Framing Dependencies Introduced by Underground Commoditization

Kurt Thomas[◇]   Danny Yuxing Huang[†]   David Wang[◇]   Elie Bursztein[◇]   Chris Grier[□]
Thomas J. Holt[*]   Christopher Kruegel[§]   Damon McCoy^{‡▽○}   Stefan Savage[†]   Giovanni Vigna[§]

◇Google   †University of California, San Diego   §University of California, Santa Barbara
▽University of California, Berkeley   ○International Computer Science Institute
□Databricks   ‡George Mason University   *Michigan State University

Abstract

Internet crime has become increasingly dependent on the *underground economy*: a loose federation of specialists selling capabilities, services, and resources explicitly tailored to the abuse ecosystem. Through these emerging markets, modern criminal entrepreneurs piece together dozens of *à la carte* components into entirely new criminal endeavors. From an abuse fighting perspective, criminal reliance on this black market introduces fragile dependencies that, if disrupted, undermine entire operations that as a composite appear intractable to protect against. However, without a clear framework for examining the costs and infrastructure behind Internet crime, it becomes impossible to evaluate the effectiveness of novel intervention strategies.

In this paper, we survey a wealth of existing research in order to systematize the community's understanding of the underground economy. In the process, we develop a taxonomy of *profit centers* and *support centers* for reasoning about the flow of capital (and thus dependencies) within the black market. Profit centers represent activities that transfer money from victims and institutions into the underground. These activities range from selling products to unwitting customers (in the case of spamvertised products) to outright theft from victims (in case of financial fraud). Support centers provide critical resources that other miscreants request to streamline abuse. These include exploit kits, compromised credentials, and even human services (e.g., manual CAPTCHA solvers) that have no credible non-criminal applications. We use this framework to contextualize the latest intervention strategies and their effectiveness. In the end, we champion a drastic departure from solely focusing on protecting users and systems (tantamount to a fire fight) and argue security practitioners must also strategically focus on disrupting frail underground relationships that underpin the entire for-profit abuse ecosystem—including actors, infrastructure, and access to capital.

1   Introduction

Over the last two decades, attacks on computer systems have transitioned from rare incidents to ubiquitous events. Part of this transformation has been driven by technology.

Indeed, the combination of universal Internet connectivity and fragile homogeneous software systems provided fertile ground for the development of large-scale host infections with a centralized command and control infrastructure (*c.f.* the DDoS botnets of the early 2000s). However, the more significant evolution—taking place almost entirely in the last decade—has been around the motivation and structure of these attacks. In particular, the rise of e-commerce, both monetized directly through sales and indirectly via advertising, engendered Internet-attached hosts with latent value that could then be monetized via abuse. The confluence of these two factors—the ease with which hosts could be compromised at scale and the fact that each such hosts could be monetized for profit—fueled a bloom in criminal entrepreneurship that underlies most threats we experience today online.

Starting with early partnerships between malware authors and e-mail spammers (largely focused on the simple problem of laundering MTA origin), miscreant innovators soon identified a broad range of monetization strategies and associated technical needs. Through their actions, today we understand that a compromised host can encapsulate a broad range of extractable value: both through its commodity technical resources (e.g., its bandwidth and IP address for sending spam, its CPU for mining crypto-currencies, its storage for hosting content for some scam) and through its unique data resources (e.g., account usernames and passwords entered, PageRank of site, credit card numbers, social network membership, and so on).

Extracting all of this value can be complex and require a range of specialized knowledge and capabilities. Indeed, it would be challenging for any single actor to operate the myriad components making up a modern scam. Instead, the emergence of underground marketplaces has allowed individual actors to specialize in particular capabilities, services, or resources types—without needing to own the entire value chain. Thus, a criminal entrepreneur today will use their own seed capital to purchase individual resources or capabilities *à la carte* (e.g., compromised accounts, CAPTCHA solving, or malware) and combine them in new ways. It is this emergence of markets that is the final component of the modern abuse ecosystem and has served both to rapidly distribute new

business models and to reduce costs through economies of scale. However, migration to this market introduces visible, cost-sensitive dependencies that, if disrupted, undermine entire criminal profit-generating schemes that as a composite otherwise appear intractable to defeat.

While individual elements of the abuse ecosystem have been covered to various degrees in the academic literature, none captures the rich fabric of this underground economy in its full breadth nor provides the context required to understand the structure and inter-dependencies between individual elements. It is this broader perspective that motivates our survey paper. Indeed, it is our contention that a systematized understanding of underground relationships is critical to developing effective, long-lasting countermeasures. We champion that research and industry must make a drastic departure from solely focusing on protecting users and systems (tantamount to a fire fight) and strategically pursue disruptions of the brittle dependencies that underpin the entire for-profit abuse ecosystem—including actors, resources, and capital flow. To this end, our paper makes four contributions:

- *Underground Structure.* We define a framework for structuring underground assets based on the role they play in the monetization process: profit-creating activities (scams), cost centers (infrastructure services and markets), and value realization (internal and external cash-out services).

- *Classification.* For most of the best-known scams, services, and capabilities, we explain how they have been specialized, how they fit into our structure, the kinds of business models that they naturally express, and the dependencies they produce.

- *Interventions.* We examine various techniques—both proposed and explored—for intervening in different parts of abuse markets with an eye for evaluating how different actions impact miscreant profitability.

- *Standing Challenges.* We stratify the breadth of methodologies thus far for studying cybercrime and identify key challenges that will shape the field moving forward.

Finally, pursuing research into the abuse ecosystem requires a great deal of domain knowledge and context for which there are few good sources. We have spent a decade working in this field and we hope in documenting our experience that this paper can serve as an effective stepping stone for new researchers to build upon.

2 Organization Within the Underground

The current cybercrime landscape stems from a rich history of which black market commoditization is only a recent innovation. We explore the fitness function driving this evolution:

profit. In the process, we capture the stratified roles and their inter dependencies into a taxonomy of underground organization. These roles place an increased importance on open communication and self-policing between criminal communities, the consequences of which open criminal activities to the research community at-large.

2.1 What is the Black Market?

Computer-based crime and abuse has a long history, with well-documented cases of computer fraud dating back to the 1970s.[1] Personal computers provided a common substrate for would-be actors, giving birth to the first widespread viruses in early 1980s, and the Internet provided a broad transmission vector allowing the first network worms to emerge in the late 1980s. However, it is only in the 21st century that this activity morphed from the independent actions of a small number of motivated individuals, to a burgeoning set of cooperative enterprises, shared business models, stratified service offerings, and ever increasing degrees of specialization.

The core of this transformation is the emergence of a "black market" economy, built around *for profit* cybercrime, in which a large number of geographically distributed actors trade in data, knowledge and services [5, 6].[2] Absent such a structure, early miscreants needed to operate every facet of their business.[3] By contrast, the same scams today may involve a dozen different parties each responsible for some particular piece of the operation. This is possible because a shared marketplace allows for economies of scale, and encourages specialization and competition (and hence efficiency). We find evidence of this specialization within underground forums that sell *a la carte* access to virtually every part of the criminal "value chain" including compromised hosts, fraudulent accounts, stolen credit cards, and even human laborers. Thus, it is possible for a criminal entrepreneur to outsource these parts of their business and combine them in innovative ways to support new value creation strategies (typically scams based on defrauding consumers, businesses or both). However, whether this commoditization has yet achieved wide-spread adoption within the criminal community remains an open research question.

Commoditization directly influences the kinds of business structures and labor agreements that drive recent cybercrime. For example, the *affiliate marketing* business model, which is endemic to the spam, concisely encapsulates the drive to specialization. In this model, an entrepreneurial group es-

---

[1]For example, in the early 1970s the Union Dime Savings Bank lost over $1M due to computer-based fraud directed by their head teller. [161]

[2]Our interpretation of cybercrime is centered around crimes unique to electronic networks, thus we omit anonymous markets which primarily exist to support traditional crime using online distribution channels. [21]

[3]Thus a spammer needed to own the means of sending e-mail, acquire mailing lists, create storefront web sites, contract with Web hosting, register domains, purchase and warehouse products, accept payments, provide customer service and so on.
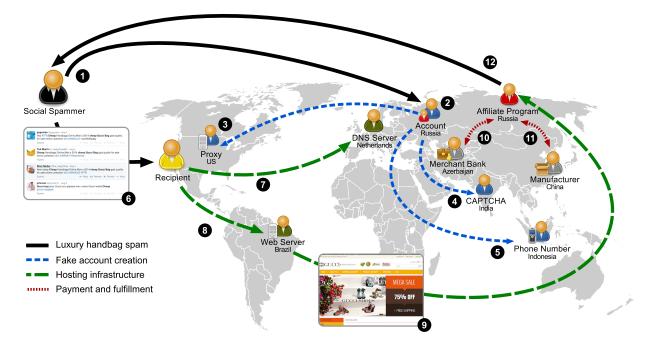
Figure 1: Specialized roles in the underground economy that underpin extracting wealth from victims. This represents just a single abuse monetization value chain to serve as a motivating example.

tablishes a core line of business in some niche (e.g., selling counterfeit pharmaceuticals). This group, called the affiliate program, typically provides all aspects of the business *except* acquiring new revenue. Revenue generation is outsourced to "affiliates"—independent contractors paid on a commission basis for each sale they bring in. This division of labor respects the value of specialization (e.g., affiliates can just focus on how to advertise products, and need not understand anything about fulfillment, payment processing, domain registration, hosting, etc.) and also provides natural risk transfer between both parties. To wit, since an affiliate has no sunk costs in concrete goods, they are free to move between affiliate programs (e.g., if one fails), or advertise for multiple programs, across a variety of independent niches.[4] Similarly, the affiliate program is insulated from the risk of engaging poorly performing affiliates because they only pay commissions on new revenue (e.g., a successful Viagra sale). Thus, an affiliate program will typically engage many hundreds or even thousands of affiliates—most of whom may be completely ineffective—yet will only pay the small subset who develop effective advertising techniques. [69, 96, 133]

Such structures, which reflect the dynamic nature of the

[4]Indeed, there is substantial evidence that this behavior is commonplace. John et al. document large-scale botnets advertising for a range of different goods and storefronts [63]. Similarly, an analysis of the data from McCoy et al. shows a range of affiliates operating in multiple pharmaceutical programs, and Stone-Gross et al. echo this finding in their analysis of several fake anti-virus programs [96, 133]. The same behavior is seen in the SEO vector, where *Wang et al.*'s analysis of the GR botnet demonstrates the spammer advertising for multiple niches simultaneously and explicitly switching product categories in response to market conditions. [154]

ecosystem, are ubiquitous in the underground economy. Indeed, this combination of an on-demand labor force, minimal sunk or indirect costs, and no regulatory limitations, creates a "pure" form of capitalism that naturally encourages rapid innovation. New business ideas are constantly proposed, deployed and tested (e.g., we are aware of multiple programs trying to monetize plagiarized term papers as a service today). Most of these fail, but those new ideas that generate significant revenue (e.g., fake anti-virus, ransomware) attract competition and become "commodity crimeware".

