Crime Pays If You Are Just an Average Hacker*

Allodi, Luca
University of Trento
Italy
Email: allodi@disi.unitn.it

Massacci, Fabio
University of Trento
Italy
Email: massacci@disi.unitn.it

Shim, Woohyun
University of Trento
Italy
Email: woohyun@disi.unitn.it

November 7, 2012

Abstract

This study investigates the effects of incentive and deterrence strategies that might turn a security researcher into a malware writer, or vice versa. By using a simple game theoretic model, we illustrate how hackers maximize their expected utility. Furthermore, our simulation models show how hackers' malicious activities are affected by changes in strategies employed by defenders. Our results indicate that, despite the manipulation of strategies, average-skilled hackers have incentives to participate in malicious activities, whereas highly skilled hackers who have high probability of getting maximum payoffs from legal activities are more likely to participate in legitimate ones. Lastly, according on our findings, we found that reactive strategies are more effective than proactive strategies in discouraging hackers' malicious activities.

I Introduction

The past few years have witnessed the remarkable increase of cyber attacks [2]. As the Internet technology develops rapidly and facilitates world-wide connections, cyber attacks are occurring at a higher rate and are becoming a significant problem for the society. For example, as identified in the study of Kanich et al. [17], businesses using spam campaigns and botnets have been flourishing. However, spam and botnets are only a small fraction of criminal businesses related with information security: there are a wide variety of business models including identity theft and stealing credit card and bank account information. The prevalence of this phenomenon led many to consider at least some kind of means should be employed to redress malicious cyber activities. Government agencies, international organizations and security vendors have therefore made a concerted effort to prevent cyber attacks by developing and introducing several policies and strategies. For example, several countries have enacted laws related to cyber-security (e.g., U.S. security breach notification laws and data protection laws) and have made cyber-security related agreements and guidelines (e.g. Seoul-Melbourne Anti-Spam Agreement and OECD Security Guidelines). In addition, various extra mechanisms that intend to reduce these wrongdoings have also proposed by several researchers [28, 29]. While a range of policy tools and strategies continue to be developed to deal with this issue, most of them tend to be adopted without ascertaining the effectiveness. Moreover, few countermeasures are currently addressing the ever increasing issue of cybercrime markets [17, 24].

In this study, we analyze various possible policies and strategies which may be able to mitigate cyber perpetrators' malicious activities; we focus mainly on exploit markets in which tools, exploits and means to automatize cyber attacks are traded. Specifically, we use a scenario which features two players: a hacker, who needs to choose between legal activities (i.e., selling exploits to legitimate security vendors) and illegal activities (i.e. writing and selling an exploit kit in a black market), and a defender (i.e., a software vendor and a policy-maker) who needs to develop policies to mitigate hackers' illegal activities. In the analysis, we use a simple game theoretic model. We believe an exploit market is an appropriate target for the application of game theory, since it can assist in increasing our understanding of the effects of implemented security strategies on the decision making process of a hacker. The proposed game is a deterrence game, in which defenders make decisions based on what they know about exploit markets in hope of deterring hackers from producing, spreading and selling exploits.

In this article we develop a preliminary study to measure how changes in key security strategies of a defender can affect the decision making process of potential or practicing hackers. More specifically, this article's purpose is to:

1. Lay the basis for future work to use noncooperative game theory as a possible approach for analysis of hacker behaviour; we aim at explaining why illegal hacking behaviour is preferred to lawfully conforming behaviour.

2. Study how hacking technologies affect and are affected by the game.

3. Investigate possibly effective strategies and policies to be enforced by government agencies and security vendors to deter hacker malicious activities.

Our results show that, interestingly, only hackers with an average skill are prone to participate in malicious cyber-activities; on the

*This work is partly supported by the projects EU-IST-NOE-NESSOS and EU-SEC-CP-SECONOMICS.

other hand, highly skilled hackers are more likely to engage in legitimate activities and disregard criminal ones. We also show that, of an array of potentially effective strategic alternatives, directly reducing the returns from malicious activities is the only effective strategy for hackers both with a low - medium skill or a high skill. Furthermore, our results confirm that security defenders should put more effort into reactive strategies than proactive strategies to mitigate hackers' malicious activities [2].

However, we should note that this study is only a first step toward a more complete modeling of cyber-perpetrators' actions and incentives for a variety of decision-making situations. The results presented in this paper are not to be intended as definitive, but rather as a starting point for more complete and articulated models for cybercrime. Nevertheless, we think our work provides interesting insights into the cyber-security environment, including interesting observations on which defensive actions are effective against strategic cyber-attackers. We also expect more empirical work to arise, hopefully, from our present discussion.

The remainder of the article is organized as follows: the next Section reviews the previous literature. Section **III** develops a game-theoretic model and Section **IV** presents the results of our simulations. Lastly, discussion and limitations of this article are presented in Section **V**.

