

Timing of cyber conflict

Robert Axelrod¹ and Rumén Iliev

Ford School of Public Policy, University of Michigan, Ann Arbor, MI 48109

Contributed by Robert Axelrod, December 6, 2013 (sent for review October 26, 2013)

Nations are accumulating cyber resources in the form of stockpiles of zero-day exploits as well as other novel methods of engaging in future cyber conflict against selected targets. This paper analyzes the optimal timing for the use of such cyber resources. A simple mathematical model is offered to clarify how the timing of such a choice can depend on the stakes involved in the present situation, as well as the characteristics of the resource for exploitation. The model deals with the question of when the resource should be used given that its use today may well prevent it from being available for use later. The analysis provides concepts, theory, applications, and distinctions to promote the understanding strategy aspects of cyber conflict. Case studies include the Stuxnet attack on Iran's nuclear program, the Iranian cyber attack on the energy firm Saudi Aramco, the persistent cyber espionage carried out by the Chinese military, and an analogous case of economic coercion by China in a dispute with Japan. The effects of the rapidly expanding market for zero-day exploits are also analyzed. The goal of the paper is to promote the understanding of this domain of cyber conflict to mitigate the harm it can do, and harness the capabilities it can provide.

computer security | international conflict | stealth | persistence

The world's economy and international security have come to depend on what the Council on Foreign Relations called "an open, global, secure, and resilient Internet" (1). From an American point of view, cyber security is essential for the health of the nation's economy and national security. In fact, the Director of National Security, James R. Clapper, listed cyber security first among the threats facing America today (2). The risks include financial loss, loss of privacy, loss of intellectual property, breaches of national security through cyber espionage, and potential large-scale damage in a war involving cyber sabotage.

This paper seeks to contribute to our understanding of cyber conflict by developing ways to analyze the issues, and concepts with which to do the analysis.

This paper focuses on one aspect of the problem: the timing of a cyber conflict, either in the form of espionage or disruption. The paper takes the point of view of an actor who has a resource to exploit a vulnerability in a target's computer system, and a choice of just when to use that resource. The best time to use such a resource depends on the stakes involved in the present situation, as well as the characteristics of the resource for exploitation. A mathematical model is offered to help analyze such a choice. Our model deals with the question of when the resource should be used by the attacker, knowing that its use today may well prevent it from being effective later. The heart of our model is the trade-off between waiting until the stakes of the present situation are high enough to warrant the use of the resource, but not waiting so long that the vulnerability the resource exploits might be discovered and patched even if the resource is never used. The question of when to use a resource to exploit a vulnerability in the target's computer network is ultimately a matter of human judgment, as Milevski (3) says. The intent of our model is to help in making informed choices about the trade-offs involved in such a judgment.

Others studies have clearly recognized that a cyber weapon has a strong tendency to depreciate once used (e.g., ref. 4) The implication has often been more or less explicitly drawn that it may pay to wait for an appropriate moment to deploy such a resource.

The next step in the logic has also not escaped the notice of some previous studies, namely that the longer one waits the more likely the target will have recognized and fixed the vulnerability one's resource is meant to exploit. What is unique here is a formal model that specifies how to conceptualize the variables that are implicit in this logic of the timing of cyber conflict, how to specify the decision problem inherent in the trade-offs among these variables, and how to solve the resulting decision problem.

For the present paper, we need a term to describe the means to exploit a specific vulnerability in a given target's computer system. The term that will be used here is "cyber resource," or just "resource." A resource need not be a weapon in the sense of something that can cause damage by itself. It might instead be used for espionage in which case its use is not necessarily an attack. A resource might use one or more zero-day exploits, which are ways to take advantage of hitherto-unknown vulnerabilities in computer software. In addition, a resource to exploit a target's vulnerability might include nontechnical means such as social engineering or an insider, either on their own or in conjunction with technical means of intrusion.

The question of timing is analogous to the question of when to use a double agent to mislead the enemy, where it may be worth waiting for an important event but waiting too long may mean the double agent has been discovered by the target and become useless. The present model is an adaption and extension of the model developed to study "the rational timing of surprise" (5).

Our model is presented from the perspective of the offense: when should a cyber resource be used to exploit a vulnerability in a target's computer network. The results, however, are equally relevant to a defender who wants to estimate how high the stakes have to be in order for the offense to exploit an unknown vulnerability.