2.2 Bird's-Eye View of a Value Chain

We present an example of a complex value chain capturing the flow of capital between actors in the black market in Figure 1. In our example, a spammer seeks to monetize user interest in trademarked products on Twitter by selling knock-off replica handbags (❶). In order to engage with Twitter users, the spammer first requires a multitude of fake accounts to post messages. This is satisfied by a subset of the underground that coordinates all of the components required to bulk register accounts in return for a fee (❷). This includes paying parties with access to dynamic proxy infrastructures to evade IP blacklisting (❸); human workers solving CAPTCHAs (❹); and SMS verification challenge farms reliant on foreign SIMs (❺). With the accounts in hand, the spammer posts links to Twitter, which ultimately land in a legitimate user's timeline (❻). When the victim clicks on the URL, an entirely independent set of components is required to provide domain resolution (❼) and Internet hosting (❽). In turn, the victim

Product

Spamvertised goods | Scareware | Clickfraud | Financial fraud | Banking theft

Theft

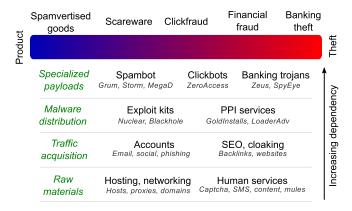|			
Specialized payloads	Spambot *Grum, Storm, MegaD*	Clickbots *ZeroAccess*	Banking trojans *Zeus, SpyEye*
Malware distribution	Exploit kits *Nuclear, Blackhole*		PPI services *GoldInstalls, LoaderAdv*
Traffic acquisition	Accounts *Email, social, phishing*		SEO, cloaking *Backlinks, websites*
Raw materials	Hosting, networking *Hosts, proxies, domains*		Human services *Captcha, SMS, content, mules*

Increasing dependency

Figure 2: Taxonomy of underground actors. Profit centers supply the revenue for all abuse, while support centers provide critical resources that streamline defrauding victims.

is handed off to an affiliate program (**9**) that handles order placement, including processing the victim's credit card for payment (**10**) and coordinating delivery from illegal manufacturers (**11**). Ultimately, the spammer receives a cut of this sale from the affiliate program (**12**), while the victim receives the intended replica handbag.

2.3 Defining a Black Market Taxonomy

Our previous example represents just one configuration for how criminals leverage underground resources to monetize fraud and abuse. Regardless of the composition, we argue there is always a *profit center* through which victims transfer new capital into the underground. In our example, revenue originates solely from victims buying trademark-infringing products. As we will discuss, criminals could adapt this strategy to sell illegal pharmaceuticals, fake anti-virus, or a multitude of other payout mechanisms. From there, any number of specialized *support centers* can facilitate abuse for a fee. These support centers have no credible non-criminal applications (e.g., CAPTCHA solvers, managers of compromised hosts, provider of exploits), and thus cannot exist without some eventual external payout due to operational costs.

Figure 2 captures these relationships into a taxonomy of abuse. Profit centers, shown on top of the figure, form a spectrum between selling products to unwitting victims to outright theft from victims. Within this spectrum, we highlight spamvertised products, scareware, click fraud, financial fraud, and liquidating funds from stolen bank accounts. A medley of alternatives such as dating scams, call-center scams, premium SMS fraud, DDoS extortion, or even stealing and re-selling gaming assets all fall within this spectrum and ultimately derive a payout from victims outside the underground. Monetization schemes that fall closer to theft are more likely to see recourse on behalf of victims (e.g., disputed credit charges, denied ad payments, reported extortion attempts). We provide an in-depth analysis of the best understood payout mechanisms in Section 3.

These profit centers are propped up by an ecosystem of support infrastructure (shown in the bottom of the figure), which are comprised of increasingly critical infrastructure that can be configured arbitrarily (and optionally) by criminals per their requirements—at a cost. We decompose the graph of relationships between criminals to stratify these cost centers based on *dependency*.[5] In our taxonomy, we argue that compromised hosts and basic human services form the foundation of all abuse irrespective of profit center. Criminals use these raw materials to bootstrap increasingly sophisticated activities within the underground that garner traffic (e.g., spammed URLs, compromised sites, phone calls). This traffic can either directly fuel profit centers or feed into more involved abuse such as malware distribution. The pinnacle of this hierarchy captures specialized botnets and trojans that depend entirely on the underground ecosystem for distribution while simultaneously tapping directly into profit centers that inject new revenue into the market.

Our list of profit and support centers is by no means exhaustive. The future of black market research lays in developing a better understanding of the tenuous connections between the actors in this space. Nevertheless, as we will show, this taxonomy proves fruitful for analyzing historical abuse and reasoning about threats on the horizon. Our primary argument is that criminals reliant on commoditization expose themselves to a range of new intervention strategies. As such, its critical we provide the security community a framework for understanding how criminals realize a profit from value chains and the fragile relationships involved that are ripe for disruption.

2.4 Underground Communities

Specialization within the underground hinges on open communication between criminals who advertise goods and services as well as potential buyers. Popular mediums include forums, Internet chats, web storefronts, and freelance labor; each supporting Chinese, Russian, German, and English black markets [108, 117, 170]. These communication channels belie the term "underground". With the exception of invite-only markets, the underground transparently conducts business transactions with (a perceived) impunity towards law enforcement. This same transparency exposes the broad range of criminal activities to researchers at large.[6]

Contact between researchers and underground communi-

---

[5]We loosely define dependency as the ratio between an asset's outdegree (e.g., external requirements) versus indegree (e.g., criminal use cases). Human services such as CAPTCHA farms play a critical role in account creation and certain spam varietals and can operate without any additional resource requirements from the underground beyond capital. In contrast, while exploit kits are fundamental to distributing malware families, the kits are obsolete without access to vulnerabilities and web traffic.

[6]We caution that risk-averse criminals who avoid forums—or vertically integrated crime that exists in isolation—biases the types of underground activities observable by researchers.

ties date back to at least the mid 2000s [45,146]. For example, in 2007 Franklin et al. passively monitored criminals selling stolen credit cards in public IRC channels [45]. Miscreants would advertise credit cards numbers, CVV codes, and account balances for dumps of credentials in their possession to anyone who bothered to connect to the IRC network. Similar investigations have yielded a wealth of information on services, pricing, and criminal business transactions posted to to public black market forums [40, 57, 168] and freelance job listings [81, 107, 156]. Each of these studies provide us with an unprecedented view into the division of labor between criminals that we explore later in Section 4.

While the existence of underground communities simplifies new actors starting a career in abuse, honor among thieves is sorely lacking. Accordingly, black markets cope with fraud by self-policing. For forums in particular, a range of roles exist to vet new buyers and sellers. This includes trusted reviewers who scrutinize newly listed products and services and verify their authenticity for the remainder of the general community who purchase illicit goods [117,169]. Furthermore, it is increasingly common for black market communities, particularly those in niches under external stress, to exert some due diligence before accepting a new members—conducting an online interview, requests for underground references, and any documentation of past illicit activities to demonstrate true affiliation with the underground world. What emerges is a vibrant social network between criminals that expunges known scammers and their multiple identities [108,167].

However, like any system founded on transitive trust, the ecosystem is vulnerable to infiltration and sybil collusion. Multiple cybercrime rings have fallen to law enforcement agents who posed as vetted conspirators, even taking control of entire forums to gather evidence for prosecution [117, 119, 120]. In response, underground markets have become increasingly insular. Forums including *darkode* and *blackseo* allegedly go as far as watermarking content in order to detect accounts involved in leaks [76]. Miscreants have also shed insecure VPN services in favor of anonymous communication over Tor in response to law enforcement actions against a popular credit card forum Shadowcrew [58]. The impact of this lockdown has yet to be felt by researchers, but will likely blind some windows into the criminal zeitgeist.

2.5 Alternative Profit Sources

The underground places no restrictions on who can participate or how actors leverage marketed resources. While we focus on the black market as a vehicle for illegally deriving wealth from victims, actors can adapt the same services for equally nefarious ends. We briefly highlight the role of illegal services in reputation gaming, anti-competitive practices, and politics.

Fame & Notoriety: If all that distinguishes an obscure video from a viral sensation is the perception of popularity, then a perverse incentive emerges for artists, musicians, and other public personalities to inflate their notoriety through fake views, followers, and subscriptions. More than mere fiction, in 2012 more than 2 billion YouTube views were removed from music videos controlled by Sony and Universal [56]. Similarly, Facebook was forced to purge nearly 1% of its "likes" as they originated from coercion and abuse [51]. The practice of buying and selling synthetic engagement impacts any service reliant on crowd-sourced reputation to guide consumer choice. Examples include Yelp restaurant ratings; Amazon product reviews; Android Play Store and iPhone App Store star ratings; Facebook likes; and YouTube views.