II Literature Review

In recent years, there has been increased discussion about cyber-attacks and cyber-defenses. While plenty of research has studied the harmful effects of various types of cyber attacks on organizations, industries and society, surprisingly little research has focused on the effectiveness of policy tools and strategies for coping with malicious cyber-activities. Particularly, many studies have recognized and have addressed the harmful effects of cyber-perpetrators' wrongdoings. Few, however, studied policies and strategies that can mitigate cyber-perpetrators' malicious activities. Accordingly, a growing number of strategies and policies related to cyber-crime have been employed in recent years, without enough consideration of the effects of these on cyber-attackers. Furthermore, most of the studies that suggested measures for preventing security incidents have been concerned about potential victims' prevention activities rather than investigating solutions to mitigate cyber-perpetrators' criminal activities. In this section, we first discuss cyber black market economics that initially motivated this study, then explore studies related to the redress of malicious cyber-activities.

A Cyber Black Market Economics

A first analysis of black market economics was addressed in [9] by Franklin et al. They analyzed the amount of credit card numbers, banking information, and Social Serial Numbers (SSNs) circulating in Internet Relay Chat (IRC)[1] markets for a period of 7 months. According to their estimations the market is worth, overall, about 100 Million USD. Moreover, they show that about 5 percent of the logged data concerns trading of compromised hosts.

However, Herley et al. are skeptical about the reliability of these results [14]. They show that IRC markets feature all the characteristics of a typical "market for lemons" [1]: the vendor has no drawbacks in scamming the buyer because of the absence of a unique-ID and of a reputation system. Moreover, the buyer cannot in any way assess the quality of the good (i.e. the validity of the credit card and the amount of credit available) beforehand. On a folkloristic note, indeed, IRC markets are well known, in the underground community, to be markets for "newbie" and wanna-be scammers [14]. There are underground markets other than IRC ones; Savage et al. [24] analysed the private messages exchanged in six underground forums. Most interestingly, their analysis shows that these markets feature the characteristics typical of a regular market: sellers do re-use the same ID, the transactions are moderated, and reputation systems are in place and seem to work properly.

Dealing with criminals and illegal underground activities can be not only difficult and prone to error, but interpretation of experimental results can also be tricky and sometimes misleading [14, 16]. Moreover, Anderson et al. in [2] showed that, when it comes to new crimes perpetrated through and thanks to the Internet, the investment to defend against them surpasses the gains for the attacker of one order of magnitude: traditional technical countermeasures and strict business-internal policies proved to be extremely expensive and unfruitful. This suggests that more efficient and practical policies and "reactive" practices should be considered when dealing with cybercrime (e.g. increasing the cost of attacks by putting the bad guys in jail).

In regards with these new forms of cybercrime, we are mainly interested in Exploit Kits: these are tools that, once deployed, attack the victim systems that try to connect to them. They are widely used by cybercriminals to, for example, build botnets. These attack techniques are very well explored in a foundational study from Provos et al. [25]: the attacker usually compromises a popular website, and puts an i-frame in it[2]. Once loaded, the i-frame points the victim's browser toward a second domain under the control of the attacker, whose tool can now attack the victim. These dynamics are enforced in online advertisement and porn networks as well [31]. Moreover, successful business models are arising from the trade of users' "network traffic": traffic brokers [25] re-sell users' connections to their client, that will eventually provide content to the final user (advertisement or, in our case, cyber attacks).

The economic returns for the attacker have been studied in literature as well. Kanich et al. analyzed the return on investment for three spam campaigns [17] launched by the Storm botnet, and show that the conversion rate (i.e. number of times the victim "clicks" on the spammed link and goes through the trade process to buy the product) are extremely low. This low success rate is taken into consideration by Herley in [13] as well; he observes that attackers pay the cost of "false positives" as well (e.g. users that are accounted as victims but are not). As a result, he shows that the cost for the

[1] IRC used to be a very popular channel for quasi-anonymous instantaneous interactions between users.

[2] An i-frame is an HTML page within-the-page that can feature no printed content and can therefore be invisible to the user

attacker steadily increases as the density of "vulnerable" users decreases. Therefore, to economize the attack process, the attacker needs to choose carefully the population of victims she is going to attack. With spam campaign, it is very hard to understand which user is a true positive and which one a true negative. However, in the case of cyber attacks, criminals have at disposal a number of technological resources to assess, rather accurately, which user they should attack and which not [25]. This, alongside with cryptographic techniques and tool differentiation, allows the attackers to minimize the number of false positives, dramatically decreasing their cost to attack. For example, less unsuccessful attacks (false positives) mean less visibility, which means that the attackers can minimize the chance of having the police knocking on their doors.

B Redress of Malicious Cyber Activities

Since Becker's seminal study [6], a vast literature has been published dealing with strategies for coping with individual criminal behavior. The objective in the early era of research on criminal behavior was to increase the understanding of criminal behavior and develop corresponding effective countermeasures, primarily from the classical expected utility theory. While the literature focused mostly on analyzing a general model of criminal behavior, Cornish & Clarke [8] started to study a crime-specific model. They argue that people's choice to participate in criminal activities might be very different according to what specific goal and act are taken into account.

More recently, as the field of information security emerges, many studies have started to apply the previous models and findings to malicious behaviors in cyber-space. Of these studies, the most referred policies for mitigating illegal activities in cyber-space were the legal system. According to Lipton [20], in spite of several deficiencies, relying on criminal laws which intend to penalize cyber-perpetrators could be the most effective way to deal with many malicious activities in cyber-space. He, however, pointed out that in order to be effective for redressing malicious activities, criminal laws that deal particularly with malicious activities should clearly state what constitutes cyber-crimes and avoid relying on an approach from a pre-Internet era.