Section 1 provides a model that expresses the value of a resource for exploiting a vulnerability in the target's computer system, and then calculates when best to use that resource. The development of our model provides some useful concepts including the Stealth and Persistence of a resource for exploiting a cyber vulnerability. Section 2 applies our model to illuminate

Significance

The world's economy and international security have come to depend upon a secure Internet. International rivalries and conflicts have already provided challenges to Internet security in the form of espionage, sabotage, and denial of service. New vulnerabilities in computer systems are constantly being discovered. When an individual, group, or nation has access to means of exploiting such vulnerabilities in a rival's computer systems, it faces a decision of whether to exploit its capacity immediately or wait for a more propitious time. This paper introduces a simple mathematical model applied to four case studies to promote the understanding of the new domain of cyber conflict.

Author contributions: R.A. and R.I. designed research; R.A. and R.I. performed research; R.A. and R.I. contributed new reagents/analytic tools; R.A. and R.I. analyzed data; and R.A. wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: axe@umich.edu.

some recent cases including the Stuxnet attack on the Iranian nuclear program, the Iranian cyber attack on Saudi Aramco, the persistent cyber espionage carried out by the Chinese military, and Chinese economic warfare against Japan. Section 3 analyzes the effects of the flourishing market for zero-day exploits, i.e., ways to take advantage of hitherto-unknown vulnerabilities in computer software or hardware. Section 4 concludes with a review of concepts, a list of some future avenues for research, and the implications of our model for the future of cyber conflict.

1. When to Use a Resource to Exploit a Vulnerability

1.1. The Model and Its Assumptions. In determining when to use a resource to exploit a vulnerability in a target's computer system (hereinafter "resource"), you should take into account what is at stake in the current situation. For example, in one year you may be at war with the target, making the stakes very high. In another year, you may be at peace but you may have just discovered that the target has some new technology you would like to be able to steal, so the stakes would be moderate. In still another year, you may have no problems with the target and the stakes would be low.

1.1.1. Assumption 1. Stakes. You know the current stakes. You do not know what the stakes will be at any future point, although you do know the distribution of stakes over time.

The assumption about stakes means you may know that the stakes are low today, and you may be able to estimate the likelihood of various possible stakes in the future, but you do not know when—if ever—the stakes associated with a particular event will occur. (An alternative assumption is that the stakes are path dependent. Path dependence of the stakes will not be explored in this paper, although it would be appropriate to consider for very short time intervals.) Later, we will look at the implications of several distributions on the likelihood of various stakes—such as how likely high-stakes events are compared with routine low-stakes events.

1.1.2. Assumption 2. Resource characteristics. For a given resource, you can estimate two parameters that determine whether the resource will be available next time, say 1 y later. These are the Stealth and the Persistence of the resource. The Stealth of a resource is the probability that if you use it now it will still be usable in the next time period. The Persistence of a resource is the probability that if you refrain from using it now, it will still be useable in the next time period.

The characteristics of a resource can be stated more formally in terms of conditional probabilities as follows:

$$S = \text{Stealth} = \Pr(\text{resource survives} \mid \text{use it}), \text{ and}$$

$$P = \text{Persistence} = \Pr(\text{resource survives} \mid \text{not use it}).$$

For estimating Stealth of a given resource, a relevant benchmark is that the average duration of a zero-day attack is 312 d (6). Another benchmark is provided by the Conficker worm that infected ~370,000 machines without being detected over more than 2 mo (6).

For estimating Persistence of a given resource, a relevant benchmark is that in a 3-y period, only about 3–5% of the hundreds of vulnerabilities found in the Chrome and Firefox browsers were independently rediscovered (6). Therefore, a resource designed to take advantage of these vulnerabilities would have had P close to 1.0.

Both Stealth and Persistence depend not only on the resource itself, but also on the capacity and vigilance of the intended target. The Stealth of resource used against a well-protected target is likely to be less than the Stealth of the same resource against a target that is not particularly security conscientious. Likewise, a resource will typically have less Persistence against a target that keeps up-to-date on security patches than one that does not. In the case of a distributed denial of service, the

effectiveness of the attack depends on the current capacity of the target to handle massive inputs, whereas the ability of the attacker to repeat the attack (i.e., Stealth) depends on the target's subsequent attainment of sufficient capacity to handle another such attack.

