

Dissecting One Click Frauds

Nicolas Christin, Sally S. Yanagihara, Keisuke Kamataki

April, 23, 2010

CMU-CyLab-10-011

CyLab
Carnegie Mellon University
Pittsburgh, PA 15213

Dissecting One Click Frauds*

Nicolas Christin
Carnegie Mellon
INI/CyLab
nicolasc@cmu.edu

Sally S. Yanagihara
Carnegie Mellon
INI/CyLab Japan
sallyy@cmu.edu

Keisuke Kamataki
Carnegie Mellon
LTI/CS
keisuke@cs.cmu.edu

Carnegie Mellon University Technical Report CMU-CyLab-10-011

April 23, 2010

Abstract

“One Click Fraud” is an online confidence scam that has been plaguing an increasing number of Japanese Internet users, in spite of new laws and the mobilization of police task forces. In this scam, the victim clicks on a link presented to them, only to be informed that they just entered a binding contract and are required to pay a registration fee for a service. Even though no money is legally owed, a large number of users prefer to pay up, because of potential embarrassment due to the type of service “requested” (e.g., pornographic goods).

Using public reports of fraudulent websites as a source of data, we analyze over 2,000 reported One Click Frauds incidents. By correlating several attributes (WHOIS data, bank accounts, phone numbers, malware installed...), we discover that a few fraudsters are seemingly responsible for a majority of the scams, and evidence a number of loopholes these miscreants exploit. We further show that, while some of these sites may also be engaging in other illicit activities such as spamming, the connection between different types of scams is much more tenuous than expected. Last, we show that the rise in the number of these frauds is fueled by high expected monetary gains in return for very little risk. The quantitative data obtained gives us an interesting window on the economic dynamics of some online criminal syndicates.

Keywords: Security Economics, Measurements, Web Frauds, Online Crime

*This research was supported in part by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by the National Science Foundation under ITR award CCF-0424422 (TRUST).

1 Introduction

In the family apartment in Tokyo, Ken is sitting at his computer, casually browsing the free section of a mildly erotic website. Suddenly, a window pops up, telling him,

Thank you for your patronage! You successfully registered for our premium online services, at an incredible price of 50,000 JPY.¹ Please promptly send your payment by bank transfer to ABC Ltd at Ginko Bank, Account 1234567. Questions? Please contact us at 080-1234-1234.

Your IP address is 10.1.2.3, you run Firefox 3.5 over Windows XP, and you are connecting from Tokyo.

Failure to send your payment promptly will force us to mail you a postcard reminder to your home address. Customers refusing to pay will be prosecuted to the fullest extent of the law.

Once again, thank you for your patronage!

A sample postcard reminder is shown on the screen, and consists of a scantily clad woman in a provocative pose. Ken has a sudden panic attack: He is married, and, if his wife were to find out about his browsing habits, his marriage would be in trouble, possibly ending in divorce, and public shame. In his frenzied state of mind, Ken also fears that, if anybody at his company heard about this, he could possibly lose his job. Obviously, those website operators know who he is and where he lives, and could make his life very difficult. Now, 50,000 JPY (\approx USD 500) seems like a small price to pay to make all of this go away. Ken immediately jots down the contact information, goes to the nearest bank, and acquits himself of his supposed debt.

Ken has just been the victim of a relatively common online scam perpetrated in Japan, called “*One Click Fraud*.” In this fraud, the “customer,” i.e., the victim, does not enter any legally binding agreement, and the perpetrators only have marginal information about the client that connected to their website (IP address, User-Agent string), which does not reveal much about the user.² However, facing a display of authority stressed by the language used, including the notion that they are monitored, and a sense of shame from browsing sites with questionable contents, most victims do not realize they are part of an extortion scam. Some victims even call up the phone numbers provided, and, in hopes of resolving the situation, disclose private information, such as name or address, to their tormentors, which makes them even more vulnerable to blackmail.

As a result, One Click Frauds have been very successful in Japan. Annual police reports show that the estimated amount of monetary damages stemming from One Click Frauds and related confidence scams are roughly 26 billion JPY per year (i.e., USD 260 million/year). The rapid rise of such frauds has led to new laws being passed [17, 19], to the deployment of police task forces specifically put in charge of solving these frauds, and to specialized help desks [15]. On the other hand, there have been so far, on average only 657 arrests and 2,859 solved cases per year [14]. Considering the lack of technical sophistication needed to set up such extortion schemes, and the fact that bank accounts and phone numbers can be made very hard

¹100 JPY \approx 1 USD.

²The Panoptick project [11] aims to show that browser fingerprints can divulge considerable information about the user. However, none of the One Click Fraud websites we investigated appears to use any advanced browser fingerprinting techniques.

to trace in Japan (discussed in Section 5), it appears that One Click Frauds offer easy money to aspiring criminals.

From a research point of view, One Click Frauds offer a unique opportunity for a case study in online crime economics. First, One Click Frauds, are extremely localized. We have not seen any instance of this specific fraud outside of Japan, presumably due to the fact that the scammers prey on unique Japanese cultural characteristics, such as respect of authority, or embarrassment at the idea of causing trouble. Rather than limiting our research, we view this local aspect as an opportunity. Indeed, because One Click Frauds are contained to Japan, and are almost exclusively used in Japanese-language websites, we can obtain a relatively exhaustive view of how these frauds are deployed and the characteristics they share, without the need for a complex measurement infrastructure. Furthermore, One Click Frauds, albeit unique, are very closely related to the body of scareware scams (e.g., windows popping up trying to entice users to buy fake anti-virus software) that are becoming increasingly pervasive, and we believe One Click Frauds offer an interesting window in this type of criminal enterprise.

The first contribution of this paper is to collect and analyze a corpus of over 2,000 reported One Click Fraud incidents, and to use the data collected to paint a relatively clear picture of the economic dynamics of an instance of online criminal activity.

Specifically, we set out to answer the following questions: Who is committing these frauds? Are One Click Frauds the product of organized criminal activities, or are they more artisanal in nature, with many aspiring crooks trying their luck? How much do criminals stand to gain? Are there vulnerabilities in the network infrastructure (e.g., weak registration processes, easy access to compromised accounts, etc...) that are easily exploited by the miscreants?

In the process of finding answers to these questions, a second important contribution of our paper is to provide monetary amounts, which help dimension the size of the One Click Fraud market, and the profits potentially made. While economic modeling of cybercrime has been an increasingly active research topic (see for instance [12, 16, 22, 24, 35]), we believe that obtaining additional measurements and clearly detailing the methodology we employ should help enrich our knowledge of how online criminal syndicates operate, while assuaging concerns on the validity of the estimates formulated [13]. Furthermore, some of the measurement methodology we employ is likely to be applicable beyond the context of One Click Frauds.

Third, our analysis methodology can be helpful to law enforcement agencies. Law enforcement typically assigns different fraud instances to specific agents, who mostly operate independently of each other. However, we show that, by analyzing a relatively large dataset of frauds, we can extract characteristics that are not visible when considering smaller subsets of these frauds. These characteristics can in turn be useful to law enforcement in identifying, and perhaps prosecuting, the miscreants behind these crimes [29].

The rest of this paper is organized as follows. We review the related work in Section 2. We then explain our data collection methodology in Section 3. We turn to analyzing the data collected in Section 4. We use this analysis to study economic incentives for the perpetrators in Section 5, before discussing the implications of our study, and drawing brief conclusions in Section 6.

2 Related work

Moore et al. contend that online crime has become economically significant since around 2004 [22]. Hence, research in measuring the economic impact of online crime is a relatively recent field, but a number of papers on the topic have been published in the past couple of years. We present several important advances in the field in this section, and refer the reader to Moore et al. [22] for a more exhaustive treatment of the literature.

Some of the pioneering work in the field has tried to quantify the value of fraudulent financial credentials as well as that of compromised hosts (“bots”), through passive observations. In particular, Thomas and Martin [30], and Franklin et al. [12] monitored IRC channels where miscreants attempt to sell commodities as diverse as stolen credit card numbers, bank account credentials, or email databases for spam, to obtain estimates of the value of these items as well as the volumes of the market associated with exchanges of such goods.

