Hack for Hire: Exploring the Emerging Market for Account Hijacking

Ariana Mirian
University of California, San Diego
amirian@cs.ucsd.edu

Joe DeBlasio*
University of California, San Diego
jdeblasio@cs.ucsd.edu

Stefan Savage
University of California, San Diego
savage@cs.ucsd.edu

Geoffrey M. Voelker
University of California, San Diego
voelker@cs.ucsd.edu

Kurt Thomas
Google
kurtthomas@google.com

ABSTRACT

Email accounts represent an enticing target for attackers, both for the information they contain and the root of trust they provide to other connected web services. While defense-in-depth approaches such as phishing detection, risk analysis, and two-factor authentication help to stem large-scale hijackings, targeted attacks remain a potent threat due to the customization and effort involved. In this paper, we study a segment of targeted attackers known as "hack for hire" services to understand the playbook that attackers use to gain access to victim accounts. Posing as buyers, we interacted with 27 English, Russian, and Chinese blackmarket services, only five of which succeeded in attacking synthetic (though realistic) identities we controlled. Attackers primarily relied on tailored phishing messages, with enough sophistication to bypass SMS two-factor authentication. However, despite the ability to successfully deliver account access, the market exhibited low volume, poor customer service, and had multiple scammers. As such, we surmise that retail email hijacking has yet to mature to the level of other criminal market segments.

KEYWORDS

email security; hacking; phishing; account compromise

1 INTRODUCTION

It has long been understood that email accounts are the cornerstone upon which much of online identity is built. They implicitly provide a root of trust when registering for new services and serve as the backstop when the passwords for those services must be reset. As such, the theft of email credentials can have an outsized impact—exposing their owners to fraud across a panoply of online accounts.

Unsurprisingly, attackers have developed (and sell) a broad range of techniques for compromising email credentials, including exploiting password reuse, access token theft, password reset fraud and phishing among others. While most of these attacks have a low success rate, when applied automatically and at scale, they can be quite effective in harvesting thousands if not millions of accounts [27]. In turn, email providers now deploy a broad range of defenses to address such threats—including challenge questions to protect password reset actions, mail scanning to filter out clear phishing lures, and two-factor authentication mechanisms to protect accounts against password theft [7–9]. Indeed, while few would claim that email account theft is a solved problem, modern defenses have dramatically increased the costs incurred by attackers and thus reduce the scale of such attacks.

However, while these defenses have been particularly valuable against large-scale attacks, targeted attacks remain a more potent problem. Whereas attackers operating at scale expect to extract small amounts of value from each of a large number of accounts, targeted attackers expect to extract large amounts of value from a small number of accounts. This shift in economics in turn drives an entirely different set of operational dynamics. Since targeted attackers focus on specific email accounts, they can curate their attacks accordingly to be uniquely effective against those individuals. Moreover, since such attackers are unconcerned with scale, they can afford to be far nimbler in adapting to and evading the defenses used by a particular target. Indeed, targeted email attacks—including via spear-phishing and malware—have been implicated in a wide variety of high-profile data breaches against government, industry, NGOs and universities alike [10, 12, 13, 31].

While such targeted attacks are typically regarded as the domain of sophisticated adversaries with significant resources (e.g., state actors, or well-organized criminal groups with specific domain knowledge), it is unclear whether that still remains the case. There is a long history of new attack components being developed as vertically integrated capabilities within individual groups and then evolving into commoditized retail service offerings over time (e.g., malware authoring, malware distribution, bulk account registration, AV testing, etc. [27]). This transition to commoditization is commonly driven by both a broad demand for a given capability and the ability for specialists to reduce the costs in offering it at scale.

In this paper, we present the first characterization of the *retail* email account hacking market. We identified dozens of underground "hack for hire" services offered online (with prices ranging from $100 to $500 per account) that purport to provide targeted attacks to all buyers on a retail basis. Using unique online buyer personas, we engaged directly with 27 such account hacking service providers and tasked them with compromising victim accounts of our choosing. These victims in turn were "honey pot" Gmail accounts, operated in coordination with Google, and allowed us to record key interactions with the victim as well as with other fabricated aspects of their online persona that we created (e.g., business web servers, email addresses

*Author DeBlasio has since joined Google.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

ACM ISBN 978-1-4503-6674-8/19/05.
https://doi.org/10.1145/3308558.3313489

of friends or partner). Along with longitudinal pricing data, our study provides a broad picture of how such services operate—both in their interactions with buyers and the mechanisms they use (and do not use) to compromise victims.

We confirm that such hack for hire services predominantly rely on social engineering via targeted phishing email messages, though one service attempted to deploy a remote access trojan. The attackers customized their phishing lures to incorporate details of our fabricated business entities and associates, which they acquired either by scraping our victim persona's website or by requesting the details during negotiations with our buyer persona. We also found evidence of re-usable email templates that spoofed sources of authority (Google, government agencies, banks) to create a sense of urgency and to engage victims. To bypass two-factor authentication, the most sophisticated attackers redirected our victim personas to a spoofed Google login page that harvested both passwords as well as SMS codes, checking the validity of both in real time. However, we found that two-factor authentication still proved an obstacle: attackers doubled their price upon learning an account had 2FA enabled. Increasing protections also appear to present a deterrent, with prices for Gmail accounts at one service steadily increasing from $125 in 2017 to $400 today.

As a whole, however, we find that the commercialized account hijacking ecosystem is far from mature. Just five of the services we contacted delivered on their promise to attack our victim personas. The others declined, saying they could not cover Gmail, or were outright scams. We frequently encountered poor customer service, slow responses, and inaccurate advertisements for pricing. Further, the current techniques for bypassing 2FA can be mitigated with the adoption of U2F security keys. We surmise from our findings, including evidence about the volume of real targets, that the commercial account hijacking market remains quite small and niche. With prices commonly in excess of $300, it does not yet threaten to make targeted attacks a mass market threat.

2 METHODOLOGY

In this section we describe our methodology for creating realistic, but synthetic, victims to use as targets, the infrastructure we used to monitor attacker activity, and the services we engaged with to hack into our victim email accounts. We also discuss the associated legal and ethical issues and how we addressed them in our work.

2.1 Victims

We created a unique victim persona to serve as the target of each negotiation with a hack for hire service. We never re-used victim personas between services, allowing us to attribute any attacks deployed against the persona back to the service we hired. In creating victim personas, we spent considerable effort to achieve three goals:

- *Victim verisimilitude*. We created synthetic victims that appeared sufficiently real that the hacking services we hired would treat them no differently from other accounts that they are typically hired to hack into.
- *Account non-attributability*. We took explicit steps to prevent attackers from learning our identities while we engaged with them as buyers, when they interacted with us as victims, and even if they successfully gained access to a victim email account.
- *Range of attacker options*. We did not know a priori what methods the hacking services would use to gain access to victim email accounts. Since there are many possibilities, including brute-force password attacks, phishing attacks on the victim, and malware-based attacks on the victim's computers, we created a sufficiently rich online presence to give attackers the opportunity to employ a variety of different approaches.

The remainder of this section details the steps we took to achieve these goals when creating fictitious victims, the monitoring infrastructure we used to capture interactions with our fake personas, and the selection of "hack for hire" services we engaged with.

