Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild

Elie Bursztein[†] Borbala Benko[†] Daniel Margolis[†] Tadek Pietraszek[†]
Andy Archer[†] Allan Aquino[†] Andreas Pitsillidis[*] Stefan Savage[*]

[†]Google, Inc [*]University of California, San Diego

{elieb, bbenko, dmargolis, tadek, archer, aaquino}@google.com {apitsill, savage}@cs.ucsd.edu

ABSTRACT

Online accounts are inherently valuable resources—both for the data they contain and the reputation they accrue over time. Unsurprisingly, this value drives criminals to steal, or hijack, such accounts. In this paper we focus on manual account hijacking—account hijacking performed manually by humans instead of botnets. We describe the details of the hijacking workflow: the attack vectors, the exploitation phase, and post-hijacking remediation. Finally we share, as a large online company, which defense strategies we found effective to curb manual hijacking.

1. INTRODUCTION

Online accounts are inherently valuable resources—both for the data they contain and the reputation they accrue over time. With the advent of the cloud, the most intimate details of our lives are contained on remote servers in a single account. This makes account theft, or *account hijacking*, a lucrative monetization vector for miscreants. Criminals leverage millions of hijacked credentials to send spam [5, 10], tap into the social connections of victim's to compromise additional accounts [24], or alternatively liquidate a victim's financial assets using malware such as Zeus or SpyEye [12].

While significant research attention has focused on wide-spread automated hijacking facilitated via botnets, we explore a second class of attacks we refer to as *manual hijacking*. Manual hijackers spend significant non-automated effort on profiling victims and maximizing the profit—or damage—they can extract from a single credential. In contrast to automated hijacking, manual hijacking is exceedingly rare. We observe an average of 9 incidents *per million* Google users per day. However, the damage manual hijackers incur is far more severe and distressing to users [22] and can result in significant financial loss [4]. These needle-in-a-haystack attacks are very challenging and represent an ongoing threat to Internet users.

IMC'14, November 5–7, 2014, Vancouver, BC, Canada.

ACM 978-1-4503-3213-2/14/11 ...$15.00.
http://dx.doi.org/10.1145/2663716.2663749.

In this paper, we explore the manual hijacking lifecycle as gleaned from incidents that occurred at Google between 2011–2014. Our study consists of three components: we explore how criminals acquire a victim's credentials (Section 4); examine how criminals monetize account credentials (Section 5); and finally highlight the process Google used to ultimately restore control back to the victim (Section 6). Based on our findings, we offer a set of best practices for others to defend against manual hijacking and discuss many of the corner cases that exacerbate the problem (Section 8).

In particular, we link manual hijacking with phishing and provide evidence supporting the hypothesis that phishing is the main way manual hijackers steal user credentials. We found that phishing requests target victims' email (35%) and banking institutions (21%) accounts, as well as their app stores and social networking credentials. Of the hijacking case samples we analyzed, we found that most of the hijackers appear to originate from five main countries: China, Ivory Coast, Malaysia, Nigeria, and South Africa.

Injecting decoy credentials in phishing pages targeting Google users reveals that criminals' response time is surprisingly fast. We found that criminals attempted to access 20% of the accounts within 30 minutes. Looking at real hijacking cases, we observed that, once logged in, manual hijackers profile the victim's account and spend an average of 3 minutes to assess the value of the account before exploiting it or abandoning the process. This step entails searching through the victim's email history for banking details or messages that the victim had previously flagged as important. We also see attackers scanning through email contacts which are then either solicited for funds or targeted with a salvo of targeted phishing emails.

Restoring a victim's account access is a non-trivial problem. We found that SMS is the most reliable out-of-band channel, where users that provided a phone number recover their account 81% of the time. Providing a secondary email address is also fruitful, succeeding 75% of the time. Absent these two mechanisms, we must rely on secret questions or manual review where our success rate falls to 14%.

Ethics We acknowledge that most of the results in the paper depend on proprietary data. As a result, we attempted to clearly explain how we reached our conclusions while balancing our need to protect the security and privacy of our users.

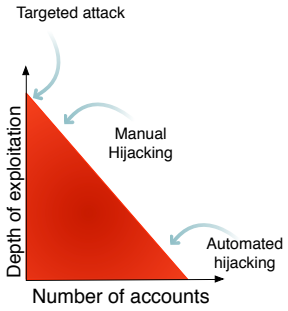

Figure 1: Hijacking trade-off: the depth of exploitation versus the number of accounts compromised.

2. HIJACKER TAXONOMY

Account hijacking stems from a multitude of attack vectors and underlying incentives. At Google, we categorize hijacking campaigns based on the depth of exploitation (e.g. the damage incurred to each victim) and the number of accounts impacted, as depicted in Figure 1. We currently observe three classes of attacks:

Automated hijacking: An automated hijacking attack attempts to compromise large quantities of accounts using botnets or other professional spamming infrastructure [11]. This type of attack is carried out entirely by automated tools that monetize the most common resources across all compromised accounts (e.g. spamming via a victim's email).

Targeted attacks: Targeted attacks [23] include industrial espionage and state-sponsored break-ins. In our experience, these attacks are carried out by highly sophisticated parties who have the resources to extensively profile targets and launch tailored attacks. These include using dedicated 0-day exploits [20], less expensive readily-available exploits, or highly targeted phishing campaigns.

Manual hijacking: Manual hijacking consist of attack that opportunistically select victims with the intent of monetizing the victim's contacts or personal data; any sufficiently lucrative credential will suffice. These attacks are carried manually rather than automatically. As a result, we observe orders of magnitude fewer instances compared to automated hijacking. However, criminals heavily abuse victim's information, making such attacks highly distressing to the impacted parties [22].

In this paper, we choose to focus on manual hijacking. While automated hijacking impacts the most users, manual hijacking incurs significantly more damage to the victims involved and represents an ongoing pertinent threat to Google. We restrict our analysis of manual hijacking to cases where the attacker does not know the victim personally. Similarly, we also exclude hijacking instances where an attacker leverages physical access to a victim's devices. These personalized attacks represent their own area of focus [2, 17], and while they are of serious importance, combating such threats relies on individual insights that benefit little from large-scale measurement.

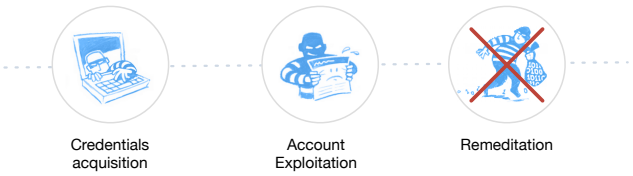

Figure 2: Overview of the account hijacking cycle.

3. METHODOLOGY

We provide a high-level workflow of the account hijacking lifecycle, depicted in Figure 2, along with our methodology for studying each component.

In total, we rely on 14 distinct datasets collected between 2011–2014, outlined in Table 1. Each dataset originates from various system logs that we aggregate via map-reduce computation. For privacy and storage reasons, Google sanitizes or entirely erases many authentication-related logs within a short time window. Consequently, some of our datasets are drawn over a period of only a few weeks despite the three year span of our study. Similarly, as many of our datasets rely on user reports, we are forced to manually curate data points sampled from a much larger, noisy source to have precise ground truth. This problem originates from the fact that both computers and humans alike are imprecise at distinguishing phishing or similar attacks from scams and other bulk spam.