Because stakes are not under your control, your best policy is to wait until the stakes are high enough to risk losing the resource because of its limited stealth. This means your best policy can be expressed in terms of T , the Threshold of stakes that will cause you to use the resource. For example, with linear stakes a policy of using the resource only on a roll of 5 or 6, gives average short-term gain, $G(T)$, to be $G(5) = (5 + 6)/2 = 5.5$. Note that the lower the Threshold you choose, the more often you will use the resource, but the lower the average gain will be when you do use it. This illustrates the fundamental problem: you want to use the resources as often as possible (T low), but you also want to preserve it for the times when the stakes will be large (T high).

The last consideration is the discount rate, a reflection of the fact that a given payoff is less a year from now than it is today. A typical discount rate, w , is about 0.9 meaning that a dollar today is worth only 90 cents a year from now.

The policy question is how to choose T to maximize the value of the resource. The value of the resource when used today is the expected gain from this use (which depends on the Threshold) plus the expected future value that depends on the discount rate, and its Stealth, namely $G(T) + wSV$. The value if not used is the discounted value of the chance it will survive, namely wPV .

The chance the resource will be used at a given time is the probability that the current stakes, s , is at least as great as the Threshold, namely $\Pr(s \geq T)$. The chance that the resource will not be used at a given time is the complement of this, namely $1 - \Pr(s \geq T)$. This gives the discounted expected value of the resource, V , as follows:

$$V = \Pr(s \geq T)(G(T) + wSV) + (1 - \Pr(s \geq T))wPV. \quad [1]$$

1.1.3. Assumption 3. Value of a resource. The value to its owner of a resource to exploit a target's vulnerability depends on its Persistence and Stealth, and the distribution of future stakes as specified in Eq. 1.

Solving for V in Eq. 1 gives the following:

$$V = \Pr(s \geq T)G(T) / [(1 - wP) + \Pr(s \geq T)w(P - S)]. \quad [2]$$

Now that we have an equation for the value of a resource for exploiting a target's vulnerability, we can evaluate what that resource is worth. Even more useful is that we can calculate the best way to use the resource in terms of the optimal Threshold, T_{opt} , specifying how large the stakes have to be to make it worthwhile to use the resource and take the added risk that it will no longer be available.

1.2. Optimal Timing. To illustrate the determination of T_{opt} , we return to the simple case where the distribution of stakes is linear, meaning that in each year there is an equal chance that the stakes will be 1, 2, 3, 4, 5, or 6. We focus first on how the optimal Threshold depends on the Persistence of the resource. In this example, we use a discount rate of $w = 0.9$. For convenience, we assume that using the resource cuts in half the chance that the vulnerability will continue for a year, making the Stealth half of the Persistence, $S = P/2$. Using these values for the distribution of stakes, w , and S , we can use Eq. 2 to calculate the value of the resource for different levels of Persistence and different policies for the Threshold. The results are shown in Table 1.

Table 1 shows the effect of Persistence on the value of a resource as a function of the Threshold, T . Notice that in the first column, where Persistence is only 10%, the highest Value comes from the policy of always using the resource ($T = 1$). So our

Table 1. The effect of Persistence

T	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
6	1.09	1.20	1.33	1.49	1.70	1.98	2.37	2.94	3.88	5.71
5	1.98	2.16	2.37	2.62	2.93	3.33	3.86	4.58	5.64	7.33
4	2.68	2.89	3.13	3.42	3.77	4.20	4.74	5.43	6.37	7.69
3	3.19	3.41	3.66	3.95	4.29	4.69	5.17	5.77	6.52	7.50
2	3.52	3.72	3.96	4.22	4.52	4.87	5.27	5.75	6.32	7.02
1	3.66	3.85	4.05	4.27	4.52	4.79	5.11	5.47	5.88	6.36

The value of a resource is shown as a function of the chosen Threshold (rows) and the resource's Persistence (columns). Stealth is set to one-half the Persistence, and the discount rate, w , is 0.9. The highest value for each level of Persistence is indicated in bold type. Note that the higher the Persistence, the higher the optimal Threshold for use.

model indicates that if your resource has very low Persistence you should use it right away, even if the stakes are very low. This makes sense because low Persistence implies that the resource has only a small chance of surviving until next time even if you do not use it now. In general, as Persistence increases, the optimal Threshold increases.