Victim Identities. Each victim profile consisted of an email address, a strong randomly-generated password, and a name. While each of our victims 'lived' in the United States, in most cases we chose popular first and last names for them in the native language of the hacking service, such as "Natasha Belkin" when hiring a Russian-language service.[1] The email address for the victim was always a Gmail address related to the victim name to further reinforce that the email account was related to the victim (e.g., `natasha.r.belkin@gmail.com`). We loaded each email account with a subset of messages from the Enron email corpus to give the impression that the email accounts were in use [5]. We changed names and domains in the Enron messages to match those of our victim and the victim's web site domain (described below), and also changed the dates of the email messages to be in this year.

Each victim Gmail account used SMS-based 2-Factor Authentication (2FA) linked to a unique phone number.[2] As Gmail encourages users to enable some form of 2FA, and SMS-based 2FA is the most utilized form, configuring the accounts accordingly enabled us to explore whether SMS-based 2FA was an obstacle for retail attackers who advertise on underground markets [1] (in short, yes, as discussed in detail in Section 3.4).

Online Presence. For each victim, we created a unique web site to enhance the fidelity of their online identity. These sites also provided an opportunity for attackers to attempt to compromise the web server as a component of targeting the associated victim (server attacks did not take place). Each victim's web site represented either a fictitious small business, a non-governmental organization (NGO), or a blog. The sites included content appropriate for its purported function, but also explicitly provided contact information (name and email address) of the victim and their associates (described shortly). We hosted each site on its own server (hosted via third-party service providers unaffiliated with our group) named via a unique domain name. We purchased these domain names at auction to ensure that each had an established registration history (at least one year old) and the registration was privacy-protected to prevent post-sale attribution to us (privacy protection is a common practice; one recent study showed that 20% of .com domains are registered in this fashion [17]). The sites were configured to allow third-party crawling, and we validated that their content had been incorporated into popular search engine indexes before we contracted for any hacking services. Finally, we also established a passive Facebook

[1] These example profile details are from a profile that we created, but in the end did not need to use in the study.

[2] These phone numbers, acquired via prepaid SIM cards for AT&T's cellular service, were also non-attributable and included numbers in a range of California area codes.

profile for each victim in roughly the style of Cristofaro et al. [3]. These profiles were marked 'private' except for the "About Me" section, which contained a link to the victim's web site.[3]

Associate Identity. In addition to the victim identity, we also created a unique identity of an associate to the victim such as a spouse or co-worker. The goal with creating an associate was to determine whether the hacking services would impersonate the associate when attacking the victim (and some did, as detailed in Section 3.2) or whether they would use the associate email account as a stepping stone for compromising the victim email account (they did not). Similar to victim names, we chose common first and last names in the native language of the hacking service. Each victim's web site also listed the name and a Gmail address of the associate so that attackers could readily discover the associate's identity and email address if they tried (interestingly, most did not try as discussed in Section 3.2). Finally, if the victim owned their company, we also included a company email address on the site (only one attack used the company email address in a phishing lure).

Buyer Identity. We interacted anonymously with each hack for hire service using a unique buyer persona. When hiring the same service more than once for different victims, we used distinct buyer personas so that each interaction started from scratch and was completely independent. In this role, we solely interacted with the hacking services via email (exclusively using Gmail), translating our messages into the native languages of the hacking service when necessary.

Many hacking services requested additional information about the victim from our buyers, such as names of associates, to be able to complete the contract. Since we made this information available on the victim web sites, we resisted any additional requests for information to see if the services would make the effort to discover this information themselves, or if services would be unable to complete the contract without it (Section 3.1).

2.2 Monitoring Infrastructure

Email Monitoring. For each Gmail account, we monitored activity on the account by using a modified version of a custom Apps Script shared by Onaolapo et al. [23]. This script logged any activity that occurs within the account, such as sending or deleting email messages, changing account settings, and so on (Section 3.6 details what attackers did after gaining access to accounts). The script then uploaded all logged activity to a service running in Google's public cloud service (Google App Engine) as another level-of-indirection to hide our infrastructure from potential exposure to attackers. Since the monitoring script runs from within the Gmail account, it is possible in principle for an attacker to discover the script and learn where the script is reporting activity to, though only after a successful attack. We found no evidence that our scripts were detected.

Login Monitoring. In addition to monitoring activity from within the accounts, the accounts were also monitored for login activity by Google's system-wide logging mechanisms. Google's monitoring, shared with us, reported on login attempts and whether they were successful, when attackers were presented with a 2FA challenge, and whether they were able to successfully respond to the challenge (Section 3.4). These monitoring logs also include the infrastructure and devices used to make login attempts, which Google used to identify other Gmail accounts attacked by these services (Section 4.1).

Phone Monitoring. As described earlier, each victim account was associated with a unique cell number (used only for this purpose)

[3]None of the service providers we contracted with appeared to take advantage of the Facebook profile, either by visiting the victim's web site via this link or communicating with the victim via their Facebook page.

Service	Price	Lang	Prepay	Payment	Respond	Attack
A.1	$229	RU	50%	Qiwi	Yes	Yes
A.2	$229	RU	50%	Qiwi	Yes	Yes
A.3	$458	RU	50%	Qiwi	Yes	Yes
B.1	$380	RU	No	Webmoney, Yandex	Yes	Yes
B.2	$380	RU	No	Webmoney, Yandex	Yes	Yes
C.1	$91	RU	No	Bitcoin	Yes	Yes
C.2	$91	RU	No	–	Yes	Yes
D.1	$76	RU	No	–	Yes	Yes
E.1	$122	RU	No	–	Yes	Yes
E.2	$122	RU	No	–	Yes	No
D.2	$76	RU	No	–	Yes	No
F	$91	RU	No	–	Yes	No
G	$91	RU	No	–	Yes	No
H.1	$152	RU	No	Webmoney	Yes	No
H.2	$152	RU	No	Webmoney	Yes	No
J	–	EN	–	–	Yes	No
K	$200–300	EN	Yes	Bitcoin	Yes	No
L	$152	RU	No	–	Yes	No
M	$84	RU	No	–	Yes	No
N	$69	RU	No	Webmoney, Yandex	Yes	No
O	–	RU	No	Webmoney, Yandex	Yes	No
P	$305	RU	No	–	Yes	No
Q	$46	RU	Yes†	–	Yes	No
R	$100	EN	No	–	No	No
S	$400–500	EN	50%	–	No	No
T	$95 or 113	EN	No	Bitcoin, Credit Card	No	No
U	$98	RU	No	Webmoney	No	No
V	$152	RU	No	Webmoney, Yandex, Qiwi	No	No
W	$152	RU	No	–	No	No
X	$152	RU	No	Webmoney, Yandex	No	No
Y	$23 – $46	RU	No	–	No	No
Z	$61	RU	No	–	No	No
AA	$46	RU	No	–	Yes	No
BB	–	CN	–	–	No	No

Table 1: We contacted 27 hacking services attempting to hire them to hack 34 different victim Gmail accounts. We communicated with the services in the language in which they advertised, translating when necessary. The prices they advertised were in their native currency, and we have normalized them to USD for ease of comparison. (Yes† : for first-time customers.)