To offer some perspective on sample size, we observed an average of 9 case of manual hijacking per million active users per day in 2012 and 2013. This number includes those our abuse detection mechanisms detected before the hijacker fully exploits the account, and those for which the victim had to file a recovery claim. An account is considered active if it has been accessed within the past 30 days. In term of phishing page, SafeBrowsing detected between 16,000 and 25,000 page per week on the Internet for the period 2012–2013 [9] [Dataset 2 and 3 in Table 1].

Credential acquisition [Section 4]: Account hijacking begins when a hijacker steals a user's credentials. This can occur in a multitude of ways: phishing a user's credentials; installing malware on the victim's machine to steal credentials; or guessing a victim's password. Overall we find corroborating evidence throughout our measurements that phishing is likely to be the main way hijackers compromise user accounts. We argue that phishing is the attack vector of choice for manual hijackers as it is easier and cheaper to perform than other mean to compromise accounts: e.g using 0-day exploits to install malwares. This is why for this part of the study dedicated to the attack vectors used by manual hijackers we decided to focus on phishing emails and phishing pages targeting Google users.

Datasets used: In this section we use a manually curated sample of 100 phishing emails selected from a random sample of 5000 emails reported by our users to understand how email phishing is structured and what type of accounts are targeted[1] [Dataset 1 in Table 1]. Next we use a manually reviewed random sample of 100 phishing pages detected by the SafeBrowsing anti-phishing pipeline [26] while indexing the web for our search engine to understand what type of accounts are phished [Dataset 2 in Table 1].

[1]we only have access to the emails the users manually reported as spam and phishing.

Id	Data type	Samples	Date	Used in section
1	Phishing emails	100	Jan 2014	4.1
2	Phishing pages detected by SafeBrowsing	100	Jan 2014	4.1
3	Google Forms taken down for phishing	100	Nov 2012	4.2
4	Decoy credentials injected in phishing pages	200	Nov 2012	5.1
5	Login attempts from IPs belonging to hijackers	300 IPs/day	Nov 2012	5.1
6	Keywords searched by hijackers	Top 10	Jan 2014	5.2
7	High-confidence hijhacked accounts	575	Nov 2012	5.2
8	Mail sent from hijhacked accounts from dataset 7 and reported as spam by users	200	Nov 2012	5.3
9	Hijacked account contacts and active users random sample hijacking rate	3000 / 3000	Nov 2011	5.3
10	High-confidence hijhacked accounts	600	Oct 2011	5.4
11	Hijacked accounts recovery cases	5000	Nov 2012	6.2
12	Accounts recovery	1 month	Feb 2013	6.3
13	IP used by hijackers	3000 hijacking cases	Jan 2014	7
14	Phone numbers used by hijackers	300	Mid-2012	7

Table 1: List of dataset used throughout this study.

The manual review was used to identify what type of account are targeted by phishers (e.g Bank, Social Network). Finally we look at the HTTP logs of a random sample of 100 Google Drive forms that were used as phishing pages before we detected and took them down [Dataset 3 in Table 1]. This dataset is used to understand the conversion rate of phishing pages, the victims targeted, and the methods for luring victims to the pages.

Account exploitation [Section 5]: If the hijacker is able to log into the victim's account, then a second phase begins that we call the exploitation phase. We observe that this phase consists of two distinct parts: The *account profiling* part where the hijacker decides on a concrete exploitation plan and the *exploitation* itself.

The existence of this profiling phase is one of the most surprising insights we gained by fighting manual hijackers. Instead of blindly exploiting every account, hijackers take on average 3 minutes to assess the value of the account before deciding to proceed. This profiling phase and the fact that hijackers decide to not exploit certain accounts is consistent with our claim that this type of hijacking is manual rather that automated. Similarly, we observe that most manual hijackers use semi-personalized scams as the main exploitation vector, e.g. to trick the victim's contacts into transferring money to the hijacker. We have also observed other more opportunistic approaches such as holding the account for ransom. We note that this type of exploitation is very different from the one we observe for automated hijacking, which mainly focuses on abusing the hijacked account's good reputation to send email spam that will bypass spam filters.

Datasets used: In this section we start by using a dataset of 200 fake credentials injected into various phishing pages targeting Google to measure a hijacker's responsiveness [Dataset 4 in Table 1]. Those pages were detected while indexing the web by the SafeBrowsing anti-phishing pipeline. Then we look at 200 IPs that were used to access stolen accounts to understand hijackers' access patterns [Dataset 5 in Table 1]. Next we analyze a set of 575 accounts that were hijacked to characterize how hijackers are exploiting accounts [Dataset 6 in Table 1]. Those accounts were selected based on their account recovery claims, which clearly indicate that they were manually hijacked.

Next we look at 200 phishing emails sent by those accounts during the hijacking period to understand what type of attack the hijackers conducted [Dataset 7 in Table 1]. We also examine the likelihood that a victim's contacts will become compromised compared to a random sample of users, using a sample of 3,000 contacts and 3,000 random accounts [Dataset 8 in Table 1]. Our analysis confirms that a contact's victims are heavily targeted by hijackers as their hijacking rate is 36 times higher than our random sample. Finally we study how hijackers' tactics to retain control of the hijacked accounts evolved over time by comparing the tactics used on our 2012 dataset with the tactics used on a set of hijacked accounts from 2011 [Dataset 9 in Table 1]. The accounts used in datasets 8 and 9 have no overlap.

Hijacking remediation [Section 6]: This last phase of the cycle starts when the exploitation is over or is interrupted by our defense mechanisms. The goal of this phase is to give the account back to its real owner and revert all the account changes made by the hijacker.

Datasets used: We analyze a set of 5000 hijacked accounts that were successfully recovered to estimate how long hijackers retained control of each account [Dataset 11 in Table 1]. We note that there is no overlap between dataset 7 and dataset 11. We next examine a month's worth of account recovery claims (Feb. 2013) to determine the most successful method for recovering account ownership [Dataset 12 in Table 1].

Hijacking attribution[Section 7] Stepping back from the manual hijacking workflow, we shed light on the origin of hijackers. Of the hijacking case samples we analyzed, we found that most of the hijackers appears to originate from five main countries: China, Ivory Coast, Malaysia, Nigeria and South Africa.

Datasets used: Our analysis relies on the geolocation of IPs used to access 3000 hijacked accounts selected at random in January 2014. Independently, we examine the country code of a set of 300 phones that hijackers used in a attempt to lock out their victims by turning on two-factor authentication for hijacked accounts in 2012. We lack more recent data on phone numbers as hijackers abandoned this tactic after realizing it was unsuccessful.[Dataset 13 and 14 in Table 1].

Account type	Phishing emails	Phishing pages
Mail	35	27
Bank	21	25
App Store	16	17
Social network	14	15
Other	14	15

Table 2: Number of phishing emails and pages targeting a specific type of account.

4. ATTACK VECTORS

Manual hijacking begins with acquiring a victim's credentials.

Phishing is a natural attack vector of choice for manual hijackers. It requires less infrastructure than operating a botnet and allows attackers to target specific victims, especially an existing victim's contacts.

Although phishing as an attack mechanism is well studied [1, 6, 16, 27, 13], little is known about its victims or effectiveness.

To understand what types of accounts are targeted by phishers and the relation of phishing to manual hijacking, we looked at what type of accounts are targeted by email-based phishing and website-based phishing. We start by analyzing email-based phishing attacks, in which phishers send email that pretends to come from a legitimate source and requests user credentials under a false pretext, such as impending account deactivation. We then turn to web-based phishing, in which an attacker sets up a webpage which looks like a traditional sign-in page for a web service and attempts to get users to input their login credentials. Finally, to understand phishing effectiveness and how users are lured to phishing pages, we look at the traffic of a random sample of phishing pages that were inadvertently hosted on Google Drive.

