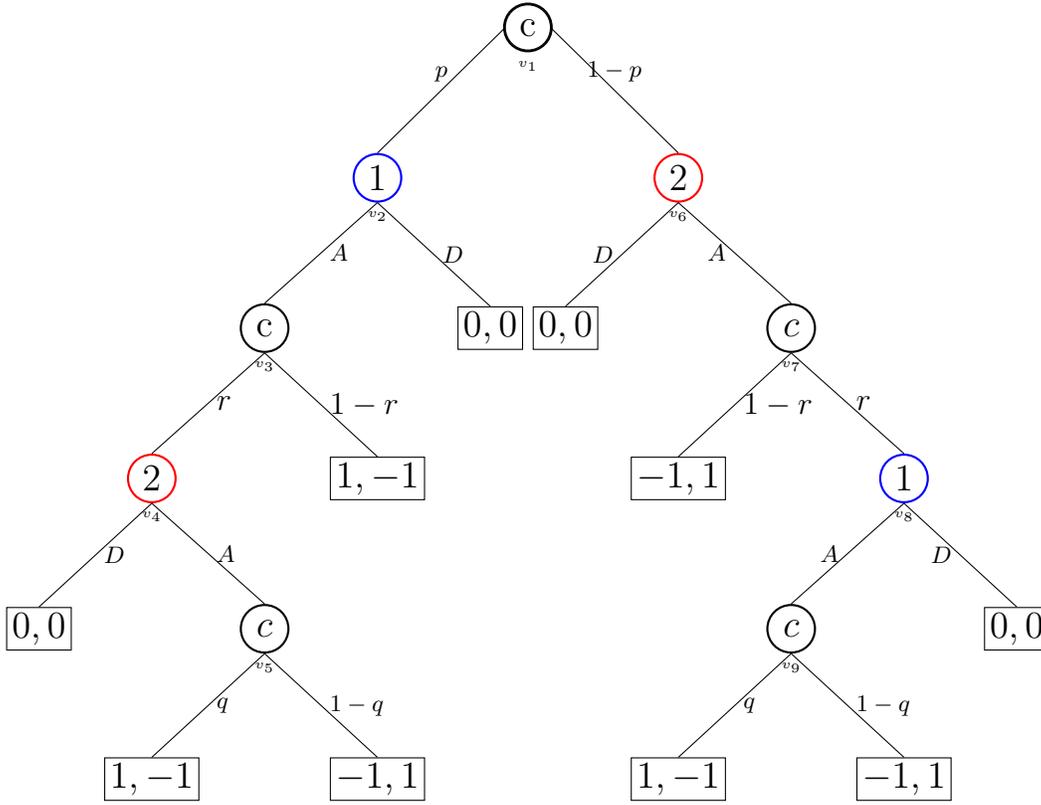


Figure 3: Cyber-hawk game. Similar to the vulnerability-stockpiles game, but when both players have the vulnerability, the winner of a zero-sum game is determined by the aggressiveness parameter q .

5.2 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, it must hold that $u_1(A, A) \geq u_1(D, A)$ 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 $u_1(A, D) \geq u_1(D, D)$ 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.

5.3 Discussion

For two evenly matched 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 tech-

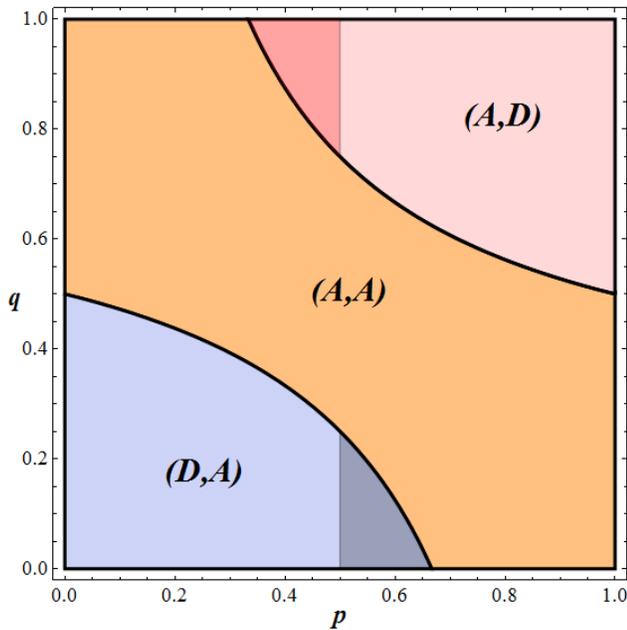


Figure 4: Different equilibria are possible for differing values of p and q in the cyber-hawk game.

nical 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.