Table 2 shows the effect of Stealth. Notice that, although large Persistence and Stealth are both valuable attributes of a cyber resource, they have opposite effects on policy. As Stealth increases, the optimal Threshold for use decreases, whereas we saw the opposite effect for Persistence.

Having considered settings in which the distribution of stakes is linear, we can now look at the effects of different distributions. A simple setting is where the stakes are constant over time. For example, the stakes will be constant for criminals whose payoff comes from exploiting stolen credit card information. Likewise, a terrorist organization that is bent on causing harm may see the stakes as relatively constant. When stakes are constant the optimum policy is to use a cyber resource as soon and as long as possible.

In international affairs, by contrast, the most important events are really important, far more so than typical events. An example of a distribution in which the stakes are highly skewed is the exponential distribution. An exponential distribution of stakes can be illustrated in the die roll game by making the payoffs 1, 2, 4, 8, 16, and 32, rather than linear as in the original game with stakes of 1, 2, 3, 4, 5, and 6. In terms of the model, the more skewed the distribution of stakes, the higher the optimal Threshold. In addition, the higher the Threshold, the longer the expected wait until the stakes exceed that Threshold.

The takeaway from this analysis is that there are three factors that make it pay to wait for high stakes before using a cyber resource to exploit a vulnerability in a target's computer system. The three factors that favor patience are low Stealth, high Persistence, and large stakes in rare events.

Table 2. The effect of Stealth

T	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
6	2.60	2.70	2.82	2.94	3.08	3.23	3.39	3.57
5	3.74	3.99	4.26	4.58	4.95	5.39	5.91	6.55
4	4.20	4.55	4.95	5.43	6.02	6.76	7.69	8.93
3	4.29	4.69	5.17	5.77	6.52	7.50	8.82	10.71
2	4.14	4.57	5.09	5.75	6.60	7.75	9.39	11.90
1	3.85	4.27	4.79	5.47	6.36	7.61	9.46	12.50

The value of a resource is shown with different levels of Stealth, and constant Persistence ($P = 0.8$). The highest value for each level of Stealth is in bold type. Note that the optimal Threshold goes down as the Stealth increases.

2. Application to Case Histories

2.1. Stuxnet: The Stealthy Attack That Got Away. The basic story of Stuxnet is familiar: an extraordinarily sophisticated computer worm infected the control system of Iran's nuclear enrichment plant at Natanz, and temporarily disabled 1,000 of the 5,000 centrifuges there (7).

In the terms of this paper, the worm probably had low Persistence because it used four different zero-day exploits to accomplish three functions: loading the malware from a flash drive, spreading the malware to other machines sharing a printer, and escalating the attacker's privileges on a machine and giving full control of it (8) (also see the Symantec official blog, www.symantec.com/connect/blogs/stuxnet-using-three-additional-zero-day-vulnerabilities, accessed Dec. 24, 2013). The low Persistence must therefore have put pressure on the attackers to use their resource sooner rather than later.

Stuxnet's designers took great care to make it Stealthy and succeeded in avoiding detection for 17 mo (9). Instead of simply destroying the centrifuges, it caused some to speed up in short bursts that damaged but did not destroy them. In addition, Stuxnet masked the change in speed by preventing the control panel from revealing what was happening. After completing its mission, Stuxnet erased itself.

The stakes involved in the Stuxnet attack were to delay Iran's ability to attain enough enriched uranium for nuclear weapons. Once Iran did attain nuclear weapons capability, further delays in its enrichment program would be much less valuable than delays before it achieved that capability. In terms of our model, the attacker's view of the current stakes must have been very high, for both the United States and Israel.

Our model predicts that a resource like Stuxnet that was expected to have poor Persistence and comparatively good Stealth would be used as soon as possible, and certainly in a high-stakes situation. This is apparently just what happened (7).

The delay achieved was probably about 6 mo because only about a fifth of the centrifuges were damaged, and even that damage seems to have been repaired within 2 mo of the attack. How much was gained from this delay remains to be seen.