Anti-competitive Practices: Rather than using the black market as a tool for extracting additional wealth, anti-competitive businesses can rely on the underground as a tool for extinguishing the finite budgets or resources of their competitors. Examples range from relying on click fraud to deplete a company's advertising budget with fake traffic; launching denial of service attacks to take down a competitor's web presence; or directly disparaging the brand of another party with negative reviews. In these cases, the black market profits from companies seeking a competitive advantage.

Political Propaganda, Censorship, and Espionage: The underground's capability to compromise hosts, disrupt network access, and bulk generate spam can also be deployed by governments and political institutions. Along these veins, there is a growing market for government surveillance software that leverages exploits as a mechanism for taking hold of a victim's machine [93]. Other capabilities such as denial of service can serve to disrupt access to critical resources or objectionable news media—similar to the attacks allegedly launched against Georgia prior to Russian bombings in 2013 [94]. Even social media accounts can serve as a mechanism for governments to control political dialogues. In 2012, over 25,000 fake accounts were used to drown out political discourse surrounding purported fraud in Russia's parliamentary elections [141], with similar evidence of political tampering by unknown parties appearing in American Senate races and economic policy debates [91, 124].

3 Criminal Profit Centers

Profit centers reflect the amalgam of tens of interconnected specialized criminal communities working in cohort towards a final payout. The abuse surface area involved at first glance is overwhelming and rapidly expanding. To help guide our discussion, we provide a breakdown of some of the most lucrative criminal business models in Table 1. Prolific examples to date include spam-based advertising, scareware, click fraud, financial fraud, and credit card theft. We stress this

represents only a fraction of for-profit fraud and abuse payout mechanisms; those research best understands.

We highlight commonalities between monetization strategies, discuss scenarios where external pressure has forced cyber criminals to seek alternative profit sources, and explore open challenges ahead. We note that research into many of these profit centers is sorely lacking: we rely on industry and government estimates of profit in the absence of methodologically sound measurements. While we acknowledge these values may be overblown as explored by Anderson et al. [6], at the same time we argue they serve as qualitative assessments of the scale of abuse and illustrate that our current perspective of criminal monetization is largely incomplete despite intense efforts outside of research to disrupt criminals.

3.1 Spamvertised Products

Spam-based product advertising is among the oldest forms of economically driven Internet abuse. Criminals solicit consumer purchases of counterfeit physical goods (pharmaceuticals, luxury goods, apparel, electronics, cigarettes), pirated digital goods (software, videos, music, e-books), high-risk services (pornography, gambling, fraudulent dating) and quasi-fraudulent goods and services (certain nutraceuticals, male-enhancement products and work-from-home offers) [86]. Consumers willingly transfer their wealth into the underground, all be it unaware their purchases fuel the development of support infrastructure (e.g., malware, botnets, CAPTCHA solving, etc.) that enabled a far wider range of criminal activity.

The early spam era of the 1990s was a vertically integrated affair where spammers handled email, product selection, warehousing, and shipping—a "soup-to-nuts" operation [98]. The modern spam advertising ecosystem is both specialized and stratified, forced in part by technical countermeasures (e.g., IP blacklisting) and criminalization (e.g., the CAN-SPAM act in the United States). Central to this organization is the affiliate marketing business model, discussed previously in Section 2. Spam affiliate programs operate consumer-facing web storefronts and handle credit card payment processing, fulfillment, customer service, and frequently, domain service and hosting. Independent affiliates are left to focus entirely on driving user traffic to the affiliate storefronts through whatever means possible (e.g., email spam [63, 79, 165], social network spam [47, 50], and search engine optimization [84, 105, 154]), earning a 30–50% commission for each sale.

Spamvertised revenue taps into latent demand from consumers that is unsatisfied in the open market due to high costs imposed by brands or regulatory overhead. For example, Chachra et al. demonstrated that between 20-40% of revenue from email advertised pharmaceuticals could be attributed to customers who *explicitly* visited their Junk/Spam folder to click on links in email spam messages [17]. Moreover, re-

searchers found that over 30% of spamvertised revenue originated from repeat customers who established an ongoing relationship with the spammed storefronts [96].

To reach these consumers, spammers send billions of messages with abysmal click-through rates: 0.003%-–0.006% for email [67] and 0.13% for social networks [50]. Despite this, the spam revenue model has proven immensely successful. Studies have documented annualized gross revenues over $65M for *a single pharmaceutical organization* [96]. Similarly, a recent study of SEO-based spamming of luxury goods demonstrated a single supplier delivering over 28K items per month to customers (given a typical price of $200 each, that would suggest an annualized gross revenue of $68M *for one supplier*) [152]. Today, spam remains a critical source of underground profit that continues to thrive despite major botnet takedowns and ubiquitous spam filtering.

3.2 Scareware & Ransomware

Where spam hinges on consumer interest in off-market products, scareware relies on socially engineering victims under duress into buying ineffectual goods or services. The most prominent example of scareware is *fake anti-virus*. The scheme relies on software prompting victims with a warning that their machine is infected with a multitude of malware variants that can be cleaned up by purchasing a one-year subscription for $60–80 (paid via credit card) for an anti-virus service that in fact does nothing. Fake anti-virus emerged as the goto cash cow for malware installations in the underground from 2008–2011, accounting for 15% of all the malware detected by Google Safe Browsing [123].

As with spam, criminals organized into affiliate programs to streamline abuse and expand the pool of victims impacted. Operators like Gagarincash, Gizmo, and BestAV handled credit card processing and provided their tailored fake anti-virus binary to affiliates [74]. Independent affiliates would obtain installs through whatever means possible: social engineering, drive-by downloads, or pay-per-install (discussed later in Section 4.1). Stone-Gross et al. executed a detailed investigation of logs obtained from some of the largest fake anti-virus affiliate programs and found that 2.2% of victims prompted with warning dialogues subsequently purchased fake anti-virus software, grossing the criminals involved upwards of $130 million over the course of 2008–2010 [133]. These programs abruptly ceased in 2011 when law enforcement, banks, and security researchers froze financial transactions to fake anti-virus affiliate programs and dismantled the Conficker botnet distributing fake anti-virus software [73,74].

Following in fake anti-virus' stead, ransomware emerged as the predominant successor in 2012–2013 [82]. The most notable variant of its type was CryptoLocker. Once installed, the malware would encrypt all of a victim's files and extort victims to pay $100–400 via Bitcoins or pre-paid cash vouchers in return for the decryption key. Reports from the me-

Profit Center	Strategy	Estimated Revenue	Time Frame
Spamvertised products	Pharamcuticals [97]	$12–92 million	2007–2010
	Luxury knock-offs [152]	$68 million	2013–2014
Scareware & Ransomware	Fake anti-virus [133]	$130 million	2008–2010
	CryptoLocker [159]*	$3 million	2013–2014
Clickfraud	ZeroAccess [115]	$36 million	2013
	DNS Changer [149]*	$14 million	2007–2011
Financial Scams	Pump and dump [150]*	$120 million	2008–2013
	419 scammers [8]*	$200 million	2006
Credit Card Theft	ATM withdrawl scam [118]*	$45 million	1 day
	Zeus banking trojan [9]*	$70 million	2009–2010
	Re-selling stolen cards [35]*	$300 million	?–2013

Table 1: Estimated revenue from a multitude of profit strategies (irrespective of operating costs). These strategies span the spectrum of cybercrime: from selling illegal products to outright credit theft. We annotate all industry and government estimates of criminal revenue with an asterisk to emphasize an unknown collection methodology. We caution these values may be overestimates.

dia claim the scheme impacted over 500,000 users, 1.3% of whom paid out an estimated $3 million before law enforcement and security researchers intervened [159]. A slew of copycats have followed that include BitCrypt [26] and CryptoLocker 2.0 [72].

3.3 Click Fraud

Online advertising is a multi-billion dollar market that fuels much of the Internet. Criminals have tapped into this revenue stream: masquerading as publishers to profit from fraudulent ad traffic [135]. Click fraud schemes rely on a two-step process. Attackers will register with ad networks as a publisher, sourcing ads from the network's catalog. They then drive traffic—simulated or hijacked—in order to receive payment from advertisers. Criminals have honed their techniques over the years to include hired manual workers who click on URLs [92]; automated bots like Clickbot.A, 7cy, and Fiesta that simulate user clicks [32, 101]; and malware variants such as TDL and ZeroAccess that redirect legitimate user clicks and searchers to ads controlled by criminals [115,129]. A variant of this scheme relies on monetizing traffic via ad-based URL shorteners [112, 142].

Ad networks protect advertisers from synthetic click traffic through a variety of technical measures that include signature and anomaly-based detection [80, 100]. Under normal circumstances when no fraud is detected, publishers are paid upwards of 70% of the revenue generated from a user clicking on an advertisement [34]. The ad network takes the remaining cut. Conversely, if an ad network deems a click fraudulent then the publisher involved goes unpaid. Dave et al. conducted a measurement of several ad networks, setting up a series of bluff advertisements to catch click fraud [33, 52]. They found that 22% of all clicks on their bluff ads were synthetic, of which 10–25% went uncaught by each ad networks' fraud detection.

Small per-click revenue—reduced even further by fraud detection—yield a threat landscape where the most vested click fraud outfits operate botnets comprising millions of infected hosts to turn a profit. Major players include the ZeroAccess botnet which comprised an estimated 1.9 million hosts that grossed the criminals involved $100,000 a day in ad revenue before Microsoft's takedown effort [115]. Similarly, DNS Changer infected nearly 4 million machines, pulling in $14 million via click fraud between 2007 and 2011 [149]. While the two players are no longer active, click fraud remains as a highly lucrative abuse profit center.

3.4 Financial Fraud

Social engineering plays a significant role in criminal profit. We explore two cases where miscreants convince unwitting victims to willingly transfer their funds in return for a "promise" of future high yield returns. The financial fraud is either indirect (e.g., via the stock market as in the case of pump and dump schemes), or direct (e.g., through checks as in the case of 419 scams). Invariably, the rewards will never materialize, as the miscreants cease all subsequent communications and disappear with the victims' money.