6. 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. In both games, there is no equilibrium that would predict that

both the United States and its adversary would pursue a policy of actively patching newly discovered vulnerabilities. In the first game, if social cost is not built into the manager’s decision-making process, both parties will play aggressively. In the second game, the closer the rival in technical prowess and bellicosity, the more likely we are to both seek to attack the other.

The above results may be pessimistic with respect to increased information security, but these models can also inform the policy process. The findings of the game theory models predict how rational cybersecurity managers and commanders will behave under certain conditions, but they can also be used to explain observed behavior and make predictions about future behavior. How might a state actor move towards a world where sharing vulnerability data to improve domestic security is more common?

For a given technical level in the stockpiling game, incorporating an understanding of the social cost, the externalities of insecure systems, can increase the number of players willing to play a defensive strategy. This can be accomplished through education, management or organizational reorientation. Both games also allow the cyber war team to dominate the other player through technical competence. Increasing the technical advantage that a player wields over another increases the likelihood that she will adopt a defensive strategy. Of course, greater technical competence also rids our own cybersecurity teams from worrying have to consider a defensive posture.

6.1 Modeling Policy Interventions

The game paradigm in this paper is deliberately built to capture a single strategic choice, with a limited set of parameters for better tractability. Players have a symmetric strategy set. Yet one benefit of a parsimonious model is that it can be easily modified to explore new ideas. Can the models predict greater security through policy interventions? What steps can Player 1—the US government—take to encourage an approach to cybersecurity with a greater emphasis on defending its computer systems? Figures 5 and 6 explore the effects of a national network detection system and a regime of mandated transparency of harms through cyber attack.

In the above models, the likelihood of rediscovery of a vulnerability in possession of ones adversary is a function of technical prowess. There is little a country can do to discover the specific software flaw that another country already knows, apart from increase its total prowess. If, however, the adversary is attempting to test the effectiveness if its store of attacks without actually attacking, we can take certain steps to detect these trial runs. The United States is under a constant barrage of probing attacks, some of them quite mundane but others almost certainly comprise exploratory attempts to test new vulnerabilities. With the right resources, network security experts can detect these signs. It requires immense quantities of data and careful processing, and it may not even work. It might be able to, however, increase the rate of rediscovery.

Figure 5 demonstrates the effect of doubling the American ability to rediscover an opponent’s vulnerability through

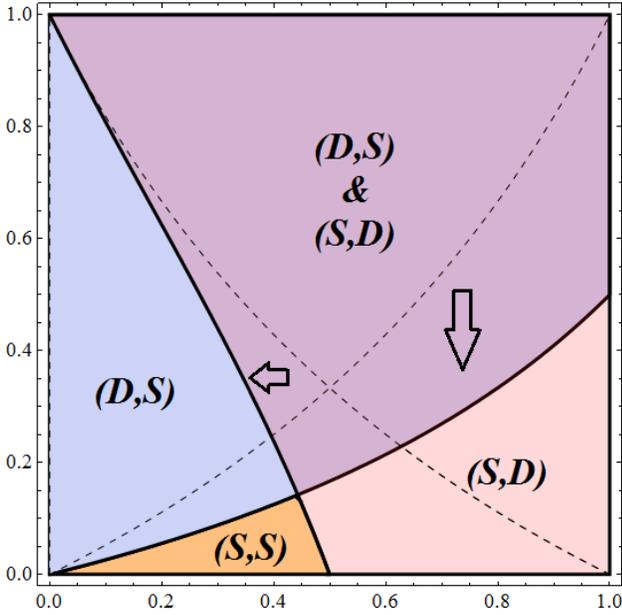


Figure 5: Policy Interventions in the Stockpiling Game. On the left, network detection efforts double Player 1’s likelihood of rediscovery. On the right, transparency about losses from doubles the cost of a potential attack on Player 1.

network detection. This capacity translates to a more defensive posture from our adversary. The dotted line from upper left to lower right denotes the equilibrium line for Player 2 from Section 4.2. Pivoting that line to the left increases the set of conditions where Player 2 would consider sharing vulnerability information. Note that most of this sweep represents a change when the social cost is quite low. This allows for the American cybersecurity organization to continue to play an efficient, aggressive strategy absent a high social cost, and yet patching is still more likely.

Similarly, in the Hawk game, our adversary also plays more defensively in response to doubling our rediscovery rate through network detection. In Figure 6, at the upper right, Player 2’s zone of defensiveness shifts left and down. Notably, defensive play is now more likely with a better-matched adversary.

