The Underground Economy of Fake Antivirus Software

Brett Stone-Gross, Ryan Abman, Richard A. Kemmerer, Christopher Kruegel, Douglas G. Steigerwald, and Giovanni Vigna

Abstract Fake antivirus (AV) programs have been utilized to defraud millions of computer users into paying as much as one hundred dollars for a phony software license. As a result, fake AV software has evolved into one of the most lucrative criminal operations on the Internet. In this paper, we examine the operations of three large-scale fake AV businesses, lasting from three months to more than two years. More precisely, we present the results of our analysis on a trove of data obtained from several backend servers that the cybercriminals used to drive their scam operations. Our investigations reveal that these three fake AV businesses had earned a combined revenue of more than $130 million dollars. A particular focus of our analysis is on the financial and economic aspects of the scam, which involves legitimate credit card networks as well as more dubious payment processors. In particular, we present an economic model that demonstrates that fake AV companies are actively monitoring the refunds (chargebacks) that customers demand from their credit card providers. When the number of chargebacks increases in a short interval, the fake AV companies react to customer complaints by granting more refunds. This lowers the rate of chargebacks and ensures that a fake AV company can stay in business for a longer period of time. However, this behavior also leads to unusual patterns in chargebacks, which can potentially be leveraged by vigilant payment processors and credit card companies to identify and ban fraudulent firms.

B. Stone-Gross (✉) • R.A. Kemmerer • C. Kruegel • G. Vigna
Department of Computer Science, University of California, Santa Barbara, CA, USA
e-mail: bstone@cs.ucsb.edu; kemm@cs.ucsb.edu; chris@cs.ucsb.edu; vigna@cs.ucsb.edu

R. Abman • D.G. Steigerwald
Department of Economics, University of California, Santa Barbara, CA, USA
e-mail: ryan@econ.ucsb.edu; doug@econ.ucsb.edu

B. Schneier (ed.), *Economics of Information Security and Privacy III*,
DOI 10.1007/978-1-4614-1981-5_4,

1 Introduction

Over the past few years, electronic crimes revolving around a class of malware known as *scareware* have become extremely lucrative ventures. The concept is simple; design a ploy through social engineering that exploits a computer user's fear of revealing sensitive information, losing important data, and/or causing irreversible hardware damage. The most common form of scareware is *fake antivirus* (AV) software, also known as "rogue security software." More specifically, a fake AV program impersonates an antivirus scanner and displays misleading or fraudulent alerts in an attempt to dupe a victim into purchasing a license for a commercial version that is capable of removing nonexistent security threats. Some fake AV programs may also lock down system functionality to prevent victims from accessing files or web sites or from creating new processes, such as Windows Explorer, Task Manager, and a Command Prompt under the false pretense that it is for the victim's own protection. In addition, we have observed fake AV software that contains hidden backdoor capabilities, enabling the program to be used for other malicious purposes, such as launching distributed denial-of-service (DDoS) attacks against adversaries.

Over the past year, we have been able to acquire backend servers for several multi-million dollar criminal operations selling fake AV products. These fake AV businesses are run out of Eastern Europe and utilize affiliate networks known as *partnerka* to distribute the rogue software [32]. These partnerka networks use various pseudonyms, and operate by recruiting affiliates to install their software on as many computers as possible. In exchange, the affiliates receive a commission for driving traffic to landing pages, malware installations (also known as *loads*), and fake AV sales. Moreover, some partnerka offer additional incentives to the most successful affiliates with prizes including expensive cars, computers, and cell phones [18].

Since we have access to the servers used by these criminal organizations, we are able to directly analyze the tools that are used to create the fake AV products, including programs that assist perpetrators in controlling the malware's behavior and brand names, as well as custom *packers* that obfuscate the malware to evade detection by legitimate antivirus products. Some fake AV groups even make use of third-party commercial services to track the detection rates by the most popular antivirus vendors (e.g., McAfee, Symantec, and Trend Micro) [19], and they tweak their obfuscation algorithms until a low detection rate is achieved. We also have access to the instruments that are used to direct traffic to fake AV web sites, the infrastructure that prolongs the longevity of the operations, and a very detailed view of the financial profits that fuel these illicit enterprises. Interestingly, the miscreants behind fake AV products even offer refunds to victims who are persistent, in order to reduce the amount of credit card chargebacks, which we will discuss in more detail later.

Although various aspects of fake AV software have been studied, there are many facets of these operations that are not well understood, including the modus operandi of the criminals, the amount of money involved, the victims who purchase the software, the affiliate networks that promote the campaigns, and the flow of money from the victims' credit cards, to the payment processors, to the bank accounts controlled by the criminals. In this paper, we attempt to fill this void by presenting the analysis of several criminal organizations that sell fake AV products. More specifically, we make the following contributions:

- We provide an in-depth analysis of fake AV operations and present detailed statistics based on the analysis of more than a dozen servers belonging to several criminal organizations. This is the most comprehensive, large-scale study of fake AV campaigns that highlights different aspects of their operations from the infection process, to the financial complexities of maintaining a fraudulent business.
- We examine how fake AV campaigns are managed and orchestrated, from the ringleaders' point of view. We discuss the software infrastructure that is utilized, the functionality it provides, and its role in the underground economy.
- We present an economic model that encapsulates financial patterns that are indicative of fake AV ventures. Our intent is to formalize the essential factors of these operations and to identify potential weaknesses that can be exploited to increase the criminals' functional and operational costs.

2 Technical Background

Before we present the financial logistics, we first discuss the methods that are utilized to infect machines with fake AV software and the infrastructure behind the process. In addition, we present details about three particular criminal operations running fake AV businesses. To protect ongoing law enforcement investigations, we refer to these three ventures as AV_1, AV_2, and AV_3. Note that we currently see ongoing activity (e.g., new malware samples, installations and online advertisements) from all three fake AV operations.

2.1 Infection Methods

There are three primary infection methods used by fake AV distributors to propagate their malware: social engineering, drive-by-download attacks, and botnets. In this section, we present how these strategies are used to infect as many computers as possible with fake AV malware.

One of the most popular infection methods uses social engineering techniques to convince a victim to voluntarily install the fake AV. To launch this attack, a malicious web page displays a window in the browser (e.g., via JavaScript or Adobe Flash) that pretends that the machine has been infected with malware. An example

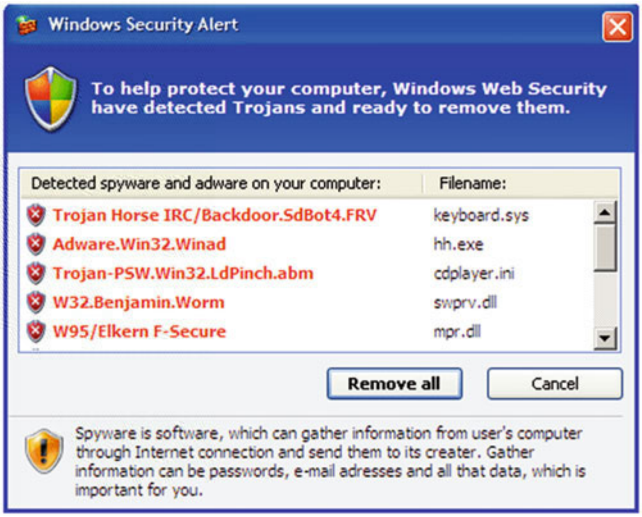

To help protect your computer, Windows Web Security have detected Trojans and ready to remove them.