3.4.1 Pump and Dump

Miscreants manipulate stock prices belonging to legitimate commercial entities via "pump and dump" spam campaigns [147]. Criminals will buy low-cost stocks and then entice victims into investing by sending messages that state a small company with a low stock price is on the cusp of becoming a hot commodity due to the development of a product or idea with substantive growth potential. Criminals target stocks from "over-the-counter" markets with limited requirements on publicly-available information, simplifying the deception involved [147]. Once the perceived value of a stock rises due to investor interest, criminals cash out their shares at

a profit and cause the stock price to deflate at the expense of the victims involved. Unlike spamvertised products and fake anti-virus, this scheme requires no direct financial transaction (e.g., credit card charge) between criminals and victims; markets act as the middlemen.

Limited research suggests that pump-and-dump spam messages may constitute as much as 3% of all spam email in a given year [11]. Spammers purchase stocks in advance of the pump and dump scheme and sell within a few days of the onset of campaigns when stocks achieve their peak inflated price [53]. Selling at this time ensures criminals the greatest possible rate of return on their investment. In fact, Frieder et al. suggest that spammers can generate a 4% rate of return on their initial investment, while victims lose at least 5% within a two day period [46]. According to an FBI press release, the largest known pump and dump crime ring impacted over 35 countries and netted the criminals involved over $120 million in fraudulent stock sales from 2008 until their arrest in 2013 [150].

3.4.2 Advanced Free Fraud

Advance fee fraud scams, more commonly referred to as "Nigerian scams" or "419 scams," fleece victims by offering a large future reward if victims send the criminal a smaller initial down payment. Examples include lottery awards, outstanding inheritances that are held up by "legal fees," romantic affairs that require flight tickets to rendezvous, or fake apartment rental deposits [61, 114]. Once victims pay out, the criminals cease communication and move on to the next victim. While originally rooted in physical mail and email, 419 scams have evolved to target users of classifieds, dating, real estate, and other websites [114].

Advanced fee fraud schemes consist of three stages: (1) an initial salvo of automated scam messages; (2) a manual relationship-building phase where criminals interact with victims to engender trust; and (3) an irreversible exchange of funds between the victim and the criminal (e.g., wire transfer, shipment of physical goods). Park et al. developed an automated honeypot that would interact with scam artists to study their operations [114]. They found that 70% of scammers provided physical addresses located in Nigeria for delivering goods. Furthermore, they found evidence of a largely manual work force that would respond to scams within 1–2 days during peak work hours in Nigeria. A follow on study by Jones et al. identified that scammers relied on support networks within the United States to mail fake checks and process money transfers via Western Union and Money Gram [66]. According to reports from the news media, such schemes have persisted for nearly a decade, netting criminals over $200 million in 2006 [8].

3.5 Credit Card and Online Banking Theft

Criminals leverage stolen credit and debit cards to defraud victims and corporations of their financial assets. Direct access to capital obviates any requirement for revenue generated via consumer products or scams like all of the profit centers discussed so far. However, the reversibility of fraudulent transactions forces criminals to launder stolen cash through unwitting victims, physical goods, and digital services. While in the United States consumers are insulated from credit fraud, the negative cost is still reflected in transaction fees and merchant charge backs.

3.5.1 Initial Theft of Credit Card Data

Credit card fraud hinges on a complex interplay between underground miscreants that begins with stealing credit data. Criminals acquire credit card numbers, CVV codes, and address details through (1) point-of-sale malware and skimming; (2) payment databases containing copies of card holder information and account numbers; or (3) phishing and malware that siphon credit card details directly from victims.

Point-of-Sale Malware & Skimmers: Point-of-sale (POS) malware and credit card skimmers steal card details directly from the physical cards used with card-present transactions. In a recent breach at Target, attackers infected POS systems with malware that scanned the active memory for credit card information [90]. When victims swiped their cards at POS terminals, the POS software temporarily stored the unencrypted information which the malware extracted and sent to the attackers. In this manner, criminals harvested over 70 million credit cards. This same attack impacted over 1,000 United States retailers in 2014 [128]. Alternatively, criminals will install physical devices on ATMs and gas station terminals that "skim" credit details that customers provide that are then remotely delivered to criminals [13].

Payment Storage Breaches: Businesses store credit card data for accounting, customer convenience, and recurring charges. The systems housing this data represent ripe targets for compromise. In 2003, criminals stole data tied to over 45 million credit and debit cards from TJX [71]. A similar breach of Sony's Play Station payment network in 2011 leaked names, addresses, and allegedly credit card details for over 77 million victims [87]. If card processors are Payment Card Industry (PCI) Data Security Standard (DSS) compliant, they cannot store the track data or the CVV code after a transaction has been authorized. This restricts attackers to credit card numbers, card holder names, expiration dates, and addresses.

Phishing & Malware: Organized phishing efforts persuade users to hand over their credit card details. Moore et al. analyzed thousands of live phishing sites hosted on compromised Web sites and estimated that on average 30 victims

are phished over the lifetime of banking phishing sites [103]. Client machines are equally vulnerable to harvesting card details. Stone-Gross et al. observed thousands of credit card details harvested from zombie machines tied to the Torpig botnet over a 10 day window [134]. This process is streamlined by software-as-a-service banking trojans such as Zeus and SpyEye that sold to the underground at large for approximately $1,000–10,000 [148, 151]. The two trojans collectively infected over 5 million victims before Microsoft and law enforcement intervened [10,151]. Tajalizadehkhoob et al. examined roughly 15,000 banking URLs that criminals configured as targets for Zeus [140]. They found 32% of domains belonged to financial institutions in the United States, Germany, United Kingdom, Spain, and Australia, many of which were attacked for more than six months at a time.

3.5.2 Cashing Out Stolen Cards

We find preliminary evidence that the parties responsible for stealing credit cards (and developing the malware involved) are independent from underground miscreants who specialize in laundering stolen assets [45]. These cash-out vectors currently remain opaque to the research community. Anecdotally, criminals funnel non-reversible transactions through businesses and (un)witting accomplices. For example, cybercriminals provided seven New York-based criminals duplicate credit cards and associated PIN numbers that were used to withdraw over $45 million in a matter of hours [118]. An unknown cut of these funds was slated to be paid back to the cybercriminals for the groundwork of obtaining credit cards.

A consistent theme of cashouts is the reliance of criminals on "money mules" who either act as drop points in the United States for expensive physical goods that are re-shipped abroad, or alternatively wire transfer stolen funds to criminals [43, 44]. Reportedly, many of these middlemen are in fact victims netted by work from home programs promising salaries that are too good to be true [12]. Apart from physical goods and cash, criminals can also use stolen credit cards to purchase in-app goods (e.g., virtual game resources) that are re-sold—in effect creating a spamvertised revenue stream backed by stolen credit cards. Similarly, criminals can purchase cloud computing resources to mine crypto currencies. While fees on the stolen cards in either case are reversed, the digital asset persists. Due to the lack of visibility into cash-out vectors, we rely on estimates from news reports which claim this confluence of techniques grosses criminals $70–300 million per major attack [9, 35].

3.6 Standing Challenges

Methodologically speaking, research thus far into profit center revenue overwhelmingly relies on data dumps of financial records that "serendipitously" appear due to criminal leaks or seizures by law enforcement.[7] The exception to this rule are revenue estimation techniques based on leaky side channels (e.g., incremental payment order ids [68]) or botnet infiltration. This challenge is reflected in profit centers we previously listed in Table 1. Of eleven revenue estimates, we aggregate seven from law enforcement, two from research of database dumps, and two from estimation techniques. Given the unknown methodology that law enforcement uses to arrive at these revenue estimates, we argue that our current perspective of criminal monetization is largely incomplete—in part due to sensitive financial data only accessible to advertisers and banking institutions.

An open question remains as how to formalize revenue estimate methodologies across profit centers to obviate the need for chance data dumps. A sister challenge to this is coverage. Our perspective of the underground thus far is a collection of observations largely drawn from Russian and Eastern European cybercrime. Language barriers create natural competing markets outside these regions, leading to a question of representativeness of global criminal activities. This challenge is exacerbated as criminal communities become more insular to the detriment of infiltration and monitoring which provide key data insights. Finally, there is a growing data divide between research and opaque abuse epicenters like clickfraud and banking theft. We must bridge each of these gaps in order to produce profit-driven countermeasures to cybercrime.

4 Ecosystem Supporting Abuse

Criminals weave a vibrant tapestry of abuse by combining threads of compromised hosts, human labor, networking and hosting, and accounts and engagement—support infrastructure that streamlines an eventual payout from victims. We focus in particular on commoditization. "Everything" is available for a price, as captured in Table 2, though whether criminals actively buy into these services remains an open question for research. Within this space, we explore the socio-economic factors that influence abuse prevention as well as black market pricing as a reflection of the industry's performance. We argue these concepts enable the community to reason about fraud and abuse as a *financial* battle. Profit centers pinned up by cost-ineffective resources will eventually crumble.

4.1 Compromised Machinery

Compromised machinery is the lifeblood of fraud and abuse. Subverted systems encapsulate a broad range of value to the underground: from raw computing and network power to unfettered access to a victim's data. Criminals who compromise hosts require an attack vector and a delivery mechanism. Social engineering, drive-by downloads, and malicious attach-

---

[7]Indeed, the long standing feud between ChronoPay and GlavMed [78] yielded a wealth of information about underground profit.

