Dialing Back Abuse on Phone Verified Accounts

Kurt Thomas[◇] Dmytro Iatskiv[◇] Elie Bursztein[◇]
Tadek Pietraszek[◇] Chris Grier[†*] Damon McCoy[‡]

[◇]Google, Inc [†]University of California, Berkeley
[*]International Computer Science Institute [‡]George Mason University

{kurtthomas, diatskiv, elieb, tadek}@google.com grier@cs.berkeley.edu mccoy@cs.gmu.edu

ABSTRACT

In the past decade the increase of for-profit cybercrime has given rise to an entire underground ecosystem supporting large-scale abuse, a facet of which encompasses the bulk registration of fraudulent accounts. In this paper, we present a 10 month longitudinal study of the underlying technical and financial capabilities of criminals who register phone verified accounts (PVA). To carry out our study, we purchase 4,695 Google PVA as well as pull a random sample of 300,000 Google PVA that Google disabled for abuse. We find that miscreants rampantly abuse free VOIP services to circumvent the intended cost of acquiring phone numbers, in effect undermining phone verification. Combined with short lived phone numbers from India and Indonesia that we suspect are tied to human verification farms, this confluence of factors correlates with a market-wide price drop of 30–40% for Google PVA until Google penalized verifications from frequently abused carriers. We distill our findings into a set of recommendations for any services performing phone verification as well as highlight open challenges related to PVA abuse moving forward.

Categories and Subject Descriptors

K.4.1 [**Public Policy Issues**]: Abuse and crime involving computers

Keywords

Account abuse; phone verification; underground economies

1. INTRODUCTION

In the past decade the increase of for-profit cybercrime has given rise to an entire underground ecosystem supporting large-scale abuse, a facet of which encompasses the bulk registration of fraudulent accounts. Miscreants leverage this market to obtain cheap email addresses and social network credentials for as little as 0.50¢ an account [26], in turn fueling spam and abuse at the expense of millions of users. The deluge of messages that follow seek to monetize victims in a variety of manners: from spamvertised products [14] to phishing and malware attacks that perpetrate software scams [22], clickfraud [7], banking theft [23], or convert infected victims into assets for the pay-per-install market [5].

Web services attempt to rate limit this torrent of automatically generated accounts through CAPTCHAs, email verification, and most recently phone verification. While CAPTCHAs and email accounts are trivially available from the underground for relatively low prices [15, 26], ideally phone numbers represent a scarce resource for criminals that are otherwise globally accessible to legitimate users [16, 17]. Consequently, when Google deployed phone verification as a signup protection, prices on the underground surged from $30 per 1K to over $500. Yet there are signs that criminals have streamlined the circumvention of phone verification. Prices for Google accounts have declined to as low as $85 per 1K at the time of this study.

In this paper, we present a longitudinal study of the underlying technical and financial factors influencing the diminishing effectiveness of phone verification. To conduct our study, we track phone verified account (PVA) abuse on Google over a 10 month period from July, 2013–April, 2014. Our perspective includes 4,695 accounts purchased from a cross-section of 14 *account merchants* selling Google PVA on blackmarket forums and storefronts as well as a sample of 300,000 PVA disabled by Google for abuse. We rely on these dual datasets to monitor the pricing and organization of the market for phone verified accounts; evaluate how miscreants circumvent the intended cost of phone verification; and identify a set of recommendations for preserving the long term viability of phone verification.

We find that merchants are capable of registering a steady stream of thousands of PVA that subsequently sell for $85–$500 per 1K to the underground. This wildly different price range reflects both the financial barrier imposed by phone verification (where some merchants advertise real SIM cards from a variety of regions and prices) as well as the influence of account resellers operating in a similar fashion to spam affiliate programs [14]. Merchants fulfill orders for fully functioning, phone verified accounts within 24–48 hours, though the lifetime of accounts is dubious; 68% of the PVA we purchase are disabled within a month of their changing hands despite laying dormant, likely due to re-used infrastructure.

We analyze the registration process tied to abusive PVA to understand the root source of phone numbers and the most frequently abused carriers. We find that 24% of PVA dis-

CCS'14, November 3–7, 2014, Scottsdale, Arizona, USA.
ACM 978-1-4503-2957-6/14/11.
http://dx.doi.org/10.1145/2660267.2660321.

abled by Google are verified with free VOIP numbers from Bandwidth.com (which services Google Voice, Pinger, and other providers [3]), effectively allowing miscreants to circumvent the cost of acquiring SIM cards. The remaining accounts in our dataset are verified with phone numbers tied to a variety of mobile carriers, the most popular of which originate from India and Indonesia. We find evidence these regions are traditionally related to CAPTCHA farms and underground hired labor, suggesting that abusive phone verification may be a manual endeavor. Combined with regular re-use of short lived phone numbers, this confluence of factors correlates with a market-wide price drop of 30–40% for Google PVA from November, 2013–February, 2014 until Google penalized verifications from frequently abused carriers.

Based on our findings, we produce a set of recommendations and best practices for services that rely on phone verification. In particular, we propose a carrier reputation system that automatically penalizes SMS and VOIP providers consistently associated with abusive accounts. Alternative approaches—such as blacklisting phone numbers upon abuse detection to prevent re-use—are too slow compared to the velocity that phone numbers appear and disappear. Ultimately, we argue that a global phone number reputation similar to IP and domain reputation systems [1, 12] is required to prevent miscreants from amortizing the cost of abusive SIMs and VOIP numbers across multiple services, as well as to prepare for the potential of compromised phones—already a challenge for banking two-factor authentication [27]—serving as a platform for verifying accounts in the future.

In summary, we frame our contributions as follows:

- We conduct a 10-month longitudinal study of the financial and technical challenges related to phone verified abuse.
- We find an increased reliance on VOIP numbers and inexpensive SIMs from India and Indonesia—likely tied to manual verification farms—correlate with a price drop of 30–40% for Google PVA.
- We evaluate a number of underground practices including phone re-use, phone access durations, and preferred carriers.
- We distill our findings into a set of recommendations and best practices for services that rely on phone verification.

2. BACKGROUND

Phone verification is a single iteration in a long evolution of abuse safeguards that aim to prevent the bulk registration of accounts. We provide an overview of the process behind phone verification, how the underground market has undermined prior protections, and privacy and ethical standards we obied by when studying phone verified abuse.

2.1 Phone Verification Process

Phone verification serves as both an initial *signup protection* as well as an *abuse escalation* where services prevent suspicious account from conducting further actions until after verification. To start the verification process, a client provides a number they wish to associate with their account. The server then sends a challenge PIN via SMS or voice to that number which the client must correctly enter into a web form to prove receipt and complete the process. Phone verification is currently employed by Google, Facebook, Twitter, LinkedIn, and Craigslist among other services to combat abuse as well as for security and account recovery purposes.