Detected spyware and adware on your computer:	Filename:
Trojan Horse IRC/Backdoor.SdBot4.FRV	keyboard.sys
Adware.Win32.Winad	hh.exe
Trojan-PSW.Win32.LdPinch.abm	cdplayer.ini
W32.Benjamin.Worm	swprv.dll
W95/Elkern F-Secure	mpr.dll

Spyware is software, which can gather information from user's computer through Internet connection and send them to its creater. Gather information can be passwords, e-mail adresses and all that data, which is important for you.

Fig. 1 Example alert from a fake antivirus advertisement displayed in a user’s web browser

is shown in Fig. 1. To fix the security problem, the window also contains a link to a program that presumably helps to clean up the infection. Of course, this program is the fake AV software that attackers aim to install.

A second technique to install fake AV software is via drive-by download attacks. In a drive-by download attack, a web site is prepared with malicious scripts that exploit vulnerabilities in the web browser or one of its plugins. When the exploit is successful, the fake AV malware is installed automatically, without the user’s knowledge or consent.

Both in the case of fake alerts and drive-by downloads, the initial goal of the attacker is to drive as many web visitors to their malicious web pages (sometimes called landing pages) as possible. In order to achieve this objective, attackers often make use of *blackhat search engine optimization* (SEO). Their intention is to poison search engine results by creating landing pages that contain popular search phrases. Many of these campaigns target current events such as the death of a celebrity, natural disasters, and holidays. Blackhat SEO relies on the fact that when search engine crawlers index a web site they identify themselves through the HTTP `User-Agent` field (e.g., googlebot). Thus, a site under an attacker’s control can serve content that contains popular keywords that a search engine will use in the computation of the page rank. If the process is done correctly, the landing page is ranked high in the search engine’s results for these popular keywords.

When a user clicks on a search engine result that leads to a blackhat SEO landing page, the server analyzes the user’s web browser (via the `User-Agent` header), and the referring web site (through the HTTP `Referer` field). The tools that are used to manage these SEO campaigns are known in the underground economy as a *traffic direction system* (TDS). These TDSs can leverage the header information to distinguish between search engine bots and web browsers. In order to avoid detection, TDSs often take additional countermeasures such as resolving the visitor's IP address to a geographic location and recording the number of accesses. Once the TDS has verified the traffic, a user is redirected a number of times to a landing page. This landing page will then launch a social engineering or drive-by download attack, as described previously.

Note that most TDSs also define a *time-to-live* (TTL) value that specifies how long a particular redirection URL will remain active. Most TTL values are very short, which makes it more difficult for security researchers to track active campaigns.

An alternative approach to using blackhat SEO techniques for traffic generation is to exploit the distribution systems and ubiquity of online ad networks. An attacker may compromise a legitimate ad network, or sign up as an advertiser to display malicious advertisements disguised as free pornography, missing audio/video codecs, or virus scans that perform similar social engineering attacks to con visitors into installing their malware. Online ad networks are also frequently used in conjunction with drive-by-download attacks, known collectively as *malvertisements*, to covertly install the fake AV software (without user interaction or permission).

A third infection method is through *botnets*, a collection of compromised computers under the control of an attacker. Several large botnets, such as Koobface, Conficker, and Bredolab, have been known to distribute fake AV software to machines under their control, which is believed to be one of their top sources of revenue [17,27,38].

Once fake AV software has been installed on the victim's machine (either voluntarily through social engineering or involuntarily through a drive-by attack or botnet), intrusive nags will be shown continuously to the victim, warning of "malware infections" or "intrusion attempts" that pose a risk to the user's system. At this point, the fake AV software usually advertises itself as a free trial version with limited functionality (i.e., detection only). If a victim wants to remove the malware infections, they must upgrade to a commercial version by purchasing a license key. When a victim clicks the software's purchase button, they are taken to one of the fake AV company's web sites. After a victim enters their personal information and credit card, they are sent a license key (e.g., through email) that essentially deactivates the bogus malware alerts, providing the user with a sense that their purchase was valuable.

2.2 Infrastructure

Similar to any other legitimate online business, when a fake AV company's servers are down, they lose potential revenue streams. Therefore, there are a number of measures that these organizations take to ensure the availability of their infrastructure. The first strategy is to deploy an array of proxy servers that are publicly visible. The sole purpose of these proxies is to relay content to one or more backend servers

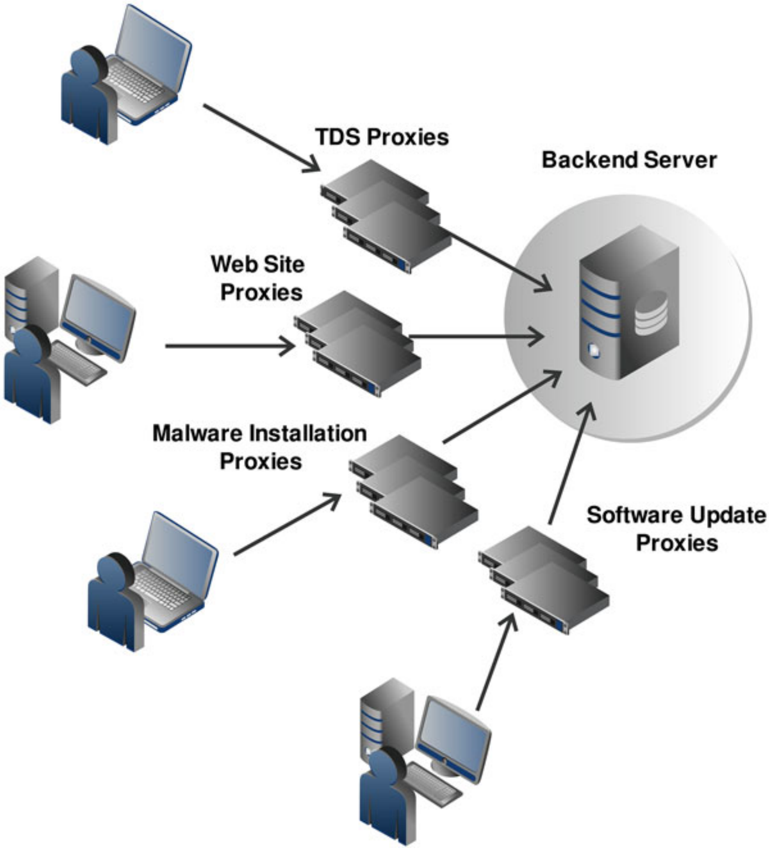

Fig. 2 Tiered infrastructure for many online criminal operations including fake antivirus businesses. We were able to obtain copies of three different fake AV organization's backend servers (in the *shaded circle* above) that control the entire operation

as shown in Fig. 2. More specifically, these machines communicate directly with users that are redirected to a landing page or infected hosts that purchase a license. The proxy servers are typically partitioned depending on the specific role that they fulfill (e.g., TDS servers are not reused for relaying sales information). The main purpose of the front-end servers is to thwart mitigation efforts. Hence, taking down one, or even several, of these machines often has little impact, since the domain name address records that point to these servers can be changed quickly and easily. These front-end servers are designed to be lightweight and expendable, and typically have an automated deployment program that accelerates the process of creating new proxy nodes.

The main drawback of proxies (from an attacker's point of view) is that when a defender obtains access to one of these front-end servers (or monitors their ingress and egress network traffic), she can learn the location of the backend infrastructure.