9

Support Center	Resource	Estimated Cost	Volume or Period
Compromised Hosts	Blackhole exploit kit [27]	$1,500	1 year
	Nuclear exploit kit [27]	$1,500	1 year
	Neutrino exploit kit [27]	$450	1 month
	Phoenix exploit kit [27]	$1,000–1,500	1 month
	Pay-per-install: US/UK [15]	$100–180	1,000
	Pay-per-install: Europe [15]	$20–160	1,000
	Pay-per-install: Other [15]	<$10	1,000
Human Services	CAPTCHAs [106]	$1–2	1,000
	SMS challenge [143]	$200	1,000
	Mobile SIMs [143]	$140–420	1,000
	English blog content [107]	$2–4	1
	Chinese blog content [156]	$0.25	1
Networking & Hosting	Proxy: 150 IPs	$25	1 month
	Proxy: 15,000–30,000 IPs	$250	1 month
	DDoS: 800 Mbps [70]	$10	1 month
	DDoS: 100 Gbps [30]	$200	1 day
Accounts & Engagement	Hotmail account [145]	$4–30	1,000
	Yahoo account [145]	$6–15	1,000
	Twitter account [145]	$1–20	1,000
	Facebook PVA [145]	$80–400	1,000
	Google PVA [145]	$80–500	1,000
	Twitter followers [136]	$4–20	1,000
	Twitter retweets [136]	$79–550	1,000
	Facebook likes [36]	$15–70	1,000

Table 2: Estimated cost of goods and services rooted in the black market as aggregated from underground studies and our own investigations.

ments are all examples of attack vectors; spam, ad networks, and compromised websites are examples of delivery mechanisms. For example, a spam email can contain a malicious PDF in the hope that the recipient will be careless or curious as to the contents. Once opened, the victim's machine is compromised and the attacker gains some degree of control.

Within this space, we focus on a dramatic shift in the underground that now decouples host *monetization* from host *compromise*. Specialized criminals abstract away the complexity required to uncover exploits and compromise victims at-scale. These criminals then sell access to compromised hosts to other miscreants in the underground. Here, we focus on two markets: *exploit-as-a-service* and *pay-per-install*. Each approach frees buyers to focus their energy on maximizing the profit derived from subverted devices.

4.1.1 Exploit-as-a-Service

The exploit-as-a-service ecosystem relies on drive-by downloads, an attack that leverages a vulnerability in a web browser to deliver malware. For a drive-by download to be successful, it needs three things: 1) victims to visit the website that serves an exploit, 2) a browser exploit, and 3) a malware payload to deliver [37]. Exploit as a service decouples these requirements and enables miscreants to specialize

and trade in each. The result: traffic is bought, exploits are "rented," and malware is broadly installed.

Grier et al. studied the exploit-as-a-service ecosystem and found a small number of *exploit kits* (pre-packaged software containing a suite of browser exploits) fueling the majority of drive-by downloads [49]. At the time, the Blackhole exploit kit was the leader in the market and even advertised the costs. Renting a server with Blackhole setup would cost $50/day, while you could license it at $700 for three months. Grier et al. found that exploit kits in the wild were responsible for delivering some of the most prominent malware families at the time including ZeroAccess, SpyEye, and TDL—all popular click fraud and banking trojans previously discussed in Section 3 that netted the criminals involved millions.

4.1.2 Pay-Per-Install

Pay-per-install is the natural extension of exploit-as-a-service that provides yet another distribution channel for malware payloads. In the pay-per-install marketplace, criminals sell access to infected machines by the thousands. This is enabled by a class of malware commonly called *droppers*. On their own, droppers are binaries that criminals configure to fetch arbitrary remote software for installation, in effect acting as a delivery mechanism for future malware. Many *malware*

platforms such as TDL and Zeus also include dropper functionality, often blurring this distinction. These systems allow malware authors to acquire an install base without going through the effort of developing exploits or worrying about distribution.

Caballero et al. found that the pay-per-install marketplace was responsible for distributing at least 60% of the most prevalent malware families in 2011 [15]. Since most of the popular malware families do not have a spreading mechanism built in, pay-per-install and exploit-as-a-service provide an alternative to infecting victims via spam or social engineering. Caballero et al. also showed how the pay-per-install economy allows country-level granularity for customers purchasing hosts. Geographic specialization lets malware authors focus their efforts to particular regions: for instance, only installing banking trojans in the US or Europe where credit card data is more lucrative. Criminal demand for specific regions is reflected in the cost of hosts: $100–180 per thousand in the United States and United Kingdom, $20–160 in Europe, and less than $10 in the rest of the world.

4.2 Human Services

Criminals supplement their computing and network cycles with an array of human services. In the following, we explore a few examples of these services, including CAPTCHA solving [48, 106], SMS verification, and content generation. We note that multiple others exist: human click farms [92], manual workers tasked with phishing [14], and even re-shippers who handle illegally-sourced packages [75]. These services highlight a critical socio-economic component of abuse: wage disparities between countries can create a secondary market for tasks and materials that would otherwise be too difficult for criminals to automate or procure.

4.2.1 CAPTCHA Solving

Websites rely on CAPTCHAs as a first line of defense against automated bulk account registration, posting, and friend requests. Two underground services exist for breaking these visual and audio challenges: (1) automated software solvers and (2) human laborers that manually solve CAPTCHAs at-scale. An example of software solvers is *spamvilla.com*, which advertised Hotmail and Yahoo CAPTCHA breakers with 25–30% and 48–50% accuracy respectively. As Motoyama et al. argues, however, the adversarial nature of CAPTCHAs makes maintaining the software more expensive than relying on human labor [106]. Indeed, services like *anti-gate.com* advertise 1,000 human-solved CAPTCHAs for as little as $1 with an accuracy rate over 98%.

Motoyama et al. identified at least 8 human CAPTCHA farms with prices ranging from $1–20 per thousand CAPTCHAs. In this model, miscreants provide CAPTCHA im-

ages to solution services via an API. These services act as middlemen who farm the CAPTCHAs out to manual laborers in China, India, Indonesia, Pakistan, Ukraine, and Vietnam. Motoyama et al. estimate the most active laborers in this market make as little as $47.32 a month. Consequently, the willingness of low-wage workers to solve CAPTCHAs offers miscreants in the underground a mechanism to side-step the technical challenge of CAPTCHAs. What remains is only an economic disincentive. Nevertheless, we argue that CAPTCHAs still serve two important roles: (1) They rate limit the velocity of abuse to the capacity of human CAPTCHA farms; and (2) they prevent criminals from engaging in low-return automated behaviors that are economically infeasible due the price of CAPTCHAs.

4.2.2 Phone Verification

Phone verification is the latest deterrent against bulk automated tasks. In this model, web services force miscreants to provide phone numbers and subsequently verify ownership via the receipt of SMS challenges. Unlike CAPTCHAs, phone verifications tie abuse to the cost of a physical SIM card rather than a digital resource. Miscreants rely on two approaches to pass SMS challenges at-scale: (1) circumventing SIM costs with free VOIP services; or (2) relying on inexpensive SIM cards from socio-economically disadvantaged countries.

Thomas et al. found that miscreants used VOIP numbers to solve 24% of Google SMS verification challenges [143]. The remaining 76% were satisfied with mobile phone numbers sourced from India, Indonesia, Nigeria, Bangladesh, Pakistan, and Vietnam, with many of the same regions providing the labor for CAPTCHA farms. The costs of SIMs from these regions are likely less than Chinese, Russian, or Ukranian SIMs that we currently see being sold on underground forums for $140–420 per thousand cards. A second layer of specialization exists within this space, where we observe middle men such as *sms-area.org* and *sms.xudan123.com* who offer SMS verification as a service for as little as $0.20 per challenge to a mobile number. Whether manual laborers operate physical phones to respond to challenges is currently unknown, though Thomas et al. found advertisements for modified hardware to simplify the task of workers swapping SIM cards. For the time being, SMS challenges remain an effective defense despite a burgeoning market for cheaply sourced phone numbers.

4.2.3 Content Generation

Criminals source realistic and grammatically correct spam templates, blogs, forum posts, and microblog posts from manual freelance laborers. While these markets are not explicitly malicious, they nevertheless satisfy a requirement among criminals for non-automatable tasks. Motoyama et al. examined one year worth of job requests on *freelancer.com*

and found numerous miscreants requesting 250–500 words of English content that included specific keywords [107]. Workers would write batches of 10–50 articles at a time, charging $2–4 an article. Wang et al. observed a similar market for over 170,000 Chinese laborers on Zhubajie and Sandaha who would write blog content for roughly $0.25 a task [156]. These "dirty" jobs allow criminals to outsource even the most menial tasks in exchange for paltry fees.

4.3 Networking and Web Hosting

The underground leverages compromised Internet networks and hosts to reach victims from across the globe. Criminals prop up stolen infrastructure to host illegal storefronts, serve exploits, and drive traffic to profit centers, all the while contending with persistent blacklisting and takedowns.

4.3.1 Domains & Hosting

Internet hosting is a necessity for connecting victims to criminals. We focus on two particular components: domain registration and IP hosting. Rather than robust, these raw materials must be cheap enough for criminals to renew in response to continuous blacklisting. Hao et al. examined over 130,000 spam domains in popular blacklists and found 70% were concentrated on 10 registrars, the most popular of which were eNom and Moniker [54]. Thomas et al. observed spammers would avoid registration costs altogether and abuse free subdomains provided by *co.cc* and *dot.tk* [142]. Automation prevention on each service consisted only of a CAPTCHA—far less expensive than domain registration. Regardless the source of domains, criminals must quickly cycle through infrastructure to stay abreast of the security industry: Anderson et al. observed that 40% of scam domains persisted for less than 120 hours [4]. Levchenko et al. tracked this phenomenon as reflected in spam campaigns for over 3 months [86]. They found a quickly narrowing funnel of hosting: criminals distributed over 346 million URLs hosted by 54,000 domains which in turn served only 968 HTML templates, the breadth of which were provided by 30 spam affiliate programs.

Where criminals rely on porous channels to register abusive domains, hosting subsists almost entirely on compromised hosts. Zhang et al. examined over 160 million blacklisted IPs flagged for spam and phishing and observed a strong correlation between poorly managed networks and the likelihood of blacklisting (e.g., compromise). The consequences are reflected in multiple studies of criminal hosting. Miscreants serve 75% of phishing pages from compromised hosts and 17% from free hosting [104]. Similarly, criminals hijacked over 180,000 websites from 2007–2008 to serve drive-by downloads [122]. High bandwidth and proximity to victims becomes a necessity for successful hosting: Anderson et al. found 57% of spam domains were hosted in the United States, followed by a long tail of other countries including

China, Canada, and much of Europe [4]. These are more expensive hosts from the perspective of the pay-per-install market. Criminals bolster the reliability of compromised machines through fast flux networks. Holz et al. found 30% of spam domains in 2008 would rapidly update DNS records every 5 minutes with a fresh set of compromised IPs [59]. Paired with domain generation algorithms that evade sinkholing with randomized, rapidly cycled naming [166], criminals have thoroughly adapted to surviving in a hostile environment.