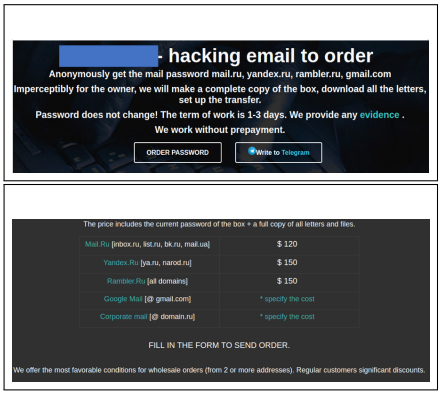

Figure 1: An online advertisement for Gmail hacking services. We remove any identifiable information and translate the page from Russian to English.

Service	Advertised	Discussed	Final
A.1	$230	$230	$307
A.2	$230	$230 - $307	Failed
A.3	$460	$460	$460
B.1	$383	$383	Failed
B.2	$383	$383	$383
C.1	$92	$102	$100
C.2	$92	–	Failed
D.1	$77	$184	Failed
D.2	$77	$184	Failed
E.1	$123	$383 - $690	$383
E.2	$123	$690	Failed

Table 2: The changes in negotiated prices when advertised, when initially hired, and when finally successful at hacking into victim Gmail accounts. All prices were originally in rubles, but are converted to USD for easier comparison.

which was configured in Gmail to be the contact number for SMS-based 2FA. To capture attacks against these phone numbers or notifications from Google (e.g., for 2FA challenges or notification of account resets) we logged each SMS message or phone call received.

Web Site Monitoring. To monitor activity on the web sites associated with the victims, we recorded HTTP access logs (which included timestamp, client IP, user agent, referrer information, and path requested). For completeness, we also recorded full packet traces of all incoming traffic to the target server machines in case there was evidence of attacker activity outside of HTTP (e.g., attempts to compromise the site via SSH). Overall, we found no evidence of attackers targeting our web sites.

2.3 Hacking Services

Recruitment. We identified hacking services through several mechanisms: browsing popular underground forums, searching for hacking services using Google search, and contacting the abuse teams of several large Internet companies. We looked for services that specifically advertised the ability to hack into Gmail accounts. While we preferred services that explicitly promised the passwords of targeted accounts, we also engaged with services that could instead provide an archive of the victim's account. Figure 1 shows an example service advertisement (one we did not purchase from).

When hiring these services, we followed their instructions for how to contact them. Typically, interactions with the services consisted of a negotiation period, focused on a discussion of what they would provide, their price, and a method of payment. The majority of the services were non-English speaking. In these cases, we used a native speaker as a translator when needed. We always asked whether they could obtain the password of the account in question as the objective, and always offered to pay in Bitcoin. If the sellers did not want to use Bitcoin, we used online conversion services to convert into their desired currency (the minority of cases). Interestingly, only a handful of services advertised Bitcoin as a possible payment vector, though many services were generally receptive towards using Bitcoin when we mentioned it.

Table 1 summarizes the characteristics of all services that we contacted, which we anonymize so that our work does not advertise merchants or serve as a performance benchmark. In total, we reached out to 27 different services and attempted to hire them to hack 34 unique victim Gmail accounts. When a service successfully hacked into an account, we later hired them again (via another unique buyer persona) with a different victim to see if their methods changed over time (we denote different purchases from the same service by appending a number after the letter used to name the service).

Service reliability. Of the twenty-seven services engaged, ten refused to respond to our inquiries. Another twelve responded to our initial request, but the interactions did not lead to any attempt on the victim account. Of these twelve, nine refused up front to take the contract for various reasons, such as claiming that they no longer hacked Gmail accounts contrary to their contemporary advertisements. The remaining three appear to be pure scams (i.e., they were happy to take payment, but did not perform any service in return). One service provided a web-based interface where we entered the email address we wanted hacked into a form. This form triggered a loading bar that showed that the "hacking" was *in progress* with a Matrix cinematic-style background. Once the bar reached "100%", the site reported that the password was captured, but we would need to pay money to decrypt the (fictional) UFD2 hash of the password.[4] Another service advertised payment on delivery, but after our initial inquiry, explained that they required full prepayment for first-time customers. After payment, they responded saying that they had attempted to get into the account but could not bypass the 2FA SMS code without further payment. They suggested that they could break into the mobile carrier, intercept the SMS code, and thus break into the Gmail account. We paid them, and, after following up a few times, heard nothing further from them. During this entire exchange, we did not see a single login attempt on the victim's Gmail account from the hacking service. The third site similarly required pre-payment and performed no actions that we could discern.

Finally, five of the services made clear attempts (some successful, some unsuccessful) to hack into eleven victim accounts. We focus on these services going forwards.

[4]We did not pay them since our monitoring showed that they had made no attempts on the victim's Gmail account and hence we would learn nothing more by paying.

Service	Method	Lure	Inbox or Spam	Promised goods	Requested	Success
A.1	Phishing	A, G, S	Inbox	Archive	–	Y
A.2	Phishing	A, G, S	Inbox	Archive	Victim and associate name, phone number	N
A.3	Phishing	A, G, S	Inbox	Archive	Victim and associate name, phone number	Y
B.1	Phishing	B	Inbox, Spam	Password	–	N
B.2	Phishing	A, G, V	Inbox, Spam	Password	Victim name, associate name/email, phone number*	Y
C.1	Phishing	G	Inbox	Password	–	Y
C.2	Phishing	G	Inbox, Spam	Password	–	N
D.1	Malware	V	Spam	Password	Victim name and occupation	N
E.1	Phishing	G, V	Inbox, Spam	Password	–	Y

Table 3: Overview of attack scenarios per service. Lure emails include impersonating an associate (A), bank (B), Google (G), government (V), or a stranger (S). In the event a service indicated they could not succeed without additional information, we indicate what details they requested. In one case (marked *), this was only for the second attempt.

Pricing. The cost for hiring the hacking services often varied significantly between the advertised price and the final amount we paid. Table 2 shows a breakdown of the price differences during engagement with the hacking services we successfully hired. The table shows the service, the purported price for that service from their online advertisement, the initially agreed upon price for their services, and then any price increase that may have incurred during the attack period. When services failed to hack into the account, they did not request payment. Several factors influenced the changes in prices, in particular the use of 2FA on the accounts (Section 6).

As a rule, we always paid the services, even when they requested additional money, and even when we strongly suspected that they might not be able to deliver when they asked for payment up front.[5] Our goal was to ultimately discover what each service would actually do when paid.

2.4 Legal and Ethical Issues

Any methodology involving direct engagement with criminal entities is potentially fraught with sensitivities, both legal and ethical. We discuss both here and how we addressed them.

There are two legal issues at hand in this study: unauthorized access and the terms of service for account creation and use. Obtaining unauthorized access to third-party email accounts is unlawful activity in most countries and in the United States is covered under 18 USC 1030, the Computer Fraud and Abuse Act (CFAA). Contracting for such services, as we did in this study, could constitute aiding and abetting or conspiracy if the access was, in fact, unauthorized. However, in this study, the email accounts in question are directly under our control (i.e., we registered them), and since we are acting in coordination with the account provider (Google), our involvement in any accesses was explicitly authorized. The other potential legal issue is that this research could violate Google's terms of service in a number of ways (e.g., creating fake Gmail accounts). We addressed this issue by performing our study with Google's explicit permission (including a written agreement to this effect). Both our institution's general counsel and Google's legal staff were appraised of the study, its goals, and the methods employed before the research began.