To address this problem and to further hide the location of the backend, the miscreants of fake AV operations may use multiple tiers of proxy servers. However, each extra tier will introduce additional network delay that could make a user who is purchasing a fake AV product more suspicious. In our experience, most fake AV operations use only one tier of proxy nodes. Thus, we were able to locate the backend infrastructure by tracking the network traffic from an infected host to a proxy node to the backend servers. By taking down the backend servers, the entire fake AV operation is disrupted (i.e., servers relaying sales, malware installations, and TDS become inoperable).

A second, important strategy is to register a large number of domain names. The domain names fulfill several purposes. First, it makes the fake AV web site look more legitimate (e.g., the domains are usually related to antivirus or security keywords). Second, the large number of domains makes takedown efforts more difficult, since the DNS records can be changed to point to any of their proxy servers. In addition, the reputation of a fake AV domain will decline as more people are defrauded, and many of the domains will become blacklisted. As a result, domain registrars may ultimately suspend some of the fake AV domains. Overall, the AV_1 crew purchased 276 domains, 17 front-end servers, and one back-end server. Similarly the AV_2 operation registered at least 188 domains, managed 16 front-end servers, and two back-end servers. We did not have complete visibility over the total number of domains used by AV_3, but from our observations, the infrastructure was similar to the others with a large number of free domains registered through the `co.cc` top-level domain (TLD), and approximately 20 front-end servers, and one back-end server.

3 Data Collection

In the following section, we describe the process that facilitated our efforts in obtaining access to these fake antivirus backend servers and the data we collected. The main tool that we utilized to analyze the fake AV malware was ANUBIS, a system that dynamically analyzes binary programs via runtime analysis [15]. ANUBIS runs a Windows executable and documents the program's behavior, including system modifications, processes creation, and network activity. ANUBIS is able to process on the order of tens of thousands of samples per day, providing us with a comprehensive view of the current malware landscape [1].

By searching through the network connections logged in the ANUBIS database, we were able to identify a number of unique network signatures commonly used by fake antivirus software. More specifically, when fake AV is installed, it often *phones home*, by connecting back to servers under the control of the fake AV criminal organization. For example, infected machines made an HTTP request similar to `GET/install.php?aff_id=151&p=34&s=7&ip=192.168.1.3&cn=US`, to notify the criminals of the installation and to credit the affiliate responsible for the infection. The parameters *p* and *s* provided details about the type and name of the malware.

After observing network signatures associated with these fake AVs, we contacted the hosting providers whose servers were being used for controlling these operations. We provided them with network traces, malware samples, and other evidence that revealed the location of the servers that were situated within their network. The hosting providers responded by taking these servers down, and they provided us with direct access to the information stored on them. Note that we had previously collaborated with a number of these vigilant ISPs in the U.S. and abroad through FIRE [34], our network reputation service that tracks where malicious content resides on the Internet.

In total, we were able to get a complete snapshot of 21 servers: 17 of which were proxy nodes, and 4 of which were backend servers. The information that we collected from these servers included data for AV_1 for approximately 3 months from January through April 2010, 16 months from January 2009 through May 2010 for AV_2, and from March 2008 through August 2010 for AV_3. From these data sources, we have a view of nearly the entire operation including web site source code, samples of the fake AV malware, and databases. The most interesting information is contained in the database records, which document everything from malware installations, fake AV sales, refunds, technical support conversations to the TDSs controlling the fake AV landing pages.

4 Following the Money Trail

Now that we have provided a summary of the fake AV infrastructure and our data sources, we will focus on the financial aspects that drive the sales of fake AV software. In particular, we analyze the flow of money from a victim to the criminals and their affiliates. In addition, we examine the ways in which the fake AV groups manage to stay under the radar when interacting with credit card payment processors.

4.1 Transaction Process

Before we present the detailed statistics of sales, revenue, chargebacks and refunds, we introduce an overview of the various entities involved in a fake antivirus business. The transaction process, as shown in Fig. 3, begins when a victim purchases the rogue AV software. This purchase is done through the fake AV company's web site (Step 1), where the victim enters her credit card information. The fake AV business (i.e., the merchant) then submits the credit card data to a third-party payment processor (Step 2). The payment processor forwards the information through one of the major credit card companies (Step 3), who requests authorization from the credit card issuer (Step 4). If the credit card issuer (i.e., a bank) approves the transaction, the victim's credit card is charged (Step 5), and the credit card company

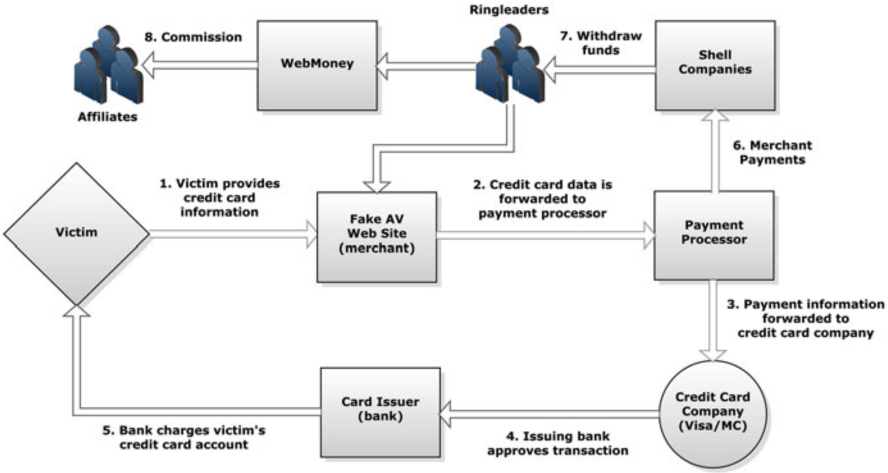
Ringleaders
8. Commission
WebMoney
7. Withdraw funds
Shell Companies
Affiliates
6. Merchant Payments
1. Victim provides credit card information
2. Credit card data is forwarded to payment processor
Victim
Fake AV Web Site (merchant)
Payment Processor
3. Payment information forwarded to credit card company
Card Issuer (bank)
Credit Card Company (Visa/MC)
5. Bank charges victim's credit card account
4. Issuing bank approves transaction

Fig. 3 High-level overview of the transaction process for fake antivirus businesses

notifies the payment processor of the successful sale. Periodically (e.g., biweekly or monthly), the payment processor deposits funds into bank accounts set up by the fake AV businesses (Step 6). The ringleaders of the fake AV operation then withdraw the funds (Step 7) and pay a commission to their affiliates (Step 8). We will provide more details about this process in the following sections.

4.2 Sales

There are a number of factors that contribute to whether a victim purchases a license, such as the aggressiveness of the fake AV software (e.g., frequency of alerts, type of threats, and whether system performance is affected). In addition, the price and subscription models offered by most fake antivirus products play an interesting role, with subscriptions that range from 6-month licenses to lifetime licenses. The AV_1 operation offered licenses for 6-months at $49.95, 1-year at $59.95, and 2-years at $69.95. These options were purchased almost uniformly with rates of 34.8%, 32.9%, and 32.3%, respectively. The AV_2 company's products also offered 6-month licenses at $49.95, 1-year at $69.95, and a lifetime license at $89.95. The 6-month option was the most popular (61.9%), followed by the lifetime license (24.6%) and the 1-year license (13.5%). The products sold by AV_3 were priced at $59.95 for a 1-year license and $79.95 for a lifetime license. All of AV_3's products were also bundled with a mandatory $19.95 fee for 24 x 7 customer support services, bringing the total price to $79.90 for the yearly license (purchased by 83.2% of victims) and $99.90 (purchased by 16.8% of the victims) for the lifetime license.