Along the same lines, Zhuge et al. [38] describe how criminals operate in Chinese underground markets, and provide some insight regarding some of the goods exchanged in these markets by monitoring web forums where such items are advertised; a particularly popular item appears to be forged or stolen online video game currency.

Passive monitoring, as discussed above, has been criticized for only measuring advertised prices, as opposed to actual realized sales [13]. A more recent line of research addresses this concern by actively participating in the online exchanges, by essentially “hijacking” some of the infrastructure used by miscreants. Kanich et al. [16] infiltrate a large botnet, and modify some of the spam traffic generated by the botnet to redirect purchases to a server under their control, thereby acquiring data on spam conversion rates. The key insight is that 350 million spam messages sent result in 28 sales. In other words, spamming is highly ineffective, but, considering its cost is negligible, and that bots can be used to generate other forms of profit (e.g., phishing site hosting), it remains a relatively popular form of crime. A similar botnet take-over is described in Stone-Gross et al. [28], who monitor the type of services supported by the Torpig botnet. In a similar vein, Wondracek et al. [35] focus on the economics of the distribution of pornographic materials, by covertly operating an adult web server.

Perhaps closer in spirit to the methodology we employ in this paper, a significant body of literature [20,21,23] is devoted to quantifying the economic impact of phishing scams, as well as the modus operandi of the attackers. One of the main outcomes of this line of research is to evidence that a large proportion of phishing activities are carried out by a modest number of phishing gangs, who use relatively sophisticated “fast-flux” techniques, where phishing sites are used as disposable commodities hosted on compromised hosts. This body of research puts the number of phishing websites at around 116,000. Likewise, Moore and Edelman [24] gather a large corpus of data on typosquatting domains, to quantify the economic value of such sites, as well as potential disincentives for advertisement networks to intervene forcefully. Finally, Provos et al. [25] provide a comprehensive overview of web-based malware distribution, showing that there are in excess of 3 million web pages attempting to infect their visitors.

The work we present here is complementary to the existing literature. To our knowledge, this is the first time One Click Frauds are analyzed from a quantitative and technical perspective. (Research on the legal

ramifications and countermeasures, on the other hand, has been active [17].) More generally, we focus on a relatively contained instance of a scam, and attempt to gather both economic data (potential revenue, deployment costs to the perpetrators...) and structural data (number of players, relationships between players, ...) through a systematic measurement methodology, which we believe can be useful beyond the study of this specific fraud.

3 Data collection

To find instances of One Click Frauds, we rely on public forums which report fraud incidents, along with details regarding the website being used, the amount of money extorted, as well as fraudster contact information (bank account number, phone number, ...). We collect data from three public forums:

2 Channel BBS [5]: The largest bulletin board in Japan, which provides discussion threads on various topics ranging from Japanese anime to sports. Several ongoing threads are dedicated to denouncing One Click Frauds.

Koguma-neko Teikoku [2]: Privately owned website providing help to solve consumer problems related with online activities. A section of the site is devoted to describing One Click Frauds, including information about scam incidents.

Wan-Cli Zukan [7]: Privately owned website solely devoted to exposing websites partaking in One Click Frauds.

Collection methodology. We gather data posted on the three websites we polled over a period of roughly three years (2006-2009). Specifically, we collect 2 Channel posts made between March 6, 2006 and October 26, 2009, Koguma-neko Teikoku posts made between August 24, 2006 and August 14, 2009, and Wan-Cli Zukan posts made between September 6, 2006 and October 26, 2009. All in all, we gathered 2,140 incident reports. While it is difficult to determine how exhaustive our data collection is compared to the total number of frauds perpetrated, we note that there is a significant overlap in the data provided by all three sites, which indicates we probably captured the most successful frauds, and that our coverage is likely adequate.

Initially, we also tried to extract information from data provided by the Japanese police [31]. Unfortunately, these reports are in image format, and do not contain much actionable information, even after using optical character recognition. Indeed, most of the interesting details (e.g., bank account used) are often not divulged.

Data parsing. The three sites present marked differences. 2 Channel is based on anonymous postings from various users who have encountered a One Click Fraud website and post notifications to other users. Due to its popularity and large user base, 2 Channel contains a wealth of information. However, because 2 Channel threads are essentially an open discussion, parsing the data for useful input is challenging. Koguma-neko Teikoku is also a collaborative web space, but, different from 2 Channel, Koguma-neko Teikoku requires posters to input information about One Click Frauds in a specified format. Finally, Wan-Cli Zukan is closer to a blog, where the website owner periodically posts notifications in a fixed format. Outsiders cannot directly post to the site.

Parsing the data from Wan-Cli Zukan and Koguma-neko Teikoku is straightforward, as both sites use a predetermined format for all posts relevant to One Click Frauds. Likewise, the data is of high quality. Randomly sampling the reported sites, we did not find a single instance of slander – i.e., benign sites which would be reported as fraudulent. Either the sites reported were, indeed, hosting One Click Frauds, or they had been recently taken down.

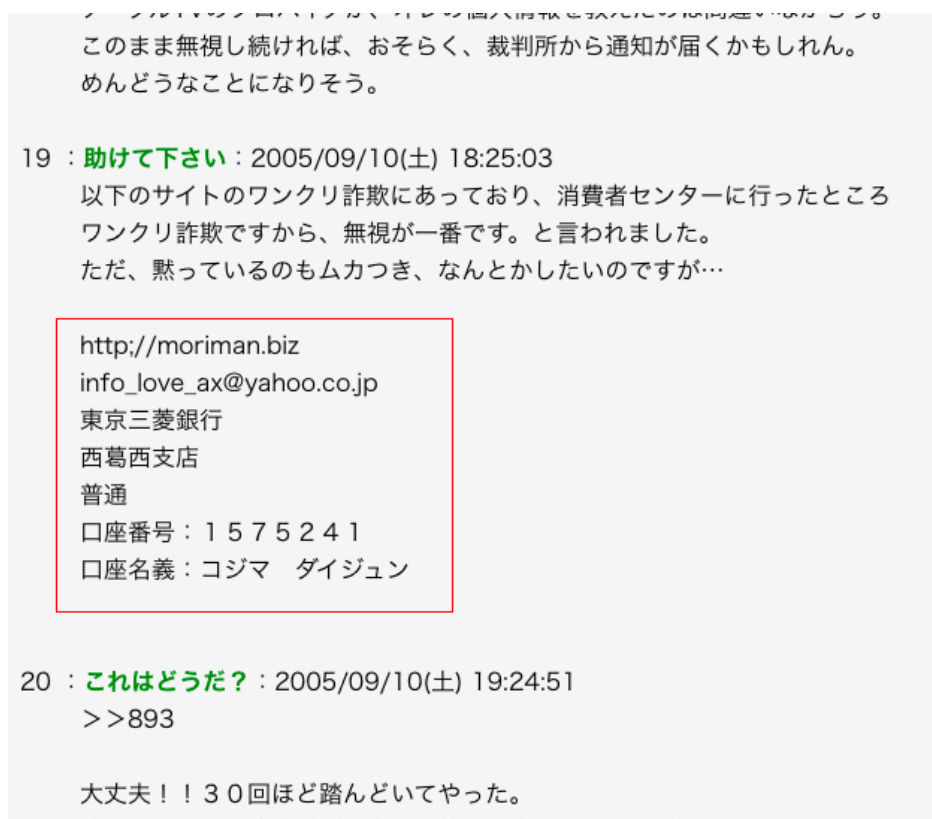


Figure 1: **Example of 2 Channel posting.** The lack of structure makes parsing slightly complex. Here, usable information is only contained in the squared area; oftentimes, formatting is even more haphazard.

On the other hand, 2 Channel is considerably more challenging to use. An example of 2 Channel posting is shown in Fig. 1. Since the data is not formatted according to a specific standard, we have to perform some basic text segmentation. Adding to the difficulty, is the fact that the Japanese language has comparatively few punctuation rules. For instance, words are rarely separated by spaces, there is no capitalization, and most characters have a couple of different readings depending on the context. Furthermore, in forums like 2 Channel, specialized slang and abbreviations are the norm, rather than the exception. To perform text extraction, we use the MeCab parser and data analyzer [4]. MeCab extracts Japanese characters into semantic units. This allows us to obtain a list of tokens, such as “ginko” (bank), “Sumitomo” (name), “hutsuu” (usual), “1234567.” We can then perform context-dependent parsing. The previous series of token tells us that it is highly likely that the poster talks about a “usual” checking account number 1234567 and Sumitomo bank.