This study is not considered human subjects research by our Institutional Review Board because, among other factors, it focuses on measuring organizational behaviors and not those of individuals. Nevertheless, outside traditional human subjects protections, there are other ethical considerations that informed our approach. First, by strictly using fictitious victims, associates and web sites, we minimized the risk to any real person resulting from the account hacking contracted for in this study. Second, to avoid indirect harms resulting from implicitly advertising for such services (at least the effective ones), we made the choice to anonymize the names of each service. Finally, to minimize our financial contributions to a potentially criminal ecosystem, we limited the number of purchases to those needed to establish that a service "worked" and, if so, that its modus operandi was consistent over time.

3 HACK FOR HIRE PLAYBOOK

Our study characterizes the operational methods that hack for hire services employ when making a credible attempt to hijack our victim personas. We limit our analysis exclusively to the five services where the attackers made a detectable attempt to gain access to our victim account. We note that the ultimate "success" of these attacks is partially dependent on our experimental protocol: in some cases, we supplied 2FA SMS codes to phishing attacks or installed a provided executable, while in other cases, we avoided such actions to see if the attackers would adapt.

3.1 Attacks Overview

We present a high-level breakdown of each hack for hire service's playbook in Table 3. Four of the five services we contacted relied on phishing, while just one relied on malware. In all cases, attacks began with an email message to our victim persona's Gmail address. We never observed brute force login attempts, communication with a victim's Facebook account, or communication to our associate personas of any kind.[6] On average, attackers would send roughly 10 email messages over the course of 1 to 25 days—effectively a persistent attack until success. All of the services but one were able to bypass Gmail spam filtering (though to varying degrees of

[5]The one exception to this rule is the aforementioned service whose automated web site immediately told us they had hacked the site when all evidence was to the contrary.

[6]In practice, a victim's password may be exposed in a third-party data breach. Our use of synthetic identities prevents this as a potential attack vector.

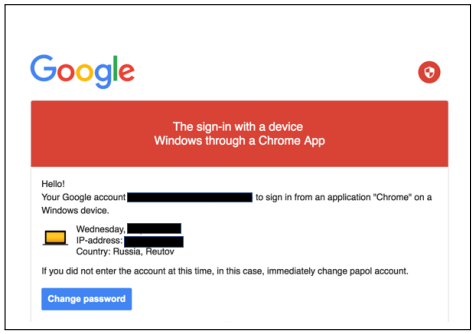

Figure 2: An example Google lure mimicking a real warning that Gmail will send to users. Identifying information removed and translated to English.

success) such that at least one of their messages appeared in our victim's inbox. However, this outcome is expected: since these are targeted attackers with more focused motivation, they have strong incentives to continue to adapt to phishing and spam defenses to ensure that their messages arrive in the victim's inbox. For example, attackers can create honeypot accounts of their own to test and modify their techniques, thereby ensuring a higher success rate; unlike their high-volume counterparts, targeted attackers only produce a modest number of examples and thus may pass "under the radar" of defenses designed to recognize and adapt to new large-scale attacks.

3.2 Email Lures

Each email message contained a lure whereby the attackers impersonated a trusted associate or other source of authority to coerce prospective victims into clicking on a link. Over the course of our study, we observed five different types of lures: those impersonating an associate persona, a stranger, a bank, Google, or a government authority. The associate lures attempted to get the user to click on an "image" for the victim's associate (using the personal connection as a sense of safety), while the Google, bank, and government lures conveyed some sense of urgency that would cause a user to click on the link. Figure 2 shows a sample Google lure that mimics a real warning used by Google about new device sign-ins. Such lures highlight the challenge of distinguishing authentic communication from service providers, whereby attackers repurpose potentially common experiences to deceive victims into taking an unsafe action.

Attackers cycled through multiple lures over time in an apparent attempt to find any message that would entice a prospective victim into clicking on a link. Figure 3 shows the elapsed time since attackers sent their first email message to our victim account, the type of lure they used for each message, and when we clicked on the lure acting as a victim (potentially halting further attempts). Each row corresponds to one attack on a victim, and the x-axis counts the number of days since the service sent their first message to the victim. The numbers on the right y-axis show the total number of messages sent by the service to the victim. The most popular lure mimicked Google, followed by associates and then lures from strangers.

Of the five services, two relied on personalized messages when communicating with four victim personas. In three of these cases, the service asked for additional details upfront about the victim persona during negotiation. Only service A.1 was able to construct personal

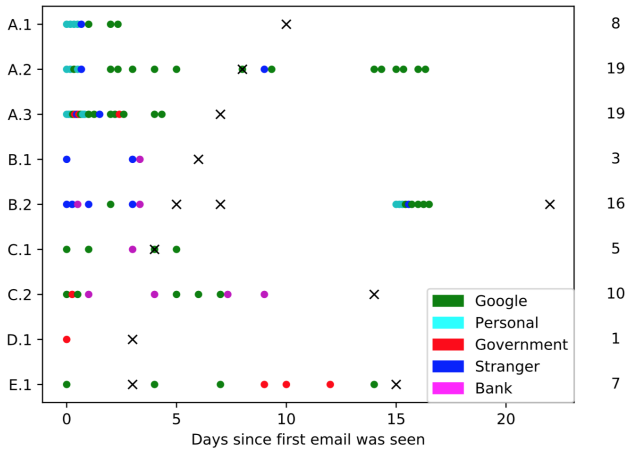

Figure 3: The different types of lures used by services that attempted to access the victim account. An 'X' marks when we clicked on a link in a message sent to a victim. Numbers on the right denote the total number of emails sent by that service.

Service	'google' in URL?	HTTPS	# redirects to phishing page
A.1	Yes	Yes	2
A.2	Yes	Yes	2
A.3	Yes	Yes	2
B.1	Yes	No	1
B.2.1	Yes	No	1
B.2.2	Yes	No	1
B.2.3	Yes	Yes	2
C.1	No	No	0
C.2	NA	NA	NA
D.1	NA	NA	NA
E.1.1	Yes	Yes	1
E.1.2	Yes	Yes	2

Table 4: For services that attempted to hack a victim account, we show whether Google was used in the phishing URL, whether the phishing page used HTTPS, and the number of redirects to the phishing page. We include separate rows for the services that sent multiple messages (services B and E).

lures without requesting assistance from the buyer, finding the details from the victim persona's website. The extent of personalization was limited, though, consisting either of mimicking the victim persona's company or their associate's personal email address. No additional branding was lifted from our web sites.

3.3 Phishing Landing Pages

All services but one relied on phishing as their attack vector. Once we clicked on the links sent to the victim personas, we were redirected to a spoofed Google login page that requested the credentials from the victim. Table 4 lists the different attack attempts and the degree to which attackers tried to spoof a Google domain, use HTTPS, or mask URLs from a crawler via multiple redirects. All services but one used "combo" domain name squatting [14] with the keyword 'google' in the URL, presumably to trick the victim into thinking that

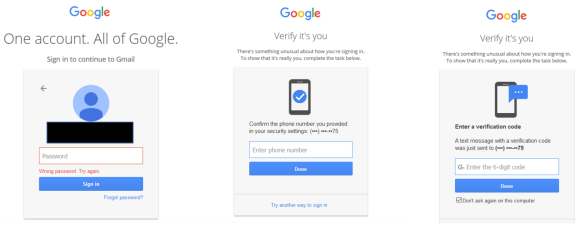

Figure 4: A service phishing flow, with identifiable information redacted. The flow is purposefully designed to mimic Gmail to trick the user into trusting the site.

the URL was a real Google subdomain. Services A.2 and B.2 used the same fully qualified domain name for the phishing landing page, suggesting that they share a business relationship (i.e., they may both be value-added resellers for the same phishing page service). Long-lived, reused domains suggest that they are valuable and perhaps relatively costly to acquire.