Phone verification imposes a cost on both criminals and services. For criminals, a single number typically has a hard limit on the quantity of accounts it can be associated with. Re-use also exposes bulk accounts to clustering where one abuse violation can trigger a cascade of deactivations across correlated accounts. Consequently, miscreants require a constant stream of fresh numbers to seed registrations. Conversely, services employing phone verification as a defense incur a fee for each SMS or voice challenge. This exposes services to typical operational costs as well as resource exhaustion attacks where miscreants request SMS verification for a deluge of numbers tied to expensive carriers to incur exorbitant SMS fees, a threat we discuss further in Section 6.

2.2 Evolution of Abuse Safeguards

Phone verification builds on a history of defenses that includes IP reputation, CAPTCHAs, and email verification. Ideally, these are scarce resources for criminals to acquire that otherwise exert little friction on legitimate users. In practice, many of these components are readily available from the underground.

IP Addresses: Services can leverage IP addresses as a weak identity tied to newly registered accounts. When thousands of accounts are registered from a single IP, there is a strong likelihood of abuse. To circumvent detection, criminals rely on compromised hosts and proxy services to acquire access to tens of thousands of IPs [26]. Anecdotally, we see advertisements for proxies as low as $250/mo for 15,000–30,000 IPs on blackmarket forums.

CAPTCHAs: CAPTCHAs—intended as human-solvable tasks that prevent automation—have become a staple of the underground economy [15]. Services such as *http://spamvilla.com* advertise automated CAPTCHA solvers with 50% accuracy for $30/mo, while human CAPTCHA farms such as *http://antigate.com* outsource CAPTCHA solving to an array of laborers operating out of India, Pakistan, Ukraine, Russia, Vietnam, and Indonesia for $0.70 per 1K solutions. The availability of manual solvers undermines the feasibility of CAPTCHAs (though such services are not free).

Email Verification: Email verification serves to tie the rate miscreants can create accounts to the rate they can acquire email addresses. This effectively outsources abuse prevention to email providers who in turn must rely on alternative signals. In response, email addresses have become a fundamental resource of the underground. Hotmail.com and Yahoo.com accounts are available from merchants for as low as $5 per 1K [26].

Each of these scenarios highlight how the underground evolves over time to respond to new defenses. While bleak from a defenders perspective, each successive protection increases the cost of accounts, cutting into the bottom line of spam and abuse.

2.3 Privacy and Ethical Considerations

Part of our study requires interacting with underground merchants selling Google phone verified accounts as well as

analyzing registration data tied to abusive signups. We build on the guidelines originally discussed by Thomas et al. [26] for interacting with the account underground. Prior to our study, we worked with the authors respective institutions as well as Google to set down a policy for purchasing accounts. We conduct all purchases (which would otherwise violate Google's Terms of Service) with Google's express permission. Furthermore, even though merchants provide us with passwords, we never access accounts. Finally, we restrict our analysis to merchants who publicly advertise Google phone verified accounts; we do not purchase accounts beyond this scope nor attempt to deceive or coerce the merchants involved.

3. CAPTURING ABUSIVE ACCOUNTS

To conduct our study, we rely on two sources of phone verified accounts (PVA): *purchased accounts* acquired from a cross section of the underground economy and a sample of *abusive accounts* disabled by Google for Terms of Service violations related to spam and abuse. We combine these two datasets to provide insights into the pricing of phone verified accounts as well as to understand the scope of Google phone verified abuse.

3.1 Purchased Accounts

Our purchased account dataset consists of 2,217 PVA that we buy in July 2013 at the onset of our study and a second set of 2,478 we purchase at the conclusion of our study in April 2014. We rely on purchasing to validate the authenticity of merchants as well as to understand the market organization for phone verified accounts. We provide an overview of how we identify account merchants, the prices they charge, and the duration merchants stockpile accounts. We find that 68% of the accounts we purchase in July are disabled by Google's infrastructure within one month. Given this high coverage, we elected not to conduct regular repurchases and instead concentrate our analysis on PVA disabled by Google throughout our study. We believe this minimizes our financing of underground merchants without sacrificing access to a representative sample of PVA abuse. We rely on our second purchase in April 2014 to understand how the market has adapted, providing a detailed comparison in Section 5. We restrict the remainder of our discussion in this section to our first purchase set.

3.1.1 Merchants

We identify a cross section of 14 merchants advertising access to Google accounts (among other services) on web storefronts, blackhat forums, and freelance labor pages. For operational concerns we refrain from documenting the identities of the merchants we solicit. Advertisements range in sophistication from automatically generated accounts with no profile information to "manually generated" accounts with "real SIM cards" from Eastern Europe which cost substantially more. From this bazaar, we elect to purchase 2,217 PVA split across 3 merchants on blackhat forums and 4 merchants operating their own storefronts in July, 2013. Merchants fulfilled all orders in 24-48 hours with working, phone verified accounts. We provide a summary of these purchases in Table 1 which we reference throughout this section.

As an extension of purchasing, we also track the pricing of Google PVA (and non-PVA) based on public listings advertised by 6 of the 14 merchants we identify.[1] We poll this data on a weekly basis from July, 2013 until the conclusion of our study in April, 2014. We rely on this historical pricing data to understand the stability of the PVA marketplace and to understand how price correlates with adaptations in phone verification techniques. When available, we also track the prices of Facebook and Twitter PVA accounts which we use to understand the burden imposed by phone verification as a general technique.

Asset	Price/1K	Volume	Disabled
Google PVA	$85	105	77%
Google PVA	$100	1,000	89%
Google PVA	$172	168	100%
Google PVA	$200	100	0%
Google PVA	$300	103	11%
Google+ PVA	$135	81	100%
YouTube PVA	$95	220	100%
YouTube PVA	$153	98	5%
YouTube PVA	$276	192	0%
YouTube PVA	$300	100	28%
YouTube PVA	$500	50	0%

Table 1: List of assets we purchase, the associated price per 1K, the volume we purchase, and whether the accounts are eventually disabled.

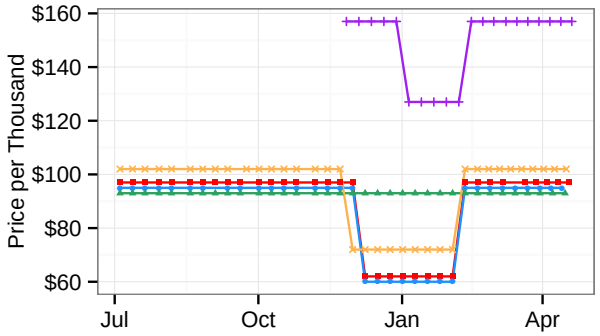

Figure 1: Historical pricing data for Google PVA merchants from July, 2013–April, 2014. A market wide price decrease of 30–40% is visible from November until February for PVA.

3.1.2 Account Pricing

Prices for the accounts we purchase in July, 2013 range from $85 per 1K at the lowest and $500 at the highest, as detailed in Table 1. We provide a breakdown of the historical prices these and other merchants charge throughout our study in Figure 1, with prices over $250 omitted for clarity. Prices in this upper bracket were $250, $300, $350, $500, and $600 per 1K accounts. These prices never changed throughout our monitoring.

Despite a large pool of competing storefronts, we find no evidence of merchants attempting to undercut one another by lowering prices. Instead, the cost of Google PVA remain fixed throughout our study, with the exception of a single

[1]Other merchants rely on email or Skype conversations to determine up-to-date pricing which precludes passive monitoring.