Data source	N. records
2 Channel	1077
Unambiguous, w/ URL	353
Unambiguous, w/o URL	174
Ambiguous, w/ URL	218
Ambiguous, w/o URL	332
Koguma-neko	372
Unambiguous, w/ URL	2
Ambiguous, w/ URL	362
Ambiguous, w/o URL	8
Wan-Cli Zukan	691
Unambiguous, w/ URL	632
Unambiguous, w/o URL	59

Table 1: Number of extracted fraud incidents from each source.

Despite the added difficulty of parsing 2 Channel threads, we did not encounter any erroneous or libelous postings. While a certain level of verbal abuse is present in such an unmoderated forum, it is easily separated from relevant information, once text segmentation is performed.

Extracted attributes. For all three sources of data, after having parsed and cleaned the input, we extract the following attributes: URL of the fraudulent website; Bank account number given by the fraudster for remittance of funds; Bank name; Branch name; Account holder name; Contact phone number; Registration fee requested (i.e., amount of the fraud).

Some of these attributes (e.g., account holder name) are probably fake. However, for the fraudsters to be able to receive money, the financial information has to be genuine. Likewise, contact information (email and phone numbers), when present, is usually accurate, as miscreants provide the victims with a way of “calling back” to further their social engineering attacks.

Table 1 describes the number of fraud incidents that we extracted from each of the three sources of data. Note that, these sources are not mutually exclusive. In fact, we have significant overlap between the different sources.

Second, the raw data that we extract needs to be cleaned up significantly. A number of records are incomplete to some extent (missing phone number for example); out of these incomplete records, particularly problematic are records that do not contain a URL field, as the existence of a fraud cannot be verified. In addition, a number of records are ambiguous in that some fields have more than one entry. For instance, a given post may contain one URL but several phone numbers. Most of these records need to be scrutinized manually, to check whether parsing was done correctly and whether the record indeed contains several fields.

All results are stored in a MySQL database. We illustrate a typical entry in our database in Table 2. To help make our experiments reproducible, we make a copy of a dump of our database available (in gzip’ed

form) at <http://arima.ini.cmu.edu/public/oneclick.sql.gz>.

After having cleaned up our original dataset, we use the extracted URL to infer the domain hosting the fraud, and, whenever possible, to obtain registrar (“WHOIS”) information. We also periodically download the corresponding website (using `wget` in recursive mode) to check for potential changes indicating a take-down, and to check for the presence of malware on the fraudulent site.

Note that, due to the time interval (three years) over which incident reports we collect were reported, a significant amount of data points to sites that have long been taken down. Thus, a number of records are hard to verify and may be missing information. For example, we have a smaller number of entries with valid DNS domain information, compared to the number of incidents reported.

ID	Date posted	URL
370	2007-02-03 00:00:00	http://www.331164.com/
IP	Email	Bank
69.64.155.122	info@331164.com	Sumitomo Mitsui
Branch	Acct. Type	Acct. Number
Koenji	Checking	7184701
Acct. Holder	Phone #	Fee
Takahashi, Mizuki	080-5182-7956	JPY 88,000

Table 2: **Example of database entry.** Note that not all attributes can always be extracted.

4 Data analysis

We next turn to analyzing the relatively large corpus of One Click Frauds we gathered. Specifically, we seek loopholes in the infrastructure that attackers are able to exploit. We then try to assess some of the characteristics of the online criminal market behind One Click Frauds, and in particular, investigate whether some miscreants are responsible for several frauds. Last, we try to determine whether fraudsters engage in other illicit activities beyond One Click Fraud.

4.1 Infrastructural loopholes

We start by checking our data corpus for evidence of repeating patterns in the attributes we collected. The idea is that, departures from usual patterns may indicate evidence of vulnerabilities in the infrastructure. If, for instance, a disproportionate number of frauds uses bank accounts at Bank *X*, one can reasonably infer that Bank *X* processes for identity verification and account establishment are flawed.

We first consider the phone numbers used by fraudsters. We are able to identify the phone number used for callback in 516 separate incidents. We check the phone numbers in our database against the phone number ownership list published by the Japanese Ministry of Internal Affairs and Communications. We find that 38.6% of the phone numbers used in One Click Frauds are *au* cellular lines, 23.3% are Softbank cellular lines, and 16.5% are local landlines for the Tokyo area. A 2009 report about the Japanese cellular phone

Registrar	Market share [34]	Registrar	Prop. of frauds
Go Daddy	29.08%	ENom	47.56%
ENom	8.30%	GMO Internet	17.48%
Tucows	6.82%	Above	5.14%
Network Solutions	6.06%	Go Daddy	4.11%
Schlund + Partners	4.38%	Tucows	3.34%
Melbourne IT	4.34%	Key Systems	2.83%
Wild West Domains	2.89%	New Dream Network	1.54%
Moniker	2.43%	Abdomainations	1.29%
Register.com	2.40%	All Earth Domains	1.03%
Public Domain Registry	2.17%	Dotster	1.03%
		Moniker	1.03%
		OTHER	13.62%

Table 3: **Registrars used in One Click Frauds.** The left table shows the market share of the 10 most popular registrars (and thus does not sum to 100%). The bottom right shows the 11 most frequently used registrars in One Click Frauds, along with the respective proportion of frauds they host. Numbers are obtained out of the 389 fraudulent domains for which WHOIS information was accessible.

market [33, pp. 78–79], tells us that *au* has about 25% of the market share, while NTT Docomo represents about 50% of the market. So, it appears as though *au* cellular lines are disproportionately targeted by fraudsters. Whether this is due to lax registration practices, or to easier access to compromised lines, is unclear. The large proportion of Tokyo numbers may be explained not only by the amount of population in that area, but also by the fact that most forwarding services, which redirect all incoming calls to a different number while preserving the anonymity of the callee, use Tokyo numbers.

Next, we identify the bank used in 803 separate incidents. We observe that online banks tend to be slightly more represented in these frauds than their actual market share should warrant. For example, 116 frauds (14.44%) use the online Seven Bank. On the other hand, this bank does not have a significant market share in Japan [33, pp. 82–83]. eBank and JapanNet Bank are two other examples of online banks used at a notable level (around 3%) in One Click Frauds, while holding only a small market share. A possible explanation is that online banks do not require a physical interaction to let individuals open an account, and are thus potentially more prone to abusive registrations.

We observe even more interesting patterns, when we look at DNS registrars and resellers that are abused by fraudsters. Table 3 compares the market share, as of Nov. 2009, of the top domain name registrars [34] with the proportion of registrars used by domain names participating in One Click Frauds. We are able to identify the registrars used in 389 incidents. We note that ENom is apparently much more victim of abuse than other top domain name registrars. Likewise, Above and GMO Internet rank highly with fraudsters, while not having significant market share. This result can indicate one of two possibilities: Either these registrars have very lenient registration policies, or they entrust some of their domains to resellers that are

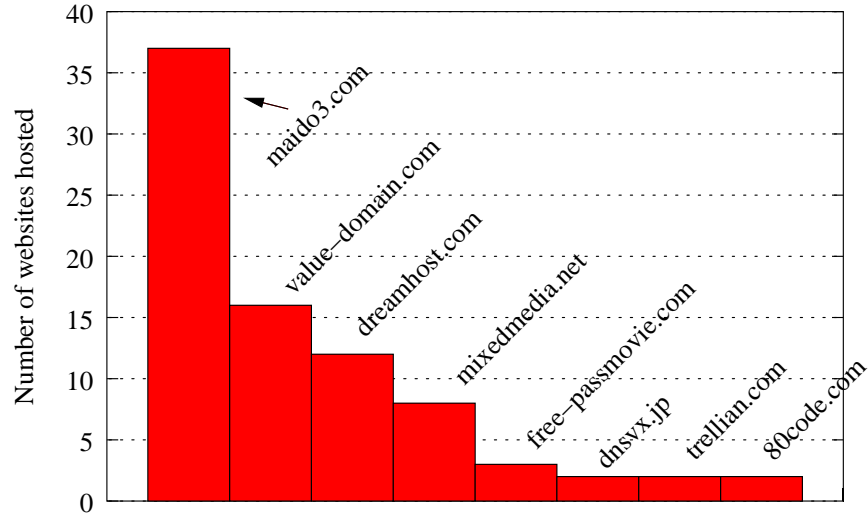


Figure 2: **Number of websites hosted by the resellers used in more than one incident.**

not perfectly scrupulous.