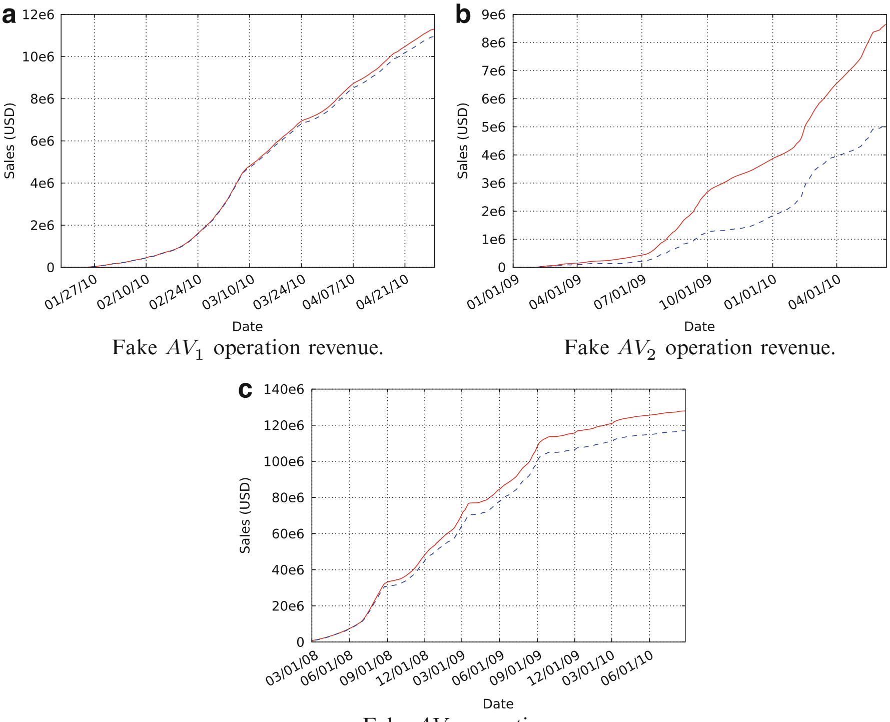

Fake AV_1 operation revenue.

Fake AV_2 operation revenue.

Fake AV_3 operation revenue.

Fig. 4 Three criminal organizations' revenue from fake antivirus sales. The *solid line* displays the total revenue, while the *dotted line* displays the revenue after chargebacks and refunds

In total, AV_1 "trial" products were installed 8,403,008 times, which resulted in 189,342 sales, or upgrades to the "commercial" version (a conversion rate of 2.4%) in only 3 months. Likewise, AV_2's programs were installed 6,624,508 times, with 137,219 victims that purchased the fake antivirus over 16 months. That is a conversion rate of approximately 2.1%. The AV_3 business sold 1,969,953 licenses out of 91,305,640 installations from March 2008 through August 2010 (a conversion rate of approximately 2.2%).

The total victim loss from the three fake AV operations was $11,303,494, $5,046,508, and $116,941,854 from AV_1, AV_2, and AV_3, respectively. Figure 4 shows the cumulative daily revenue for each of these fake antivirus operations. If we extrapolate these profits over one year, the AV_1 crew was on track to earn more than $45 million dollars per year, while the AV_2 group earned approximately $3.8 million per year. The largest and most profitable operation was AV_3, which raked in an average of $48.4 million dollars per year.

As we will discuss in Sect. 4.4, some credit card transactions were reported to be fraudulent and were credited back to the victim. Interestingly, victim complaints force these illegitimate firms into a complex position with their payment processors, as we will discuss in the following sections.

4.3 Payment Processors

An interesting facet of fake AV sales is the process in which credit card transactions are handled. In particular, payment processors (also known as payment service providers) are an integral part of every sale. Without these processors, fake AV operations would not be able to accept credit card payments. This would make it not only harder for a victim to purchase the product (i.e., they would have to use an alternative form of payment, such as cash, check, or money order), but it would also likely raise red flags that the software may be fraudulent. Note that payment processors must maintain a degree of legitimacy, or they risk losing the ability to accept major credit cards. For instance, a payment processor known as ePassporte lost the rights to accept Visa credit cards, due to a large amount of fraudulent transactions, money laundering, and other questionable activities [20]. Note that the AV_2 crew at one point set up an ePassporte merchant account for processing credit card transactions.

Perhaps the most notorious payment service provider is Chronopay, which is headquartered in the Netherlands and operated by Russian businessmen. Chronopay has long been associated with processing transactions for various forms of online criminal organizations [24]. However, Chronopay also provides legitimate services to large organizations such as Electronic Arts, Kaspersky, and charities including the World Wildlife Federation, Greenpeace, and UNICEF. Because the volume of legitimate transactions from these businesses may far outweigh the fraudulent activities, major credit card companies may be hesitant to sever ties with Chronopay. Note that all three fake AV businesses that we analyzed used Chronopay's credit card payment services.

There were several other, smaller payment processors that the fake AV operations used for credit card transactions. Interestingly, we found communications between one of these small payment processors and the fake AV perpetrators that revealed that the payment service provider was well aware of the fake AV business and even offered advice to help the group sell more products. There are a number of tricks that some of these dishonest payment service providers perform in order to benefit from fraudulent transactions. First, payment processors may offer *high-risk merchant accounts*, where the processor may earn close to 15% for each transaction. These are typically for questionable businesses that have significant problems with customer complaints (e.g., online pharmacies or pornography). Second, we observed that some of these payment processors allow an illicit company to create multiple merchant accounts in which transactions are periodically rotated (approximately every 30–45 days) through each account, such that a single account is never flagged for fraudulent activities, since the transactions are distributed over all of the accounts.

4.4 Chargebacks and Refunds

Interestingly, all three fake antivirus groups that we studied offered a certain number of refunds to individuals who requested them. At first, it may seem counter-intuitive for a criminal operation that is selling fraudulent products to provide refunds to victims. However, it is important to keep in mind that these criminal organizations have to use legitimate (or semi-legitimate) credit card payment processors for every transaction. In addition, payment processors are required by statutory (federal regulations) and contractual obligations (PCI) to provide various levels of consumer protection against theft and fraudulent purchases. When a victim reports a fraudulent transaction to their credit card issuer, they are issued a credit, which is known as a *chargeback*. If a business receives too many chargeback complaints, the payment processor may sever ties with the company and prohibit further credit card transactions. Therefore, it is important to minimize the number of chargebacks, which has the effect of extending the lifetime of the fake AV operation.

Overall, AV_1 granted 5,669 refunds (3% of sales) at a cost of $346,039 (in addition to 1,544 chargebacks worth $94,963). In comparison, AV_2 issued 11,681 refunds (or 8.5% of sales) at a cost of $759,666 (in addition to 3,024 chargebacks valued at $183,107). AV_3 refunded 151,553 (7.1% of sales) for a total of $10,951,191 (with 30,743 chargebacks valued at $2,225,430). Note that the primary credit card processor for AV_3 temporarily froze AV_3's merchant account for approximately one month in March 2009, due to a high number of chargebacks. After this incident, AV_3 offered more refunds, and the number of chargebacks dropped accordingly.

Another important factor that has an impact on chargebacks and refunds is how frequently a fake AV business changes the name of their product. This is due to the fact that after a short interval (typically 3–7 days), victim complaints start appearing on consumer web forums that are in turn indexed by search engines. Thus, a victim may perform a Google search for the name of the fake AV and find that other users have similar grievances and complaints. Interestingly, we found that AV_2 had significant server problems and maintained the same product names for an extended period of time. As a result, they had the highest chargeback and refund rates.