Service	Reg. Cost	PVA Cost	Increase
Google	$80	$100	1.25x
Youtube	$270	$349	1.29x
Youtube	$80	$150	1.875x
Google	$120	$230	1.9x
Google	$80	$500	6.25x
Facebook	$300	$600	2x
Facebook	$70	$350	5x
Facebook	$400	$1800	4.5x
Twitter	$20	$500	25x

Table 2: Price difference between phone verified and regular accounts for Google, Facebook, and Twitter based on advertisements from 3 merchants. Despite a wide range of prices, phone verification tends to impose a 1.25x–6.25x increase, with the exception of Twitter at 25x.

market-wide drop lasting November, 2013–February, 2014. During this period, almost all of the merchants we tracked (with the exception of those in the upper bracket) lowered their pricing by 30–40% before returning to their previous rate. The correlated behavior of merchants leads us to believe that many storefronts are merely resellers for the same miscreant in a similar fashion to spam affiliate programs [14].

3.1.3 Cost of Phone Verification

The primary goal of phone verification is to throttle the rate miscreants can register fraudulent accounts, and as a byproduct, increase the cost of credentials. While we cannot determine the fees that merchants pay to acquire fresh phone numbers, we can measure how merchants pass these costs on to blackmarket consumers. Of the merchants we track, three simultaneously advertised non-PVA and PVA equivalents for Google as well as Facebook, while one merchant advertised access to Twitter non-PVA and PVA. (We were unable to find merchants advertising both LinkedIn or Craigslist PVA and non-PVA.) We use these merchants to measure the *price increase* imposed by phone verification. Assuming that merchants rely on the same infrastructure to register PVA and non-PVA, this allows us to isolate the impact of phone verification from variable merchant sophistication and registration safeguards across services.

Table 2 shows the relative price increase underground merchants charge for phone verification. While the base price of accounts are wildly different between merchants, this increase is relatively fixed: 1.25x–6.25x for Google and 2–5x for Facebook. The 25x increase for Twitter is likely a result of merchants not yet adapting to phone verification on the service, with PVA accounts emerging only at the end of March, 2014 (a month before our study concluded). We observed a similar drastic price difference with the initial release of Google PVA in 2012, where prices were 17x their non-PVA equivalent. We note that a direct comparison between PVA multipliers is difficult due to varying service-level policies on the number of accounts that can be associated with a single phone or whether certain phone numbers are prohibited as verification endpoints.

We caution there is no indication whether blackmarket consumers are willing to bare the fees charged by account merchants. Equally opaque is whether the price differential between non-PVA and PVA accounts is grounded in the scarcity of phone numbers, demand, or consumer naivety. Consequently, we explore the relation between phone verification techniques and market price, particularly during the market-wide price reduction, further in Section 4.

3.1.4 Stockpiling

The freshness of accounts is an important metric for whether merchants conduct real-time bulk registrations or instead rely on outdated stockpiles. We measure the age of the 2,217 accounts we purchase as the delta between the time we order accounts versus the time merchants registered the accounts. Accounts range in age from 1–164 days, with an average age of roughly 27 days. Our results indicate that merchants are not reliant on old stockpiles, but instead have access to recently registered accounts. Paired with stable pricing throughout our analysis, this suggests that merchants have a regular supply of phone numbers at their disposal.

3.1.5 Disable Rate

Inactive accounts that merchants stockpile are not immune to abuse detection. We measure the volume of accounts per merchant that Google disables (independent of our purchasing and analysis), shown in Table 1. Overall, Google disables roughly 68% of the accounts we acquire within one month of their purchase. We find that cheaper accounts are more frequently correlated with being caught and deactivated by Google, indicating that price may have some bearing on the effort account merchants put into bulk registering accounts (e.g. limiting the reuse of infrastructure to avoid clustering). For the purposes of our study, the high recall rate allows us to rely on sampling abusive accounts disabled by Google without risk of omitting a large market segment of PVA abuse. We note however that without regular repurchases, we cannot guarantee the detection rate remains stable throughout our analysis.

3.2 Abusive Accounts

The bulk of our analysis relies on a retroactive random sample of 300,000 Google PVA created and disabled for spam and abuse between July, 2013–April, 2014. No account information ever leaves Google datacenters or is accessed in non-aggregate form by external researchers. For each of these accounts (as well as our purchased account dataset), we have access to the registration IP, registration phone number, and other signals tied to the registration process. We note that due to potential delays in abuse detection, we may underestimate the volume of abuse towards the tail end of our collection period. We consider this limitation whenever we discuss trends in the volume of abuse over time or changes in registration behaviors.

4. ANALYZING ABUSIVE ACCOUNTS

A fundamental question of our investigation is the sustainability of phone verification as a defense against bulk account creation. We find evidence that phone verified abuse is a persistent threat. To dissect this problem, we analyze the origin of abusive number and techniques miscreants use to maximize the value they garner from a single phone number. We relate these technical measurements that capture the complexity of creating phone verified accounts to the prices merchants charge. Finally, we analyze the effectiveness of other registration safeguards including IP reputation, CAPTCHAs, and secondary email addresses.

4.1 Origin of Abusive Phone Numbers

We examine multiple facets tied to the origin of phone numbers including the country of origin, the carrier providing service, and whether fraudulent accounts are registered with collocated IPs and phone numbers. We acquire these phone signals from an MSISDN[2] database used by Google to map phone numbers to carrier data (including whether the number is VOIP). We note that similar databases are publicly available, though typically for some fee.

4.1.1 Breakdown by Country

We examine the country code of each phone number associated with our abusive and purchased accounts to capture which regions serve as the most popular verification endpoints. We find that the United States is the single largest origin of phone numbers, accounting for 27% of abusive PVA in our dataset. This is followed in popularity by India (22%), Indonesia (12%), Nigeria (4%), South Africa (4%), and Bangladesh (4%), with other regions accounting for 28% of abuse. We note that receiving an SMS in all of these top countries other than the United States is free.

Bulk access to phone numbers in these regions appears to be a variable process. Figure 2 shows a weekly breakdown of the top six countries serving as verification endpoints throughout our study. Phones from India, while prevalent at the onset of our measurement (contributing nearly 40% of new PVA), has fallen off in favor of Indonesia. In contrast, phones from the United States dominate 60% of new PVA registrations from October–February. This period overlaps with the drastic price reduction we observe from November–February, a phenomenon we explore further in the next section.

For the accounts we purchased at the onset of our study, 97% were verified with phone numbers from the United States while 3% were associated with numbers from Ukraine. We find that only one of the 7 merchants we solicit rely on non-US numbers. If we examine pricing based on the region that phones numbers originate from, merchants appear to charge arbitrarily for accounts verified with US numbers. Such accounts range $85–300 per 1K accounts, while the sole merchant verifying accounts from Ukraine charged $500 per 1K. Our findings indicate that the origin of phone numbers alone cannot explain the cost of an account or why certain merchants are more likely to have their stockpiles disabled.