4.1 Phishing targets

We generated a manually curated sample of the phishing emails from January, 2014 [Dataset 1 in Table 1] This dataset was constructed by first extracting a random sample of 5,000 phishing emails and then manually reviewing the sample to find 100 emails that explicitly phish for users credentials or point to phishing webpages. This manual curation is necessary to reliably distinguish spam emails from phishing emails.

We found that *62* of those emails contain URLs, which were likely to point to phishing pages designed to impersonate well known site login pages in order to trick users into submitting their credentials. The remaining *38* did not contain URLs and instead asked users to reply to the email with their credentials. By manually categorizing the type of account each phishing email targets, we found that email account credentials were by far the most popular target, followed by bank credentials (Table 2). The result of a similar analysis for phishing pages reported on the same table reveals a mirrored distribution in terms of which type of credentials were phished for. This analysis was conducted by manually reviewing 100 phishing pages detected while indexing the web [26].

The target distribution consistency between our two distinct datasets suggests that phishing is mainly used by the group of attackers that focuses on acquiring email and financial institution credentials.

This focus is consistent with manual hijackers priorities, as this group primarily monetize victim's accounts by running emails and financial scams as detailed in section 5.

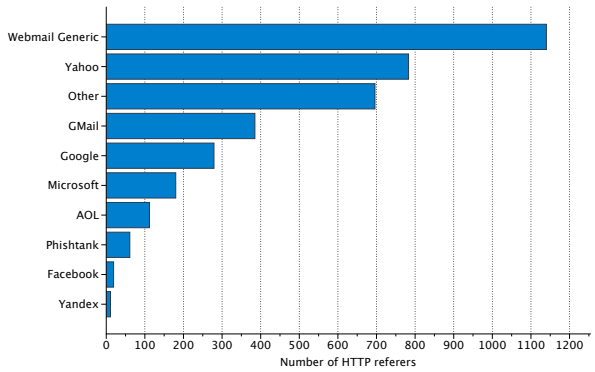

Figure 3: Non-Blank HTTP referrers breakdown.

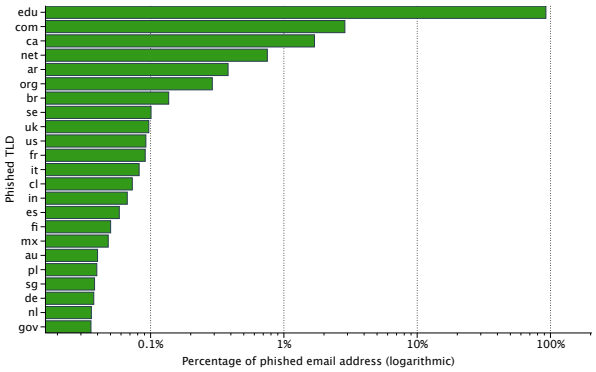

Figure 4: Phished email TLD breakdown

4.2 Phishing effectiveness

Using a separate dataset, we analyzed traffic to a random sample of 100 Google Forms that Google later flagged as phishing and disabled for abuse [Dataset 3 in Table 1]. We used this dataset to understand how users were lured to phishing pages, what domains were targeted most often by phishers, how many credentials hijackers could collect per hour/page, and what success rates phishing pages had. We caution that the behaviors we observed on Google Forms may not be representative of all phishing pages, but no other studies exist on the topic to offer comparison.

Origin of Phished Users: We looked at the HTTP referer of the requests made to phishing pages to understand how users are lured onto those pages. We were surprised to observe that above 99% of those referrers were blank. We hypothesize that it is due to the fact that most traffic toward phishing pages is driven via emails as this source of traffic don't have HTTP referer. Traffic from desktop email clients don't have HTTP referer as it come from an application, and major webmails, including Gmail, ensure the HTTP referer is not set by opening links in a new tab. The hypothesis that most victims are lured via emails is further supported by the fact that most of the remaining 1% of visitors arrived from various webmail providers (Figure 3). The Gmail referrers oddity can be traced back to one of our old HTML frontends used by legacy phones.

The TLDs of the email addresses phished, reported in Figure 4, reveal that the vast majority ($> 99\%$) of the emails address phished come from .edu domains.

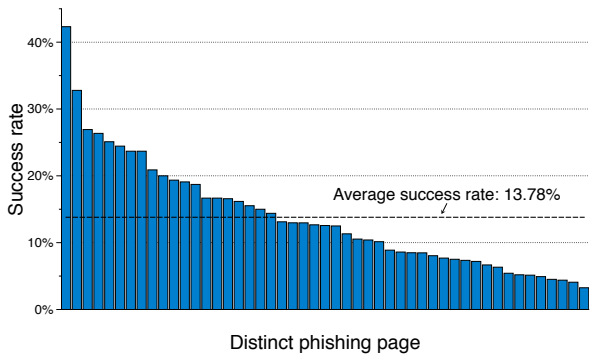

Figure 5: Credential submission rate for phishing pages. Ratio between the number of submitted credential and the number of page views

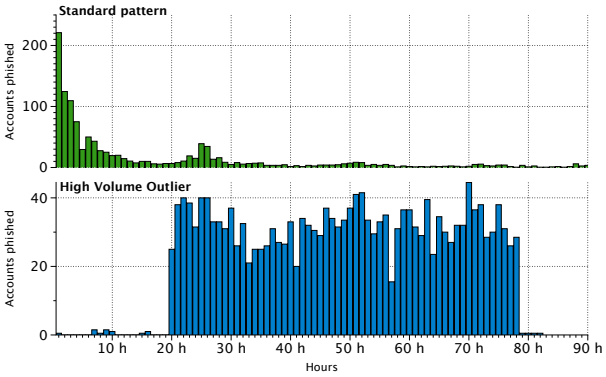

Figure 6: Average number of submitted credentials over time.

A possible explanation, behind the overwhelming prevalence of .edu emails address is that self-hosted emails, e.g by universities, have less robust spam filters than large mail providers and social networks. This explanation is supported by previous work [14] that observed that the spam delivery rate was 10 times higher for webmail protected by commodity spam filtering than by large mail provider such as Gmail, Yahoo, and Hotmail. This explanation is also consistent with our hypothesis that most phishing victims are lured to phishing pages via email.

Phishing Submission Rates: Using the submission rate of visitors to phishing forms, we can estimate the success rate of phishing campaigns. The phishing page success rate is computed by dividing the number of page submissions (HTTP POST requests) by the number of page views (HTTP GET requests). We observe that 13.7% of visitors complete the form (which we assume indicates users submitted accurate credentials); much higher than we anticipated. Broken down by individual page,s we observe a huge variance in success rate, with the highest page having a 45% success rate and the lowest only 3% (Figure 5).

We visually inspected a sample of pages and found that those with low submission rates were very poorly executed and contained only a form asking for a username and password.

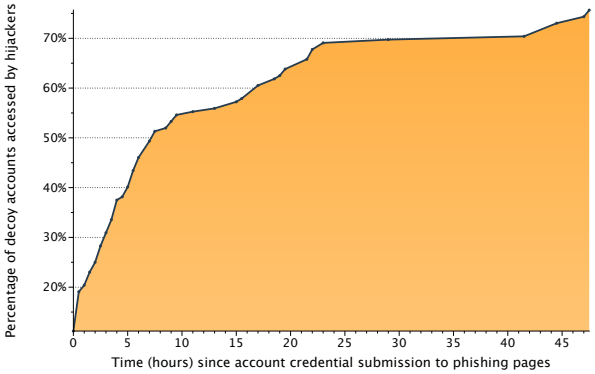

Figure 7: Speed of compromised account access.