The Stuxnet case highlights two factors that came into play only because of the inadvertent escape of the Stuxnet code. These two factors should have been included in the estimate of possible gains (and losses) from its use, $G(T)$, but might not have been given proper weight because the escape of the code was probably quite unexpected. The first factor was the negative effect of breaking the precedent against cyber attacks used for industrial sabotage. [The need to consider international norms is now included in Presidential Policy Directive 20, "U.S. Cyber Operations Policy," November 16, 2012. The document is classified TOP SECRET/NOFORN, leaked by Edward Snowden, and published by the Federation of American Scientists (10).] The second factor was the positive effect of alerting those responsible for critical infrastructure in the United States and elsewhere that security breaches of industrial control systems such as SCADA were not just hypothetical, but had to be taken seriously.

2.2. Iranian Attack on Saudi Aramco. Shortly after Stuxnet was revealed, Iran launched a new wave of cyber attacks against Saudi Arabia and the United States with the aim of destroying data and manipulating machinery such as oil pipelines (11). The attack spread to 30,000 Aramco workstations but did not escape to harm the company's critical infrastructure because those systems were on isolated networks (12).

The effect of the attack was modest, although US Secretary of Defense Leon Panetta declared it to be the most destructive cyber assault the private sector has seen to date (11). In terms of our model, the attack was not very Stealthy. Indeed, it was promptly

noticed and stopped, and within 4 d cleanup efforts were completed (9). Mistakes in the attacking program suggested that the attack was prepared in haste. From the point of view of the stakes involved, the Iranians presumably felt that haste was needed to demonstrate to both domestic and international audiences that they were not passive. With sufficiently high stakes, the model predicts immediate use of a cyber resource, which is just what happened.

2.3. Everyday Chinese Cyber Espionage. The Chinese Army has for years been deploying cyber resources for espionage against defense and industrial targets in the United States and elsewhere (13, 14). Their cyber espionage often has only moderate Stealth against vigilant targets, so it is frequently discovered. It is able to continue because many of the targets have not maintained state-of-the-art defenses, or even kept up-to-date on patches for known vulnerabilities.

A result of widely detected industrial espionage was hostility against the Chinese government. American officials acknowledge that all countries spy on each other, but they say China is unique in its theft of foreign technology (15). The problem has become so serious that in the context 2013 Sino-American summit meeting, Obama's National Security Adviser, Tom Donilon, said that resolving cyber security issues would be "key to the future" of the relationship between China and the United States (16). [China denies being the source cyber espionage and insists it is a major victim of cyber attacks, including from the United States (17). Denial of espionage has long been so common that any country's denial of espionage has the status of a "polite fiction" (18).]

In terms of our model, one might well ask why the Chinese are deploying their resources for cyber exploitation now when the stakes are not particularly high, rather than wait for a time when the stakes are much higher? In other words, why might the Chinese be operating with a low Threshold? One possibility is that Chinese might have thought that the resource they were deploying had a low shelf life (low P). Another reason might be that China expected high S against at least some targets because it has taken outliers several years to even detect that they have been compromised (19).

2.4. Premature Chinese Use of a High-Persistence, Low-Stealth Resource. The three cases considered so far, all illustrate the expected timing for using the relevant resource based on the Stealth, Persistence, and stakes involved. To illustrate a situation in which the timing of the employment of a resource does not seem optimal, consider the case of the Chinese halt of its rare-earth exports pressure to provide strong economic pressure against Japan.

The incident started on September 7, 2010, when a Chinese fishing trawler collided with a Japanese patrol ship near some disputed islands. The Japanese took the trawler to Japan. On September 9 and again on the 12th, the Chinese demanded that the captain and crew be released. The next day, Japan released the crew, but continued to detain the captain. Tension continued to escalate, and on September 21, China abruptly halted its exports of rare-earth materials. Rare-earth materials are needed for the production of electronics, automobiles, and much else. Because China controlled 97% of the world's exports of rare earths, and Japan imported one-half of that supply, the effects on Japan of the cutoff were immediate and drastic. Japan complained that this was economic warfare and released the captain within 3 d. China waited a month to restore exports to most of the world, and 2 mo before restoring exports to Japan (20, 21).