4.1.2 Breakdown by Carrier

We further subdivide countries based on the abused carriers operating in each region, the results of which we show in Table 3. Bandwidth.com—a VOIP provider in the US tied to multiple free telephony services including Pinger and Google Voice [3]—represents the single largest gateway for abuse. This is followed in popularity by a multitude of mobile carriers predominantly operating out of India and Indonesia. We evaluate each of these verification approaches separately.

VOIP Abuse: VOIP in particular poses a significant threat to the intended cost of phone verification. Services such as Pinger [18] and TextPlus [24] allow new customers to register for a free, SMS-receivable number in exchange for solving a CAPTCHA or email verification challenge. Such resources

[2]An MSISDN is the unique international representation of a phone number which is associated with a SIM card.

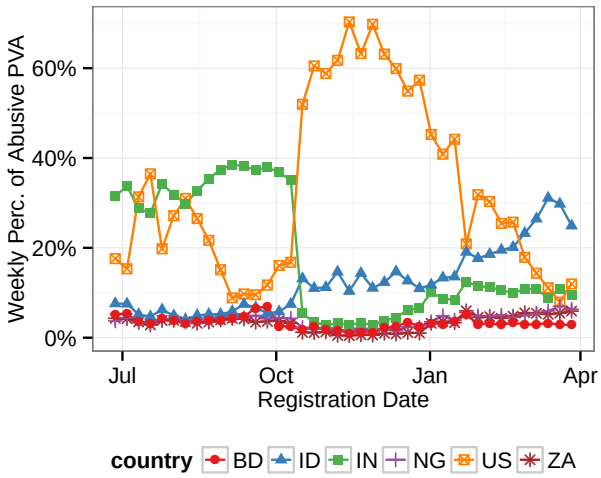

Figure 2: Weekly breakdown of the top 6 country codes associated with abused phone numbers. The most popular origins of numbers are the United States (US), India (IN), Indonesia (ID), Nigeria (NG), South Africa (ZA), and Bangladesh (BD).

are cheaply available from the underground as we previously discussed in Section 2. Similarly, services such as Google Voice allow miscreants to convert an existing phone number (including US VOIP numbers) into multiple new phone numbers. This creates an *abuse multiplier* that allows miscreants to amortize the cost of the original phone number seed as well as mask the original carrier. All of these services are available online, opening up the possibility for miscreants to scrape page content to automate SMS verification challenges. In total, 24% of all abusive PVA in our dataset were verified with VOIP numbers.

The merchants we solicit readily exploit cheap VOIP numbers to circumvent the intended cost of phone verification. Of the accounts we purchased, 97% were verified via numbers tied to a mixture of VOIP providers including Bandwidth.com, Level 3, and Telengy. This trend is also represented in Figure 2 where 94% of all US numbers used to verify accounts between October–January were VOIP. The decrease in US phone numbers after January is the result of Google penalizing new registrations tied to frequently abused US VOIP providers. This confluence of events correlates with the 30–40% price drop in accounts that we observe from November–February after which prices returned to their normal levels. While we cannot provide definitive proof, our results suggest that market prices can serve as an indicator of the underlying performance of abuse safeguards.

Mobile carriers: VOIP numbers alone do not explain the entire phone verified abuse ecosystem; a second substantial component is fueled by mobile carriers tied to India and Indonesia including PT, Bharti, and Vodafone. Our understanding of how miscreants acquire phone numbers from these regions and subsequently respond to SMS challenges is less clear than VOIP. Anecdotally, when we conducted our search to identify merchants selling PVA, we also encountered an underground market segment surrounding *verification as a service*. Sites such as *http://sms-area.org* advertise automated APIs for phone verifying Vkontakte, Google, and Facebook accounts. Prices for these services are as low

Rank	Carrier	Country	Popularity
1	Bandwidth.Com	US	19.91%
2	Pt	ID	7.29%
3	Bharti	IN	5.31%
4	Vodafone	IN	4.04%
5	Mtn	NG	2.99%
6	Idea	IN	2.79%
7	Telekomunikasi	ID	2.23%
8	Aircel	IN	2.11%
9	Tata	IN	1.87%
10	Viettel	VN	1.71%
11	Reliance	IN	1.71%
12	Mtn	ZA	1.52%
13	Gramenphone	BD	1.51%
14	Vodacom	ZA	1.29%
15	Bsnl	IN	1.28%
16	Excelcom	ID	1.16%
17	Hutchison	ID	0.95%
18	Level 3	US	0.86%
19	Cell	ZA	0.84%
20	Telengy	US	0.81%
-	Other	-	37.80%

Table 3: Top 20 carriers used by abusive phone verified accounts. Verification challenges are predominantly sent over VOIP or to a concentrated set of carriers in India and Indonesia.

as $140 per 1K verification codes for mobile (non-VOIP) numbers originating from Russia, Kazakhstan, and Belarus. Similarly, miscreants can acquire SIM cards in bulk. We see resellers advertising prices of $140–420 per 1K SIM cards for Russian carriers such as Beeline, MTS, and MegaFon. While we can only speculate, discussions we observe on underground forums suggest that workers manually respond to verification challenges using modified cell phones to simplify cycling through SIM cards. This reflects related strategies for CAPTCHA solving that rely on manual laborers operating out of India and Indonesia as discussed in Section 2.

4.1.3 *Collocation of Phones and IP Addresses*

One potential measurement of the validity of a newly registered PVA is the collocation of a phone's country of origin and the geolocation of the IP address a miscreant uses to register the account. Figure 3 shows a weekly breakdown of the six most popular IP geolocations tied to signups at a country granularity. We find that trends in IP geolocation nearly mirror that of phone origins (previously presented in Figure 2). The exception to this trend is the decreased popularity of US IP addresses, which may be a direct consequence of the difficulty of purchasing hosts in the US from the pay-per-install market [5] or the respective cost of freelance labor in the US compared to other regions. Quantitatively, 60% of abusive accounts in our dataset share the same IP and phone origin, while the same is true for only 33% of purchased accounts. This deviation for purchased accounts is a consequence of the majority of phone numbers coming from the United States while IPs originate in India, a trend that is also reflected in Figure 3 for abusive accounts back in July. Our results indicate that miscreants bulk registering PVA take care to mimic the expected behavior of legitimate registrations beyond avoiding clustering.

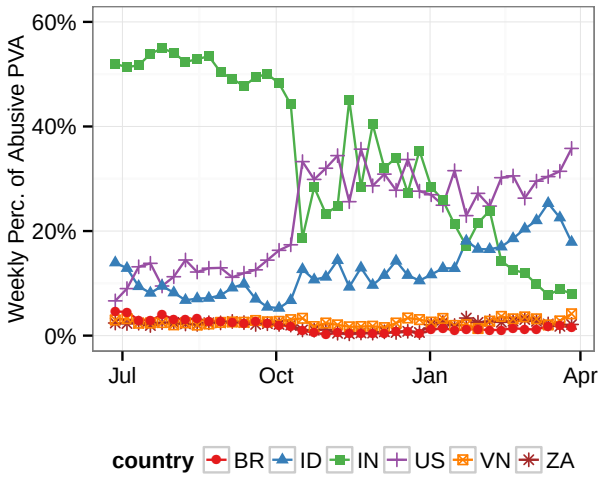

Figure 3: Geolocation of IPs related to abusive signups. The most frequent source of abusive IPs mirrors that of abusive sources of phone numbers.