Arrival Rate of Phishing Victims: Finally we looked at the average hourly volume of credential submissions for a phishing page, calculated from the time when the page was first visited (first entry in the HTTP request log) until it was taken down (Figure 6) as determined by our logs. The pattern exhibits a clear decay, from the moment the webpage receives its first visitors until it is taken down. This pattern is consistent with a mass mailed email, with clicks centered around the initial delivery time.

We note that while analyzing our phishing page set, we found a single outlier (Figure 6, bottom) that received a huge number of submissions after a step function following a gentle diurnal pattern through several days. This is a prime example of a successful large-scale phishing campaign—starting when the first email went out and ending abruptly when the phishing page was taken down. The initial 15 hour quiet period is, upon a closer investigation, best explained by the attackers testing the page themselves before launching the campaign.

Overall our analysis supports our hypothesis that phishing is a key vector of attack used by manual hijackers and that email is the primary vector by which victims are phished or lured to phishing pages.

5. ACCOUNT EXPLOITATION

Once hijackers obtain access to a victim's login credentials, we observe a multitude of monetization vectors. We find that criminal activities are well-structured, efficient, and savvy at taking advantage of human psychology. A typical hijacker's method adheres to the following playbook: access the account, assess its value, exploit it, and make efforts to delay account recovery in order to increase the chances of successful exploitation. We present our findings about each of these steps.

5.1 Logging Into Accounts

Rapid response time: In order to measure the hijacker response time—the delta between a victim submitting an account's login credentials to a hijacker and the account being accessed by the hijacker—we manually submitted 200 fake credentials into a random sample of 200 phishing pages that explicitly ask for Google credentials (each credential was submitted to exactly one phishing page)[Dataset 4 in Table 1]. We recorded the time when each credential was submitted to a phishing page, and used our logs to observe when the hijacker first attempted to access each account. We found that 20% of the decoy accounts were accessed within 30 minutes of credential submission, and 50% within 7 hours.

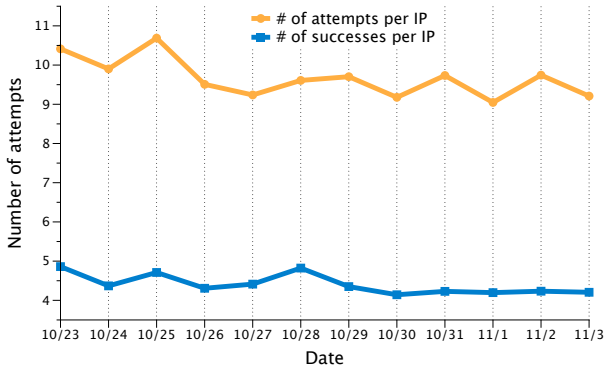

Figure 8: Average hijacker activity per IP.

Figure 7 describes this relationship in more detail. Although not all of the decoy accounts were accessed, possibly due to the suspension of either the phishing website or the email account used by the hijacker to collect credentials [19], the speed with which decoy accounts were accessed was astonishing. Our findings suggest that in order to prevent a hijacked account's exploitation, the reaction time to the credential compromise needs to be even faster than previously thought [18]. Hijhackers' extreme reactivity emphasizes the need to perform an accurate, real-time login risk analysis (Sec 8.2) at the time of log-in in order to detect hijacking attempts since it is unreasonable to expect both the user and the service provider to reliably react within a time window as short as 30 minutes. How we implemented such real time hijacking detection system is discussed in Section 8.

Efforts to Blend in With Organic Traffic: Our next surprising finding is the systematic effort the manual hijackers make in order to avoid detection. Studying the log-in activity of a random sample of 300 IPs used by hijackers, selected daily over a period of two weeks in October/November 2012, reveals that on average, the hijackers attempted to access only *9.6* distinct accounts from each IP, which makes their activity extremely difficult to distinguish from organic traffic [Dataset 5 in Table 1].

The average number of account access attempts per IP is consistently under 10 during the entire two week period studied (Figure 8), suggesting that the manual hijackers may have established guidelines to avoid detection. We observe that hijackers have the correct password for an account 75% of the time (including retries with trivial variants).

5.2 Account Value Assessment

A surprising finding is that hijackers spend 3 minutes on average assessing the value of the account and will not attempt to exploit accounts that they deem not valuable enough. This systematic assessment phase and the fact that certain accounts are not exploited suggest that manual hijackers are "professional" and follow a well established playbook designed to maximize profits. The assessment phase ends when the hijacker performs an active action (e.g sending an email) or logs out of the account. The remaining results in this section are based on a study of 575 randomly sampled accounts that were compromised by manual hijackers [Dataset 7 in Table 1].

We analyze which features of Gmail were accessed by hijackers and find that the most common way for them to assess an account's value is via the Gmail search feature. We also observe that beside familiarizing themselves with the contents of the victim's account via the search terms, hijackers often also open content folders of special significance, such as *Starred* (viewed by 16% of hijackers) *Drafts* (11% of hijackers), *Sent Mail* (5% of hijackers), or *Trash* (less than 1%). The hijackers also look at the user contacts to estimate the number of potential scam and phishing victims.

To better understand how the search feature is abused, we set up a temporary experiment that collected and analyzed the search terms hijackers used for exploring the contents of the victim's mailbox. We found out that hijackers mainly look for financial data (including the victim's financial status and images of signatures to be used for future impersonation), linked account credentials (e.g., usernames and passwords for the victim's other accounts), and personal material that might be sold or used for blackmail (e.g., adult pictures). Table 3 summarizes the frequency of hijackers' top 10 search keywords [Dataset 6 in Table 1]. We find that searches are overwhelmingly for financial data as opposed to other account credentials or content.

The fact that some searches were performed in Spanish and Chinese suggests that certain hijackers target specific regions the world. This observation is consistent with our hijacking attribution analysis (section 7), which shows that besides Africa, hijackers seem to come from Asia (China, Malaysia) and South America (Venezula).

We believe few hijackers look for account credentials because most websites will not send them in clear and instead favor a password reset process. We note that searches for credentials for e-commerce and gaming sites are most prominent among the account searches. Confirming how any of these searches yield an actual financial attack or blackmail is beyond the capabilities of our datasets.

Although, as discussed in Section 8, detection of account hijacking at the initial log-in stage is the most effective way to prevent harm due to hijacking, the systematic approach hijackers use to assess the value of accounts opens another possibility for rapid account hijacking detection. Namely, an approach that models manual hijacker initial activity on hijacked accounts and compares a logged-in user's activity to this model in order to flag those that exhibit excessive similarity to hijacker activity.

5.3 Account's contacts exploitation

If the brief account value exploration using the above techniques yields promising results, the hijackers spend an additional 15 to 20 minutes per account sifting through emails and finding ways to monetize the account. We now discuss how hijackers use this time to exploit the victim's contacts through scams and phishing.

Scam Schemes: Among the exploitations that we can observe, the most common one employed by the manual hijackers is to scam the contacts of the hijacked account owner. A typical scam consists of an email describing a reasonably credible story of how the account owner got into a difficult situation and a plea for money to help get out of the situation. We observe that scams are semi-personalized, e.g take into account the victim gender and location, appeal to human emotions, and systematically exploit known psychological principles [3] to maximize their success rate as discussed below.