As we will discuss in Sect. 6, the amount and timing of refunds follows an interesting pattern, which indicates that the criminals maximize their profits by refunding just enough sales to remain under a payment processors chargeback limit.

4.5 Affiliate Programs

The financial incentives for cybercrime play an important role both in the type and amount of fraud. In order to infect as many machines as possible and therefore maximize sales, fake AV businesses rely upon affiliate networks based primarily in Eastern Europe known as *partnerka*. The backend servers that we obtained contained payment records to these partners. The profits for some of the affiliates are immense, with members earning as much as 30–80% commission from sales leads. Remarkably, the top affiliate of AV_1 made more than $1.8 million dollars in approximately two months. Over the course of these two months, there were a total of 44 affiliates who were paid (out of 140 that enrolled), with four earning more than $500,000, 11 in excess of $100,000, and 15 more than $50,000. The average affiliate income was approximately $60,000 per month. In comparison, AV_2 had 98 active affiliates out of 167 total registered, and stored records for 9 months of payments to these affiliates. Overall, five of these affiliates made more than $300,000, 16 earned more than $100,000, and 22 earned more than $50,000. The AV_3 operation had a total of 1,107 affiliates with 541 who were active. The top AV_3 affiliate earned $3.86 million, and three others made more than $1 million. There were 15 AV_3 affiliates that earned over $100,000, and 23 that were paid more than $50,000.

By comparing the affiliate email addresses across the three different fake AV partnerka, we were able to determine that 70 affiliate members were involved in multiple groups. Interestingly, there was one affiliate who was associated with all three fake AV businesses.

The affiliate payments were made through WebMoney, a virtual electronic currency. There are several advantages that WebMoney provides for criminal activities. In particular, all transactions are anonymous and irreversible. That is, once a transfer has occurred it cannot be voided, regardless of whether it was fraudulent. Other benefits include a very low transaction fee (0.8%), and a large number of places, especially in Eastern Europe, that will exchange WebMoney for local currencies.

4.6 Shell Companies

One of the most important parts of the financial system from a fake AV company's perspective is the ability to *cash out* earned funds. Thus, a fake AV company must open one or more bank accounts to receive merchant remittances from their payment processors. These accounts are typically set up and registered to fictitious *shell companies*. We observed accounts registered primarily in Europe and Asia, including the Czech Republic, Finland, Cypress, and Israel. Once money is deposited into a shell account, the ringleaders can directly withdraw the funds. However, criminals who are more cautious may opt to use the services of *money mules*. A money mule is a person who is recruited (usually under the pretense of a work from home job) to accept a bank deposit, withdraw the funds, and wire the money (minus a service fee) back to the criminals. This greatly minimizes the risk that a criminal will be apprehended when receiving funds. Unfortunately, we were not able to determine the precise method used by these three fake AV groups to withdraw funds. Nevertheless, we believe the money was probably picked up directly by the ringleaders (or one of their close associates), based on the geographic locations of the bank accounts.

5 Victims

In this section, we analyze the victims that purchased fake AV software. In particular, we will study various characteristics of victims including: geographic location, operating systems, and institutions. In addition, we will examine the technical support and customer service provided by the three fake AV businesses.

The largest concentration of victims (by far) was in the U.S. (76.9%) followed by the U.K., Canada, and Australia. This is likely due to the fact that the fake antivirus products are primarily written for English speakers (only a few of them had been translated to other languages). The most popular, compromised operating systems were Windows XP (54.2%), Windows Vista (30.8%), and Windows 7 (14.8%). Internet Explorer 7 was the most commonly used browser (65.6%). The most frequently used email addresses of customers of fake AV products were Yahoo, Hotmail, AOL, Gmail, and Comcast. Other residential ISPs placed in the top 10 including AT&T, SBC Global, Verizon, and Bellsouth. This indicates that most victims probably purchased the fake AV software for their personal computers at home. However, there were a number of sales from victims at commercial, government, and military institutions.

All three of the fake AV companies offered various forms of customer service and technical support. Customer service for fraudulent products may seem contradictory, but its purpose is clear: to reduce the number of refunds and victim complaints. Overall, the fake AV groups offered two types of support systems. The first was an online system where victims could open tickets describing their problems, and technical support representatives would periodically reply to these tickets. The second type of support system was an interactive, live chat service, where a victim would talk in real-time with technical support personnel.

We were able to observe the communications in many of these support systems, and analyze how operators responded to questions, and how they handled irate customers. For the most part, victims were upset, realized that the fake AV software was a scam, and requested instructions for removing the malware from their system. The fake AV representatives typically responded with removal directions, but they warned users that their computer was still infected and made claims that competitors (i.e., legitimate antivirus vendors) were slandering their products.

We also performed automated data mining techniques to determine the relationship between complaints, sales, chargebacks, and refunds. To this end, we queried the fake AV groups' internal databases for patterns such as credit card numbers, unique identifiers (e.g., orders), email addresses, and various keywords (e.g., *fraud, scam, refund*, etc) that were relevant to disgruntled customer reactions. By correlating these database records, we examined whether a victim who purchased a fake AV product later filed a complaint through any of the support forums, and if a refund or chargeback was issued. Overall, only a small percentage (less than 10%) of victims actually sought refunds, and those who were issued refunds received their credit within 7 days on average. Note that the low rates of victim complaints that we discovered are similar to those reported by the computer security news investigation web site, KrebsOnSecurity [21].

6 Economic Model

In this section, we utilize the data that we have collected to identify behavior that is representative of a fake AV business. We then propose an economic model based on a key observation of refunds that may be used to detect other businesses that are engaged in illegal activities.

6.1 Refund Patterns

Fake antivirus software firms (hereafter, firms) act to maximize profits. To do so, the firms rely not only on the systematic transfer of funds to their accounts, but also on a return flow of refunds that mimics the behavior of legitimate providers. As this flow of refunds provides a clear pattern of behavior, we model the refund flow with consideration toward using it to detect and punish firms.

The flow of funds, and refunds, depends on two key players that act as intermediaries between the buyer of the fake software and the firm. As outlined in Fig. 3, the payment processor is a key player that serves to transmit credit information from the buyer to the credit card network. The second key player is the credit card network, which incorporates both the actual card company (e.g. Visa) and the bank that issues the card (and thereby hosts the buyer's account). The payment flow is from the buyer, through the payment processor and then the credit card network, to the firm.

The trigger for a refund is a request, made by a purchaser, for return of payment upon discovery that the software is fake (or not what they expected). The purchaser may then issue a request for a refund at any point after the sale. To construct a model of requests, we let s denote the number of sales in a given period and let rq denote the number of refund requests that result from s. We model requests in period t as a Poisson random variable:

$$rq_t = \lambda s_{t-1},$$

where λ captures the expected portion of buyers from period $t-1$ who will issue a request for a refund in period t. Given the speed at which information is received and decisions are made, we are primarily concerned with periods corresponding to individual days.

When a refund request has been made, the firm can either ignore the request or grant a refund. If the firm ignores the request, then the buyer may contact the credit card network to obtain a refund. When the credit card network grants a refund to the buyer, the network must collect the funds from the firm by reversing the charge, hence refunds of this type are called chargebacks. This pattern is born out in the data as, for each of the firms under study, the average time to receive a chargeback is substantially longer than the average time to receive a refund (for AV_1, chargebacks average 23.7 days longer to process than refunds; the comparable numbers for the other firms are 21.4 days for AV_2 and 10.6 days for AV_3). For AV_1 and AV_2, 35–37% of all refunds occur within three days of sales. In contrast, only 1–6% of all chargebacks for AV_1 and AV_2 occur within three days of sales. For AV_3, only 12% of refunds occur within 3 days of sales but less than 1% of chargebacks occur within that same time.