The current framing of the games above equate the harms of a security advantage to a threat to the United States as roughly equivalent to one we could inflict on others. While this is the standard approach for modeling conflict [23], one policy option is to prioritize an attack against the United States as more damaging. Indeed, some suggest this is the case already [5].

There are other reasons to suggest this apart from pure America-centrism. Clarke suggests that we, as a country, are more dependent on our IT infrastructure than any of our potential adversaries for the foreseeable future. There is also an explicit policy paradigm that would increase the visible costs of a threatened or actual cyber attack: transparency.

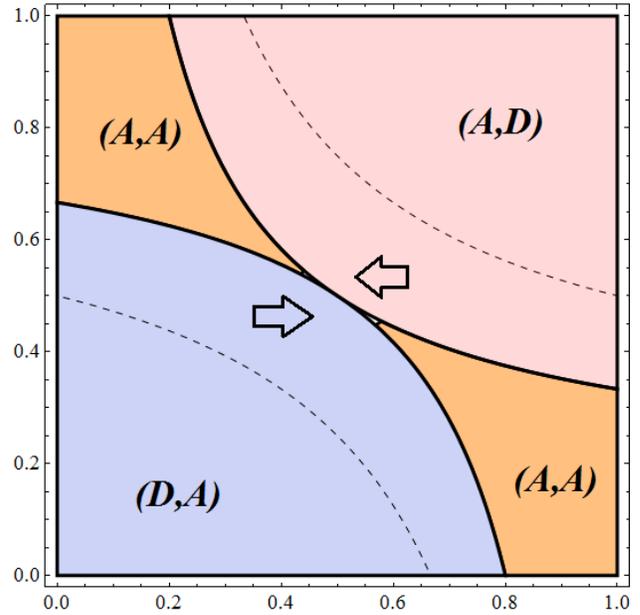


Figure 6: Policy Interventions in the Cyber-Hawk Game. On the left, transparency about losses from doubles the cost of a potential attack on Player 1. On the right, network detection efforts double Player 1’s likelihood of rediscovery.

Transparency about the loss of a database, a trade secret or negotiating position make the costs of an attack more salient to key decision makers. Transparency can not only drive better security behavior for privacy actors [BREACH06], but can push public actors to address a growing issue.

Introducing a transparency regime into the models is simply a matter of increasing the harms suffered by the US from a real or feared attack. In this case, we double it relative to our adversary. This is similar to increasing δ , except it only applies to a player one. This time, it is the best strategy for Player 1 that changes.

The right side of Figure 5 captures the changed equilibrium of doubling the relative cost of Player 1 (USA) being at a disadvantage. As the line starting at the origin and moving up and right shifts down, the US is less likely to stockpile and more likely to patch. Social cost becomes more important, particularly as the US approaches technical dominance. This paints a picture of a powerful state that will nonetheless prefer to have a safe environment.

Unsurprisingly, increasing the social cost of an actual attack in the Cyberhawk game also increases the range of conditions when a state will play defensively (lower-left side of Figure 6). The bias is still towards aggressive play when at a technical or behavioral advantage.

6.2 Understanding the National and International Policy Landscape

How do these findings compare to current events in cyber international relations? 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. There are also more subtle dynamics. Perversely, this model predicts an increase in global government attention to cyber defenses if it increases δ itself by raising the general social threat. That is, a rise in cybercrime might lead to better security. This is not to suggest that we cease fighting online crime, but it is important to understand the second-order effects of decisions. Take, for example, the tacit support Russian cybergangs receive from the Russian government [11]. While frustrating from an American law enforcement perspective, lax law enforcement in Moscow can increase the social cost for all players, actually leading to greater levels of vulnerability patching.

In the cyber hawk game, the perceived readiness and willingness to use a weapon factors into a player's decision process. 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 [15]. 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)).

Another side effect of the militarization of cyber space is an increased level of secrecy. Computer security, particularly with its roots in cryptography, disdains secrecy as "security through obscurity." Beyond slogans, however, is this increased 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 [15], 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.

7. 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.

Similarly, the game rests on many classic game theory components such as perfect information and a common understanding of the value of key parameters. Assuming similarly estimated values—or any reliable estimates at all—for some components, such as relative technical sophistication and aggressiveness would be dangerous. Nevertheless, a common framework of interaction such as the above games can actually lead to cooperation and data sharing for the purpose of conflict avoidance.

Finally, 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.

8. CONCLUSION AND RECOMMENDATIONS

Understanding governance issues in a complex policy environment requires a thorough understanding of incentives, and the interactions and effects of the decision set. This paper examined the management challenge created by unifying attack and defense capabilities for cybersecurity. We

presented two games that 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 model also allows us to make several recommendations to promote better investment in securing information systems.