4.2 Lifecycle of a Phone Number

Miscreants have a multitude of strategies for how they leverage SMS access once acquired. We measure the number of accounts that miscreants associate with the same phone number, estimate the duration miscreants control phone numbers, and identify whether miscreants opt for SMS or voice challenges.

4.2.1 *Phone Reuse*

Miscreants can reuse a phone number multiple times to amortize the cost of acquiring access to a VOIP number or SIM card. For each phone numbers in our abusive account sample we calculate the total number of *all* registered accounts that share the same number, regardless if Google disabled those accounts for abuse. We repeat this process again for phone numbers tied to purchased accounts. To serve as a comparison, we also obtain re-use statistics tied to a random sample of 300,000 benign phone numbers (provided by Google). Figure 4 shows a summary of our results.

We find that 36% of numbers associated with abusive accounts are unique. This results in a skewed distribution, where the top 10% of phone numbers (ranked by popularity) are used to register 23% of accounts. A similar result appears for numbers tied to purchased accounts where only 13% of numbers are unique and the top 10% of numbers are used to create 25% of accounts. In total, phone clusters of size 5 or greater contain 58% of abusive accounts. As a result, a safeguard that restricts the number of accounts miscreants can associate with a single phone number can have a substantial impact on the volume of phone numbers required to sustain PVA abuse while having little impact on legitimate users.

4.2.2 *Phone Access Lifetime*

During our investigation of the underground we found that blackmarket merchants offering SMS as a service frequently advertised a limited window of availability—clients would receive access to a number for 30-90 days, after which that number would no longer be accessible. We estimate the life-

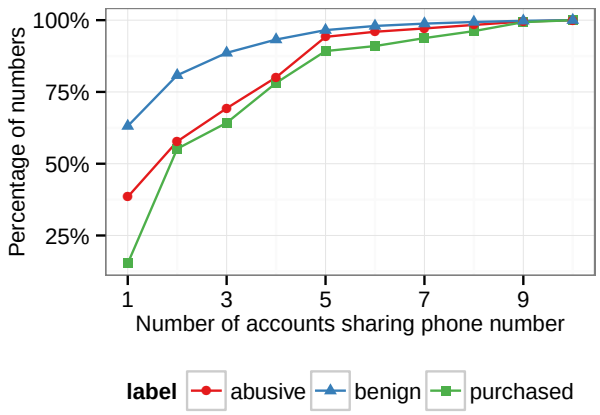

Figure 4: CDF of the size of account clusters all verified with the same phone number.

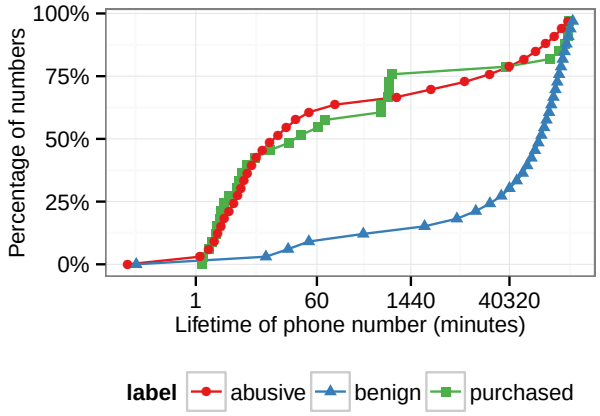

Figure 5: CDF of time between successive account registrations sharing the same phone number.

time that miscreants control a number as the time between its first use to its last use within our measurement window, restricting our analysis to numbers used at least five times. This timestamp is determined based on the creation time of an account. Our results are shown in Figure 5.

We find that 62% of phone numbers tied to abusive PVA have a lifetime of only 1 hour. Similarly, 55% of phones tied to purchased accounts have a lifetime of only 1 hour and 76% less than a day. This indicates that once miscreants acquire a phone number, there is a high velocity period of abuse, after which the phone number is never used again. Consequently, blacklisting an individual phone number upon an abuse report is ineffective at preventing future abuse if the delay between an account's creation and deactivation is greater than a day.

Defenders can leverage this short lifetime to their advantage. Miscreants that purchase accounts from the underground have no means to re-verify access to a phone number once the account transfers hands. Similarly, merchants that bulk generate thousands of accounts would need to catalog and retain the VOIP number or SIM card used to verify each account, something that may be impossible. We explore how services can perform re-verification without increasing friction on legitimate users further in Section 6.

4.2.3 *Verification Challenge Type*

When miscreants verify fraudulent PVA they can opt to receive challenge codes via SMS or voice. We find that 90% of abusive accounts rely on SMS codes, while the same is true for 85% of purchased accounts. This holds for benign accounts as well where users verify their phone 94% via SMS. We find that the verification method miscreants use is independent of whether they rely on VOIP or mobile numbers. The exception to this rule is PT in Indonesia; 31% of verification codes served through this carrier were conducted over voice. This further suggests that human verifiers may respond to challenges from this region, though we note that voice transcription software is an alternative explanation. We cannot draw any conclusion as to how SMS codes from mobile phones are recovered, though we note that VOIP services such as Google Voice digitize text messages that miscreants can scrape.

4.3 Alternative Account Challenges

Phone verification is one layer in a set of successive challenges miscreants must pass in order to register a new account. These other challenges include IP signals, CAPTCHA solving, and providing a secondary email. We briefly examine how miscreants circumvent each of these measures. When possible, we compare our results to those of the Twitter account market examined by Thomas et al. [26] to determine whether merchant techniques for bulk registration generalize across services.

4.3.1 *IP Diversity*

Miscreants rely on IP addresses from India (32%), the United States (19%), Indonesia (11%), Vietnam (2%), and a range of other countries to register abusive PVA. Our purchased accounts were registered via IPs exclusively from India (70%) and the United States (30%). In contrast, miscreants targeting Twitter relied on a much more diverse range of countries; India, the most popular region, accounted for only 8% of abuse [26]. We believe this difference stems directly from miscreants relying on IP addresses collocated with phones, as previously discussed in Section 4.1.

Once miscreants have access to an IP they register tens to hundreds of accounts from that portal. Figure 6 shows a breakdown of IP reuse amongst abusive, purchased, and a random sample of benign accounts. The merchants we purchase from clearly restrict the number of accounts they register from a single IP. Other miscreants bulk registering accounts are not so cautious. Nevertheless, IP reuse is not a foolproof signal; benign accounts are frequently registered via the same IP due to NATing and mobile traffic.