Finance		Account		Content	
wire transfer	14.4%	password	.6%	jpg	.2%
bank	11.9%	amazon	.4%	mov	.2%
transfer	6.2%	dropbox	.1%	mp4	.2%
bank transfer	5.2%	paypal	.3%	3gp	.1%
wire	4.7%	match	.1%	passport	.1%
transferencia	4.6%	ftp	.1%	sex	.1%
investment	3.4%	facebook	.1%	filename:(jpg or jpeg or png)	.1%
banco	3.0%	skype	.1%	is:starred	.1%
账单 (account statement)	1.9%	username	.1%	zip	.1%

Table 3: Top search terms used by hijackers for the various type of information searched.

A prominent example of such a scam scheme is the *Mugged-In-"City"* scheme, in which, according to the story told by the hijacker, the account owner was robbed during his or her trip to a faraway city and is asking for a temporary emergency loan to help settle bills and get out of the situation:

...My family and I came down here to West Midlands, UK for a short vacation...

...we were mugged last night in an alley by a gang of thugs on our way back from shopping, one of them had a knife poking my neck for almost two minutes and everything we had on us including my cell phone, credit cards were all stolen, quite honestly it was beyond a dreadful experience...

...I'm urgently in need of some money to pay for my hotel bills and my flight ticket home, will payback as soon as i get back home...

Besides the *Mugged-In-"City"* scheme, there is a large variety of scam schemes with different stories that appeal to the same human emotions and exploit the same psychological principles. For example, the following excerpt presents an example where the reason for the plea is a sick relative with a sudden need for a medical procedure:

Sorry to bother you with this. I am presently in Spain with my ill Cousin. She's suffering from a kidney disease and must undergo Kidney Transplant to save her life. ...

Over time, we realized that scam schemes share a set of core principles that we were then able to formalize as follow:

- A story with credible details to limit the victim suspicion.
- Words or phrases that evoke sympathy and aim to persuade. E.g apologizing and providing distressing details such as "had a knife poking my neck for almost two minutes".
- An appearance of limited financial risk for the plea recipient, as financial requests are typically requests for a loan with concrete promises of speedy repayment.
- Language that discourages the plea recipient from trying to verify the story by contacting the victim through another means of communication, often through claims that the victim's phone was stolen.
- An untraceable, fast and hard-to-revoke yet safe-looking money transfer mechanism. The payment also needs to be picked up anywhere and somewhat anonymously as the scammers might not be from the country they claim the victim what mugged in.

  For example, a request to transfer money to the victim by name via Western Union/MoneyGram helps the request appear credible, and enables the recipient to reclaim money anywhere in the world. [28].

Thus, despite the appearance of simplicity, in reality, the scam emails are well-formed and thought-out in a way to maximize efficiency by preying on known human physiological principles [3]. The potentially high level of financial and psychological distress due to scam emails [22], explains why detecting and filtering out such emails is a high priority for us.

The volume of scam emails sent from hijacked accounts is relatively low and supports our earlier claim that they are likely a result of manual work. For 65% of the victims, the hijacker sends at most five messages, each with a high number of recipients We further analyze the 6% of the cases in which the number of recipients is less than 10, and conclude that those tend to contain a more customized message.

We hypothesize that although message customization comes at an additional (time) cost for the hijackers, certain hijacker groups anticipate higher returns on those messages, thus making the trade-off worthwhile.

Phishing: Another common exploitation pattern is to use the hijacked account for email-based phishing. We reached this conclusion by observing 2 key facts: First the number of outgoing emails for our sample accounts was only 25% higher on average on the suspected day of the hijacking compared with the previous day's volume, the number of distinct recipients of that traffic was 630% higher than on the previous day. This suggests that hijackers generate traffic by sending the same message to many recipients.

Secondly the traffic generated on the suspected day of the hijacking received 39% more user spam/phishing reports than the previous day's traffic. This sharp increase of reports confirms that malicious emails were sent from stolen account and corroborate our hypothesis that manual hijackers attempt to scam and phish the victim's contacts. To further confirm this, we performed a manual analysis of 200 randomly selected phishing messages sent from one of the hijacked accounts on the day of the suspected hijacking [Dataset 8 in Table 1]. This analysis reveals that 35% of them were phishing messages and the remaining 65% were attempts to scam the recipients, as described earlier in this section. These results also support our hypothesis that hijacked accounts are used for phishing.

We were able to confirm that hijackers favor the use of the victim's contacts to select their next set of phishing victims and to improve the attack's chances of success using a second measurement. For this second analysis we measured the number of account hijacked for two account sets: The first set consisted of random sample of 3000 accounts that belonged to the contact list of accounts that were hijacked [Dataset 9 in Table 1]. The second set consisted of 3000 randomly sampled 7-day active users[Dataset 9 in Table 1] We found that the number of manual hijackings over the next 60 days among the users of the first set was 36 times higher than among the users of the second set. This supports our hypothesis that hijackers use victim's contacts as their next targets.

These results are consistent with our first experiment observations and support the hypothesis that phishing is manual hijackers' main way to compromised accounts. We hypothesize that the rationale behind phishing people in the social circle of previous victims is that hijackers try to leverage the sometimes more lenient and trusting treatment given by automated mail classifiers and humans to emails originating from a person's regular contact. This observation is consistent with previous studies on the subject [24].

5.4 Account Retention Tactics

In order for the scam attempts to succeed, the hijacker needs to control the account for a sufficiently long period of time. For example, the *Mugged-In-"City"* scams typically require at least two rounds of emails – one with the call for help and one with the money transfer details, meaning that even the shortest process may take one or two days, depending on the responsiveness of the contact. In this section, we present our analysis of the account-level tactics deployed by the hijackers in order to increase the chances of their scam's success. This longitudinal study was performed by comparing the retention tactics used for 600 high confidence hijacking cases from October 2011 [Dataset 10 in Table 1] with the retention tactics used for the 575 high confidence hijacking cases from November 2012 that are used in throughout this section [Dataset 7 in Table 1]. The most popular tactics during this period were: *locking out the victim from his account* and *delaying account recovery*, *minimizing the chances of account hijacking discovery*, and *redirecting future communications from the plea recipients to a "doppelganger" account*. These tactics are often combined with each other or alternated between.

Locking out the Victim and Delaying Recovery: The intent behind locking out the account owner is to separate him from his contacts. Most people do not have a copy of their contact list offline, so by losing access to the account the victim also loses the ability to warn his contacts about the scam. The lockout can be as simple as changing the account's password. Most victims are not experts in manual hijacking, so even if they realize that they cannot log in to their account they may not conclude that their contacts are at risk of being scammed. As a result, they are unlikely to warn their contacts of a potential scam via an alternate channel. Locking out the victim is, however, a double-edged sword. While it helps the hijacker in delaying the victim's counter-actions, it also provides a very clear signal for abuse detection and account recovery purposes.

Besides locking out the user, the hijackers also often invest effort in delaying account recovery by changing the account's recovery options, including the secondary email, recovery phone number, and secret question. Finally, they often delete the user's emails and contact lists, so even after the recovery, the victim cannot easily warn his contacts.

Acting in the Shadow: Hijackers sometimes find it more beneficial to remain unnoticed and carry out the scam while the victim remains in control of the account. However, operating in the background is risky for the hijacker: his direct actions (e.g., marking a mail as read) as well as the service provider's countermeasures (e.g., suspicious activity notifications) can easily reveal his presence. In order to remain undetected, hijackers try to separate their emails from the victim's communications. A common tactic for doing so is to set up an email filter and redirect all hijacker-initiated communication to the Trash or to the Spam folder.