All but one service tried to obscure the URL to their phishing page with at least one layer of redirection. (The exception was the link in the phishing message from C.2, which redirected to an error page on a Russian hosting service indicating that the page had been taken down.) The redirection URLs seemed to be one-time use URLs, since we were not able to visit them after the attack executed and did not see repeat redirection URLs in any of the attacks. One-time use URLs are attractive for attackers because they can greatly complicate investigating attacks after the fact or sharing attack information among organizations.

Figure 4 shows an example page flow used by one hacking service. We always entered the Gmail credentials of the victim to see how the hacking attempt would progress. After collecting the password, all but one of the hacking services would redirect to a new screen which asked for the 2FA code that the victim had just received on their phone from Google.

Six of the nine hacking attempts captured the password from the phishing page and then immediately tried to use it to login to the victim's account (as verified with our Gmail access logging). Due to the similar behavior and speed at which these logins occurred, we believe that most of these services used an automated tool, similar to Evilginx [6], for this step.

Moreover, three of five of these attacks captured the necessary information in one session visiting the phishing pages. This degree of sophistication suggests that attackers can readily adapt any additional information requested by Google as a secondary factor. Since our study, Google launched additional protections at login to prevent automated access attempts [26]. However, hardware security keys remain the best protection mechanism against phishing for users.

3.4 Live Adaptation

Services B.2 and E.1 exhibited phishing attacks that adapted over time to overcome obstacles. These services, once realizing that the account used 2FA, sent new phishing email messages with a different structure than the ones they sent previously. Service E.1, for example, initially used a phishing attack that only captured the Gmail password. When the service attempted to login, they were blocked by the 2FA prompt. The service then contacted our buyer persona asking for the victim's phone number. The victim's email account subsequently received more phishing messages in their inbox. Clicking on the link in the phishing messages led to a page that requested the 2FA code that was sent to the victim's phone. When we entered the 2FA code into the phishing page, the service was able to successfully login. This behavior indicates live testing of password validity, as the attackers were able to determine if the account had 2FA.

Service B.2 was similar to service E.1, but when they were blocked by the 2FA challenge they switched to phishing messages that looked exactly like the messages from service A. Upon collecting the password and the 2FA code that was sent to the phone number for the victim, the service was able to login.

3.5 Malware Attachments

Service D was the only service that attempted to hijack our victim account using malware. The attacker in this case sent just one email message to our victim persona—flagged as spam—that contained a link to a `rar` archive download. The archive contained a sole executable file. The attackers most likely concealed the executable in a `rar` to impede scanning because Gmail forbids executable attachments by default. We unpacked and ran the executable in an isolated environment, but to no effect. According to VirusTotal [32], the executable is a variant of TeamViewer (a commercial tool for remote system access) which would have enabled the attacker to hijack any existing web browsing sessions.

After no further visible activity, the service eventually contacted our buyer persona to say that they could not gain access to our victim account. We decided to hire them again via a different contract (and different buyer and victim personas) to see if the seller would adapt to Gmail's defenses. However, we observed no email messages from the attacker the second time around, even in our spam folder. The seller eventually responded stating that they could not gain access to our second persona's account. While this malware vector proved unsuccessful, the presence of remote access tools poses a significant risk for adaptation, as session hijacking would enable an attacker to bypass any form of two-factor authentication.

3.6 Post Compromise

For those services that did obtain our victims' credentials and 2FA codes, the attackers proceeded to sign in to each account and immediately removed all Google email notifications (both from the inbox and then trash) related to a new device sign-in. None changed the account password. We also observed that services A, B, and E removed the 2FA authentication and the recovery number from our victim accounts as well. Presumably they took these steps to regain access to the account at a later time if needed without having to phish an SMS code again, but we did not see any service log back into the accounts after their initial login. However, these changes to the account settings could alert a real victim that their account had been hijacked, a discovery which the attackers are willing to risk.

Once accessed, all but one of the services abused a portability feature in Google services (Takeout) to download our victim account's email content and then provided this parcel to our buyer persona. One advantage of this approach is that it acquires the contracted deliverable in one step, thus removing risks associated with subsequent

credentials changes, improvements in defenses, or buyer repudiation. Only service C avoided logging into our victim account and only provided the buyer persona with a password.[7] These findings highlight an emerging risk with data portability and regulations around streamlining access to user data. While intended for users, such capabilities also increase the ease with which a single account hijacking incident can expose all of a user's data to attackers. Since our study, Google has added additional step-up verification on sensitive account actions.

4 REAL VICTIMS & MARKET ACTIVITY

Based on our findings from the hack for hire process, we returned to the forums of the most successful attackers to understand their pricing for other services and how they attract buyers. Additionally, we present an estimate of the number of real victims affected by these services based on login traces from Google. Our findings suggest that the hack for hire market is quite niche, with few merchants providing hijacking capabilities beyond a handful of email providers.

4.1 Victims Over Time

Of the 27 initial services we contacted, only three—services A, E, and B—could successfully login to our honeypot accounts. As part of our collaboration with Google, they examined metadata associated with each login attempt and found that all three services rely on an *identical* automation process for determining password validity, bypassing any security check such as producing an SMS challenge, and downloading our honey account's email history. Whereas the email messages from the services had varied senders and delivery paths for each contracted campaign, this automation infrastructure remained stable despite eight months between our successive purchases. This stability in turn allowed Google to develop a signature allowing the retrospective analysis of all such login attempts from the three services in aggregate.

Over a seven-month period from March 16 to October 15, 2018, Google identified 372 accounts targeted by services A, B, and E. Figure 5 shows a weekly breakdown of activity. On an average week, these services attacked 13 targets, peaking at 35 distinct accounts per week. We caution these estimates are likely only lower bounds on compromise *attempts* as we cannot observe users who received a phishing URL, but did not click it (or otherwise did not enter their password on the landing page). Despite these limitations, we see that the volume of activity from these hack for hire services is quite limited when compared to off-the-shelf phishing kits which impact over 12 million users a year [29]. Thus, we surmise that the targeted account hacking market is likely small when compared to other hacking markets, e.g., for malware distribution [11]. While the damage from these commercialized hacking services may be more potent, they are only attractive to attackers with particular needs.

Apart from the volume of these attacks, we also examine the sophistication involved. As part of its authentication process, Google may trigger a "challenge" for sign-in attempts from previously unseen devices or network addresses [20]. All of the hack for hire attempts triggered this detection. In 68% of cases, the attacker was forced to solve an SMS challenge, while in 19% of cases the attacker

[7]The service demanded additional payment to defeat the 2FA, which we paid, at which point they stopped responding to our requests.