4.3.2 *CAPTCHA Solving*

We find that CAPTCHAs are a minor roadblock in the account creation process. In total, 56% of abusive accounts solved a CAPTCHA. Miscreants solved these CAPTCHAs correctly 96% of the time. This accuracy is a strong indication of human solvers based on the results of Motoyama et al. [15], though it may be possible that OCR CAPTCHA solving has vastly improved since then. Our findings differ from Thomas et al. [26] where the majority of Twitter account merchants relied on what appeared to be automated solvers with roughly 7% accuracy. We cannot directly measure the time it takes miscreants to solve a CAPTCHA. However, we can estimate the overall time it takes miscre-

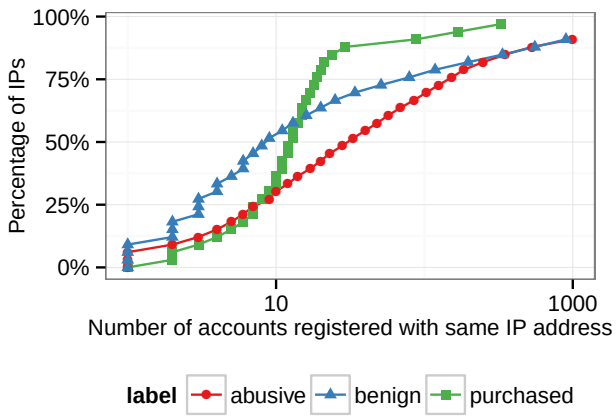

Figure 6: Reuse of IP addresses for registering accounts.

ants to register a new account. We find it takes a median of 1.86 minutes from the time Google displays a signup form to the time miscreants submit a response. In contrast, benign registrations take a median of 3.36 minutes. This overall timing yields little signal into the abusiveness of a newly minted account. It is likely miscreants purposefully delay their automation speeds to avoid rudimentary timing detection.

4.3.3 *Secondary Email Address*

Google allows new accounts to associate a secondary email for recovery purposes. While not required, we find that 83% of purchased accounts and 34% of abusive accounts provide a secondary email. A breakdown of the most popular email providers tied to abusive accounts is shown in Table 4. We observe that 52% of abusive PVA list a second Gmail address as a recovery email. For purchased accounts (not shown in the Table), 57% use a secondary Gmail address. The remaining 43% of accounts use Hotmail.com, one of the cheapest emails on the underground that we observe merchants selling for $5 per 1K. Apart from an increased prevalence of Gmail addresses, the most popular email providers are identical to those abused by Twitter account merchants with the exception of rediffmail.com [26].

The frequency of Gmail recovery addresses amongst bulk generated accounts is a result of *email chaining* where miscreants specify a recovery address tied to a previously generated account they control. We build a graph of all email-secondary email pairs and find that 95% of accounts we purchased share a link with another purchased accounts. Miscreants form chains that are 2–4 accounts long before forming a cycle with the start of the chain. We note that forming a cycle is possible because secondary addresses are not validated on signup. This practice removes any cost associated with merchants purchasing secondary emails while providing a credible veneer in the event recovery emails are factored into spam analysis.

5. REVISITING THE UNDERGROUND

Following the conclusion of our market and abuse monitoring in April, 2014, we revisit the underground merchants we originally solicited to procure a fresh set of 2,478 Google PVA. We use these accounts to independently verify the long-term impact of Google's penalization of frequently abused carriers in January, 2014 (discussed in Section 4.1).

Rank	Email provider	Popularity
1	gmail.com	52%
2	rediffmail.com	10%
3	yahoo.com	10%
4	hotmail.com	6%
5	mail.com	2%
-	Other	19%

Table 4: Top 5 email providers used as recovery addresses for abusive PVA.

Of the accounts we purchase, 29% are *stale accounts* registered back in mid-October, 2013 while the remaining 71% are *fresh accounts* created in mid-April, 2014. We discuss the implications each set has on our understanding of the account blackmarket and abusive phone verification.

Stale Accounts: Merchants providing stale PVA are still liquidating stockpiles they generated during the period of rampant VOIP abuse. All of the accounts in this set were phone verified via Bandwidth.com, re-affirming our findings presented in Section 4.1. Even though old stockpiles are regularly disabled (discussed in Section 3), we find that merchants are still able to retain some operational credentials. The presence of such accounts also suggests a lack of liquidity in the market; merchants have held on to accounts for 6 months without sale. One potential explanation is consumer skepticism on the merchant's credibility or the exorbitant fees they charge.

Fresh Accounts: Merchants providing fresh PVA have adapted to the new phone verification requirements imposed by Google. Of fresh accounts, only 12% were verified via Bandwidth.com—an evolution we see mirrored at the end of our analysis in April. Instead, merchants verify 74% of PVA with a previously unobserved US carrier that is likely a VOIP provider and 9% from a carrier related to TextMe (an Android and iOS VOIP app) which we previously observed. Our results highlight the resilience of the underground to intervention. Nevertheless, we believe that if services force merchants away from VOIP it will in turn raise the operational costs of miscreants and ultimately cut into the profitability of spam and abuse.

Noticeably absent in our fresh account sample are any phones tied to carriers in India or Indonesia as we see in our abusive dataset. This may reflect a limitation in our coverage of the underground market or alternatively indicate that some spammers are vertically integrated and beyond our access. As such, while we believe that underground infiltration provides an invaluable oracle into the performance of abuse safeguards, it should be supplemented with service-side data to provide a dual perspective on abuse.

6. ADAPTING PHONE VERIFICATION

We distill our underground market analysis of PVA abuse into a set of recommendations and best practices for services reliant on phone verification. While our perspective of PVA abuse is limited to Google, we believe the threats we identify are fundamental to phone verification and thus apply outside its confines. We also take a moment to discuss open challenges for phone verification services moving forward including resource exhaustion attacks and the potential for compromised phones.

6.1 Restricting Phone Numbers

The long term validity of phone verification hinges on services enforcing the scarcity of phone numbers as an underground resource. To satisfy this requirement, we propose two solutions: a *carrier reputation* system which tracks the most frequently abused telephony providers at a coarse level and *phone reputation* that provides fine-grained abuse information related to phone numbers, similar to existing IP, domain, and social reputation systems [1, 12, 13, 28].

6.1.1 Carrier Reputation

As our analysis shows, account merchants gravitate towards free or inexpensive regional telephony carriers for the bulk of their phone numbers. If we examine the aggregate contributions of carriers ranked by popularity–shown in Figure 7–we find miscreants verify 20% of all abusive PVA from a single carrier and 50% of PVA from the top 10 carriers. Blacklisting carriers outright (with the exception of VOIP) is not an option. Table 5 shows a breakdown of the top 10 carriers tied to PVA abuse and the fraction of all accounts verified via those carriers that are considered legitimate by Google (e.g. not disabled). Only Bandwidth.com, the VOIP provider we saw affiliated with rampant abuse, has a low reputation. All other carriers are popular amongst legitimate users.

An alternative to blacklisting is to adaptively throttle the number of accounts that can be tied to a single phone number on a per-carrier basis. While there are legitimate users who rely on re-using phone numbers (discussed in Section 4.2), phone verification services can strike a balance between user friction and abuse prevention. In particular, services can restrict phones tied to frequently abused carriers to a one-to-one mapping between accounts and phone numbers, forcing miscreants to acquire more numbers. Other carriers would default to a many-to-one allowance. Carrier reputation scores can also be considered as a factor into machine learning risk evaluations tied to new registrations.