4.3.2 Search Engine Optimization & Cloaking

Beyond the raw networking resources of compromised hosts, hijacked websites have an intrinsic value reflected in their content and search ranking. Attackers harness reputation data through *search engine optimization* (SEO) techniques that drive user traffic to profit centers. Mechanistically, attackers engaging in SEO manipulate web search results by falsely promoting their (compromised) sites for specific targeted keyword queries, with the goal of acquiring targeted user traffic. In other words, contrary to the traffic from email or social networks where users may have no interest in the profit centers designed to monetize them, user traffic obtained through SEO is characterized by their implicit interest in the targeted keywords from their explicit query. Various works document the effectiveness of SEO in directing users to a litany of the profit centers described in Section 3, specifically fake anti-virus, pirated OEM software, and counterfeit luxury goods [64, 84, 85, 89, 105, 152–154].

Criminals in turn rely on network *cloaking* to maximize the value of compromised hosts while simultaneously evading detection. At its essence, cloaking allows practitioners to deliver different content to different types of users. Search engines are presented enticing content for indexing; users are redirected to profit centers; and site owners and security crawlers are presented benign content [153]. These tactics rely on fingerprinting the respective parties through HTTP headers, IP addresses, and cookies [18, 113, 157, 158, 163, 164]. While we present cloaking as a vehicle for SEO, the strategy extends to drive-by download exploits that fingerprint a victim's machine to detect specific vulnerabilities before launching attacks that might otherwise trigger security crawlers [37]. As such, cloaking has become a commodity in its own right.

4.3.3 Denial of Service

Miscreants wield the sheer bandwidth power of compromised hosts as a bludgeon in distributed denial of service (DDoS) attacks that degrade or disable a victim's network presence. The motives for this are multifarious and range from extorting technology companies for as little as $200 to cease attacks [111] or "smokescreen" tactics used to distract infrastructure security teams to conceal simultaneous large-scale banking theft [1]. The threat of DDoS is exacerbated by the

commoditization of DDoS software such as DirtJumper [7]. Buscher et al. examined botnet deployments of this software in the wild and identified tens of thousands of machines located in India, Thailand, Indonesia, and an array of non-US and non-European hosts—the cheapest hosts available from the pay-per-install market previously discussed in Section 4.1. Popular targets included online shopping, pornography, gambling, and hacking communities. Karami et al. observed that criminals would in turn rent out similar "booter" services for as little as $10 per month for a 800 Mbps attack [70]. Backend logs leaked for one of such service reveal criminals attracted over 300 customers purchasing 48,000 attacks.

4.3.4 Proxies

Attackers use network proxies to both mask their true identity when conducting criminal transactions as well as to evade blacklists that prevent bulk automation. A variety of proxy services exist and are differentiated by the level of desired anonymity, reliability, geolocation, and number of simultaneous connections. These features are ultimately reflected in the price of proxy services. We are unaware of any existing research into black market proxies. Anecdotally, we find *5socks.net* selling monthly access to 150 proxies for $25. Similarly, *spamvilla.com* sells 15,000–30,000 IPs for $250/mo (mailing prohibited). Non-criminal proxies also exist, including the Tor network [39] and *hidemyass.com* VPN service.

4.4 Accounts & Endorsements

Underground merchants reacted to the migration of surfing crowds into closed-garden web services like YouTube, Facebook, and Amazon by selling access to reputable accounts and endorsements. Popular products include bulk, automatically generated *fraudulent* accounts; *compromised* accounts hijacked from victims; and a range of synthetic followers, likes, subscribers, and reviews.

4.4.1 Fraudulent Accounts

Fraudulent accounts are a keystone for monetizing spam and fake engagement outside of email [60, 137, 142, 155]. Facebook publicly estimates that abusive accounts comprise 1.5% of its user base [121], while Twitter estimates its own problem at 5% of users [127]. As originally discussed by Thomas et al., fraudulent accounts are readily available from the underground [145]. Account merchants in this space abstract away the complexity required to evade IP blacklisting, solve CAPTCHAs, and satisfy verification challenges over email or SMS. In turn, merchants sell accounts for a fee ranging from $5–500 per thousand accounts. These prices differ based on the complexity of the security protections in place in addition to demand. Hotmail accounts cost $4–30 per thousand, Yahoo accounts $6–15, Twitter accounts $1–20, Facebook accounts $80–400, and Google accounts $100–500.

Technical protections against automated account creation are in truth financial barriers. CAPTCHAs, previously discussed in Section 4.2, add a minimal fee to account creation. Similarly, web services that require newly minted accounts to pass an email verification challenge force criminals to purchase an address for as little as $0.04.[8] Among registration challenges, only phone verification represents a significant financial hurdle for criminals to acquire SIM cards in bulk, discussed previously in Section 4.2. Thomas et al. observed that phone verification increased the cost of Google accounts from roughly $30 per thousand to $500 per thousand, though criminals have now streamlined the process of acquiring foreign SIM cards. What emerges is a protected battle between web services and account merchants where websites attempt to increase the cost of accounts and narrow a fraudulent account's window of activity such that fraud and abuse become financially unsound.

4.4.2 Compromised Accounts

Criminals rely on hijacking account credentials as an alternative strategy for gaining a foothold in registration-based websites. In particular, social network accounts are valuable as they come with established trust relationships. Miscreants obtain these credentials through password guessing, re-used passwords leaked by database dumps, phishing, and malware. However, detailed evidence of which approach poses the greatest risk remains elusive.

Measurements of large-scale compromise date back to 2010 when researchers observed over 1 million hijacked accounts spamming Facebook and Twitter [47, 50]. The problem has since grown, with evidence showing that compromise now outpaces fake accounts as the primary source of spam. Work by Cao et al. in collaboration with Facebook in 2014 found over 2 million accounts spamming, 71% of which were legitimate users who fell victim to malware and social engineering [16]. Similarly, Thomas et al. found over 14 million compromised accounts on Twitter compared to nearly 5 million fraudulent accounts during the same period [144]. The majority of these users fell victim to phishing and malware propagated within Twitter as opposed to password guessing or database dumps. So long as victims fall for social engineering and web services fail to detect anomalous logins, compromise will remain a significant threat.

4.4.3 Fake Endorsements & Engagement

Fraudulent and compromised accounts feed into a vast network of services that provide fake endorsements and engagements ranging from YouTube subscribers, Twitter followers, Facebook likes, and Amazon and Yelp reviews. These services inflate the credibility of miscreants and result in heightened visibility, much like search engine optimization dis-

---

[8]The additional cost over CAPTCHAs covers access to IP addresses, account templates, and rate limiting to avoid detection.

cussed in Section 4.3. Some of the most popular programs target social networks. Stringhini et al. identified a vibrant market for Twitter fake followers where services advertise ten thousand followers in exchange for $40–214 [136]. Retweets are also available for $79–550 per thousand. De Cristofaro et al. observed a similar market for Facebook likes priced at $150–700 per ten thousand [36]. Stringhini et al. concluded that miscreants source these engagements from over 740,000 compromised accounts. In particular, "premium" customers pay a fee to get followed, or to spread content on Twitter. "Free" users, instead, provide their credentials to the market operators in exchange for a certain number of followers. Once these accounts are under the control of the market operators, they are used to provide followers to the premium users or to spread the "premium" users' content. This scheme continues to persist, with over 3 million victims on Twitter generating fake follows and retweets through surreptitiously installed third-party applications [144]. Despite ready evidence of markets for fake engagement, the long-term payoff or effectiveness of miscreants purchasing from this space remains unknown.

4.5   Standing Challenges

The research communities understanding of support infrastructure is years ahead of profit centers, in part because it has been the dominant focus of interventions over the last decade. However, there are a number of pitfalls that remain that stymie reasoning about the costs of goods and services. The foremost is whether pricing accurately reflects the value of a resource. There are tens to hundreds of black market forums and web storefronts, many of which are scams or wildly overpriced. CAPTCHAs are a prime example: the dominant merchants in the space charge $1–2 per thousand solutions, but others charge upwards of $20. Herley et al. expressed similar skepticism over the purported costs for stolen credit cards with advertised balances far exceeding the cost of the card [55]. These inconsistencies may arrive due to arbitrage, insular markets, influid exchanges, or risky payment mechanisms (discussed later in Section 5). A significant challenge remains for how to distinguish between valid and inflated prices.

A second major challenge is understanding the return on investment generated per resource, which is reflected in a service's durability. CAPTCHAs are single use; compromised accounts last till victims wrest back control; exploit kits persist until every victim patches their system against stale attacks. This "value add" is critical for determining which interconnects of the underground are the most fragile to price increases and thus juicy targets for intervention.

5   Payment in the Black Market

One resource connects the whole underground ecosystem more than any other: payment. Banks, credit cards, digital currencies, and even crypto currencies are all highly regulated systems that make it difficult for criminals to fluidly extract and exchange funds. Consequently, payment represents a unique bottleneck for victim-to-criminal and criminal-to-criminal transactions.

Credit Card Processing: Victims transfer their personal wealth into the black market via credit cards for all product-based profit centers. While alternative digital currencies exist, these are not the norm consumers expect: of over a million customer transactions to pharmaceutical affiliate programs, McCoy et al. found 95% were paid via Visa and Master Card [96]. Mechanically, credit card processing involves four critical parties: the *cardholder* (e.g., customer); an *issuing bank* that manages the cardholder's finances; a *merchant* (e.g., criminal affiliate program); and the *acquiring bank* that manages the merchant's account [95]. The acquiring banks that process black market payments are few and far in between. McCoy et al. found only 25 international banks that supported all of the top 40 product-focused affiliate programs, the most popular of which were based in Azerbaijan, Latvia, and Mauritius [95].