1. *Bring the social cost into the decision-making equation.* Cybersecurity managers are more likely to consider securing themselves if the externalities of remaining unpatched are considered in the equation.
2. *Maintain technical superiority to increase defensive behavior in others.* Evenly matched adversaries are more likely to seek the benefits of an aggressive offensive strategy. Since the United States already has a formidable technological capacity, it should continue to cultivate this advantage.
3. *Specific civilian cybersecurity policies can, all other things equal, improve the levels of system security.* Imposing transparency laws that make the costs of the cyber threat more visible, and deploying systems to detect adversaries' offensive experimentation can shift the strategic equilibrium in favor of more defensive behavior on all sides.

Using these models, we can better understand the incentives facing states juggling the sometimes conflicting goals of cyber attack and defense, better analyze the context of domestic and international cybersecurity politics, and better guide the shaping of policies that promotes more security investment.

Acknowledgments

Discussions with Bruce Schneier, Melissa Hathaway and Roger Hurwitz inspired this article. The authors are also grateful to the participants of the CRCS Computer Science Seminar at the Harvard Computer Science department and the ECIR Cyber Policy Lunch at the Harvard Kennedy School.

9. REFERENCES

- [1] C. Albanesius. NSA confirms 'Perfect Citizen' exists, but as R&D? *PC Magazine*, July 2010. <http://www.pcmag.com/article2/0,2817,2366270,00.asp>.
- [2] R. Anderson and T. Moore. The economics of information security. *Science*, 314(5799):610–613, October 2006.
- [3] A. Arora, R. Telang, and H. Xu. Optimal policy for software vulnerability disclosure. *Management Science*, 54(4):642–656, 2008.
- [4] H. Cavusoglu, H. Cavusoglu, and S. Raghunathan. Emerging issues in responsible vulnerability disclosure. In *4th Workshop on the Economics of Information Security*, 2005.
- [5] R.A. Clarke and R.K. Knake. *Cyber War*. HarperCollins, 2010.
- [6] 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>.
- [7] R.V. Jones. *The wizard war : British scientific intelligence*. Coward, McCann & Geoghegan, 1978.
- [8] G. Keizer. Windows market share dives below 90% for first time. *Computer World*, December 2008. http://www.computerworld.com/s/article/9121938/Windows_market_share_dives_below_90_for_first_time.
- [9] G. Keizer. NSA helped with Windows 7 development. *Computerworld*, November 2009. http://www.computerworld.com/s/article/9141105/NSA_helped_with_Windows_7_development.
- [10] F. D. Kramer, S. H. Starr, and L. K. Wentz, editors. *Cyberpower and National Security*. National Defense University Press, 2009.
- [11] B. Krebs. Shadowy russian firm seen as conduit for cybercrime. *Washington Post*, October 2007. <http://www.washingtonpost.com/wp-dyn/content/article/2007/10/12/AR2007101202461.html>.
- [12] C. E. Landwehr, A. R. Bull, J. P. McDermott, and W. S. Choi. A taxonomy of computer program security flaws. *ACM Computer Survey*, 26, 1994.
- [13] J.A. Lewis. *Securing Cyberspace for the 44th Presidency*. Center for Strategic and International Studies, 2008.
- [14] M. C. Libicki. *Cyberdeterrence and CyberWar*. The Rand Corporation, 2009.
- [15] 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>.
- [16] N.A. Schwartz M.B. Donley. 2010 United States Air Force Posture Statement, 2010.
- [17] C. Miller. The legitimate vulnerability market. *Workshop on the Economics of Information Security*, June 2007.
- [18] 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>.
- [19] 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.
- [20] Netcraft. July 2010 web server survey. *Company Report*, July 2010. <http://news.netcraft.com/archives/2010/07/16/july-2010-web-server-survey-16.html>.
- [21] A. Ozment. *Vulnerability Discovery & Software Security*. PhD thesis, University of Cambridge, 2007.
- [22] 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.
- [23] T. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.
- [24] P. W. Singer. Double-hatting around the law. *Armed Forces Journal*, June 2009. <http://www.afji.com/2010/06/4605658>.
- [25] Chairman of the Joint Chiefs of Staff. The national military strategy for cyberspace operations, 2006.

- [26] H.S. Lin W.A. Owens, K.W. Dam, editor. *Technology, Policy, Law, and Ethics Regarding U.S. Acquisition and Use of Cyberattack Capabilities*. National Academies Press, 2009.
- [27] F. C. Zagare and D. M. Kilgour. Deterrence theory and the spiral model revisited. *Journal of Theoretical Politics*, 10, 1998.