In a sample from November 2012, ,15% of the accounts had hijacker-initiated email forwarding rules and 26% had a hijacker-configured Reply-To address.

"Doppelganger" Accounts: Given that the exploitation window can close rapidly, it is advantageous for the hijacker to divert the communication to a separate email account that he owns. That way the hijacker has all the time in the world to scam its victim. In order to accomplish this, the hijacker creates and uses a duplicate ("*doppelganger*") email account that looks reasonably similar from the point of view of the victims. The exact choice of the account name depends on the available usernames and on the personal taste of the hijacker. Some hijackers prefer to set up the doppelganger with the same provider as the victim's account, and introduce a difficult-to-detect typo to the username. Other hijackers prefer to host the doppelganger with a different email provider, preferably, but not necessarily, with a similar-looking domain name [15]. For example *johndoe@example.com* is a doppelganger account for *johndoe@gmail.com* that retain the same username but use a different mail provider.

We observe that the methods of choice for attackers to divert subsequent emails to their doppelganger accounts are to either populate the *Reply-To* email field with it when sending the first email or to use the *Gmail email filters* to forward all victim's email to his own accounts. To efficiently counter those doppelganger tactics it is essential during the account recovery process to have these settings reviewed by the legitimate account owner or automatically cleared.

Retention tactics evolution: It is interesting to note that hijacker's tactics are constantly evolving in response to the new defenses we put in place. For example, in October 2011, for *46%* of the manual hijacking cases in which the hijackers initiated a password-change, the user also suffered a mass deletion of their emails. We believe hijackers deleted emails in attempt to make harder for the victims to reach their contacts when they got their accounts back. By November 2012 this probability was down to *1.6%*, since hijackers realized that this tactic was no longer effective after we improved our account recovery process by allowing users to optionally restore the emails, contacts, and settings added/modified/deleted by hijackers. Following this change we also saw the ratio of hijacker-initiated recovery option changes dropping from 60% (October 2011) to 21% (November 2012).

5.5 Manual Hijacking – an Ordinary Office Job?

We present further circumstantial evidence that manual hijackers may work in organized groups that follow a playbook based on an opportunity that we had to (retrospectively) monitor the activities of five individual hijackers for a short period of time.

We did observe the following:

- The individuals seemed to work according to a tight daily schedule. They started around the same time every day, and had a synchronized, one hour lunch break. They were largely inactive over the weekends.
- All individuals followed the same daily time table, defining when to process the newly gathered password lists, and how to divide time between ongoing scams and new victims.
- They were operating from different IPs, on different victims, and in parallel with each other, but the tools and utilities they used were the same. They also shared certain resources such as phone numbers.

We note that these observations support the results of the analysis presented throughout this paper.

6. HIJACKING REMEDIATION

In this section we discuss how our remediation process works. In particular, we study how long it takes for users to recover their account after being hijacked and analyze what is the success rate of the various options (SMS, email, knowledge test) that can be used by user to recover their accounts.

6.1 Remediation Workflow

The account hijacking remediation process consists of two parts. The first part, the recovery process, typically starts when the user realizes that his account is not accessible and submits an account recovery claim. The exact trigger may be a notification (e.g. an SMS notification about the password change on the account), or the user may just notice by himself that the password does not work or that the account was disabled by our anti-abuse systems to prevent further damage. The recovery part ends with Google verifying ownership and restoring exclusive access to the account to its rightful owner. The second part of the process is the cleanup and mitigation phase. This is triggered either by the account recovery system itself, or at the user's explicit request. At the end of it, all the changes made by the hijacker are ideally reverted.

6.2 ETA until recovery

From the anti-hijacking standpoint, a core metric is *latency*, which is the time elapsed between the hijacking and when the victim regains exclusive control over the account. The lower the latency, the less opportunity the hijacker has for exploitation, and the lower the cost for the victim. To encourage users to reclaim their accounts quickly, we keep users informed about key settings changes and potential suspicious activity via notifications over independent channels (e.g. SMS or secondary email address) whenever possible.

Figure 9 shows the distribution of the end-to-end recovery latencies for a random sample of 5000 accounts that were returned to the rightful owner after hijacking in Nov 2012 [Dataset 11 in Table 1]. In 22% of the cases, the victim successfully reclaimed the account within one hour after the hijacking, and in 50% of the cases the account was returned in less than 13 hours. The recovery time is calculated by taking the delta between the time our risk analysis system flagged the account as hijacked and the time the user started the recovery process. The fastest recoveries are best explained by the proactive notifications we send, enabling the user to recognize the hijacking and act immediately. Enabling more users to react quickly in case of hijacking is one of the main reason why we ask our users to keep their contact information up to date and give us their mobile phone number.

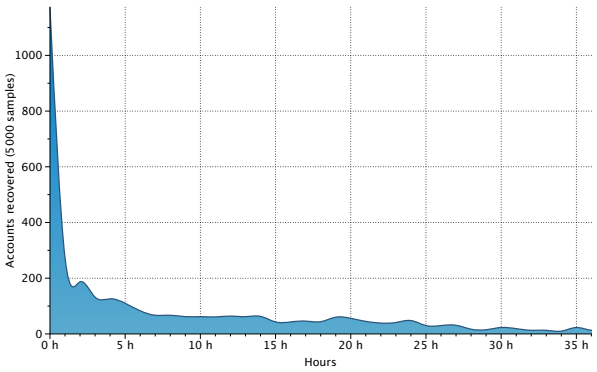

Figure 9: Hijacking recoveries by time. Sample of 5000 recoveries.

6.3 Recovery methods

Providing a secure and efficient account recovery process is a complex problem for two main reasons. First, we need to accommodate varying levels of user comfort towards sharing information for recovery (e.g. a phone number). Secondly, the recovery process needs to include contingency plans for the cases where an existing recovery option is inaccurate, outdated, or is itself compromised. This is particularly problematic for the secondary email recovery option as several providers, including Microsoft [25], expire email addresses after some period of inactivity and allows anyone, including the hijacker, to register them again. As of 2014, we estimate that 7% of the secondary emails our users provided for recovery have since been recycled. This email recycling problem makes recovery emails address a somewhat unreliable medium.

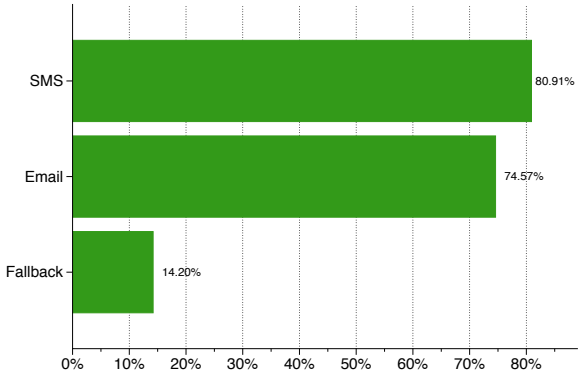

Figure 10: Success rate for various recovery methods.