These acquiring banks take on all liability (e.g., charge back due to fraud complaints) for a merchant account. Banks cope with high-risk merchants by charging higher transaction fees (10-20%), requiring upfront capital, and holding back 20% of transactions for 30–90 days before releasing funds to the merchant. Criminals are cognizant of these profit sinks: Stone-Gross et al. observed fake anti-virus affiliate programs carefully tuning their algorithms to refund customer complaints (3–8.5%) to mitigate potential charge backs to remain within each respective payment processor's "acceptable" fraud level [133]. Similarly, miscreants will purposefully label transaction as cosmetics or groceries as opposed to pharmaceuticals, which are "high risk" and might cause banks to shun their business. This concealment exposes criminals to potential legal action and asset seizure if revealed. As we discuss in Section 6, the delicate relationship between acquiring banks and criminals is one of the most fragile underground resources and ripe for intervention.

Money Exchanges & Pre-Paid Vouchers: Victim-to-criminal transactions, and to a lesser extent criminal-to-criminal transactions, rely on money exchanges (e.g., Western Union, Money Gram) and pre-paid vouchers such as Money Pak to irreversibly transfer funds in a potentially non-traceable manner. Research investigations in this area are sparse. Anecdotally, CAPTCHA solving services such as *anti-gate.com* accept MoneyGram payments. More concretely, Jones et al. found a network of mules involved in 419 scams help criminals process money transfers in the United States

and then forward the funds along [66]. The prevalence of exchange fraud and a lack of protections lead to lawsuits against both Western Union and MoneyGram [38, 65].

Digital Currencies: Criminals buying and selling within the black market set their own standards for acceptable currencies (and risk), the most common of which are digital. A litany of popular and niche non-criminal businesses experience abuse: PayPal, WebMoney, Qiwi, UKash, Skirll, PerfectMoney, CashU, Alibaba—the list goes on. In the theme of "commoditize everything" there are even currency exchanges like *24change.com* that convert dollars, rubles, and every digital currency in between. The most prolific digital currency tied to the black market was Costa Rica-based LibertyReserve. Criminals could anonymously create accounts and convert personal funds into "LR" dollars that could then be irreversibly transferred to other black market entities. Miscreants allegedly laundered *$6 billion* earned via fraud and abuse through the service before the United States government seized all the company's assets [125]. No single player has emerged since to rival the size of Liberty Reserve's fraudulent operations.

Crypto Currencies: Criminals also support crypto currencies as a decentralized payment infrastructure for criminal-to-criminal transactions. While a seemingly fertile ground for illicit transactions, crypto currencies suffer from two limitations. First, the transactions are not as anonymous as might be believed: Meiklejohn et al. demonstrated the pseudonymity of crypto currencies can be breached by monitoring public payment records [99]. Second, crypto-currencies merely push risk onto exchanges that convert fiat (e.g., government regulated) currencies to crypto equivalents. The emergence of these middle-men has not escaped attention from law enforcement: governments are increasingly placing regulations on exchanges for reporting large fund transfers and stricter proof of identity for account holders. Mt. Gox ran afoul of these requirements, which resulted in the seizure of over $5 million in assets [130]. Similarly, exchanges are ripe targets for abuse: of 40 BitCoin exchanges, Moore et al. found 18 exchanges have since closed shop and 9 suffered breaches resulting in the loss of hundreds of thousands of dollars [102]. Despite these risks, successful (traditional crime) markets such as the Silk Road were founded on crypto-currencies [21]. In our own experience however, most underground merchants who focus on electronic abuse tend to favor digital payment mechanisms such as WebMoney and PayPal.

6 Curbing the Black Market

The scope of the underground economy suggests it cannot be disrupted through the use of formal law enforcement intervention strategies alone. There are, however, numerous techniques that the security and law enforcement community can leverage in order to affect the practices of buyers, sellers, and victims whose devices serve as the infrastructure and data as a key commodity. In particular, the criminological framework of situational crime prevention may be of value in deterring and disrupting underground markets [25, 110]. This criminological perspective views offenders as rational actors who make choices to engage in crime based on their assessments of perceived risks, potential rewards, and situational factors such as environmental cues and victim behavior [28]. Situational crime prevention focuses on five categories designed to impact both offenders and victims by identify strategies to directly impact opportunities to offend by 1) making it more challenging to engage in crime, 2) increase the risk of detection, 3) reduce the rewards that may result from offending, 4) reduce provocations to offend, and 5) remove excuses for offending by affecting the behavior of targets and environmental conditions [22–24].

At the same time, offenders naturally adapt to crime prevention strategies and adjust their tactics accordingly, allowing crime to continue [19,62]. This phenomenon is referred to as displacement, recognizing that offenders may change who they target, the methods of offending, or moving to different environments in order to offend [29, 41]. Research on traditional illicit economies, including drug markets [62, 160], stolen goods [116, 132], and prostitution [88, 131], demonstrate that offenders are quick to displace their behaviors in order to continue to make a profit.

Given that the online underground economy is driven by similar economic imperatives, it is clear that there will be no way to completely deter or disrupt offender networks. Instead, innovative deterrent strategies only force behavioral change in offenders and markets. Thus, all possible strategies must considered, implemented, and revised over time. Recent approaches to disrupt the market can be situated within aspects of this theory, and examined for their value in either hardening potential targets from compromise, increasing the difficulty offenders may experience in completing a crime, hindering their ability to profit from an offense, or increasing the likelihood of arrest. It is important to note that we cannot argue in favor of any one strategy over another, as all have substantive strengths and weaknesses. Furthermore, principals from all five areas must be applied to any given crime problem in order to produce the greatest potential deterrent effect on offender behavior overall.

6.1 Protecting Users & Systems

Client and server-side security has dominated industry's response to digital abuse over the last decade. The spectrum of solutions—automated software updates, personal anti-virus, network packet scanners, firewalls, spam filters, password managers, two-factor authentication, certificates, and secure communication—all attempt to reduce the attack surface that criminals can penetrate. This strategy hinges on the belief that users can make conscious security decisions and keep

pace with the abuse arms race. In practice, managing security remains overly burdensome. Users regularly click through security warnings to access insecure HTTPS or malicious content [3]; re-use passwords to simplify accessing resources [31]; misuse encryption [162]; or fall victim to social engineering attacks that compromises a victim's system [123].

These strategies also fail to disincentive criminal fraud and abuse. As technology invariably marches forward, security professionals and users are forced into a never-ending firefight that involves shoring up system defenses and deploying weak heuristics against abuse. In turn, criminals adapt or find the subset of systems that remain vulnerable and resume operation. This reactive development cycle never affords defenders an opportunity to strike at the critical infrastructure or financial centers that underpin abuse, which might otherwise fundamentally change the war against for-profit abuse.

6.2 Exhausting Resources & Stockpiles

While the reliance of criminals on specialized support infrastructure streamlines abuse, it also exposes criminals to a range of new possible countermeasures where defenders target vulnerable infrastructure. The predominant strategy in this area has been to infiltrate botnets, identify the command & control systems, and then either sinkhole or take control of the botnet's hosting, effectively cutting off zombie machines from receiving new instructions [20, 109, 134]. Researchers have energetically debated the effectiveness of botnet take-downs due to collateral damage that sometimes ensues, such as Microsoft's takedown of No-IP [42] or the resilience of peer-to-peer systems like ZeroAccess to take-down [115]. Successful operations include the Rustock and Waldec takedown, while Mariposa and Grum's respective operators managed to regain control even after actioning. One of the largest takedown involved terminating peering agreements with McColo in 2008, a network operator tied to hosting some of the largest botnet C&Cs. Spam levels at the time dropped over 35% until spammers shifted their C&C infrastructure to other providers over the course of 5 months [139]. The presence of the pay-per-install and exploit-as-a-service market also impede takedown effectiveness, allowing botnet authors to recuperate lost infrastructure for a price. When the MegaD botnet was taken down it returned to its original operation within one month [20]. So long as compromised hosts are not a resource bottleneck, take-downs offer only a temporary reprieve.

An alternative strategy in this space is to target resource bottlenecks within the underground. For example, Twitter proactively targeted the account black market and disabled several million fakes before merchants could sell them to spammers [145]. While initially effective, Twitter made no changes to prevent new fake registrations. Within two weeks, merchants were selling accounts again. Google took an alternative route and targeted a bottleneck for cheap phone numbers that lead to throttling certain cell carriers and outright blocking commonly abused free VOIP providers [143]. This positively increased the cost of accounts by 30–40%, but did not outright defeat the market for phone verified accounts.

Similarly recent research explored the efficacy of disrupting abusive advertising through constraining another critical resource from the attacker—their domain names [17, 152]. For example, the work from Wang et al. examined the effectiveness of intervention efforts commonly used by luxury brand holders, where the brand holders use litigation to seize the domain names of storefronts selling knock off merchandise on the basis of trademark infringement. Interestingly, the authors find that despite the fact that a handful of luxury brands have seized, at a minimum, tens of thousands of domains over the last couple of years, SEO campaigns are largely unaffected as their sales of counterfeit goods continue to flourish. This is primarily due to two reasons. First, these efforts do not comprehensively seize all domains belonging to SEO campaigns. Second, these efforts typically take two months before the seizure goes into effect, thereby giving SEO campaigns a large window of opportunity to continue selling goods. Furthermore, evidence suggests attackers have already developed countermeasures to domain name seizures by stockpiling fresh domains and cycling them in response to any seizures.

While none of the strategies discussed outright defeat fraud and abuse, we believe this space offers a number of critical opportunities moving forward for disrupting fragile connections within the underground.