Recent literature suggests several additional mechanisms that could prevent cyber-perpetrators' wrongdoings. Some researchers argue that developing morality or the intrinsic motivation to do the "right thing" would be useful to reduce malicious activities. The question however is how to foster morality in online activities. According to Lipton [20] and Broadhurst [7], education or training can be used to develop morality since it could lead users to behave in a socially acceptable manner by creating an internal sense of guilt and decreasing moral satisfaction. Several studies have also addressed effects of policy tools and mechanisms which are developed to promote positive cyber-activities. According to these studies, the most commonly suggested tools are to increase rewards and reputation from positive cyber-activities. For example, researchers including Hennig-Thurau & Walsh [11], Kwok & Gao [19], and Liu et al. [21] argued that monetary and economic rewards are one of the most important mechanisms that promote users' well behavior. Wang et al. [30] further stated that users are likely to decide their online activities based on rewards that are offered by the system. They therefore concluded that the existence of the reward system which allows users to converts their activities into monetary rewards might increase their positive cyber-conduct.

III Game Theoretic Model for a Hacker's Behavior

Alongside with the literature review proposed in Section **II**, we base our model on our direct observation of the black markets. With the purpose of getting a more detailed and precise idea of how *blackhat* trades and tools work, we monitored the activities of many black markets for more than 4 months. In this work, in particular, we are interested in one of the kinds of tools traded in these markets: Exploit Kits.These tools are usually licensed over a one-year period; prices may vary in between 1,500 USD and 2,500 USD per year. In our model, cyber-attackers act as utility maximizers evaluating various factors including penalties and rewards in perpetrating cyber-crimes. In particular we consider a utility function that allows cyber-offenders to allocate their time to illegal cyber-activities while considering potential benefits and costs resulting from their wrongdoings.

A The Basic Model

We consider two types of players in the study: a hacker who can sell an exploit kit which includes various vulnerabilities, or can sell the vulnerabilities to legitimate vendors (e.g., Google's bug bounty program, tipping point initiative or exposing them in a black-hat conference to be hired as a penetration tester) and a defender (e.g., a policy-maker or a security vendor). We regard a hacker as a single decision making entity no matter who is an individual hacker or a hacking group and, throughout, we use he for a hacker. He faces uncertain situations and needs to make a choice from a set of available actions. Each of these actions has a different probability of yielding an outcome. A hacker decides which action he will make based on his belief in the utility. We assume that a hacker will choose the action that is likely to produce the highest utility. Actual outcomes are then assumed to be the result of the interplay between the decisions made by a hacker and a defender.

In the game theory, a game can take either a cooperative form or a noncooperative form. Since exploit-kit markets consists of players with competing and conflicting interests, however, this study models a game as a noncooperative form, and hence assumes that the players make an effort to maximize individual payoffs. In order to investigate the game between a hacker and a defender, we adopt and extend the framework of traditional game theoretic models used in the studies of Mesquita & Cohen [22] and Krebs et al. [18]. According to them, the game theoretic model allows us to determine a player's possible action based on his preferences for particular outcomes and his willingness to take risks to maximize the utility. The game we propose here posits that a hacker's poten-

Activity type	Variable	Meaning
General	T	hacker's total time
	t	time for detection and neutralization of criminal activity
	p	probability of obtaining maximum benefit from legal activities
	1-p	probability of obtaining only minimum benefit from legal activities
	q	probability of detection of the criminal activity
	q-1	probability of non-detection of the criminal activity
Legal	L	fraction of time the hacker devotes to legal activities
	B	maximum benefit gained from a legal activity
	S	minimum benefit gained from a legal activity
Criminal	I	fraction of time the hacker devotes to criminal activities
	Z	maximum benefit gained from a criminal activity
	C	cost for the hacker in perpetrating criminal activities

Table 1: Map of variables and their meaning in the model

tial to make efforts to develop and market exploits kits is a function of the expected payoffs from the exploit kits and the opportunity cost from committing these malicious activities. In contrast, defenders try to formulate strategies which make hackers' malicious activities less likely. They make choices based on what they know about hackers and exploit kit markets to deter hackers from producing, spreading and selling their exploit kits. There might be two types of strategies which can be employed by defenders: incentive and deterrence strategies. Deterrence strategies can be defined as the strategies which target reducing hackers' malicious activities by decreasing the returns on exploit kit related investments, and/or increasing the possibility of detection of exploits and conviction of hackers. On the other hand, incentive strategies mean the strategies which directly aim at encouraging hackers to involve more in legitimate activities by increasing the returns to these activities.

We follow a discussion on the variables considered in the model. Table 1 reports a sum-up of the variables and their respective meaning. First, we consider a hacker. He has total time, T, and is assumed to participate in only two activities, defined as malicious activities such as producing and selling exploit kits in black markets, which are harmful to sound cyber environment, and normal activities including the development of legitimate software and selling vulnerabilities to legitimate vendors, that are socially acceptable or even some-times increase social benefits. Therefore, he chooses the optimal allocation of his total time spent online between these two activities at the beginning of a given period.[3] We denote a fraction of a hacker's total time devoted to normal activities as L and a fraction of his total time spent on malicious activities as I (i.e., $L = T - I$).