If the firm ceases operations prior to a collection by the payment processor, then the processor must absorb the cost of the chargeback. Because a firm with a large number of sales in a period may decide to cease operations, leaving the processor at risk of absorbing a large number of chargebacks, the payment processor has an incentive to identify illegitimate firms and sever ties with them.

To model the interplay of requests, refunds (which are made directly by the firm to the buyer) and chargebacks, we must specify how payment processors monitor chargebacks to limit their risk. Let $\overline{cb}$ be a threshold, above which the credit card company denies all future transactions. In determining how many requests to refund, a firm that wishes to continue operations must balance the loss in current revenue from granting refunds against the loss of future revenue from being denied access to the credit card network. The number of refunds in a given period, rf, is thus an increasing function of the number of requests and a decreasing function of the number of chargebacks, cb,

$$rf = g\,(rq, cb)\,.$$

Let the threshold $\overline{cb}$ apply to the sum of accumulated chargebacks over T periods. The decision rule of the credit card network is to sever ties with a firm if $\sum_{s=1}^{t} cb_s > \overline{cb}$, for any period $t \in 1, \ldots, T$. As a consequence, a firm will increase the rate of refunds as the sum of accumulated chargebacks approaches the threshold $\overline{cb}$. That is, refunds follow the pattern

$$rf_t = \alpha \cdot rq_t + \beta \cdot rq_t \cdot \left\{ \overline{cb} - \sum_{s=1}^{t} cb_s < D \right\}, \qquad (1)$$

where $\{A\}$ takes the value 1 if the event A occurs and is 0 otherwise.

The desire to avoid crossing the threshold $\overline{cb}$ leads to a distinctive pattern of refunds and chargebacks. For a payment processor, (1) provides several patterns to distinguish these firms from legitimate software providers. For example, refunds from firms may increase at the periodic interval corresponding to T or may increase in reaction to an increase in chargebacks. Also, refunds should increase as the cumulated chargeback sum approaches $\overline{cb}$. For legitimate providers, no such dynamic pattern of refunds should emerge.

To understand the difference in the dynamic refund pattern between legitimate providers and fraudulent firms, note that in contrast to (1), refunds for legitimate providers follow the pattern

$$rf_t = \alpha \cdot rq_t \qquad (2)$$

Because refunds are not a function of chargebacks in (2), refunds should depend only on requests for legitimate providers.

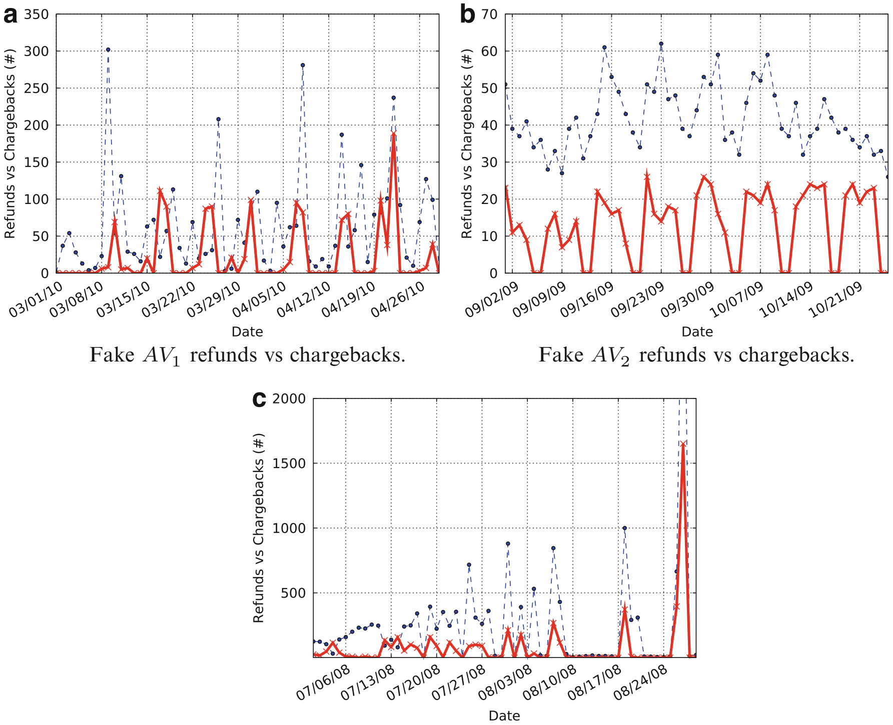

Fake AV_1 refunds vs chargebacks.

Fake AV_2 refunds vs chargebacks.

Fake AV_3 refunds vs chargebacks.

Fig. 5 Daily refunds and chargebacks from fake AV sales. The *dashed line* displays the number of refunds per day, while the *solid line* displays the number of chargebacks per day

To provide evidence that a firm's refunds respond to chargebacks, we display daily refunds and chargebacks for the firms in Fig. 5. For each of the firms, surges in daily chargebacks are closely followed by (or occur simultaneously with) surges in refunds. The only exceptions appear to be at the latter part of Fig. 5b.

While the figures reveal a dynamic pattern of refunds and chargebacks that is consistent with (1), isolating the impact of chargebacks on refunds requires that we control for the level of sales. We must do so because refunds are positively related to sales, so it is possible that sustained increases in sales could lead to increases in both chargebacks and refunds. To estimate the isolated impact of chargebacks, we construct the ordinary least squares estimates of the coefficients in

$$rf_t = \beta_0 + \beta_1 cb_t + \beta_2 cb_{t-1} + \beta_3 \overline{s_t} + u_t. \tag{3}$$

The coefficients β_1 and β_2 capture the increase in refunds on day t brought about by an increase in chargebacks on day t and day $t-1$, holding previous sales constant. The coefficient β_3 captures the increase in refunds due to an increase in average

Table 1 Coefficient estimates for (3)

AV_1 - Refunds	(I)	(II)	
Chargebacks	0.64	0.52	
	(0.24)*	(0.24)*	
Lagged Chargebacks	–	0.55	
		(0.21)*	
3-day Average Sales	0.008	0.009	
	(0.008)	(0.008)	

AV_2 - Refunds	(I)	(II)	(III)
Chargebacks	1.23	1.16	1.17
	(0.14)*	(0.15)*	(0.14)*
Lagged Chargebacks	–	0.26	0.25
		(0.12)*	(0.12)*
3-day Average Sales	0.043	0.041	0.041
	(0.004)*	(0.004)*	(0.004)*

AV_3 - Refunds	(I)	(II)	(III)
Chargebacks	0.72	0.71	0.72
	(0.24)*	(0.23)*	(0.23)*
Lagged Chargebacks	–	0.089	0.088
		(0.073)	(0.080)
3-day Average Sales	0.031	0.030	0.030
	(0.004)*	(0.004)*	(0.004)*

Note: Heteroskedasticity-robust standard errors are reported in parenthesis
Our results are not sensitive to the choice of a 3-day average sales window
*indicates significance at the 5% level

sales over the past three days ($\overline{s_t}$). As we do not observe the number of refund requests each day, we use $\overline{s_t}$ as a proxy. The quantity u_t is a random error that encompasses all other factors that influence refunds on that day.