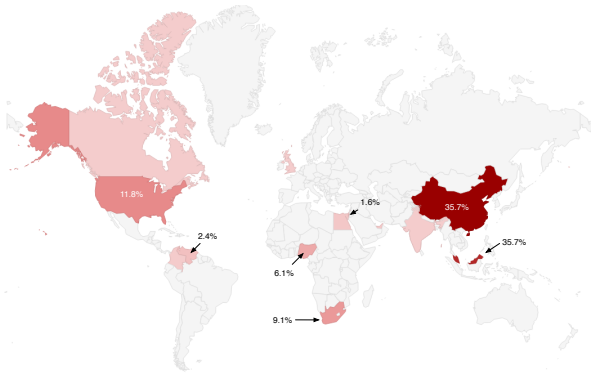

Figure 11: Top countries for the IPs involved in hijacking.

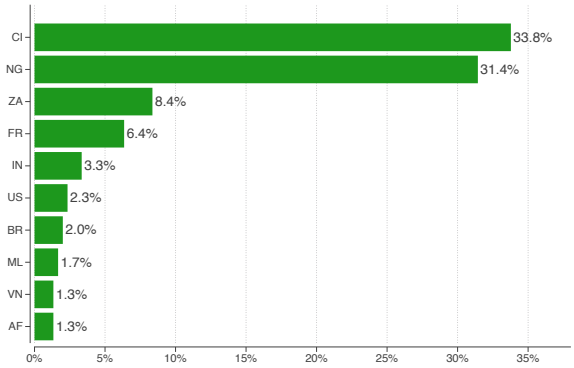

Figure 12: Top countries for the phone numbers involved in hijacking.

Figure 10 depicts the recovery success rate, based on a random sample of a month worth of account recovery claims, broken down by the method used to recover the account. Given the low volume of hijacking cases, this sample contains all the claims that were made in Feb. 2013 to avoid sample bias issues [Dataset 12 in Table 1].

We note that, while we always offer multiple recovery options to our users, the choice that we offer depends on our risk analysis and the recovery options that the account owner selected.

SMS: SMS verification, which has an over 80% success rate, is the most reliable recovery option for multiple reasons. First, users tend to keep their phone number up-to-date, which make non-existent phone numbers a non-issue. Secondly, it provides a very good user experience. Users find it easy to type in the code they receive via SMS. Finally, it is hard to fake. Failures can be traced back to the unreliability of SMS gateways in certain countries, and to confused users who did not really mean to use this option.

Email: Email is our most popular account recovery option and has a success rate of 74.57%. Email based recovery provides a good cross-device user experience as the user simply has to click on the link contained in the email. However as we explained earlier, email is a weaker recovery channel than SMS both in terms of security and reliability. The main source of failures is that users mistype their recovery emails, causing the password reset link to go to a wrong address or to a non-existing email address. We saw email bounces in approximately 5% of the cases. We are actively mitigating this issue by encouraging our users to verify their recovery addresses, but we do not enforce this step. Secondly, as also discussed above, the ownership of the secondary email is easier to lose, and users tend to not keep this recovery option up-to-date. We encourage our users to update their recovery email by showing reminder when they use our services. Furthermore to mitigate the risk of returning the account to an impostor, we do not offer this option if there is any indicatios that the secondary email address has been recycled.

Fallback options: We offer industry standard fallback options (security questions, knowledge tests and manual review) to our users when the account has no other recovery option available. For these fallback options, both the user experience and the success rate are significantly worse than for email or SMS.

We are constantly encouraging our users to upgrade to other options to improve their experience. Sadly a large portion of our users still hasn't upgraded, which force us to rely upon those less reliable options when they need to recover their accounts [21].

We also deemed security questions insecure and unreliable, and stopped allowing new users to use them as they have poor user recall and would-be hijackers may succeed by guessing the answer. However secret questions answering is still the only available method for a non-negligible portion of our user base. This is why, we only offer the ability to recover an account via security questions under certain limited circumstances, and repeatedly ask our users to give us better means to help them restore access to their accounts if needed. We only support knowledge-based options as a last resort.

6.4 Remission

After hijacking, regardless of whether the victim was locked out or has retained access, hijacker-initiated changes must to be reverted. This is very important because, as discussed in section 5.4, hijackers used to change the account settings and delete users' emails and contacts before we added this step to our account recovery flow. The remission process include restoring hijacker-deleted content, removing the hijacker-added content, and resetting all account options to their original state.
We found out by experimenting with various alternatives, that making content recovery an optional last step rather than having a fully automated process was user preferred behavior.

7. HIJACKERS' ORIGIN

Attributing accounts hijacking to specific groups of actors is very difficult as we have at best incidental proof. This section discuss hijacking attribution based on a sample of IPs used by hijackers in January 2014 [Dataset 13 in Table 1] and a set phone numbers used by hijackers sometime in 2012 [Dataset 14 in Table 1].

Figure 11 summarizes the geo-location of the IPs involved in 3000 hijacking cases that occurred in January 2014. As visible on the map, most of the traffic comes from China and Malaysia. We don't know if this traffic come from proxies or represent the true origin of the hijackers. However we do note that having Chinese traffic is consistent with the fact that hijackers search for Chinese terms.[2] Similarly having traffic originating from south america, mainly Venezuela, is consistent with hijackers search in spanish.

[2]While some Google services such as YouTube are completely blocked in China, Gmail on the other hand is not systematically blocked

The other source of data we have is a set of phone numbers that hijackers briefly used in 2012 to try to lock out users out of their accounts by enabling the two step authentication as an account retention tactic.

They quickly gave up after realizing it was not effective and we haven't seen this tactic used in the last 2 years which explain why we don't use a more recent dataset. The mapping of those 300 phone numbers to country using phone country code is summarized in figure 12. From this dataset two major groups of hijackers emerge: the Nigerian one (NG) and the Ivory Coast (CI) one. We believe those two groups to be different as their native language differs, French vs English, and they are 2000km apart. Anecdotal evidence suggest that the Ivory Coast specialize in scamming French speaking countries where as the Nigeria focus on English speaking countries. The volume of phone numbers involved in this type of attack is small enough to corroborate our hypothesis that is manual work and large enough to point to organized groups that are dedicated to monetize hijacked accounts. Finally we note that South Africa (ZA) account for 10% of both datasets which suggest that South Africa is also one of the largest home of hijackers. The fact that neither China or Malaysia show up in the phone dataset might be explained by the fact that the hijacking groups from those countries didn't try to use second factor enabling as a retention tactic.

8. DISCUSSION

In this section we discuss why it is difficult to defend against manual hijackers. In particular we summarize the defenses strategies that we found to work well, and those that did not work out so well.

8.1 Detecting manual hijackers is challenging

In our experience the greatest challenges in detecting manual hijacking is that it is extremely low volume, that hijackers are very versatile, and that it is difficult to strike the right balance between false positives (challenging legitimate users) and false negative (letting a hijacker in) when it come to detection.

Low volume: The volume of successful manual hijacking Google sees is extremely low. The exact rate varies with time, but its average is around 9 hijacks per day per million active accounts. This make the use of any large statistical model at best very difficult.

Hijacker versatility: Manual hijackers, as the name suggests, are human beings with all the flexibility and intelligence that comes with it. Our observations indicate that they are likely to have average technical ability, plus some additional knowledge of using IP cloaking services and browser plugins. Accordingly, what manual hijackers do when interacting with Google's services is not very different from what normal users do. Normal users also search their inboxes and read emails, set up email filters, and change their password and their recovery options. Thus the rules or models derived from those behaviors are not crystal clear and are certainly not high enough confidence to act upon easily. Furthermore, as we have emphasized throughout this paper, manual hijacking is an ever-moving target. The actions of the hijacker, while following certain patterns, are by no means deterministic or static.