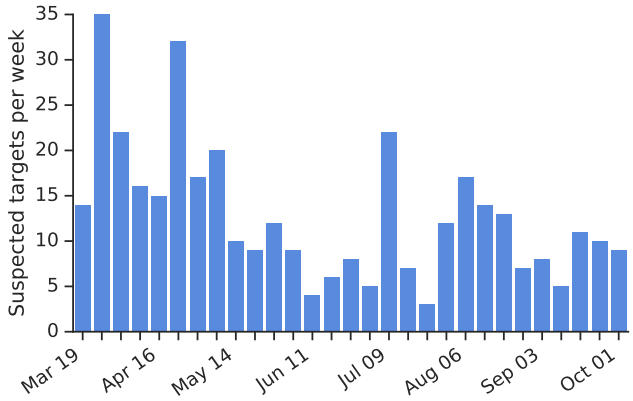

Figure 5: Weekly target accounts retroactively associated with hack for hire services.

only had to supply a victim's phone number. The remaining 13% involved a scattering of other secondary forms of authentication. This layered authentication approach provides better security when compared to passwords alone, with attackers only correctly producing a valid SMS code for 34% of accounts and a valid phone number in 52% of cases. These rates take into consideration repeated attacks: Google observed that attackers would attempt to access each account a median of seven times before they either succeeded or abandoned their efforts. As such, even though these attacks may be targeted, Google's existing account protections can still slow and sometimes stop attackers from gaining access to victim accounts.

4.2 Alternate Services and Pricing

While our investigation focused on Google—due in large part to our ethical constraints and abiding by Terms of Service requirements—the hack for hire services we engaged with also purport to break into multiple mail providers (Yahoo, Mail.ru, Yandex), social networks (Facebook, Instagram), and messaging apps (WhatsApp, ICQ, Viber). To provide a price comparison between offerings, in preparation for our study we performed a weekly crawl of the forum page or dedicated web site advertising each service starting in January 1, 2017. However, as detailed previously in Section 3, only a fraction of the services are authentic, and just three—services A, B, and C—had online prices that matched (or were close) to the final price we paid. We treat these as trusted sources of pricing information. We also include services E and D, but note their prices were higher than advertised. We exclude all other services as they failed to attack any of our victim personas.

We present a breakdown of pricing information as of October 10, 2018 in Table 5 for the five services that executed an attempt to access the accounts. Across all five services, Russian mail provider hacking (i.e., Mail.ru, Rambler and Yandex) was the cheapest, while other mail providers such as Gmail and Yahoo were more expensive. The cost of hacking a social media account falls in the middle of these two extremes.

We also note that some services have increased their prices over time. For services B and C, prices on the forums they advertise have been stable since we first began our monitoring. Only service A provided dynamic pricing, with rates increasing as shown in Figure 6.

Target	Service A	Service B	Service C	Service D*	Service E*
Mail.ru	$77	$77	$62	$54	$77
Rambler	$152	$108	$77	$77	$108
Yandex	$106	$108	$77	$77	$108
Gmail	$384	$385	$92	$77	Negotiable
Yahoo	$384	$231	$92	–	–
Facebook	$306	–	–	–	–
Instagram	$306	–	–	–	$231

Table 5: Purported price to access various accounts, based on an October 10, 2018 snapshot . All prices USD, converted from rubles. An asterisk indicates the service's advertised price was lower than the final payout requested during our buyout.

Since 2017, Gmail prices have steadily increased from $123 to $384, briefly peaking at $461 in February 2018. The advertised rates for targeting Yahoo accounts has largely tracked this same rate, while Facebook and Instagram were initially priced higher before settling at $307. We hypothesize that the price differences between services and the change in prices for a service over time are likely driven by both operational and economic factors. Thus, prices will naturally increase as the market for a specific service shrinks (reducing the ability to amortize sunk costs on back-end infrastructure for evading platform defenses) and also as specific services introduce more, or more effective, protection mechanisms that need to be bypassed (increasing the transactional cost for each hacking attempt).

4.3 Advertising & Other Buyers

As a final measure, we examined the forum advertisements each service used to attract buyers. Here, we limit our analysis to the five successful hack for hire services. Across seven underground forums, we identified two types of advertisements—pinned posts and banner ads—which require paying forum operators. Services A, B, and E, the three services that were able to bypass two-factor authentication, all had pinned posts on forums where this option was available. Only service A paid for banner advertisements on all of these forums. Together, this suggests that the services are profitable enough to continue advertising via multiple outlets. Additionally, these three services had verified accounts, indicating that a forum moderator had vetted the service stated. Further, services A, B, D, and E all stated they could work with a "guarantor", an escrow service proxying for payment between service and buyer to avoid fraud risks. By and large, feedback on the forum threads was positive, though we caution this may biased be due to the ability to delete posts and the difficulty in distinguishing between legitimate customers and virtual "shills". We avoid using forum posts as a count of purchases as most negotiation activity occurs via private messaging.

In addition to this qualitative search, we received an email advertisement from one of the services for upcoming changes to the service, which was sent to 44 other buyers as well (exposing their clientele's email addresses). The message was an announcement that the service now had a Telegram channel that was available (with a link to the channel), and to join the channel to keep up to date with relevant news. The only response to that initial email message was another customer exclaiming their excitement for this new development. Of the 44 email addresses that were leaked, 23 were accounts

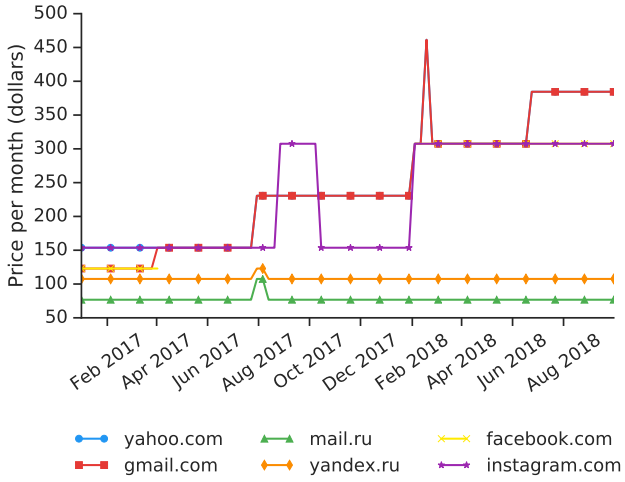

Figure 6: Monthly price that Service A charges across email and social network account providers. Over two years, the price per Gmail account increased from $123 to $384.

with .ru mail providers (`mail.ru` or `yandex.ru`), 9 were Gmail addresses, and the rest were various other providers, like Tutanota, Protonmail, or iCloud. We were unable to find these buyers online, which indicates that they did not engage in forum postings, or used a burner (one-purpose use) email address. However, the concentration of Russian mail providers suggests that interest in the market may largely be geographically limited, potentially due to language barriers or culturally-biased demands for account hacking.

5 RELATED WORK

Phishing is a well studied, yet continuing concern in the security community. Sheng et al. studied the demographic of people who are susceptible to phishing attacks, and found that users 18–25 years of age are most likely to click on phishing messages [24]. Egleman et al. studied the effectiveness of phishing warnings and found they can be successful in preventing account hijacking [4]. Following this mode of thought, there are a variety of studies on effective anti-phishing training as well as the creation of a content-based approach to detect phishing sites [15, 25, 33]; in all of these studies, the percentage of users' susceptible to phishing emails dropped. Similarly, Zhang et al. evaluted anti-phishing tools, and found that many of them are not effective on new URLs and have exploits of their own. Oest et al. also studied the phishing ecosystem via an analysis of phishing kits, and developed a URL-based scheme to detect phishing URLs [22].