Estimates of (3) are contained in Table 1. The column labeled (I) corresponds to (3) with $\beta_2 = 0$; that is, lagged chargebacks are not included (these lagged chargebacks are included in Column II). For each of the firms, chargebacks have a substantial impact on refunds after controlling for previous sales. For example, the estimate of 0.64 for firm AV_1 indicates that, after controlling for the average level of sales over the previous 3 days, an increase of 100 chargebacks leads to an increase of 64 refunds. In contrast, an increase in average sales of 100 leads to an increase of only 1 refund. The estimated standard errors describe the precision of our estimates: for this coefficient on chargebacks, the confidence interval of (0.16,1.12) indicates the range of plausible values for β_1. As the interval does not contain 0, the data is strongly supportive of a positive relationship between chargebacks and refunds.

In addition to controlling for sales, we also control for date of the month and day of the week to remove any monthly and daily trends. Column (III) in Table 1 corresponds to the coefficient estimates of (3) while controlling for monthly and weekly patterns. This was possible with AV_2 and AV_3 but not for AV_1 due to limited data.

Table 1 indicates significant correlation between chargebacks received and refunds granted while controlling for previous sales and monthly fluctuations among all three firms. Without knowing more firm-level details regarding their contracts with payment processors or restrictions from credit card networks further inference becomes difficult. However, we do interpret this as evidence that fraudulent firms seem to alter their refunds according to the chargebacks reported against them. Payment processors or credit card networks have more information and have a better understanding of the firm's chargeback constraints and may, therefore, be in a unique position to monitor these firms.

An important limitation to our analysis is that we lack comparable data for legitimate firms. Despite our findings above, we are unable to discern whether or not this pattern is distinctive to only illegitimate firms.

6.2 Detecting Fraudulent Firms

The previously described patterns in behavior could be observed by the payment processor since it knows the number of chargebacks against the firm at a particular time, the chargeback threshold faced by the firm, as well as the number of refunds the firm is offering (as these would have to pass through the payment processor). If the payment processor has an incentive to investigate its clients, the existence of this chargeback-responsive behavior could provide evidence that a particular antivirus company is fraudulent. The question is: Does the payment processor have an incentive to investigate its clients?

The payment processor (as noted in Sect. 4.3) receives a percentage of each transaction that occurs but faces a risk of losing business with a credit card company for too much fraudulent behavior. While losing a major credit card company like Visa would devastate a payment processor (as in the case of ePassporte), the credit card company may be hesitant to drop a payment processor if it does enough legitimate business (as in the case of Chronopay).

However, at any given time there is a risk that the fraudulent antivirus firm may be caught or may cease operations. In this case the firm will no longer be able to offer refunds and the payment processor will receive an increase in chargebacks from consumers who have no other way of receiving a refund. The payment processor would be forced to pay the entire amount of the chargeback (the chargeback fees as well as the entire refund amount) as it can no longer bill the firm. Depending on the volume of sales, the risk of future increases in chargebacks could be very costly. If this risk outweighs the revenue the payment processor receives from the firm's account, it may prefer to sever ties with the firm as to not be held liable for the potential chargebacks.

In the case when the firm is caught, credit card companies would have to pay the costs of the chargebacks if the payment processor is forced to shut down. The credit card companies may, therefore, be concerned if a small payment processor is serving an illegitimate firm that may be relatively large compared to the processor's overall volume. In these cases, credit card companies may have an incentive to investigate these firms if they are working with small payment processors. While the credit card company may not observe as much firm level information as the payment processor, it observes the chargebacks and refunds associated with a particular firm. Therefore, this could be a good technique for a credit card company to investigate fraudulent firms.

As mentioned above, we expect the rate of refunds offered by a fraudulent firm to vary in response to chargebacks incurred by the firm. As firms increase their sales, payment processors and credit card networks face increased risk of liability for future chargebacks if the firm ceases operations. This risk may warrant investigation of fraudulent firms using these observable patterns.

7 Ethical Considerations

The nature of the data that we collected raises a number of ethical concerns. In particular, we have a large amount of personal information for the victims who were defrauded by these three fake AV businesses. Thus, we took measures to protect the privacy and identity of the victims through the use of data encryption, automated program analysis, and by conducting our research according to established ethical principles in the field [2, 8, 12, 16]. We also obtained approval from the Institutional Review Board (IRB) at the University of California, Santa Barbara before performing our analysis. Finally, we provided all information that we obtained to U.S. law enforcement officials.

8 Related Work

In the past few years, there have been several studies that have analyzed various aspects of fraudulent businesses selling fake antivirus products. Researchers from Google described the techniques and dynamics used by cybercriminals to drive traffic to their sites via landing pages [30]. Other work analyzed the distribution and installation methods of rogue security software [10]. Various security vendors have reported on potential revenue from scareware operations based on the number of infections that they observed [4,37]. Cova et al. presented an analysis of the rogue antivirus structure and indirectly tried to measure the number of victims and profits based on poorly configured web servers used by several fake AV groups [6]. They estimated the conversion rate of infections to sales at 1.36%, which is slightly lower than the rates that we observed. We also found a similar geographic distribution of victims in the U.S., and number of domains registered by larger fake AV groups. In comparison, our data provides a much more complete view of large-scale fake AV operations, with information dating back more than two years. We also had visibility of refunds and chargebacks from fake AV sales, which has never been studied before.

Techniques to identify drive-by-download attacks have been proposed that analyze web sites for malicious content in a virtual or emulated environment to detect exploits [5, 14]. The prevalence of malicious web sites has been examined through crawler-based approaches that analyzed billions of web pages [28, 29]. Another study analyzed drive-by attacks via infiltration and provided insights into the compromised web servers used in the attacks as well as the security posture of potential victims [35].

A number of recent papers have analyzed the reasons that cause users to fall victim to phishing scams, which include lack of knowledge and attentiveness to browser and other security related cues [7, 9]. Several approaches have been proposed to detect phishing sites such as analyzing page content, layout, and other anomalies [22, 26, 31]. In addition, studies have analyzed the modus operandi of the criminal operations behind phishing [23], and the effectiveness of phishing defenses [25].

Previous work has investigated the Internet's underground economy, through advertised prices of web forums [39] and IRC chat rooms [11]. Holz et al. studied the drop zones used by botnets to store stolen information from victims [13]. Stone-Gross et al. hijacked the Torpig botnet and studied the data exfiltrated from infected computers, and estimated the value of the compromised financial information (e.g., credit card numbers and bank account credentials) [33]. The underground economy of large-scale spam operations was examined in [36]. The paper analyzed the complexity in orchestrating spam campaigns, and explored an underground forum used by spammers to exchange goods and services. Another type of scam, known as *One Click Fraud*, was studied by Christin et al. The fraud works through intimidation (similar to fake AV) by threatening unsuspecting web site visitors with potential embarrassment (e.g., the victim was browsing pornographic content) unless a payment is received for a nonexistent service. The authors presented an economic model to determine the number of users that must fall victim to the scam in order to remain economically viable, and estimated losses in the tens to hundreds of thousands of U.S. dollars [3].

9 Conclusions

In this paper, we have presented an in-depth study of how a particular type of scareware, namely fake anti-virus software, is deployed and managed. Our work is unique in that it is based on the information contained on a number of key servers that were part of the criminals' infrastructure. This unprecedented access allowed us to obtain ground truth about the type and sophistication of the techniques used to lure victims into paying for scareware, as well as the amount of transactions performed, including refunds and chargebacks.