We now consider how a hacker's expected utility is constituted. We assume that, if a hacker chooses to spend his total time only on legitimate activities, his abilities from these activities can potentially provide him with an opportunity to realize maximum benefits equal to B. We denote p as the hacker's probability on the achievement of these benefits. In contrast, with probability $1 - p$, the hacker can achieve only minimum benefits, S which is smaller than B (i.e., $B > S$). It should be noted that benefits from legitimate activities are not limited to monetary values (e.g., pecuniary income) and can include nonmonetary rewards such as the improvement of self-esteem and self-confidence. As a result, B and S can be increased not only by increasing monetary rewards from legitimate activities as suggested by Hennig-Thurau & Walsh [11], Kwok & Gao [19] and Liu et al. [21], but also by fostering morality or the intrinsic motivation to act legitimately as proposed by Liption [20], Broadhurst [7]. The levels of p and $1 - p$ are often considered to be influenced by the hacker's personal characteristics including education level and previous job experience. The hacker's expected utility from legitimate activities, therefore, can be expressed as

$$EU_N = L(pB + (1 - p)S) \tag{1}$$

where $L = T$.

We now take into account the case where a hacker chooses to participate in malicious activities (i.e., writing an exploit kit and selling it in black markets). We denote q as the probability of an exploit kit developed by the hacker being detected and disabled by defenders. The returns to the malicious activities are determined by the benefits gained from the exploit kit, Z, the timing of the detection and disablement of the exploit kit, t ($0 \leq t \leq 1$), and the costs to the hacker, C. Similarly with the benefits from legitimate activities, Z is an important factor that determines a hacker's behavior as explained by Wang et al. [30]. The costs to the hacker, C, is caused by the detection and disablement of the exploit kit, including the loss of reputation and the penalty from criminal laws considered by Lipton [20]. Three things should be noted: first, benefits and costs are not restricted to monetary payoffs and losses. These can also take the form of psychological rewards (e.g., self-esteem or self-confidence) and disappointment (e.g., a sense of sinfulness or guilt). Second, unlike the previous criminology research, since it is extremely difficult, if not impossible, to arrest a malicious hacker who develop an exploit kit [15], we assume that the hacker can still have the returns from his legitimate activities even after an exploit kit developed by him is detected and disabled by defenders. Lastly,

[3] We also assume that there is no cost for the movement between the activities.

unlike the previous literature, we include the time of the detection and disablement, t, in the model since the time has a high impact on a hacker's final payoffs. As a result, we define the returns from an exploit kit being detected as $(T-L)(Zt-C)+L(pB+(1-p)S)$. On the other hand, the probability of a hacker's exploit kit not being detected by defenders can be expressed as $(1-q)$. In this case, the returns are equal to $(T-L)Z+L(pB+(1-p)S)$. Putting it all together, a hacker's expected utility of committing malicious activities in line with the ideas of the time allocation can be denoted as

$$\begin{aligned} EU_M = q[(T-L)(Zt-C)+L(pB+(1-p)S)] \\ +(1-q)[(T-L)Z+L(pB+(1-p)S)]. \end{aligned} \quad (2)$$

As a result, if a hacker puts all of his time on malicious activities, the expected utility becomes $T(q(Zt-C)+(1-q)Z)$. From these expected utility functions, we can use a game theoretic model to investigate a hacker's decision making process.

Figure 1 depicts the course of a game that contains the possible choices of the players. In the game, a defender (e.g., policymaker and security vendor) moves first, so as to decide whether to enforce security policies and strategies against the activities related to exploit kits. A hacker then should decide whether he will involve in normal activities or malicious activities. If the hacker chooses to participate in malicious activities, the defenders again have to decide whether or not to impose additional security policies and strategies to the hacker's behavior. To solve this game theoretic model, it is important to identify the equilibria of the game. These show us under which conditions a hacker is expected to choose his involvement between socially acceptable activities and malicious activities. Briefly speaking, a hacker determines whether malicious activities or socially acceptable activities will yield a greater expected utility. He evaluates the expected utility from malicious activities which is represented by an exploit kit business and the expected utility from normal activities based on his beliefs. If he believes $EU_N \geq EU_M$, then socially acceptable activities will be selected. Otherwise, a hacker will start allocate his time to malicious activities.

B A Hacker's Response to Parameter Shifts

In this subsection, we investigate various policies and strategies of defenders which may affect a hacker's decision on participating in exploit markets. Specifically, we examine the hacker's supply shift of malicious activities in response to changes in strategies. Following Mesquita & Cohen [22] and Krebs et al. [18], several variables which may be possible remedies for malicious activities are modeled to identify which variables might be effective for reducing such activities. Following their study, we manipulate six variables: the probability of achieving maximum benefits from legitimate activities (p), the probability of an exploit kit being detected and disabled by defenders (q), the minimum benefits from legitimate activities (S), the maximum benefits from legitimate activities (B), the costs associated with the detection and disablement of the attack tool (C), and the benefits achieved from the attack tool (Z). In addition to these variables, we also propose manipulating the timing of the detection and disablement (t). This is because defenders (e.g., security vendors) can affect the value of an exploit kit by providing their customers with patches which can disable the exploit kit, or can shorten the timing of the detection of the exploit kit by monitoring exploit markets. We do not however vary a fraction of a hacker's total time devoted to normal activities L, since this variable is not likely controllable by defenders. As a result, in this study, we examine the effects of changes in key variables p, q, S, C, B, Z and t by using a simple simulation technique.