To find out which domain name resellers are used, we examine the DNS field in the WHOIS record associated with the domain used in each One Click Fraud. A large number of sites we investigate have been taken down (particularly those obtained from 2006-2007 sources). In some cases, though, the One Click Fraud sites have been replaced by domain parking pages, and thus still have valid WHOIS information, which may reveal information on the reseller used. Another large number of sites use either their own DNS server, or free DNS services like `opendns.com`, and are thus discarded from analysis. We complement investigation of the WHOIS DNS fields with a simple reverse DNS/DNS lookup script. For instance: `dig +short admovie69.com | xargs dig +short -x` yields `ikero.dreamhost.com`, which tells us the web hosting service is `dreamhost.com`.

We eventually manage to identify the resellers in 97 incidents. In Figure 2, we graph the number of at resellers that are used in more than one incident. We note that, for our successful lookups, a couple of resellers (`maido3.com`, `value-domain.com`, ...) appear to be highly represented. They tend to be cheap resellers, with lax, or absent registration checks.

Specifically, we investigated further `maido3.com`. We used a disposable email address to request rental of a specific domain name and server space. We used a fake name, address and phone number, which could have been easily spotted, as the address was that of a famous subway station. We received an immediate response with invoicing details. A simple money transfer was all that was needed; and this could be performed anonymously by using cash at a convenience store. The reseller promised the site would be up within 30 minutes of the payment being sent. No identity checks were ever performed. In other words, this specific reseller can be easily abused for fraudulent activities. While we have not attempted the same type of test with the other “popular” resellers, we venture to guess that their registration processed are likely easily abused as well.

4.2 Grouping miscreants

We next attempt to determine whether we can find evidence of organized criminal activities through miscreant behavior. In this paper, we qualify by “organized criminal activity” large scale operations involving many frauds, using a considerable number of bank accounts and stolen or abused phone lines, and possibly coordinating with other organizations to deploy malware or share stolen information. We contrast such large scale operations with those involving only a couple of sites.

Methodology. We define an undirected graph $G = (V, E)$ as follows. We create a vertex $v \in V$ for each domain name, bank account number, or phone number our database contains. Then, we connect vertices belonging to the same fraud with edges E . For instance, if our database contains an entry for a fraud hosted on `example.com`, with bank account number 1234567 and phone number 080-1234-1234, we add $(v_1, v_2, v_3) = \{\text{example.com}, 1234567, 080-1234-1234\}$ to V , and three edges $(e_1, e_2, e_3) = \{v_1 \leftrightarrow v_2, v_1 \leftrightarrow v_3, v_2 \leftrightarrow v_3\}$ to E . Note that, by building the graph G in this fashion, two nodes of the same type (domain name, bank account number, phone number) never link directly to each other, but can be connected through an intermediary. For example, two websites using the same phone number result in a connected path between the three concerned nodes. One natural question is why we chose not to link two different domain names pointing to the same IP address. The reason is, that we realized that this happened sometimes for sites that were markedly different, only because they were co-hosted on the same machine by a common reseller. In other words, both incidents shared the same web hosting provider, but were not related otherwise. We decided to err on the side of caution, and not to consider identical IP addresses as a strong indicator of identical ownership.

Parsing the entire database yields a graph with 1,341 nodes and 5,296 edges. We first isolate 26 “singletons” representing connected subgraphs containing at most three nodes, and at most one domain name, one phone number and one bank account. These singletons represent incidents that we cannot connect to any other fraud. All “one-off” scams fall in this category. Excluding the 26 singletons, the whole graph G contains 105 connected subgraphs (“clusters”), with sizes ranging from a couple of nodes and edges to a large cluster containing 179 nodes (56 domains, 112 bank accounts, 11 phone numbers), and 696 edges.

We plot the graph G in Fig. 3. While the specific domain names, phone numbers, and bank account numbers are not readable at this scale, we can clearly distinguish the connected subgraphs, which are ordered in separate boxes in the figure. The presence of large connected subgraphs (e.g., top right) indicates that some miscreants are operating a large number of sites, and are reusing phone numbers or bank accounts across several sites.

While it could, in theory, be possible that two completely different criminals share an identical bank account number, we reject this hypothesis, as the bank account numbers used have to be valid for the criminals to collect their dues. Hence, a shared bank account number means that, at the very least, the miscreants sharing the account are in a tight business relationship, and probably are the same individual or group of individuals. Likewise, because phone numbers used in these frauds are genuine, reusing the same phone number across several frauds indicates a business relationship between the perpetrators.