This same system can outright block mobile and VOIP services (if appropriate) when the threshold of abusive accounts drops below an acceptable level. While miscreants may spread their verification endpoints over multiple carriers, this forces criminals to pay higher fees. Craigslist has taken a similar policy to the extreme, blacklisting all VOIP and non-US numbers as verification endpoints [6]. This strategy has kept Craigslist PVA at a price of $3.50–$5 per account (as advertised by account merchants), though it precludes a global user bases. Carrier reputation services currently exist from Telesign and Pindrop, though the methodology or accuracy of these systems is not public.

6.1.2 Phone Reputation

Where carrier reputation can help throttle the creation rate of PVA, phone reputation can outright block abusive registrations. The primary challenge of phone reputation is the velocity of abuse. As we discussed in Section 4.2, 62% of phones used to verify abusive accounts have a lifetime under one hour. Consequently, if abuse reports cannot flag a phone number in time, the window for a reputation system to impact abuse will have passed. One option is for services to restrict the velocity of new registrations tied to the same phone number (e.g. forcing users to wait a period before a phone number is re-usable), but this offers no protection if numbers are single-use. This benefit could be expanded by services sharing abuse information which would prevent *verification as a service* merchants (discussed in Section 4.1) from registering multiple accounts cross-service with a single phone number. We also believe that phone reputation can become a predictive score that assesses the risk of a previously unseen phones tied to newly registered account. Potential features include the sequentiality of phone numbers and whether the phone number has appeared in other contexts before. We leave such a system for future work.

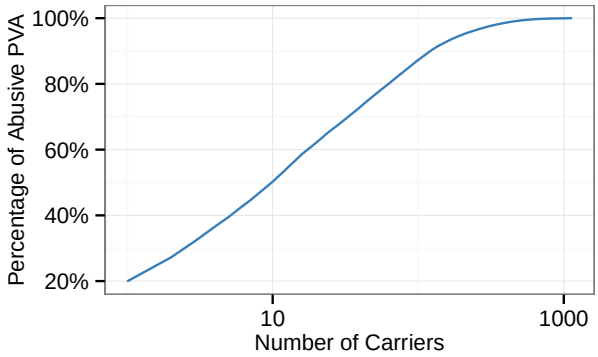

Figure 7: Aggregate contribution of fraudulent PVA of carriers ranked by popularity. Miscreants use the top carrier to verify 20% of all abusive PVA and the top 10 to verify 50% of PVA.

Rank	Carrier	Country	% Good
1	Bandwidth.com	US	41%
2	Pt	ID	91%
3	Bharti	IN	98%
4	Vodafone	IN	98%
5	Mtn	NG	97%
6	Idea	IN	98%
7	Telekomunikasi	ID	99%
8	Aircel	IN	98%
9	Tata	IN	98%
10	Viettel	VN	99%

Table 5: Top 10 carriers used by abusive phone verified accounts and their respective reputations. Blacklisting carriers outright would impact a disproportionate number of legitimate users.

6.1.3 Phone Reverification

The relatively short lifetime of phone numbers—a median of less than 1 hour in our analysis—raises the potential for defenders to re-verify phone numbers as an effective abuse escalation. By limiting re-verification as a response to suspicious activity (e.g. sending an abnormal number of email; rapidly subscribing to hundreds of YouTube channels), legitimate users will avoid the friction imposed by phone verification. Conversely, miscreants who purchased accounts will not have access to the original phone number tied to an account after it changes hands. Similarly, spammers that are vertically integrated and register their own credentials must retain access to thousands of SIM cards, something that may not be possible if miscreants rely on *verification as a service* merchants. This adds yet another dimension to the complexity of circumventing phone verification, where

duration of access becomes just as important as the quantity of phone numbers miscreants can draw from. We note that Facebook may already rely on this practice based on chatter from underground forums, but cannot confirm.

6.2 Open Challenges

6.2.1 Phone Chaining

When miscreants abuse free call forwarding services such as Google Voice, they must register a phone number to serve as the forwarding endpoint. While this initial number is intended as a safeguard—similar to email verification or a CAPTCHA—the existence of other free (virtual) numbers overcomes this protection. This is exacerbated by the potential many-to-one relationship between seed numbers and numbers handed out by forwarding telephony services. This leads to a vulnerability we call *phone chaining*. In particular, miscreants can use a single free mobile number or free VOIP number that fans out to multiple free forwarding numbers. These in turn can serve as forwarding endpoints at other services, with miscreants repeating the process ad nauseam to create an arbitrarily sized tree of phone numbers for verification purposes. Even if virtual number providers prevent chaining between their own phone number pool, miscreants can obfuscate the original identity of a number by cycling through multiple forwarding services for each level of the tree. Absent a comprehensive forwarding blacklist—which is difficult to acquire due to phone ranges constantly changing hands and limited opacity into phone ownership—there is nothing to prevent this practice. An alternative is to have telephony services charge for forwarding numbers, but global adoption is required to prevent miscreants from constantly shifting their operations to free providers.

6.2.2 Resource Exhaustion Attacks

Phone verification is not free. Services must pay a fee for each attempted SMS or phone challenge. The cost of phone verification exposes services to resource exhaustion attacks. Miscreants can request thousands of spurious verification codes for non-existent numbers or phones tied to legitimate victims (who may in turn be charged a delivery fee). Assuming a fixed daily budget, a fail-open system will allow miscreants to bypass phone verification all together; a fail-closed system will block legitimate users from registering new accounts. This challenge also exists for email verification—services that generate excessive bounce notifications expose themselves to throttling on the receiving end that can degrade the delivery rate of important messages. While services can perform a risk analysis for each new registration attempt and block suspicious registrations before a SMS challenge is sent, we view resource exhaustion attacks as an open challenge.

6.2.3 Compromised Phones

The widespread adoption of Android and iOS devices carries with it an emerging risk of mobile malware [8]. Instances already exist of the Zeus banking trojan intercepting two-factor authentication PINs for liquidating a victim's assets [27]. Other threats include mobile drive-by-downloads installing the *NotCompatible* malware family that converts a victim's phone into a mobile proxy [19]. While we find no evidence of miscreants using compromised phones as verification endpoints, there is no reason that compromised phones will not become a commodity market like compromised hosts [5]. This possibility undermines the longterm sustainability of phone verification and any protection provided by carrier-level reputation. Similarly, it forces phone reputation systems into a reactive position, effectively becoming phone blacklists much like existing IP and DNS blacklists [21].

7. RELATED WORK

Monetizing Account Access. Miscreants leverage bulk registered accounts to expose legitimate users to spam, phishing, and malware [9, 11]. At its heart, a substantial amount of spam serves to advertise products and lure people to purchase pharmaceuticals, replica goods, and counterfeit software, often with the aid of complex infrastructure managed by affiliates and affiliate programs [14]. Recent alternatives in monetization include ad syndication services and ad-based URL shortening [25]. These examples highlight the range of strategies that miscreants employ to generate a profit once they gain access to account credentials.