After this demonstration of economic coercion, Japan, the United States, and others invested in production of rare earths outside of China so as to never be subject to the same threat again. Clearly, China had the ability to stop the global supply of minerals essential for manufacturing electronics and automobiles. In terms of our model, this ability had very high Persistence

because until the Chinese actually stopped exports, other countries were happy to close down their own production in favor of the cheaper Chinese supply. Presumably, China dominance could have persisted for years. When China did use its coercive power, it tried to attain some Stealth by never acknowledging that the cutoff had any political purpose. However, the timing was so obviously connected to the Japanese detention of the trawler captain that there was little doubt that the cutoff was quite deliberate. In addition, once the Chinese did deploy this ability to coerce by cutting off exports of rare earths, they lost their ability to coerce again in the same way because importing nations woke up to the risk and took effective measures to end their total dependence on Chinese exports.

In terms of our model, China's ability to coerce others with a cutoff of rare-earth exports would have had very low Stealth because not only was the cutoff very obvious, but steps that could have been (and were in fact) taken to remove the vulnerability became salient. The resource had high Persistence because Chinese dominance had persisted for years already and would probably have persisted for many more years had the resource for coercion not been used when it was. As we have seen, a resource with both low Stealth and high Persistence has a very high optimal Threshold for use. Although the recovery of the release of the trawler captain was important to China, it is hard to see how the stakes in 2010 were greater than they would be in other situations that might arise in the not-too-distant future.

Second-guessing a nation's choice is always problematic. Nevertheless, our model strongly suggests that the Chinese would have been better off had they had the patience to wait for a situation with much higher stakes before deploying this particular low-Stealth and high-Persistence resource for economic coercion.

3. Effects of the Market in Zero-Day Exploits

The aboveboard market in zero-day exploits includes vendors such as Microsoft offering bounties for hitherto-unknown ways to exploit their products. Most of the market, however, is in a gray area: buyers and sellers preferring not to make their transaction public. Price estimate vary widely, but there are reports that the most valuable zero-day exploits can fetch over \$100,000 (11, 22, 23). [For much lower prices charged by the Russian cyber underground for basic types of hacker activity, see the study by Goncharov (24).] The US Government, through the National Security Agency and defense contractors, is the biggest buyer of all (22).

One might suppose that the market for zero-day exploits would quickly become saturated, at least for the discovery of new exploits in the most common software. Surprisingly, the evidence is that there is a very large pool of undiscovered vulnerabilities. For example, Finifter et al. (7) found that in the 3-y period of 2009–2012, there were over 400 serious problems found in the Firefox browser and more than 800 in the Chrome browser. With new versions of commonly used software being introduced at a high rate to patch recently discovered vulnerabilities and to add new features, the pool of zero-day exploits waiting to be discovered is ever renewable.

A flourishing market in zero-day exploits causes an increase in the rate of discovery of such resources, with a consequent increase the number of potential exploits actually available for use at any one point in time. Our model indicates that the more effort that is devoted to discovering hitherto-unknown vulnerabilities,

- i) the more will be the decrease in Persistence of a given resource due to the greater likelihood of its being rediscovered by someone else, and then sold to a potential target before the resource is used;
- ii) the lower Threshold for use for a given cyber resource because of its lower Persistence,
- iii) the sooner a given resource is likely to be used, and

- iv) the lower the price of a given cyber resource because supply will increase and the reduction in Persistence will make each resource less valuable.

4. Conclusion

4.1. Optimal Timing. Cyber conflict has already begun. Exploitation of vulnerabilities in computer systems has been used for both espionage and sabotage. Exploitation of vulnerabilities has also led to new ways of conducting crime and fighting crime; maintaining anonymity and destroying anonymity, resisting political authority, and reinforcing political authority. In the near future, cyber conflict will likely allow international sanctions to be more precisely targeted than economic sanctions alone and will provide powerful force multipliers for so-called kinetic warfare.

This paper clarified some of the important considerations that should be taken into account in any decision to use a method of exploiting a target's vulnerability. The focus has been on optimal timing for such use. This kind of analysis can help users make better choices and help defenders better understand what they are up against. In some situations, one may want to mitigate the potential harm from cyber conflict, and in other situations, one may want to harness the tools of cyber conflict. In some cases, one might want to do both. In any case, an important step is to understand the logic inherent in this new domain.