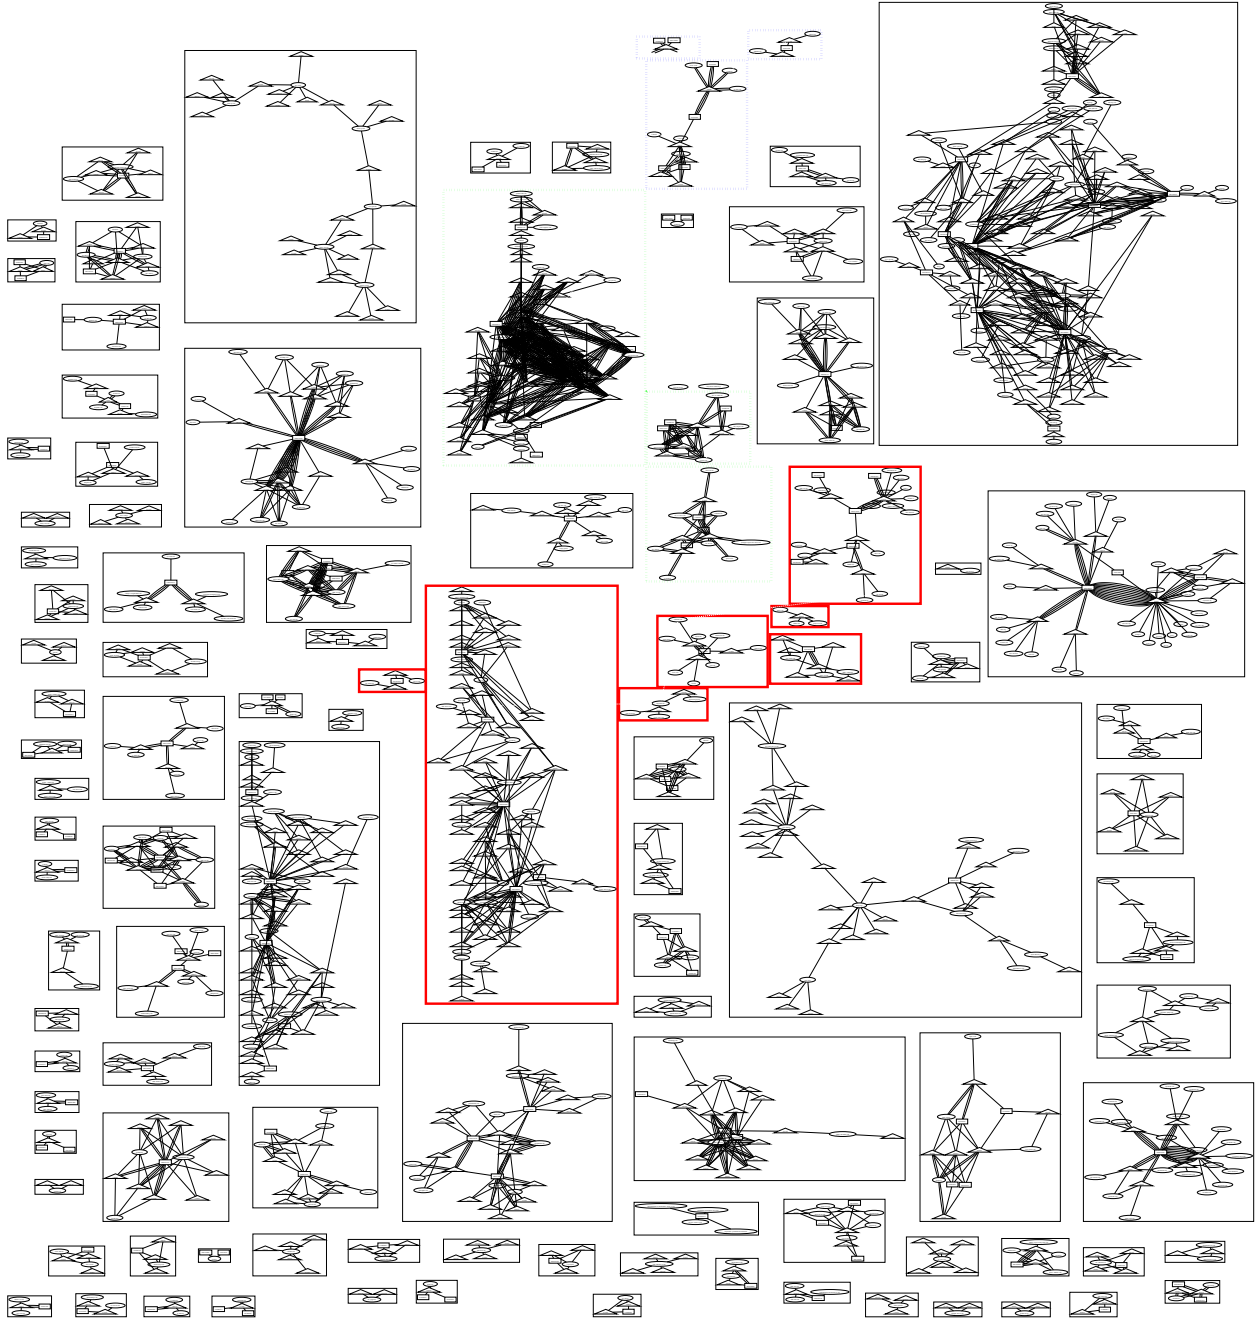


Figure 3: **Clustering of 1,341 nodes representing 481 domains (ovals), 684 bank accounts (triangles), and 176 phone numbers (small rectangles).** Links connect nodes found in the same incident. We obtain 105 clusters containing more than one node, and 26 singletons (omitted from the figure). Additional grouping by malware used allows to link seven of these clusters into two groups (denoted by dashed lines), and additional grouping by strong similarities in WHOIS registrant information allows to group another set of seven clusters (represented by thick lines).

Malware. We further noticed that websites from the second largest connected subgraph were still alive as of Nov. 2009, indicating that the fraud continues unabated. As mentioned earlier, we downloaded the entire contents of each fraudulent website listed in our database. We found that a small number (14) of One Click Fraud websites contained some malware.

Specifically, some of these sites contain a virus named `Trojan.HachiLem`. This is an MS Windows executable file which is posted on these sites as a “mandatory video viewer plug-in.” Once downloaded and executed, this virus automatically collects email addresses and contacts stored within Outlook Express or Becky! (email application popular with Japanese enterprise users), and sends the collected information to a central `hachimitsu-lemon.com` server and, possibly, to another machine as well. The collected information is in turn used to send blackmail emails to victims and notifying them they owe the website registration fees. Interestingly enough `hachimitsu-lemon.com` has not been a valid domain name since, at least, early October 2009, but a report about this virus was posted as recently as Oct. 26th, 2009 on Wan-Cli Zukan, indicating that this virus is still mildly active.

Another group of sites contains a simpler, less intrusive form of malware, which modifies the Windows registry entries to display a pop up window reminding of the registration fee due upon boot-up.

Additional clustering. An interesting feature of the sites hosting the `Trojan.HachiLem` malware is that they all share strong similarities in their WHOIS records. For instance, the technical contact phone number field contains is identical (+81-6-6241-6585). While it is likely a bogus number, this similarity, coupled with the presence of identical malware and overall similar appearance of the different websites involved strongly suggest all sites are operated by the same group. We can thus relate three connected subgraphs to each other, represented by the dashed boxes in Fig. 3.

We further look into WHOIS information details, and notice that a number of registrant entries shared similarities across seemingly different incidents. For example, we found a number of entries with identical (bogus) contact name, email information, or technical contact phone number. By grouping this entries together, we can link seven connected subgraphs, represented by the thick boxes in Fig. 3, to each other.

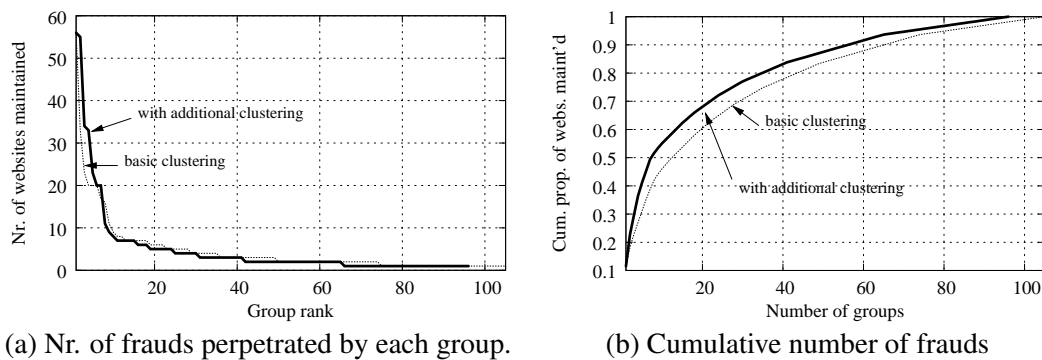


Figure 4: **Frauds distribution over the miscreant groups.** These plots exhibit considerable concentration in fraudulent activities.

Fraud distribution. With the assumption that each connected subgraph denotes a unique criminal organization, we plot, in Fig. 4, (a) the number of frauds perpetrated by each group, and (b) the cumulative number of frauds perpetrated by the most active fraudsters.

In Fig. (a), we rank criminals by the number of frauds they have been committing, and plot the number of One Click Frauds perpetrated by each group as a function of its rank. We provide data for both simple clustering (based only on phone numbers, domains, and bank accounts), as well as additional clustering (further grouping criminals by malware used, and identical WHOIS records). We note that the distribution seems to be following a Zipf law. Specifically, the curve representing a basic clustering fits very well with the function $y = 56/x^{0.76}$; $y = 80/x^{0.90}$ is a good fit for the curve obtained by considering additional clustering. The match with a power law indicates both high concentration (a few people are responsible for most of the frauds), and long tail behavior (many people participate in these scams).

We validate this observation with Fig. 4(b). Fig. 4(b) plots the same data as Fig. 4(a), but in a cumulative fashion. What we observe here, is that even with our conservative clustering strategy, the top 13 miscreant groups are responsible more than 50% of the fraud. Taking into account similarities exhibited by WHOIS records, and malware deployment, we observe that *the top 8 groups are responsible for more than half of the frauds we have collected*. At the same time, there is a large number of groups that do not seem to be involved in more than at most a couple of frauds. Including singletons, 80 out of the 112 groups we extracted using additional clustering, or about 71% of the criminals involved, appear to operate at most two sites. This number may be an overestimate, since we do not claim that we managed to establish all possible links that exist. Yet, it tends to indicate that the miscreant population is a mix of large operators, and of “artisans” running much more limited criminal enterprises.

4.3 Evidence of other illicit activities

Blacklist	Purpose	Nr. hits
cbl.abuseat.org	Open proxies, Spamware	7 (2.55%)
dnsbl.sorbs.net	Spam	22 (8%)
bl.spamcop.net	Spam	4 (1.45%)
zen.spamhaus.org	Spam, Trojans, Open proxies	23 (8.36%)
combined.njabl.org	Spam, Open relays	4 (1.45%)
l2.apews.org	Spam, Spam-friendly	90 (32.73%)
aspews.ext.sorbs.net	Spam	11 (4%)
ix.dnsbl.manitu.net	Spam	4 (1.45%)
Google Safe Browsing (URLs)	Phishing, Malware	0 (0%)
Google Safe Browsing (IPs)	Phishing, Malware	44 (16%)

Table 4: Presence of One Click Fraud domains in various blacklists.