For our simulation we adopt an approach already used in literature by a former and notable study performed by Krebs et al. [18]. In each simulation analysis, we normalize all the values of the variables to 1.00. We then fix all of the variables except for the value for the key variable being manipulated: other things being equal, the key variable whose effect is being simulated will increase from 0.05 to 1.00 by 0.05.

As pointed out by Krebs et al., while fixed values used in the previous studies might be appropriate for the purpose of each of them, some of the variables should be adjusted for the purpose of this study [18]. We fix the variables p and S at .5 and .3 respectively; these are the same values used in the studies of Mesquita & Cohen [22] and Krebs et al. [18]. We estimate the values of q, C, Z, B, t and L based on several months of explorations in the exploit markets.

The probability of an exploit kit being detected and disabled by defenders, q, may be very low. In general, delayed detection of breaches is very common in cyber-security; as an example Verizon's 2012 report on data breaches investigations [4] shows that the great majority of breaches are discovered months, if not years, after they happened. Moreover, cooperation between law forces is often difficult[4], and the rate at which an attacker can change the address of his exploit kit is way higher than its detection rate by lawful security researchers. As a result, we fix the value of q at 0.1, which is, however and most likely, an over-estimation.

The costs caused by the detection and disablement of an exploit kit, C, may also be low since arrest of a hacker is quite hard and the actual arrest rate is very low [5, 12, 15]. Recently, the U.S. Federal Bureau of Investigation and many other European agencies are increasing their effort in chasing cyber-criminals[5]. While cyber-criminals face very severe penalties when caught[6], it is certainly hard to prosecute and apprehend them since they usually stay outside the reach of law enforcement [28]. Given this situation, we fix the value of C at 0.2.

As for the benefits gained from the exploit kit, Z, and the maximum benefits from legal activities, B, we consider two cases: In one case, we fix the values of Z at 1.0 and B at 0.8 ($B > Z$). In the other case, we choose the values of Z at 0.8 and B at 1.0 ($Z > B$).

[4] http://nakedsecurity.sophos.com/2012/01/19/koobface-gang-servers-russia-police/, accessed July 05 2012

[5] http://nakedsecurity.sophos.com/koobface/, accessed July 05 2012

[6] http://www.darkreading.com/database-security/167901020/security/attabreaches/224200531/index.html, accessed July 05 2012

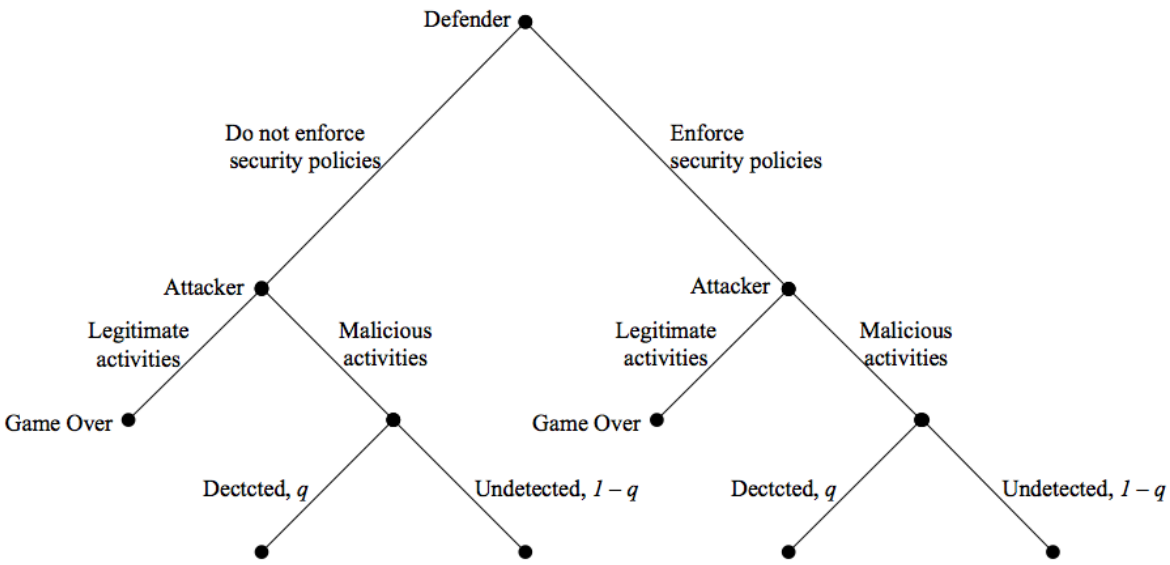

Figure 1: Defender-Attacker Game

This is to compare different types of hackers: a hacker valuing self-esteem and altruism vs. a hacker valuing sense of superiority and dominance. While indeed regular criminals often act out of need (e.g. they don't have a satisfying social status or they don't have a job), cyber-criminals are seemingly often well-educated and financially stable members of the society [15]. Hackers are indeed well-known to often act for fun or for reputation [27]. Being hackers' motivation not strictly related to their condition in the society, but rather an "emotional state", we feel that we should distinguish between the two cases in which the hacker is *a)* lawful-but-curious and *b)* criminally-minded.