Account hijacking threats represent a spectrum that ranges from financially motivated, large-scale attacks to highly-targeted incidents motivated by political, personal, or financial incentives. Thomas et al. identified billions of credentials stolen via data breaches and millions of credentials stolen by phishing kits and keyloggers, with phishing posing the largest hijacking risk [29]. Once an account was accessed, hijackers searched for financial records or used the account as a stepping stone to control connected online identities [2, 23]. While techniques such as risk-aware authentication [20] or two-factor authentication help protect against unsophisticated bulk attacks, the hack for hire outfits we studied were more dedicated,

with attackers stealing SMS two-factor codes as part of their phishing pages to bypass the additional layers of security. Security keys would prevent this attack vector.

At the other end of the spectrum, Marczak et al. investigated government actors targeting political dissidents [19]. The hijackers in these cases relied on exploits or social engineering to have victims install commercial or off-the-shelf spyware to enable long-term monitoring of the victim's activities. Email was a common delivery mechanism, where attackers customized their lures to the NGOs where employees worked or to the human rights topics they were involved with [12, 16]. Given the risks involved here, researchers have focused on how to improve the security posture of at-risk users [18]. Compared to our work, we found more generalized lures that can work for any target (e.g., your account is running out of storage space or there was a security incident), while phishing was the most popular technique for gaining one-off access to a victim's account. Pressure on the hack for hire playbook, or wider-scale adoption of security keys, may cause them to move towards malware and thus mirror government attackers.

6 DISCUSSION AND CONCLUSION

When starting this study, we had very little knowledge of what to expect in terms of attacker methods, behaviors, and ability. At a high level, we find that the commercial account hijacking ecosystem is far from mature. When such attackers are successful, they can be potentially devastating to individuals. Yet, as an overall market it is not poised to cause widespread harm.

Retail account hijacking is a niche market. Many aspects of engaging with account hijackers strongly indicate that these services are a fledgling market:

- Most telling is that only five of the 27 services we contacted were willing to take our business, a third never responded to repeated requests as buyers, and some were outright fraudulent.
- Services have inconsistent and poor customer service. For example, three of the services charged significantly higher prices than their advertised price, and two services changed their initial prices while they were executing the hack. Moreover, customer service is slow and inconsistent in their communication with the buyer, sometimes taking more than a day to respond.
- Attackers showed little initiative. Most attacks made no effort to gather information independently about their victims. Of the nine attempts, only services A.1 and A.2 discovered additional information about the victim on their web sites, such as the name of their associate. The others, including different contracts within service A, would not attempt hacking the account without explicitly requesting additional information from the buyer.

In contrast, studies on markets for CAPTCHA solving [21], Twitter spam [30], and Google phone verified accounts [28] show that those services are quick to respond, and stable in their services and pricing. This differentiation between other underground service offerings and the retail hacking market suggests that account hacking may not be the main focus of these attackers, and may simply be a "side hustle" — a method to gain opportunistic income in addition to other activities they are more fully engaged in.

Services predominantly mount social engineering attacks using targeted phishing email messages. All but one of the nine attacks used targeted email phishing to hack into our Gmail accounts. The attackers customized their phishing messages using details that we made available about the businesses and associates of our fictitious victims. To prompt engagement with a victim, the phishing messages created a sense of urgency by spoofing sources of authority (e.g., government agencies, banks, or Google itself).

These methods are a subset of those used in other targeted attack ecosystems. In particular, in *addition* to targeted phishing (frequently much more tailored than any attacks mounted by the services we studied), government-targeted attackers use malware and long-term monitoring of victim behavior to gain access to the account, requiring much more overhead than phishing alone [19]. Indeed, although these two classes of attackers are superficially similar in focusing on individual users, they are likely distinct in most other respects including the nature of the populations they target, their resource investment per target, their goals upon compromising an account, and a far greater requirement for covert operations.

Two-factor authentication creates friction. Even though phishing can still be successful with 2FA enabled, our results demonstrate that 2FA adds friction to attacks. Various services said that they could not hack into the account without the victim's phone number, had to adapt to 2FA challenges by sending new phishing messages to bypass them, and one renegotiated their price (from $307 to $690) when they discovered that the account had 2FA protection. Based on these results, we recommend major providers encourage or require their user base to use a 2FA physical token

Minimal service differentiation. Even with a variety of services advertising in the account hijacking market, they have remarkably little differentiation in their methods and infrastructure. Services A.1, A.2, A.3, and B.2 sent very similar re-usable phishing email messages to their respective victims, and all services that successfully hacked our accounts used identical automation tools for determining password validity, bypassing security checks, and downloading victim data.

Gmail as a vantage point. Overall, our study indicates that the attack space against Gmail is quite limited. Since we focused on hiring services to hack solely into Gmail accounts, it is possible that the landscape of the commercialized hacking market would look much different when deployed against native email services such as `mail.ru` or `yandex.ru`.

ACKNOWLEDGEMENTS

We would like to thank Mikhail Kolomogorov for his significant help in translating email messages, and we would similarly like to thank Kirill Levchenko, Vector Guo Li, and Ivan Mikhailin for their additional assistance in translation. Our thanks also to Shawn Loveland who provided additional data that helped us expand the set of underground hack for hire services that we considered. We also thank Elie Bursztein, Angelika Moscicki, Tadek Pietraszek, and Kashyap Puranik for their feedback on our study. We are also grateful to our anonymous reviewers for their insightful feedback and suggestions. This work was supported in part by NSF grants CNS-1629973 and CNS-1705050, DHS grant AFRL-FA8750-18-2-0087, the Irwin Mark and Joan Klein Jacobs Chair in Information and Computer Science, and by generous research, operational and/or in-kind support from Google and the UCSD Center for Networked Systems (CNS).

REFERENCES

[1] Olabode Anise and Kyle Lady. State of the Auth: Experiences and Perceptions of Multi-Factor Authentication. *Duo Security*, https://duo.com/assets/ebooks/state-of-the-auth.pdf, November 2017. Accessed: 2018-10-22.

[2] Elie Bursztein, Borbala Benko, Daniel Margolis, Tadek Pietraszek, Andy Archer, Allan Aquino, Andreas Pitsillidis, and Stefan Savage. Handcrafted Fraud and Extortion: Manual Account Hijacking in the Wild. In *Proceedings of the 2014 ACM Internet Measurement Conference (IMC)*, Vancouver, BC, Canada, November 2014.

[3] Emilano De Cristofaro, Arik Friedman, Guillaume Jourjon, Mohamed Ali Kaafar, and M. Zubair Shafiq. Paying for Likes? Understanding Facebook Like Fraud Using Honeypots. In *Proceedings of the 2014 ACM Internet Measurement Conference (IMC)*, Vancouver, BC, Canada, November 2014.

[4] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. You've been warned: An empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 1065–1074, New York, NY, USA, 2008. ACM.

[5] Enron Email Dataset. https://www.cs.cmu.edu/~enron/. Accessed: 2018-11-03.