Next, we try to see if domains engaging in One Click Fraud are also supporting other illicit activities. Out of 1608 URLs present in our database, we extract 842 domain names. These domains resolve to 275

unique IP addresses.³

We check the 275 IP addresses against a set of eight blacklisting services that flag IP addresses known for producing large amounts of spam, and/or trojans and malware, and present our results in Table 4. While some of the domains are also used for spam, the vast majority does not appear in any database. In fact, most of the hits we see (in `L2.apews.org`) are coming from sites that resolve to a parked domain IP address. That is, these sites are not active as One Click Frauds anymore, and have been reclaimed by the DNS reseller, which apparently either serves as a spam relay or has been known to be friendly to spammers. Given the results we found earlier about the concentration of frauds in some specific resellers, this outcome is not particularly surprising.

We further verify this hypothesis by checking our entries against the Google Safe Browsing database distributed as part of Firefox 3, updated as of November 15, 2009. On the one hand, none of the URLs we find in our database is listed in the Google Safe Browsing database. On the other hand, when we check IP addresses of the servers that host One Click Frauds against the IP addresses of the servers hosting pages in the Google Safe Browsing database, we see a significant (16%) number of hits. This confirms our finding that, while the websites engaging in One Click Frauds keep a relatively low profile, they are sometimes hosted on very questionable servers, which essentially turn a blind eye to what their customers are doing.

A possible reason for the relatively lack of conclusive evidence that One Click Frauds sites engage in other forms of online crime will become clear in the next section, when we look at the potential profits miscreants can make. Operating a set of One Click Fraud sites is indeed quite a profitable endeavor, and we conjecture that, engaging in other forms of fraud would only have a marginal benefit to the fraudster, while increasing the risk of getting caught.

5 Economic incentives

We next turn to looking into the profitability of One Click Frauds. We start by drafting a very simple economic model, before populating it with the measurements we obtained from our study. The objective is not to obtain very precise estimates of the actual profits that can be made, but mostly to be able to dimension the incentives (or disincentives) to engage in One Click Frauds. We first determine the break-even point in the absence of potential risks for miscreants, before looking at the impact of penalties (prison, fines) on criminals' incentives. All the numbers in this section are current as of November 2009.

5.1 Cost-benefit analysis

Costs. A miscreant interested in setting up a One Click Fraud will need a computer and Internet access. We call this startup cost, C_{init} . C_{init} is independent of the number of frauds carried out, but is actually dependent on time, as Internet access is usually provided as a monthly subscription service. Then, the miscreant will need to purchase a DNS name, possibly joint with a web hosting service. We will call this cost C_{host} . This cost, too, is time-dependent.

³A large number of domains, particularly those reported in 2006–2007, are down, and do not resolve to any IP address anymore.

The website has then to be populated with contents (e.g., pornographic images, dating databases), which will cost C_{cont} . C_{cont} can be a fixed price, or a periodic fee, depending on whether the miscreant is purchasing items in bulk, or as part of a syndicated service.

The criminal has to set up a bank account, which should not be traceable back to its real identity. We will assume the cost of this operation to be C_{bank} . This is a fixed fee. Likewise, an untraceable phone number for call-back from the victim may be purchased to make the fraud more successful. We will denote the cost associated with this purchase C_{phone} . This fee is likely to be time-dependent; for example, the forwarding service discussed in Section 4 is a monthly subscription.

Profits. Profits primarily come from money obtained through user payment P_{user} , conditioned by the number of users n paying the requested amount f for each incident. The number of users increases with time. Secondary profits, P_{resale} , may be reaped from reselling databases of victims contact information to other miscreants.

Utility. As we have discussed before, miscreants may actually set up several frauds, potentially reusing phone numbers and bank accounts across them. Let us denote by S , K , and B the number of frauds set up, phone numbers, and bank accounts purchased by a miscreant, respectively. We assume that identical contents is reused across all frauds operated by the same miscreant, which, based on casual observation, appears to be the case. We further assume that hosting costs are directly proportional to the number of sites operated, which is a very conservative assumption, as resellers tend to provide discounts based on the number of domains purchased. We also assume that the costs of phone and banking services are directly proportional to the number of phone lines and banking accounts purchased. Finally, because quantifying P_{resale} is difficult, and is likely to be small in front of the other profits, we simply consider $P_{resale} \geq 0$. Finally, we assume that each fraud has roughly the same probability of being successful, and thus, that the total number of users to fall for all frauds operated by a miscreant is simply given by multiplying S by n .

With these assumptions, we obtain, for the utility (i.e., the profit or loss miscreants can expect) U , the following inequality, in absence of any risks linked to being caught:

$$U(t) \geq Sn(t)f - C_{init}(t) - S(C_{host}(t) + C_{cont}) - BC_{bank} - KC_{phone} , \quad (1)$$

where t denotes time. Clearly, the policy maker wants $U < 0$. If, on the other hand, $U \gg 0$, there is a high incentive for people to start fraudulent operations.

5.2 Fraud profitability

We next turn to dimensioning each of the parameters in Eqn. (1). We consider $t = 1$ year in all discussions.

We plot the number of occurrences the most commonly requested amounts of money f are observed in our database in Fig. 5. The spike at 45,001-50,000 JPY (roughly USD 500) coincides with the average monthly amount of “pocket money” typical Japanese businessmen are allowed to take from the household

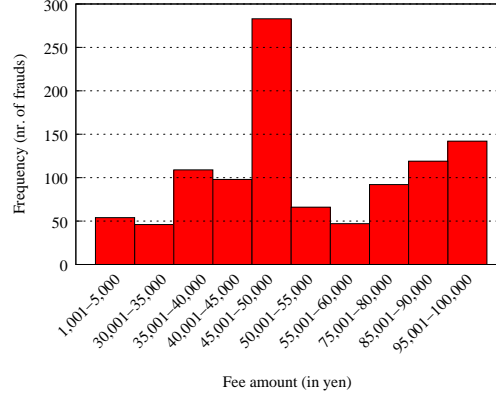


Figure 5: **Ten most common amounts of money requested.**

income (45,600 JPY in 2008, [26]).⁴ The average, over 1,268 records in our database, stands at $f = 60,462$ JPY.

Next, the initial cost C_{init} combines that of a computer, which the fraudster likely already possesses, and of an Internet subscription. If the miscreant does not already own a computer, a basic model, around 28,000 JPY (e.g., an Asus EeePC 900X), would be more than sufficient. Likewise, the miscreant probably already subscribes to an Internet service, but assuming they do not, a modest broadband access is all that they need. As an example, in 2009, Yahoo! BroadBand provides an ADSL “8 Mbps” plan for 3,904 JPY per month [8]. We get $0 \leq C_{init} \leq 74,848$ JPY, where the upper bound is given by assuming a fully dedicated machine and Internet connection for setting up One Click Frauds, and is obtained combining the prices of purchasing a new PC and subscribing for one year to the Internet service.

We evaluate C_{host} using `maido3.com`, the most popular domain reseller and rental server used by fraudsters (see Section 4). The Starter Pack Plan costs 3,675 JPY for an initial setup fee; domain registration and DNS service is free; the plan requires a three months advance payment of 22,050 JPY (monthly fee is 7,350 JPY). We gathered these numbers from our email communication with `maido3.com`. For a one-year subscription, we obtain $C_{host} = 91,875$ JPY.

Fraudulent bank accounts can be obtained for prices going between 30,000 and 50,000 JPY [1] from the black market. As previously mentioned in Section 4, bank accounts are easily forged if created in Internet banks or banks requiring only postal mail for new contracts since there is no physical interaction during the contracting process. Also, it is relatively easy to set up fraudulent accounts by taking advantage of the Japanese writing system. Bank accounts internally use a phonetic alphabet (katakana), different from the characters used for most names and nouns (kanji). It is thus possible to create ambiguous account names. For instance, both the “Baking Club of Shirai City” and “Shiraishi Mitsuko” (a person’s name) are pronounced exactly the same, and thus would have the same account holder information. By registering as a non-profit organization, the fraudster may easily bypass most identity checks, and subsequently set up fraudulent transfers using the individual name instead. Creating an account in this fashion would cost much

⁴In Japan, the wife of the household typically plays the role of “Chief Financial Officer” and dispenses monthly allowances to the husband and children.

less than 50,000 JPY. We can thus confidently assume $C_{bank} \leq 50,000$ JPY.