The implications of our model are easy to summarize: Stealth and Persistence are both desirable properties of a resource, and increase its Value (Eq. 2). However, they have opposite effects on the best time to use the resource. Persistence leads to more patience, meaning the stakes need to meet a higher Threshold before the resource is worth using. The reason is that with high Persistence you do not need to worry very much about the resource becoming obsolete before you use it. High Stealth, however, promotes use even with relatively low stakes because the resource is likely to be reusable. Moreover, in a world of exponential (rather than linear) stakes, the chance of occasional very high Gains increases the Threshold because those very high stakes are more worth waiting for. Turning the perspective around, it would be a mistake to evaluate one's own vulnerability by what one sees when the stakes are low or moderate. The potential attacker may well be waiting for an event of sufficiently high stakes to exploit the cyber resources it already has.

4.2. Useful Distinctions. Our model provides a systematic way of thinking about cyber conflict. Just as important as the model itself is the set of concepts and distinctions it clarifies.

1. A zero-day exploit vs. an effective resource for exploiting vulnerability in a target's computer systems. A given resource may require more than one zero-day exploit to be effective and may also require some nontechnical components as well.
2. Gain vs. Value. The immediate Gain to be had from using such a resource needs to be distinguished from the long-term Value of the resource based on the chance that it might be useable more than once.
3. Persistence vs. Stealth. Both terms refer to the probability that the resource will still be effective in the future (say in a year). Persistence is the probability that the resource will be

effective if it is not used now, and Stealth refers to the lower probability that it will still be effective next time if it is used now. The distinction is important because high Persistence raises the optimal Threshold for use, but high Stealth has the opposite effect.

4. Known current stakes and the future stakes with only known distribution.
5. Different distributions of stakes. For example, we have seen the distinction between constant, linear, and exponential stakes. In general, the larger the stakes are in the rare events compared with common events, the higher the optimal Threshold and the more patient it pays to be.

4.3. Future Research. An important question for empirical research is how a decision maker can estimate the parameters of our model when applied to a specific resource aimed at exploiting the vulnerabilities of specific target. Some empirical facts were offered to give a sense of how to estimate Persistence and Stealth. What is most needed is a sophisticated understanding of how to estimate the potential Gains (and losses) from actually using the resource in a particular setting. As we saw, these Gains (and losses) arise not only from the direct effects of the intrusion, but also the indirect effects such as increased vigilance if the news gets out, and also the political effects if an emerging international norm is violated or reinforced. Gains (and losses) from the employment of cyber resources will also be affected by whether or not a state of armed combat exists between the two sides.

An important question for both theoretical and empirical research is the speed at which newly revealed vulnerabilities are actually corrected by various individuals, organizations, and governments. [One recent study found that about 30% of released vulnerabilities are still not patched after 1 mo (25). The figures are Microsoft, 31%; Apple, 26%; and Linux, 35%.]

Finally, our model could be extended to explicitly deal with the interaction of two sides of a potential conflict when both sides can try to exploit the vulnerabilities of the other, and both sides can be the targets of the other. Although the present model offered some insights into cyber arms races and crisis stability, a more interactive game-theoretic treatment would allow richer considerations of mutual learning, deterrence, and the presence of more than two holders of cyber resources.

4.4. Mitigating and Harnessing. As Clarke and Knake (26) put it, "It took a decade and a half after nuclear weapons were first used before a complex strategy for employing them, and better yet, for not using them, was articulated and implemented." This paper provided some concepts, theory, applications, and distinctions to promote the understanding of this new domain of cyber conflict. The goal is to mitigate the harm cyber conflict can do, and harness the capabilities it can provide.

ACKNOWLEDGMENTS. For helpful discussions, we thank Keith Alexander, John Arquilla, Vera Axelrod, Ted Belding, John D. Ciorciari, Steve Crocker, Stephanie Forrest, Brandon Kaplowitz, Jeffrey Karrels, Herb Lin, James Morrow, Jose Nazario, Love Rönnelid, Heather Roff, Amy Saldinger, Anton Severin, Carl Simon, Peter Singer, and Allan Stam. This work was supported by Air Force Office of Scientific Research Grant FA9550-10-1-0373.