We leveraged this data to build an economic model that shows how cybercriminals are very careful in performing refunds and chargebacks in order to maintain a balanced financial posture that does not immediately reveal their criminal nature. Nonetheless, the economic model also outlines how these operations have distinct characteristics that may differentiate these criminal endeavors from legitimate business operations.

Future work will extend the current model with detection capabilities that can be directly applied to payment data streams. The goal is to develop a tool based on the model that can identify scareware operations automatically.

Acknowledgments This work was supported by the Office of Naval Research (ONR) under Grant N000140911042 and by the National Science Foundation (NSF) under grants CNS-0845559 and CNS-0905537. We would also like to thank the anonymous reviewers for their valuable suggestions and insights.

References

1. Bayer U, Habibi I, Balzarotti D, Kirda E, Kruegel C (2009) A view on current malware behaviors. In: USENIX workshop on large-scale exploits and emergent threats (LEET), 2009
2. Burstein A (2008) Conducting cybersecurity research legally and ethically. In: USENIX workshop on large-scale exploits and emergent threats (LEET), 2008
3. Christin N, Yanagihara S, Kamataki K (2010) Dissecting one click frauds. In: ACM conference on computer and communications security (CCS), 2010
4. Correll S, Corrons L (2010) The business of rogueware: analysis of the new style of online fraud. http://www.pandasecurity.com/img/enc/The%20Business%20of%20Rogueware.pdf
5. Cova M, Kruegel C, Vigna G (2010) Detection and analysis of drive-by-download attacks and malicious javascript code. In: Proceedings of the international world wide web conference (WWW), 2010
6. Cova M, Leita C, Thonnard O, Keromytis A, Dacier M (2010) An analysis of rogue AV campaigns. In: Symposium on recent advances in intrusion detection (RAID), 2010
7. Dhamija R, Tygar J, Hearst M (2006) Why phishing works. In: Conference on human factors in computing systems (CHI), 2006
8. Dittrich D, Bailey M, Dietrich S (2009) Towards community standards for ethical behavior in computer security research. Technical report 2009–1, Stevens CS, April 2009
9. Egelman S, Cranor L, Hong J (2008) You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In: Conference on human factors in computing systems (CHI), 2008
10. Fossi M, Turner D, Johnson E, Mack T, Adams T, Blackbird J, Low M, McKinney D, Dacier M, Keromytis A, Leita C, Cova M, Overton J, Thonnard O (2009) Symantec report on rogue security software. In: Whitepaper, 2009
11. Franklin J, Paxson V, Perrig A, Savage S (2007) An inquiry into the nature and causes of the wealth of internet miscreants. In: ACM conference on computer and communications security (CCS), 2007
12. Garfinkel S (2008) IRBs and security research: myths, facts and mission creep. In: Proceedings of the USENIX workshop on usability, psychology, and security, 2008

13. Holz T, Engelberth M, Freiling F (2008) Learning more about the underground economy: a case-study of keyloggers and dropzones. Reihe Informatik TR-2008–006, university of Mannheim, 2008
14. Ikinci A, Holz T, Freiling F (2008) Monkey-spider: detecting malicious websites with low-interaction honeyclients. In: Proceedings of Sicherheit, Schutz und Zuverlässigkeit, April 2008
15. International Secure Systems Lab (2010). Anubis: analyzing unknown binaries. http://anubis.iseclab.org
16. Kenneally E, Bailey M, Maughan D (2010) A framework for understanding and applying ethical principles in network and security research. In: Proceedings of the workshop on ethics in computer security research (WECSR), 2010
17. Kirk J (2010) Bredolab-infected PCs downloading fake antivirus software. http://www.pcworld.com/businesscenter/article/209031/bredolabinfected_pcs_downloading_fake_antivirus_software.html
18. Krebs B (2009) Massive profits fueling rogue antivirus market. In: Washington post, 2009
19. Krebs B (2009) Virus scanners for virus authors. http://krebsonsecurity.com/2009/12/virus-scanners-for-virus-authors/
20. Krebs B (2010) Following the money, ePassporte edition. http://krebsonsecurity.com/2010/09/following-the-money-epassporte-edition/
21. Krebs B (2010) Rogue antivirus victims seldom fight back. http://krebsonsecurity.com/2010/07/rogue-antivirus-victims-seldom-fight-back/
22. Ludl C, McAllister S, Kirda E, Kruegel C (2007) On the effectiveness of techniques to detect phishing sites. In: Proceedings of the conference on detection of intrusions and malware & vulnerability assessment (DIMVA), 2007
23. McGrath K, Gupta M (2008) Behind phishing: an examination of phisher modi operandi. In: USENIX workshop on large-scale exploits and emergent threats (LEET), 2008
24. Mick J (2010) Russian anti-spam chief caught spamming. http://www.dailytech.com/Russian+AntiSpam+Chief+Caught+Spamming/article18423.htm
25. Moore T, Clayton R (2007) An empirical analysis of the current state of phishing attack and defence. In: Workshop on the economics of information security (WEIS), 2007.
26. Pan Y, Ding X (2006) Anomaly based web phishing page detection. In: Annual computer security applications conference (ACSAC), 2006
27. Poulsen K (2009) Conficker doomsday worm sells out for $49.95. http://www.wired.com/threatlevel/2009/04/conficker-dooms/
28. Provos N, McNamee D, Mavrommatis P, Wang K, Modadugu N (2007) The ghost in the browser: analysis of web-based malware. In: USENIX workshop on hot topics in understanding botnets (HotBots), 2007
29. Provos N, Mavrommatis P, Rajab M, Monrose F (2008) All your iFRAMEs point to us. In: USENIX security symposium, 2008
30. Rajab M, Ballard L, Mavrommatis P, Provos N, Zhao X (2010) The nocebo effect on the web: an analysis of fake anti-virus distribution. In: USENIX workshop on large-scale exploits and emergent threats (LEET), 2010
31. Rosiello A, Kirda E, Kruegel C, Ferrandi F (2007) A layout-similarity-based approach for detecting phishing pages. In: Security and privacy in communication networks (SecureComm), 2007
32. Samosseiko D (2009) The Partnerka what is it, and why should you care? In: Annual virus bulletin conference, 2009
33. Stone-Gross B, Cova M, Cavallaro L, Gilbert R, Szydlowski M, Kemmerer R, Kruegel C, Vigna G (2009) Your botnet is my botnet: analysis of a botnet takeover. In: ACM conference on computer and communications security (CCS), 2009
34. Stone-Gross B, Moser A, Kruegel C, Kirda E, Almeroth K (2009) FIRE: FInding rogue nEtworks. In: Annual computer security applications conference (ACSAC), 2009
35. Stone-Gross B, Cova M, Kruegel C, Vigna G (2010) Peering through the iFrame. In: IEEE mini-conference on computer communications (INFOCOM), 2010

36. Stone-Gross B, Holz T, Stringhini G, Vigna G (2011) The underground economy of spam: a Botmasters perspective of coordinating large-scale spam campaigns. In: USENIX workshop on large-scale exploits and emergent threats (LEET), 2011
37. TrendMicro (2010) The business of cybercrime a complex business model. Technical report, 2010
38. Villeneuve N, Deibert R, Rohozinski R (2010) KOOBFACE: Inside a crimeware network. InfoWar monitor JR04–2010, The SecDev group, 2010
39. Zhuge J, Holz T, Song JGC, Han X, Zou W (2009) Studying malicious websites and the underground economy on the chinese web