A phone line is required for miscreants to have the ability to further pressure and blackmail the victim to pay money. Cell phones can be illegally purchased for approximately 35,000 JPY [1, 3], and can be made untraceable by paying the monthly subscription fees of 7,685 JPY/month [3] at a convenience store, or by simply using a prepaid phone. Forwarding and anonymizing services discussed in Section 4 can be purchased for only 840 JPY/month [6]. Assuming all these services are purchased for a year yields an upper bound $C_{phone} \leq 137,300$ JPY.

We make the modestly controversial assumption that $C_{cont} = 0$. Indeed, we suspect that most of the pornographic contents presented in One Click Fraud websites is simply copied and plagiarized from other (non-fraudulent) websites, or obtained through peer-to-peer transfers.

Last, our analysis of Section 4 tells us that, on average, $S \approx 3.7$ domains, while $B \approx 5.2$ accounts, and $K \approx 1.3$ phone lines. These numbers are obtained by looking at the graph G , and dividing, the total number of domains operated (481), bank accounts (684), and phone numbers used (176), by the number of connected subgraphs we found, including singletons (131).

Reusing these numbers in Eqn (1), we obtain a simple linear condition for the fraud to be economically viable, that is, for $U \geq 0$, we need

$$n \geq 3.8 \text{ users,}$$

to fall for each fraud, within one year. This means that, *as long as, for each scam operated, more than four people fall for the scam within a year, the miscreant turns a profit*. In other words, there is an extremely strong incentive for aspiring criminals to engage in One Click Fraud. This strong incentive is not very surprising, given the low amount of skills and equipment needed to set up One Click Frauds. What we find more interesting is how profitable One Click Fraud appears to be.

5.3 Legal aspects

The above incentive assessment ignores the probability of getting caught, and the associated penalties resulting from arrest, which we detail next.

Prosecution probability. Criminals take advantage of social norms and legal loopholes to commit One Click Frauds. First, most victims do not report these crimes. Indeed, to be able to legally charge a suspect, a victim must make a formal complaint to the court. Due to the embarrassment associated with the context of the fraud, many victims do not wish to participate in the accusations that may reveal their identity. Further complicating matters, is the fact that most DNS resellers and web hosting services used, as we have observed in our measurements, are physically situated outside of Japan. The part of the fraud infrastructure hosted in Japan, i.e., the phone line and the bank account, is equally hard to use for investigation. Indeed, police cannot obtain contact and network information from telecommunication networks unless they have an actual *arrest* warrant. Likewise, access to banking accounts is very restricted. Thus, prosecution probability is actually very low.

Date	Location	Damages ($\times 10^6$ JPY)	Victims	Ref.
2/2004 - 4/2005	Osaka	600	10,000+	[32]
8/2004 - 8/2005	Iwate	28	450+	[37]
7/2006 - 4/2007	Saitama	50	700+	[18]
7/2006 - 11/2007	Chiba	300	3,400+	[36]
7/2007 - 8/2008	Yamaguchi	240	3,500+	[10]

Table 5: **Press reports of One Click Fraud arrests.**

Penalties. In addition, sentences are relatively light. Common fraud (e.g., blackmail), in Japan, carries a sentence of up to 10 years of imprisonment. However, One Click Frauds very often do not meet the legal tests necessary for qualifying as “fraud,” as in the vast majority of cases, the victim pays up immediately, and there is no active blackmailing effort from the miscreant. One could argue that they are nothing more than a sophisticated form of panhandling. As a result, the few sentences for cases related to One Click Fraud made public [2] were rather light, with fines ranging from 300,000 JPY to 2,000,000 JPY (\approx USD 3,000 to 20,000) and prison sentences ranging from probation to 2.5 years of jail time.

5.4 Field measurements

We next use field data provided by the Japanese police [14], and compare their findings to ours. Note that the report [14] combines numbers for One Click Frauds and related confidence scams, so the numbers in this section are much more approximate than we would wish. However, they remain a useful starting point to roughly assess the economic impact of these frauds.

Number of frauds per miscreant. First, the police acknowledges, per year, 2,859 cases of frauds leading to 657 arrests. It is very unlikely that the police would inflate the number of unsolved cases in an external publication, and it thus appears that each arrested individual was responsible for an average of 4.4 frauds.⁵ This number is slightly higher than $S = 3.7$ we found earlier, but the difference is hardly surprising, as the probability of getting caught increases with the number of sites operated.

Profits made. The Japanese Police estimates, on average, between 2004 and 2008, that 26 billion JPY per year were swindled in various confidence scams, prominently including One Click Frauds. Dividing by the number of 2,859 cases the police reports, we find the average income per case to be approximately 9 million JPY over a year (\approx 90,000 USD), a quite staggering number, especially considering that each arrested individual was responsible for about four cases, potentially making around 360,000 USD. While we believe that this is likely an over-estimate, due to the unpublished number of unsolved cases, as well as some potential exaggeration in the losses incurred, profits made by miscreants appear sizable.

⁵In Japan, arrests are usually made only when the police is certain with 99% probability or more of the suspect’s guilt. The conviction rate is thus very high.

We can compare this estimate with reports gathered from the press, which we summarize in Table 5. The profits made appear to be very high, which is not overly surprising considering the number of victims involved. Recall from the previous section, that one only needs about four victims to break even. Clearly, with victim numbers in the 450–10,000 range, we can expect very significant profits, which, in Table 5 are estimated between approximately USD 28,000 and 600,000 for each group arrested. Thus, the USD 90,000 – 360,000 estimate acquired from the police reports seems to be in the right order of magnitude.

Ironically, the case reported in [37] names the suspect as a famous IT writer specializing in cybercrime and appearing in multiple TV programs to warn the public of the dangers of One Click Fraud. One cannot help but think that the suspect, being very familiar with the inner workings of One Click Fraud, must have reached a conclusion identical to that of the present paper, that is, that there is a significant economic advantage to engage in such frauds.

We also note that these arrests were likely possible by the sheer magnitude of the fraud. A criminal confining themselves to only a few dozens victims would likely be able to make a reasonable profit, while being extremely hard to catch.

6 Discussion and conclusions

This paper attempts to provide a comprehensive picture of a confidence scam used in Japan, called One Click Fraud. By gathering over 2,000 reports of incidents from vigilante websites, we were able to describe a number of potential vulnerabilities that miscreants use (lax registration checks, resellers turning a blind eye to their customers’ activities), and also showed that the market appears to be quite heavily concentrated. The top eight miscreant groups are responsible for more than half the frauds we uncovered. At the same time, a large number of individuals seem to be participating in frauds of smaller magnitude.

We showed that an important reason for these scams to flourish is the strong economic incentives miscreants have. They can break even as soon as they manage to successfully scam four victims for each fraud, and, on average can make profits in the order of USD 10,000-100,000 or more per year. Prosecution is difficult, since victims are reluctant to come forward, and the penalties that can be meted are relatively light.

Reforming the law to impose harsher penalties is not easy, as proving the mere existence of a crime is cumbersome. On the other hand, prevention, identification, and take-down seems more feasible. Our study shows that identifying perpetrators can greatly benefit from a bird’s eye view of the network. In our discussions with the police, we learned that different cases are usually assigned to separate officers, and that communication could be greatly improved. For instance, one officer may investigate a fraudulent phone number, while another may be investigating an online fraud. If the phone number is used in said fraud, clearly both officers would benefit from sharing the results of their investigations. However, such holistic approaches remain relatively rare in practice.

Future work. We believe that, while we tried to provide as comprehensive a study as possible, we could enrich the data we obtained by looking for connections with other forms of crime in more details. Our conclusions, thus far, are that domains involved in One Click Fraud do not apparently engage in other forms of fraud. However, these results are only based on a relatively simple check of known blacklists. Studying

in more details WHOIS registrations cross-listings, as well as possible connections with more “mainstream” pornographic sites could reveal further insights.