In addition to these values, we also estimate the values for t and L which were not introduced in the previous studies. As previously mentioned, the detection rate of cyber-threats is traditionally rather low. Moreover, in the market we observed a number of Exploit Kits that feature $5+$ years old vulnerabilities at the time of release. The reason why these exploits are still effective and actively used[7] is because users do not patch[12] and therefore vulnerabilities stay exploitable far longer than supposed to. Moreover, exploit kits continuously change domain, therefore tracking them down and disabling them is a very hard if not impossible task [10]. We conclude that the average time for the neutralisation of an Exploit Kit is very high. We therefore set the timing of the detection and the disablement of an exploit kit, t, to .9.

As for a fraction of a hacker's total time devoted to normal activities, L, we fix the value of L at 0.9, meaning that the fraction of time they devote to the criminal activity is low (0.1). The reason of this is because most of the hackers have regular jobs, as inferred from reports on the profiles of cyber-crooks [15]; moreover, exploit kits do not require much time or effort to be managed, once their development is complete and the final product marketed. This coincides with our starting observation that cyber-criminal activities are typically not correlated with social needs: not only many hackers have a job as programmers, they are also often young and well-educated [15]. Of course, this is very hard to prove to be the typical case; however, due to the rather high sophistication of Exploit Kits and the non-trivial exploitation of software vulnerabilities [10], it is likely that the average hacker considered in this study is at least experienced with programming and/or has at least college-equivalent preparation in the topics of Software Engineering and Information Security. Otherwise, any technical implementation of the final product would be impossible for him/her to accomplish.

[7] `http://contagiodump.blogspot.it/2011/08/targeted-attacks-against-personal-gmail.html`, accessed July 05 2012

IV Results

We now discuss the results of the simulation tests. In the simulation we let the variables p, q, S, Z change from 0.05 to 1 with 0.05 steps. When a variable doesn't change, it is fixed to the value identified above. We ran simulations for both the cases when $(Z > B)$ - i.e. when the maximum benefit from criminal activities is higher than that from legal ones - and $(B > Z)$ -i.e. *vice-versa*. Unsurprisingly, we found that most of the strategies and policies for reducing malicious activities of a hacker *do not work* as intended by defenders *when the hacker values the benefits from exploit kit development and marketing more than the benefits from legitimate*

activities ($Z > B$).[8] However, it confirms that lowering the value of Z is the only effective strategy for hindering hackers participating in malicious activities. These results correspond to those from Mesquita's foundational study from 1995 [22].

The results for the second case ($B > Z$) are reported in Table 2. The first column indicates the changes of the key variable in increment of 0.05 ranging from 0.05 to 1.00. The columns of each simulation model show the results of the comparison between the expected utilities from normal activities and malicious activities (i.e., $EU_N - EU_M$). That is, these columns display whether the changes in the variable are likely to be effective for reducing malicious activities: **succeed** indicates that the key variable might be effective whereas a blank cell means that the changes in the key variable will not be effective. Note that the models with the changes in the values of C, B and t are eliminated from the table because all the changes in the variables are not effective for mitigating malicious activities.

Table 2 indicates that, in addition to the strategies for decreasing the value of Z, several other strategies that are not effective in the previous tests become effective for reducing malicious activities if a hacker values the benefits from lawful activities more than the benefits from malicious activities. In detail, Model 1 suggests that increasing the value of p will make normal activities more attractive than malicious activities: with an increase in the value of the probability of getting maximum benefits from normal activities to 70 percent, a hacker is likely to choose normal activities over malicious activities. In contrast, this also implies that only highly skilled hackers (i.e., hackers who have a probability of getting the maximum benefits from legitimate activities higher than 70 percent) are likely to devote their resources to legitimate activities. Model 2 confirms that the increase in the value of q can be an effective strategy for reducing malicious activities: While such a scenario is unlikely, a hacker would be prone to participate in legitimate activities rather than malicious activities if the probability of detection rate rises to 55 percent. Model 3 also suggests that increasing the minimum benefits from legitimate activities to 55 percent makes a hacker participate in legitimate activities. This implies that the gap between the minimum and maximum benefits from legitimate activities should be reduced in order to mitigate hackers' malicious activities. Lastly, Model 4 indicates that reducing the value of Z makes the incentive to work in exploit markets less attractive.