1. Council on Foreign Relations (2013) Defending an open, global, secure, and resilient Internet. *Independent Task Force Report* (Council on Foreign Relations, New York), No. 70.
2. Clapper J (2013) *Statement for the Record: Worldwide Threat Assessment of the US Intelligence Community* (US Senate Select Committee on Intelligence, Washington, DC).
3. Milevski L (2011) Stuxnet and strategy: A special operation in cyberspace. *Joint Force Quarterly* 63(October):64–69.
4. Libiki M (2007) *Conquest in Cyberspace* (Cambridge Univ Press, Cambridge, UK), p 87.
5. Axelrod R (1979) The rational timing of surprise. *World Polit* 31:228–246.
6. Bilge L, Dumitras T (2012) Before we knew it: An empirical study of zero-day attacks in the real world. *Proceedings of the 2012 ACM Conference on Computer and Communications Security* (Association for Computing Machinery, New York), pp 833–844.

7. Finifter M, Akhawe D, Wagner D (2013) An Empirical Study of Vulnerability Rewards Programs. *22nd USENIX Security Symposium* (USENIX Association, Berkeley, CA).
8. Sanger DE (June 1, 2012) Obama order sped up wave of cyberattacks against Iran. *NY Times*, Section A, p 1.
9. Langner R (2013) *To Kill a Centrifuge: A Technical Analysis of What Stuxnet's Creators Tried to Achieve* (Langner Group, Arlington, VA).
10. Federation of American Scientists (2012) U.S. Cyber Operations Policy. Presidential Policy Directive 20, November 16, 2012 (Federation of American Scientists, Washington, DC). Available at www.fas.org/irp/offdocs/ppd/ppd-20.pdf. Accessed July 27, 2013.
11. Perloth N, Sanger DE (May 24, 2013) New computer attacks traced to Iran, officials say. *NY Times*, Section A, p 10.

12. Schwartz MJ (August 27, 2012) Saudi Aramco restores network after Shamoon malware attack. *InformationWeek Security*.
13. Baldor LC (October 12, 2012) US: Hackers in Iran responsible for cyberattacks. *Yahoo! News*.
14. Sanger DE (May 6, 2013) U.S. blames China's military directly for cyberattacks. *NY Times*, Section A, p 1.
15. Eckert P, Yukhananov A (July 10, 2013) U.S. opens China talks with cyber complaints, vow to boost trust. *Reuters*.
16. Donilon T (2013) Press Briefing By National Security Advisor Tom Donilon (The White House, Washington). Available at <http://www.whitehouse.gov/photos-and-video/video/2013/06/08/press-briefing-national-security-advisor-tom-donilon#transcript>. Accessed Dec. 26, 2013.
17. Xinhua (February 19, 2013) China opposes hacking allegations: FM spokesman. Available at http://news.xinhuanet.com/english/china/2013-02/19/c_132178666.htm. Accessed Dec. 24, 2013.
18. Axelrod R, Zimmerman W (1981) The Soviet press on Soviet foreign policy: A usually reliable source. *Br J Polit Sci* 11:183–200.
19. Strohm C (March 12, 2013) Cyber attackers' tactics outpace companies' responses. *Bloomberg*.
20. Bradsher K (November 19, 2010) China restarts rare earth shipments to Japan. *NY Times*, Section B, p 10.
21. Speigel R (December 7, 2010) Electronics industry braces for rare-earth-materials shortages. *EDN*.
22. Menn J (May 10, 2013) Special report—U.S. cyberwar strategy stokes fear of blowback. *Reuters*.
23. Greenberg A (March 23, 2012) Shopping for zero-days: A price list for hackers' secret software exploits. *Forbes*. Available at <http://www.forbes.com/sites/andygreenberg/2012/03/23/shopping-for-zero-days-an-price-list-for-hackers-secret-software-exploits/>. Accessed Dec. 24, 2013.
24. Goncharov M (2012) Russian underground 101. *Trend Micro Incorporated Research Paper* (Trend Micro Incorporated, Cupertino, CA).
25. Shahzad M, Zubair Shafiq M, Liu AX (2012) A large scale exploratory analysis of software vulnerability life cycles. *Proceedings of the 34th IEEE/ACM International Conference on Software Engineering (IEEE, Zurich)*.
26. Clarke RA, Knake R (2010) *Cyber War: The Next Threat to National Security* (Harper Row, New York).