Also, an article issued at the time of this writing [9] shows that One Click Fraud is now using peer-to-peer software as a propagation vector for viruses very similar to `Trojan.HachiLem`. While the propagation method is new, the psychological factors making such scams successful remain the same. Finding how to address these psychological traits is a promising area of research, quite related to ongoing works in security psychology [27].

We have also observed, in the police records that we eventually were not able to use, that 71 out of 359 scams reported were SMS-based. That is, One Click Frauds are also affecting mobile websites. Whether the perpetrators are the same, or different groups, remains unknown, and would warrant further investigation.

Despite being focused on a very specific crime, we believe that some of the methodology we used in this paper (graph-theory based modeling, and clustering techniques, for example) could very well apply to other forms of online frauds. As mentioned in the introduction, One Click Frauds are closely related to the more general body of scareware scams, and some of the techniques we use could equally apply to analyzing scareware markets. More generally, we hope this paper will help facilitate an ambitious research agenda to gather more data from online criminal activities. An improved understanding of the economics of online crime is indeed essential in determining which policies may work to curb the problem.

7 Acknowledgments

This research benefited from many discussions with Kilho Shin, Jens Grossklags, and with our colleagues at Carnegie Mellon CyLab, both at CyLab Japan and in Pittsburgh. The Hyogo Prefecture Police was instrumental in availing to us the case data used in Section 5. Ashwini Rao and Lirida Kercelli helped with some of the DNS black-listing tests. Sachi Iwata-Christin provided additional translation assistance.

References

- [1] Black market yamamoto web. In Japanese. <http://yamashita0721.web.fc2.com>. Last accessed April 16, 2010.
- [2] Koguma-neko teikoku. <http://www.kogumaneko.tk>.
- [3] Manual on how to live anonymously. In Japanese. <http://homepage3.nifty.com/nonu/tokumei1.html>. Last accessed April 16, 2010.
- [4] MeCab: Yet another part-of-speech and morphological analyzer. <http://mecab.sourceforge.net>.
- [5] Ni-channeru (Channel two). <http://www.2ch.net>.
- [6] Symphonet Services Co. In Japanese. <http://symphonet.co.jp>. Last accessed April 16, 2010.
- [7] Wan-cli zukan. <http://zukan.269g.net/>.
- [8] Yahoo! BB ADSL. <http://bbpromo.yahoo.co.jp/adsl/type2/index.html>. Last accessed April 16, 2010.

- [9] BBC News. Porn virus publishes web history of victims on the net. <http://news.bbc.co.uk/2/hi/technology/8622665.stm>. Published April 15, 2010. Last accessed April 16, 2010.
- [10] Chugoku Shimbun. Arrested for one click fraud. August 16, 2008. <http://www.chugoku-np.co.jp/News/Tn200808160061.html>. Last accessed November 21, 2009.
- [11] Electronic Frontier Foundation. Panopticlick, 2009. Available at <https://panopticlick.eff.org/>. Last accessed April 10, 2010.
- [12] 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 14th ACM Conference on Computer and Communications Security (CCS)*, pages 375–388, Alexandria, VA, October 2007.
- [13] C. Herley and D. Florêncio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Proceedings (online) of the Workshop on Economics of Information Security*, June 2009. Available from <http://weis09.infosecon.net/>.
- [14] Japan National Police Agency. Incident report and monetary damages of direct deposit frauds, 2009. Archived at http://megalodon.jp/2010-0223-1659-40/www.npa.go.jp/safetylife/seianki31/1_hurikome.files/Page386.htm. Last accessed April 16, 2010. In Japanese.
- [15] Japanese Information Technology Promotion Agency. Virus detection reports, October 2009. In Japanese. <http://www.ipa.go.jp/security/txt/2009/10outline.html>. Last accessed April 20, 2010.
- [16] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the Conference on Computer and Communications Security (CCS)*, Alexandria, VA, October 2008.
- [17] N. Kato and T. Kawabata. The advance of the fraudulent business scheme through the it and the legal measures of specified business transaction. In *Proceedings of the AIEEJ Media Computing Conference*, 2007. In Japanese. Abstract available online at http://www.jstage.jst.go.jp/browse/aieej/35/0/_contents/-char/ja/. Last accessed April 10, 2010.
- [18] Mainichi Shimbun. Fraud: Suspect arrested again for requiring a registration fee for unauthorized access. March 4, 2007. <http://blog.goo.ne.jp/alladult/e/4a53e012de3c0335a838d99161e9d8bc>. Last accessed April 16, 2010.
- [19] Ministry of Justice, Japan. Act on special provisions to the civil code concerning electronic consumer contracts and electronic acceptance notice. English translation accessible online at <http://www.japaneselawtranslation.go.jp/law/detail/?yo=&ft=2re=01&ky=&page=3>. Last accessed April 10, 2010.
- [20] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *Proceedings of the Second APWG eCrime Researcher’s Summit*, Pittsburgh, PA, October 2007.
- [21] T. Moore and R. Clayton. Evil searching: Compromise and recompromise of internet hosts for phishing. In *13th International Conference on Financial Cryptography and Data Security*, Barbados, February 2009.
- [22] T. Moore, R. Clayton, and R. Anderson. The economics of online crime. *Journal of Economic Perspectives*, 23(3):3–20, Summer 2009.
- [23] T. Moore, R. Clayton, and H. Stern. Temporal correlations between spam and phishing websites. In *2nd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET ’09)*, Boston, MA, April 2009.

- [24] T. Moore and B. Edelman. Measuring the perpetrators and funders of typosquatting. In *Proceedings of the 2010 Financial Cryptography Conference (FC'10)*, Canary Islands, Spain, January 2010.
- [25] N. Provos, P. Mavrommatis, M. Rajab, and F. Monrose. All your iFrames point to us. In *Proceedings of the 17th USENIX Security Symposium*, August 2008.
- [26] Shinsei Financial. Japanese businessmen monthly allowance, 2009. In Japanese. <http://www.shinseifinancial.co.jp/aboutus/questionnaire/kozukai2009/>. Last accessed April 16, 2010.
- [27] F. Stajano and P. Wilson. Understanding scam victims: Seven principles for systems security. Technical Report UCAM-CL-TR-754, Cambridge University, August 2009. To appear in *Communications of the ACM*.
- [28] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your botnet is my botnet: analysis of a botnet takeover. In *Proceedings of ACM CCS'09*, Chicago, IL, October 2009.
- [29] P. Swire. No Cop on the Beat: Underenforcement in E-Commerce and Cybercrime. *Journal on Telecommunications and High Technology Law*, forthcoming 2008.
- [30] R. Thomas and J. Martin. The underground economy: Priceless. *login:*, 31(6):7–16, December 2006.
- [31] Tokyo Metropolitan Police Association. Stop frivolous billing requests! In Japanese. <http://www.anzen.metro.tokyo.jp/net/attention.html>. Last accessed April 15, 2010.
- [32] Tokyo Shimbun. First arrest for one click fraud, 5 individuals arrested, website owner issued a warrant of arrest, victims said to exceed 1,000 and monetary damages of up to 600 million yen. April 13, 2005. <http://www6.big.or.jp/~beyond/akutoku/news/2005/0413-10.html>. Last accessed April 16, 2010.
- [33] Toyo Keizai Inc. *Quarterly Corporate Report Sector Map 2010*. 2009.
- [34] Webhosting.Info. Largest ICANN registrars, November 2009. <http://www.webhosting.info/registrars/top-registrars/global/>.
- [35] 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 (online) of the 9th Workshop on Economics of Information Security*, Cambridge, MA, June 2010.
- [36] Yomiuri Shimbun. Company president arrested for two click fraud. November 28, 2007. <http://www.yomiuri.co.jp/net/security/s-news/20071128nt0c.htm>. Last accessed April 16, 2010.
- [37] Yomiuri Shimbun. IT writer arrested. November 8, 2005. <http://www.yomiuri.co.jp/net/news/20051108nt03.htm>. Last accessed April 16, 2010.
- [38] J. Zhuge, T. Holz, C. Song, J. Guo, X. Han, and W. Zou. Studying malicious websites and the underground economy on the Chinese web. In *Proceedings (online) of the Seventh Workshop on the Economics of Information Security (WEIS)*, Hanover, NH, June 2008.