In sum, the simulation models suggest the followings: First, the only key variable which can be effective for hackers with either $Z > B$ or $B > Z$ is to reduce the value of Z. However, developing policies and strategies to reduce the value of Z might be difficult. For example, several researchers have argued that building legitimate markets that can substitute illegal exploit markets might be effective in giving hackers incentives to reduce malicious activities by decreasing the value of Z [3]; accordingly, many policy makers and security vendors have tried to build legitimate "markets for vulnerabilities". However these markets are not as well-activated and well-developed as originally intended [23]. Second, while shortening the timing of the detection and disablement of a security threat, which corresponds to t in our study, might be an effective tool for reducing malicious activities, it might do nothing to make hackers reduce their malicious activities. Third, it is identified that developing policies and strategies for hackers with $Z > B$ is more problematic than developing those for hackers with $B > Z$. That is, hackers who value the benefits from legitimate activities more than the benefits from malicious activities are likely to give up malicious activities by changing the values of p, q, S and Z; on the other hand, hackers who regard the benefits from malicious activities higher than the benefits from normal activities are still likely to participate in malicious activities even after the manipulation of the key variables except for Z. This result corresponds to the hackers' profiles reported in other articles and in the news [9][10] [15]: since they are relatively young, these traffic hackers are more likely to participate in malicious activities motivated by thrill-seeking, feelings of addiction, peer recognition, boredom with the educational system and lack of money [27, 26].

V Discussion and Future Research

Currently, most of the research on malware threats has been studied from a technical lens, and hence other domains such as economic and political perspectives have been largely ignored. Furthermore, the focus on the research is mostly on the targets of attacks rather than on strategies and policies that can mitigate criminal activities associated with malware. With this article we want to fill this gap in the literature by conducting a study on strategies and policies for reducing malicious cyber-activities from an economic perspective. The results of this study are therefore not to be intended as definitive: while many of our conclusions are, we believe, sound and promising for future research, more complete models are needed to design realistic and effective mitigation strategies. However, some key insights identified in this work could be interesting pointers for future work. Specifically, our results show that:

1. Only *very good programmers and professionals* who have high probability of getting maximum payoffs from legitimate activities *are not prone to engage in criminal activities*. Indeed, only when one's likelihood of getting maximum benefits from lawful activities exceeds 70% we can expect the actor not to act maliciously. This implies that it is not only true that one does not have to be a very good programmer in order to be a malicious hacker, but also true that a *very good programmer is not likely to be a malicious hacker*.

2. Good policies that can increase the likelihood of achieving maximum returns from lawful activities would prevent the very good professionals from going rogue.

[8] The table of the results is not presented here, but is available for the interested reader upon request.

[9] http://www.informationweek.com/security/management/amazoncom-ddos-attacker-busted-in-cyprus/240004073

[10] http://nakedsecurity.sophos.com/koobface/

Changes in key variable	Model 1:	Model 2:	Model 3:	Model 4:
	p changes	q changes	S changes	Z changes
0.05				**Succeed**
0.1				**Succeed**
0.15				**Succeed**
0.2				**Succeed**
0.25				**Succeed**
0.3				**Succeed**
0.35				**Succeed**
0.4				**Succeed**
0.45				**Succeed**
0.5				**Succeed**
0.55		**Succeed**	**Succeed**	**Succeed**
0.6		**Succeed**	**Succeed**	**Succeed**
0.65		**Succeed**	**Succeed**	**Succeed**
0.7	**Succeed**	**Succeed**	**Succeed**	
0.75	**Succeed**	**Succeed**	**Succeed**	
0.8	**Succeed**	**Succeed**	**Succeed**	
0.85	**Succeed**	**Succeed**	**Succeed**	
0.9	**Succeed**	**Succeed**	**Succeed**	
0.95	**Succeed**	**Succeed**	**Succeed**	
1	**Succeed**	**Succeed**	**Succeed**	

Table 2: Simulation Results when $B > Z$. Z is fixed at 0.8 and B is fixed at 1.0.

3. Policies could also be tuned to assure that only low-scale professionals are willing to "join the dark side". Accordingly, this would *decrease the quality of the attack tools* traded in black markets, and possibly their effectiveness in infecting machines and, for example, building botnets.

4. Another possible strategy is to increase the minimum benefits for a hacker ("S" in our model). This would encourage even "average skilled" hackers in joining legal activities rather than criminal ones.

Moreover, despite resulting from a completely different approach, our conclusions are in accordance with those of a recent study from Anderson et al. [2]: "response policies" is where policy makers should put more effort into: Increasing detection rates is an effective strategy to deter cyber-criminals from going rogue. We are, however, very far from achieving that goal: our model predicts a detection rate higher than 50% to be effective; in the current state of cyber-security, this is far from being accomplished. A more plausible strategy is to cleverly increase the minimum benefit for legitimate activities (S) in cooperation with higher detection rates (q): this may turn out to be an effective strategy in real-world scenarios.

In spite of the interesting findings, this study has some limitations that might offer additional avenues for future study and are important to underline here. First, one should recognize that, even with a well-crafted strategy for coping with malicious activities, its implementation might be problematic and therefore unrealistic. For example, an exploit provider may not be inside the jurisdiction where cyber-crime is committed [28]. Because the Internet can be accessed by anyone throughout the world, it might be very difficult, if not impossible, to apply strategies that are made for a specific country to other countries or to prosecute a foreign cyber-perpetrator. As a result, while this study can help in pointing policy-makers and security vendors toward theoretically supported strategies, it is clear that further investigation and additional empirical studies in the field are required. Moreover, the results of our model may change because of complementary or substitution effects between the key variables. All of these issues are very interesting points to be address in future work: we believe that the model presented in this paper can be a good candidate as a starting point for upcoming research in the field.