Reputation Systems. Reputation is commonly used to prevent or limit abuse of web services. One technique, with roots in email spam filtering, is the use of IP addresses for reputation [13]. Using IP reputation enables a web service to quickly score actions and label them as abusive or benign. IP reputation can become error prone and new techniques have been developed for evaluating IP address ranges that are populated with a mix of benign and abusive actions [13]. Approaches to reputation have also evaluated automated scoring using network-level features [12]. In addition to IP addresses, other forms of reputation are used to label domains [4], and accounts on web services by leveraging the observed actions of existing users to vet new accounts [28]. Phone reputation over voice has also been explored based on audio features introduced by telephony networks [2], though such features do not extend to numbers used purely for SMS.

Applications of Phone Verification. In addition to preventing fraudulent account registration, SMS and phone verification is used broadly to prevent account hijacking. One widely deployed use is two-factor authentication for online banking. With SMS two-factor, a one-time code is sent via SMS and used along with the username and password to login [20]. Even two-factor authentication is not without vulnerabilities, and criminals have used malware to intercept SMS messages on the target phone [10, 20]. Similar to two-factor authentication, other uses for phone verification include account recovery, and as an informational channel to alert users of changes to their account.

8. CONCLUSION

In this paper, we presented a longitudinal study of the underlying technical and financial factors influencing the diminishing effectiveness of phone verification. To conduct our study, we combined underground intelligence gleaned from 4,695 phone verified accounts purchased from 14 blackmarket merchants as well as 300,000 phone verified accounts disabled by Google for spam and abuse. We found that merchants were capable of registering a steady stream of thousands of PVA that subsequently sold for \$85–\$500 per 1K to the underground. Many of these merchants appeared to operate in a fashion similar to spam affiliate programs,

reiterating that specialization with the underground ecosystem is the norm. We found that merchants used inexpensive VOIP numbers to circumvent the intended cost of acquiring SIM cards, effectively invalidating the defense provided by phone verification. As this practice became a widespread, we observed a simultaneous market-wide price drop of 30–40% for Google PVA until Google penalized verifications from frequently abused phone carriers. Our results highlight how blackmarket monitoring can provide an invaluable oracle into the performance of abuse safeguards. This information vastly simplifies the task of defenders keeping pace with evolutions of the underground, in turn better protecting users against spam and abuse.

9. ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under grant 1237265 and 1237076, by the Office of Naval Research under MURI grant N000140911081, and by a gift from Google. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

10. REFERENCES

[1] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster. Building a dynamic reputation system for dns. In *Proceedings of the USENIX Security Symposium*, pages 273–290, 2010.

[2] V. A. Balasubramaniyan, A. Poonawalla, M. Ahamad, M. T. Hunter, and P. Traynor. Pindr0p: using single-ended audio features to determine call provenance. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 109–120. ACM, 2010.

[3] Bandwidth.com. Who we are. *http://bandwidth.com/about-us*, 2014.

[4] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi. EXPOSURE: Finding Malicious Domains Using Passive DNS Analysis. In *NDSS*, 2011.

[5] J. Caballero, C. Grier, C. Kreibich, and V. Paxson. Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Proceedings of the USENIX Security Symposium*, 2011.

[6] Craigslist. How does phone verification work? *http://www.craigslist.org/about/help/phone_verification*, 2014.

[7] V. Dave, S. Guha, and Y. Zhang. Measuring and Fingerprinting Click-Spam in Ad Networks. In *Proceedings of the ACM SIGCOMM*. ACM, 2012.

[8] A. P. Felt, M. Finifter, E. Chin, S. Hanna, and D. Wagner. A Survey of Mobile Malware in the Wild. In *Proceedings of ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, 2011.

[9] H. Gao, J. Hu, C. Wilson, Z. Li, Y. Chen, and B. Y. Zhao. Detecting and Characterizing Social Spam Campaigns. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010.

[10] L. Goddard. Researchers discover five new samples of zitmo malware for android and blackberry. *http://www.theverge.com/2012/8/8/3227638/zitmo-malware-android-blackberry-samples*, August 2012.

[11] 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.

[12] S. Hao, N. A. Syed, N. Feamster, A. G. Gray, and S. Krasser. Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine. In *USENIX Security Symposium*, volume 9, 2009.

[13] C.-Y. Hong, F. Yu, and Y. Xie. Populated IP Addresses: Classification and Applications. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, 2012.

[14] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, et al. Click trajectories: End-to-end analysis of the spam value chain. In *Proceedings of IEEE Security and Privacy*, 2011.

[15] M. Motoyama, K. Levchenko, C. Kanich, D. McCoy, G. M. Voelker, and S. Savage. Re: Captchas-understanding captcha-solving services in an economic context. In *Proceedings of the USENIX Security Symposium*, 2010.

[16] Pew Research. Emerging nations embrace internet, mobile technology. *http://www.pewglobal.org/2014/02/13/emerging-nations-embrace-internet-mobile-technology/*, 2014.

[17] Pew Research. Mobile technology fact sheet. *http://www.pewglobal.org/2014/02/13/emerging-nations-embrace-internet-mobile-technology/*, 2014.

[18] Pinger. Free unlimited texting to 35 countries from your computer. *https://www.pinger.com/content/text-from-your-computer/*, 2014.

[19] F. Ruiz. Android notcompatible looks like piece of pc botnet. *http://blogs.mcafee.com/mcafee-labs/androidnotcompatible-looks-like-piece-of-pc-botnet*, 2012.

[20] B. Schneier. Two-Factor Authentication: Too Little, Too Late. 2005.

[21] S. Sinha, M. Bailey, and F. Jahanian. Shades of grey: On the effectiveness of reputation-based "blacklists". In *Malicious and Unwanted Software*, 2008.

[22] B. Stone-Gross, R. Abman, R. A. Kemmerer, C. Kruegel, D. G. Steigerwald, and G. Vigna. The Underground Economy of Fake Antivirus Software. In *Economics of Information Security and Privacy*, 2013.

[23] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *Proceedings of the ACM CCS*, pages 635–647. ACM, 2009.

[24] Text+. Free text to anyone in the us or canada. *http://www.textplus.com/*, 2014.

[25] K. Thomas, C. Grier, V. Paxson, and D. Song. Suspended Accounts In Retrospect: An Analysis of Twitter Spam. In *Proceedings of the Internet Measurement Conference*, November 2011.

[26] K. Thomas, D. McCoy, C. Grier, A. Kolcz, and V. Paxson. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and

Abuse. In *Proceedings of the USENIX Security Symposium*, 2013.

[27] TrendMicro. Zeus now bypasses two-factor authentication. *http://blog.trendmicro.com/trendlabs-security-intelligence/zeus-now-bypasses-two-factor-authentication/*, 2013.

[28] Y. Xie, F. Yu, Q. Ke, M. Abadi, E. Gillum, K. Vitaldevaria, J. Walter, J. Huang, and Z. M. Mao. Innocent by Association: Early Recognition of Legitimate Users. In *Proceedings of the 2012 ACM conference on Computer and communications security*, 2012.