[6] Evilginx — Advanced Phishing with Two-factor Authentication Bypass. https://breakdev.org/evilginx-advanced-phishing-with-two-factor-authentication-bypass/. Accessed: 2018-10-22.

[7] Google. Add 2-Step Verification. https://support.google.com/a/answer/175197. Accessed: 2018-10-22.

[8] Google. Guard Against Targeted Attacks. https://support.google.com/a/answer/9010419. Accessed: 2018-10-22.

[9] Google. Verify a user's identity with a login challenge. https://support.google.com/a/answer/6002699. Accessed: 2018-10-22.

[10] Garrett M. Graff. DOJ Indicts 9 Iranians For Brazen Cyberattacks Against 144 US Universities. *Wired*, https://www.wired.com/story/iran-cyberattacks-us-universities-indictment/. Accessed: 2018-10-22.

[11] Chris Grier, Lucas Ballard, Juan Caballero, Neha Chachra añd Christian J. Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Da mon McCoy, Antonio Nappa, Andreas Pitsillidis, Niels Provos, M. Zub air Rafique, Moheeb Abu Rajab, Christian Rossow, Kurt Thomas, Vern Paxson, Stefan Savage, and Geoffrey M. Voelker. Manufacturing Compromise: The Emergence of Exploit-as-a-Service. In *Proceedings of the ACM Conference on Computer and Communications Security*, Raleigh, NC, October 2012.

[12] Seth Hardy, Masashi Crete-Nishihata, Katharine Kleemola, Adam Senft, Byron Sonne, Greg Wiseman, Phillipa Gill, and Ronald J Deibert. Targeted Threat Index: Characterizing and Quantifying Politically-Motivated Targeted Malware. In *Proceedings of the 23rd USENIX Security Symposium*, San Diego, CA, USA, August 2014.

[13] Ian Karambelas. Spear Phishing: The Secret Weapon Behind the Worst Cyber Attacks. *Cloudmark*, https://blog.cloudmark.com/2016/01/13/spear-phishing-secret-weapon-in-worst-cyber-attacks/, January 2016. Accessed: 2018-10-22.

[14] Panagiotis Kintis, Najmeh Miramirkhani, Charles Lever, Yizheng Chen, Roza Romero-Gońez, Nikolaos Pitropakis, Nick Nikiforakis, and Manos Antonakakis. Hiding in Plain Sight: A Longitudinal Study of Combosquatting Abuse. In *Proceedings of the 2017 ACM Conference on Computer and Communications Security (CCS)*, Dallas, TX, USA, October 2017.

[15] Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System. In *Proceedings of the 2007 Conference on Human Factors in Computing Systems (CHI)*, pages 905–914, San Jose, CA, USA, April 2007.

[16] Stevens Le Blond, Adina Uritesc, Cédric Gilbert, Zheng Leong Chua, Prateek Saxena, and Engin Kirda. A Look at Targeted Attacks Through the Lense of an NGO. In *Proceedings of the 23rd USENIX Security Symposium*, San Diego, CA, USA, August 2014.

[17] Suqi Liu, Ian Foster, Stefan Savage, Geoffrey M. Voelker, and Lawrence K. Saul. Who is .com? Learning to Parse WHOIS Records. In *Proceedings of the 2015 ACM Internet Measurement Conference (IMC)*, Tokyo, Japan, October 2015.

[18] William R Marczak and Vern Paxson. Social Engineering Attacks on Government Opponents: Target Perspectives. In *Proceedings of the 17th Privacy Enhancing Technologies Symposium (PETS)*, Minneapolis, MN, USA, July 2017.

[19] William R Marczak, John Scott-Railton, Morgan Marquis-Boire, and Vern Paxson. When Governments Hack Opponents: A Look at Actors and Technology. In *Proceedings of the 23rd USENIX Security Symposium*, San Diego, CA, USA, August 2014.

[20] Grzegorz Milka. Anatomy of Account Takeover. *Enigma*, https://www.usenix.org/node/208154, January 2018.

[21] Marti Motoyama, Kirill Levchenko, Chris Kanich, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. Re: CAPTCHAs: Understanding CAPTCHA-solving Services in an Economic Context. In *Proceedings of the 19th USENIX Security Symposium*, Washington, DC, USA, August 2010.

[22] Adam Oest, Yeganeh Safei, Adam Doupé, Gail-Joon Ahn, Brad Wardman, and Gary Warner. Inside a Phisher's Mind: Understanding the Anti-phishing Ecosystem Through Phishing Kit Analysis. In *Proceedings of the 2018 APWG Symposium on Electronic Crime Research (eCrime)*, San Diego, CA, USA, September 2018.

[23] Jeremiah Onaolapo, Enrico Mariconti, and Gianluca Stringhini. What Happens After You Are Pwnd: Understanding the Use of Leaked Webmail Credentials in the Wild. In *Proceedings of the 2016 ACM Internet Measurement Conference (IMC)*, Santa Monica, CA, USA, November 2016.

[24] Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. Who Falls for Phish?: A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions. In *Proceedings of the 2019 Conference on Human Factors in Computing Systems (CHI)*, pages 373–382, Atlanta, GA, USA, April 2010.

[25] Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, and Elizabeth Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS)*, pages 88–99, July 2007.

[26] Jonathan Skelker. Announcing some security treats to protect you from attackers' tricks. https://security.googleblog.com/2018/10/announcing-some-security-treats-to.html, October 2018.

[27] Kurt Thomas, Danny Yuxing Huang, David Wang, Elie Bursztein, Chris Grier, Tom Holt, Christopher Kruegel, Damon McCoy, Stefan Savage, and Giovanni Vigna. Framing Dependencies Introduced by Underground Commoditization. In *Proceedings of the 2015 Workshop on the Economics of Information Security (WEIS)*, Delft, The Netherlands, June 2015.

[28] Kurt Thomas, Dmytro Iatskiv, Elie Bursztein, Tadek Pietraszek, Chris Grier, and Damon McCoy. Dialing Back Abuse on Phone Verified Accounts. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security (CCS)*, pages 465–476, Scottsdale, AZ, USA, November 2014.

[29] Kurt Thomas, Frank Li, Ali Zand, Jacob Barrett, Juri Ranieri, Luca Invernizzi, Yarik Markov, Oxana Comanescu, Vijay Eranti, Angelika Moscicki, Daniel Margolis, Vern Paxson, and Elie Bursztein. Data Breaches, Phishing, or Malware?: Understanding the Risks of Stolen Credentials. In *Proceedings of the 2017 ACM Conference on Computer and Communications Security (CCS)*, Dallas, TX, USA, October 2017.

[30] Kurt Thomas, Damon McCoy, Chris Grier, Alek Kolcz, and Vern Paxson. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and Abuse. In *Proceedings of the 22nd USENIX Security Symposium*, Washington, DC, USA, August 2013.

[31] Verizon. 2018 Data Beach Investigations Report. https://www.verizonenterprise.com/resources/reports/rp_DBIR_2018_Report_en_xg.pdf. Accessed: 2018-10-22.

[32] Virus Total. https://www.virustotal.com/#/home/upload. Accessed: 2018-10-22.

[33] Yue Zhang, Jason I. Hong, and Lorrie F. Cranor. CANTINA: A Content-based Approach to Detecting Phishing Web Sites. In *Proceedings of the 16th International Conference on World Wide Web (WWW)*, pages 639–648, May 2007.