References

[1] G. A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84:pp. 488–500, 1970. A

[2] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. van Eeten, M. Levi, T. Moore, and S. Savage. Measuring the cost of cybercrime. In *Proceedings of the 11th Workshop on Economics and Information Security*, 2012. **I**, **I**, A, **V**

[3] R. Anderson and T. Moore. The economics of information security. *Science*, 314:610, 2006. **IV**

[4] W. Baker, M. Howard, A. Hutton, and C. D. Hylender. 2012 data breach investigation report. Technical report, Verizon, 2012. B

[5] J. Baltazar. More traffic, more money: Koobface draws more blood. Technical report, TrendLabs, 2011. B

[6] G. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 78:169–217, 1968. B

[7] R. Broadhurst. Developments in the global law enforcement of cyber-crime. *Policing: An International Journal of Police Strategies & Management*, 29:408–433, 2006. B, A

[8] D. B. Cornish and R. V. Clarke. Understanding crime displacement: An application of rational choice theory. *Criminology*, 25(4):933–948, 1987. B

[9] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 375–388, 2007. A

[10] C. Grier, L. Ballard, J. Caballero, N. Chachra, C. J. Dietrich, K. Levchenko, P. Mavrommatis, D. McCoy, A. Nappa, A. Pitsillidis, N. Provos, M. Z. Rafique, M. A. Rajab, C. Rossow, K. Thomas, V. Paxson, S. Savage, and G. M. Voelker. Manufacturing compromise: the emergence of exploit-as-a-service. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*, pages 821–832. ACM, 2012. B

[11] T. Hennig-Thurau and G. Walsh. Electronic word-of-mouth: Motives for and consequences of reading customer articulations on the internet. *International Journal of Electronic Commerce*, 8:51–74, 2003. B, A

[12] C. Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*, NSPW '09, pages 133–144. ACM, 2009. B

[13] C. Herley. Why do nigerian scammers say they are from nigeria? In *Proceedings of the 11th Workshop on Economics and Information Security*, 2012. A

[14] C. Herley and D. Florencio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. *Economics of Information Security and Privacy*, 2010. A

[15] G. IB. State and trends of the russian digital crime market. Technical report, Group IB, 2011. A, B, **IV**

[16] C. Kanich, N. Chachra, D. McCoy, C. Grier, D. Y. Wang, M. Motoyama, K. Levchenko, S. Savage, and G. M. Voelker. No plan survives contact: experience with cybercrime measurement. In *Proceedings of the 4th conference on Cyber security experimentation and test*, 2011. A

[17] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamalytics: an empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM Conference on Computer and Communications Security*, CCS '08, pages 3–14. ACM, 2008. **I**, **I**, A

[18] C. P. Krebs, M. Costelloe, and D. Jenks. Drug control policy and smuggling innovation: a game-theoretic analysis. *Journal of Drug Issues*, 33(1):133–160, 2003. A, B

[19] J. S. Kwok and S. Gao. Knowledge sharing community in p2p network: a study of motivational perspective. *Journal of Knowledge Management*, 8:94–102, 2004. B, A

[20] J. Lipton. What blogging might teach about cybernorms. *Akron Intell. Prop. J.*, 4:239, 2010. B, A, A

[21] S. Liu, H. Liao, and Y. Zeng. Why people blog: an expectancy theory analysis. *Issues in Information Systems*, 8(2):232–237, 2007. B, A

[22] B. B. D. Mesquita and L. E. Cohen. Self-interest, equity, and crime controls: A game-tehoretic analysis of criminal decision making. *Criminology*, 33(4):483–518, 1995. A, B, **IV**

[23] C. Miller. The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *Proceedings of the 6th Workshop on Economics and Information Security*, 2007. **IV**

[24] M. Motoyama, D. McCoy, S. Savage, and G. M. Voelker. An analysis of underground forums. In *Proceedings of the ACM 2011 Internet Measurement Conference*, 2011. **I**, A

[25] N. Provos, P. Mavrommatis, M. A. Rajab, and F. Monrose. All your iframes point to us. In *Proceedings of the 17th USENIX Security Symposium*, pages 1–15, 2008. A

[26] P. Taylor. *Hackers: crime in the digital sublime*. Psychology Press, 1999. **IV**

[27] O. Turgeman-Goldschmidt. Hackers' accounts: Hacking as a social entertainment. *Social Science Computer Review*, 23(1):8, 2005. B, **IV**

[28] M. Van Eeten and J. Bauer. Economics of malware: Security decisions, incentives and externalities. Technical report, OECD, 2008. **I**, B, **V**

[29] M. Van Eeten, J. Bauer, H. Asghari, S. Tabatabaie, and D. Rand. The role of internet service providers in botnet mitigation: An empirical analysis based on spam data. Technical report, OECD STI Working Paper, 2010. **I**

[30] X. Wang, H. Teo, and K. Wei. What mobilizes information contribution to electronic word-of-mouth system? explanations from a dual-process goal pursuit model. In *Workshop Association for Informational Systems, Oklahoma*, 2009. B, A

[31] G. Wondracek, T. Holz, C. Platzer, E. Kirda, and C. Kruegel. Is the internet for porn? an insight into the online adult industry. In *Proceedings of the 9th Workshop on Economics and Information Security*, 2010. A