Striking the right balance: A key challenge is handling the base error rate of our detection systems. Although we have invested a great deal in minimizing the error rate, our systems are not perfect. We have to carefully tune the aggressiveness of our system to balance acting upon signals that *might* indicate manual hijacking (but potentially inconveniencing legitimate users) against the risk of harm done by allowing hijackings to occur. This is especially difficult in light of the extreme low volume of manual hijacking; the vast majority of users are not hijacked this way. We concluded that a certain error rate, which we call the *false positive rate*, is a fair price to pay for greatly reducing the number of successful hijackings, e.g., temporarily not being able to access the account or being asked to answer additional verification questions.

8.2 Defense strategies

Second factor: Using a second authentication factor, such as a phone, has proven the best client-side defense against hijacking. While second factor authentication has some drawbacks, we believe that it is the best way to curb hijacking long term. The main issue with second factor authentication is that it is incompatible with legacy applications, such as mail clients. We work around this by authorizing our users to generate an application-specific password for those type of apps. However, this is far from ideal since those passwords can be phished. Consequently, we also work with vendors to move their apps to a better authentication technology, such as OAuth. An additional drawback of second factor authentication is usability. While phones provide a good user experience, we are exploring alternatives [7] for people who don't have a smartphone (e.g emerging countries) or want a separated physical device. We hope to see more research done in this space as there is a clear need of innovation in term of usability and accessibility.

Login time risk analysis is for us the best defense strategy that an identity provider can implement server-side since it stops the hijacker before getting into the account. Over the years we have built a complex login risk analysis system that assess for each login attempt whether it is the legitimate owner or not. Our system uses many signals (that we can't disclose for obvious reasons) to evaluate how anomalous a login attempt is. If the login attempt is deemed suspicious the user is redirected to an additional verification step before begin allowed to access his account: the login challenge [8]. Our login challenge asks the user to answer knowledge test questions or to verify their identity by proving he has access to the phone that was registered with the account earlier. e.g. by receiving an SMS. We view the proof of having access to the phone associated with the account as a safer challenge than knowledge question answers that the hijacker may just guess by researching the user's background. We spend a lot of time ensuring that our login challenge is easy to pass for our users, but hard for hijackers. This allows us to aggressively block hijacking in exchange for a small fraction of false positives as the "annoyance" cost associated of being mistakenly challenged is marginal for legitimate users.

Account behavioral risk analysis is important and needed, but we argue that it should be viewed as a last resort as it is already too late from the victim perspective. By that time any behavioral detector reports an anomaly the hijacker already had accessed the account data (emails / contact lists), and even might have setup means to continue his attack by setuping a "doppelganger" account for example.

User notifications is another tool that as proven to be essential to fight hijackers. Triggering notifications on critical events is very effective to thwart hijacking attempts and speed up the recovery process. We found out that is is important to being mindful about the keeping the volume of notification low so users know they are important when they see them. We notify our users upon account settings changes, blocked suspicious logins, and unusual in-product activity for which we have high confidence.

Iron tight Account recovery Finally we can't stress enough how important it is to invest into having a very secure and reliable *account recovery system*. We continuously improve our recovery process to ensure that it is easy for legitimate users to get their account back while keeping hijackers out. Developing novel ways to validate user identity both for login challenge and account recovery purpose is something that we view as critical and we would love to see more research done in this space.

9. REFERENCES

[1] APWG. Global phishing survey: Trends and domain name use in 1h2011. http://www.antiphishing.org/reports/, 2011.

[2] A. J. Aviv, K. Gibson, E. Mossop, M. Blaze, and J. M. Smith. Smudge attacks on smartphone touch screens. In *Proceedings of the 4th USENIX conference on Offensive technologies*, pages 1–7. USENIX Association, 2010.

[3] R. B. Cialdini. Influence: The psychology of persuasion. 1993.

[4] FBI. 2013 internet crime report. Technical report, FBI, 2013.

[5] H. Gao, J. Hu, C. Wilson, Z. Li, Y. Chen, and B. Y. Zhao. Detecting and characterizing social spam campaigns. In *SIGCOMM*. ACM, 2010.

[6] S. Garera, N. Provos, M. Chew, and A. D. Rubin. A framework for detection and measurement of phishing attacks. In *WORM '07*, 2007.

[7] Google. Google's internet identity research. https://sites.google.com/site/oauthgoog/gnubby.

[8] Google. Login challenge for suspicious sign-ins. https://support.google.com/a/answer/6002699?hl=en.

[9] Google. Transparency report: Safe browsing. http://www.google.com/transparencyreport/safebrowsing/?hl=en.

[10] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: The Underground on 140 Characters or Less. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2010.

[11] M. Hearn. An update on our war against account hijackers. http://bit.ly/1qfMckD, 2013.

[12] IOActive. Reversal and analysis of zeus and spyeye banking trojans. http://www.ioactive.com/pdfs/ZeusSpyEyeBankingTrojanAnalysis.pdf.

[13] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer. Social phishing. *Communications of the ACM*, 50(10):94–100, 2007.

[14] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An empirical analysis of spam marketing conversion. *Proceedings of the 15th ACM conference on Computer and communications security*, pages 3–14, 2008.

[15] J. Leyden. Typo-squatting domains can harvest corporate emails. http://bit.ly/1o8dx6d, 2011.

[16] C. Ludl, S. McAllister, E. Kirda, C. Kruegel, and C. Kruegel. On the effectiveness of techniques to detect phishing sites. In *DIMVA '07: Proceedings of the 4th International Conference on Detection of Intrusions, Malware, and Vulnerability Assessment*, pages 20–39, 2007.

[17] S. Maggi, A. Volpatto, S. Gasparini, G. Boracchi, and S. Zanero. Poster: Fast, automatic iphone shoulder surfing. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 805–808. ACM, 2011.

[18] T. Moore and R. Clayton. Examining the impact of website take-down on phishing. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, pages 1–13. ACM, 2007.

[19] T. Moore and R. Clayton. Discovering phishing dropboxes using email metadata. In *eCrime Researchers Summit (eCrime), 2012*, pages 1–9. IEEE, 2012.

[20] R. Naraine. Zdnet security blog: 'state-sponsored attackers' using ie zero-day to hijack gmail accounts, Jun. 2012.

[21] S. Schechter, A. B. Brush, and S. Egelman. It's no secret. measuring the security and reliability of authentication via "secret" questions. In *Security and Privacy*, pages 375–390. IEEE, 2009.

[22] R. Shay, I. Ion, R. W. Reeder, and S. Consolvo. "my religious aunt asked why i was trying to sell her viagra": Experiences with account hijacking. In *CHI*, 2014.

[23] Symantec. Industrial espionage: Targeted attacks and advanced persistent threats (apts). http://bit.ly/RsGkV7.

[24] K. Thomas, F. Li, C. Grier, and V. Paxson. Consequences of connectivity: Characterizing account hijacking on twitter. In *Proceedings of the 21st Annual Conference on Computer and Communications Security*, 2014.

[25] T. N. Web. Microsoft can recycle your outlook.com email address if your account becomes inactive. http://tnw.co/1sWsNAU, 2013.

[26] C. Whittaker, B. Ryner, M. Nazif, and M. Nazif. Large-scale automatic classification of phishing pages. In *NDSS*, 2010.

[27] Y. Zhang, S. Egelman, L. Cranor, J. Hong, and J. Hong. Phinding phish: Evaluating anti-phishing tools. In *NDSS*, 2007.

[28] Embassy of the Unites States warning on West African advanced fee scams, http://abidjan.usembassy.gov/art_of_scam.html.