6.3 Disrupting Payment

Disrupting the flow of money from victims-to-criminals and criminals-to-criminals can disincentivize abuse. With no payout, entire profit centers and support centers disappear. Depending on the payment mechanisms, however, such interventions may have varying degrees of effectiveness. Some payment processors, such as credit card companies, are regulated. They can shut down the accounts of miscreants on legal grounds. At the other end of the spectrum, some payment processors are semi-regulated (e.g. WebMoney); in the case of Bitcoin—a peer-to-peer crypto-currency—there is not even a centralized payment processor; disruption can only target exchanges to fiat currency. Still a nascent concept, we contextualize the efforts thus far at cutting off criminals from laundering ill-gotten gains.

Victim-to-criminal: Product-based criminal profit centers are at the mercy of credit card payments. As discussed in Section 6, high-risk and illegal merchant activities are subject to increased fees, fines, and even asset seizure. Where criminals can replace lost hosting infrastructure and domains directly from the underground, banking relationships are another matter all together. Levchenko et al. found that only three banks were responsible for accepting payments for 95% of the spam

16

URLs [86]. Brand holders impacted by fraud and trademark abuse can alert the credit card networks involved, resulting in merchant banks severing relationships with known criminals. McCoy et al. found that persistent brand holder intervention from 2011–2012 disrupted payment processing for criminals for months at a time [95]. This insight is critical for two reasons. First, there is a defenders advantage: it takes only one successful product purchase to identify abuse and notify the merchant bank involved. Second, unlike compromised hosting and domains, takedowns have a large financial impact: assets seized by the merchant bank can be in excess of $1 million [95]. In aftermath, criminals must find new merchant banks who unwittingly take on their high risk activity, the bureaucratic process of which is orders of magnitude more complicated than registering a domain.

The qualitative evidence of this intervention was recorded by monitoring underground forums, as one poetic affiliate marketer wrote (translated from the Russian) "The sun is setting on the OEM era," which was in reference to the actions on the part of Microsoft to undermine payment processing for counterfeit software. In reactions to the payment disruption efforts of a major pharmaceutical company, a leading affiliate wrote (again translated from the Russian) "Right now most affiliate programs have a mass of declines, cancels and pendings, and it doesn't depend much on the program IMHO, there is a general sad picture, fucking Visa is burning us with napalm [95]."

This payment intervention has not gone unopposed by criminal merchants. Their main responses have been an escalating effort to detect and filter test purchases and removing brand holder's products when they launch payment intervention efforts. This reinforces the notion that security does exist in a void and every intervention will be met with countermeasures.

Criminal-to-criminal: Digital currencies traded by criminals are also ripe for disruption due to international finance laws. Rather than target criminals, law enforcement has frequently taken action against currency operators. This includes the United States prosecuting e-Gold—a now defunct digital currency—for money laundering and facilitating payments between criminals [58, 117]. LibertyReserve, previously discussed in Section 6, witnessed a similar fate and was dismantled for money laundering [138]. The disruption of payment systems is not, however, a panacea for market disruption. Displaced actors can move to a range of alternative payment services both digital and crypto-based, though they lose all seized assets in the process. Similarly, introspection into this market is more difficult compared to consumer-facing payment processing. A challenge remains for security practitioners to monitor criminal chatter and identify the evolution of payment mechanisms involved.

6.4   Targeting Actors

At the end of the day, cybercrime is like any other crime, and the most lasting solution is to arrest the perpetrators. Interventions of this type are few and far in between—yet effective when preformed. Examples include a rash of FBI arrests in 2004 for 28 criminals tied to the Shadowcrew carding forum [117]. Russian law enforcement arrested the alleged Blackhole exploit kit author "Paunch" in 2013 [77]. At the time, the exploit kit was the most popular in the market, dominating 30% of drive-by download websites [49]. Similarly, the Ukrainian police arrested five miscreants involved in the development of the Zeus banking Trojan [73].

However, the infrequency of arrests posses a significant challenge. Taking down one criminal opens doors copy-cats and competitors. With Blackhole gone, a slew of new exploit kits now exist that take its place [27]. Similarly, the source for Zeus was leaked and quickly led to knock-off variants [83]. There are also failures in this space where criminals remain at-large: the miscreants behind the Koobface botnet are known but currently beyond legal action [126]. Other times the criminal outfit is gutted, but the criminals involved were never identified. Microsoft has a standing bounty of $250,000 for information leading to the arrest of the criminals behind the now defunct Conficker botnet; the FBI is offering $3 million for information on the whereabouts of alleged Zeus developer Evgeniy Mikhailovich Bogachev [2]. As such, it remains unclear whether law enforcement is more effective than merely seizing assets, which requires a lower threshold of bureaucratic process.

6.5   Standing Challenges

Intervention—technical or otherwise—is the ultimate goal of security. Moving forward, we argue that research must focus on *data-driven interventions*. Throughout this work we have developed a taxonomy for reasoning about cybercrime and its most fragile business relationships. While evidence suggests that payment and resource disruptions have a lasting impact on the underground's profit, these conclusions are preliminary. Moving forward, the challenge of future research is to measure the *long-term efficacy* of proposed legal, technical, and financial solutions. Our taxonomy naturally provides the necessary metrics: overall revenue; the pricing of support infrastructure; resource availability (including renewability post-takedown); and the durability of goods and services (e.g., the value extracted before rendered obsolete). These metrics force researchers and industry practitioners to reconsider fire-fighting behavior and instead reflect on how criminals adapt to interventions. Furthermore, the metrics provide a global standard of performance that captures an important truth that abuse fighting is a cross-institutional ecosystem, not an island.

7 On the Horizon

Maturing Underground Measurement as a Field: We believe that underground research can mature from an exploratory niche to a core component of data-driven security research. To do this, research contributions must shift from analysis based off anecdotal evidence to thorough investigations of the ecosystem. Like other observational fields, this includes in-depth and longitudinal studies, the techniques for which we have systematized throughout this work. Nevertheless, there is a challenge moving forward to gain vantage points in an increasingly decentralized criminal ecosystem that has an incentive to prevent infiltration.

Monetizing Beyond Black Markets: Our conception of monetization must keep pace with shifts towards "grey" markets. These include potentially unwanted programs that tamper with browser settings (e.g., toolbars), legitimate pay-per-install companies (e.g., *installsmonetizer.com*), and other affiliate programs that criminals can infiltrate both for profit or acquiring installs. This interplay adds yet another component to understanding abuse: the degree of user consent.

Changing Threat Landscape: Technology leads and abuse follows. Much of our study focused on monetizing compromised desktops and servers, yet mobile and cloud computing are poised to dominate the field. While fresh from a research perspective, future threats targeting these systems nevertheless fit within our taxonomy. Compromised cloud clients and mobile devices are simply new support infrastructure components that serve as proxies, network resources, or gateways to sensitive financial information. These abuse vectors do not open up new profit centers; criminals must still fundamentally rely on consumer choice, social engineering, or outright theft to make money. Of the potential new profit centers, only crypto currencies pose an interesting new outlet that allows criminals to capitalize on stolen compute cycles, in a similar vein to click fraud.

8 Conclusion

The underground economy has evolved into a complex ecosystem with commoditized services that criminals compose to realize very different verticals of abuse. Understanding the intricate relationships that compose this opaque market is crucial. Without a clear framework for how attacks monetize victims and institutions, we argue it is impossible to devise effective countermeasures that have long-lasting effects.

To satisfy this gap, we developed a comprehensive taxonomy that captures the myriad components of the underground economy. We systematized the findings of the last decade of black market into a framework of profit centers that draw revenue into the underground and support centers that streamline abuse. As a part of our analysis, we emphasized the

fragile dependencies introduced by underground commoditization that are ripe targets for disruption. We believe that researchers and industry can leverage our framework to evaluate novel approaches in undermining existing cybercrime operations and to predict future trajectories of Internet crime.

References

[1] Dell SecureWorks. `http://www.secureworks.com/assets/pdf-store/other/2012.threat.report.pdf`, 2013.

[2] Wanted by the FBI. `http://www.fbi.gov/wanted/cyber/evgeniy-mikhailovich-bogachev`, 2015.

[3] Devdatta Akhawe and Adrienne Porter Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Usenix Security*, pages 257–272, 2013.

[4] David S. Anderson, Chris Fleizach, Stefan Savage, and Geoffrey M. Voelker. Spamscatter: Characterizing internet scam hosting infrastructure. In *Proceedings of the USENIX Security Symposium*, Boston, MA, August 2007.

[5] Ross Anderson. Why information security is hard-an economic perspective. In *Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual*, 2001.

[6] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J.G. van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, 2012.

[7] M.M. Andrade and N. Vlajic. Dirt jumper: A key player in today's botnet-for-ddos market. In *Internet Security (WorldCIS)*, 2012.

[8] BBC News. Nigeria scams 'cost UK billions'. `http://news.bbc.co.uk/go/pr/fr/-/2/hi/business/6163700.stm`, 2006.

[9] BBC News. More than 100 arrests, as FBI uncovers cyber crime ring. `http://www.bbc.co.uk/news/world-us-canada-11457611`, 2010.

[10] Hamad Binsalleeh, Thomas Ormerod, Amine Boukhtouta, Prosenjit Sinha, Amr Youssef, Mourad Debbabi, and Lingyu Wang. On the analysis of the zeus botnet crimeware toolkit. In *Privacy Security and Trust (PST), 2010 Eighth Annual International Conference on*, 2010.

[11] Rainer Böhme and Thorsten Holz. The effect of stock spam on financial markets. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, 2006.

[12] Brian Krebs. A One-Stop Money Mule Fraud Shop. `http://bit.ly/1uV4oC0`, 2010.

[13] Brian Krebs. All About Skimmers. `http://krebsonsecurity.com/all-about-skimmers/`, 2014.

[14] 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 Internet Measurement Conference (IMC)*, 2014.

[15] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. Measuring pay-per-install: The commoditization of malware distribution. In *USENIX Security Symposium*, 2011.

[16] Qiang Cao, Xiaowei Yang, Jieqi Yu, and Christopher Palow. Uncovering large groups of active malicious accounts in online social networks. In *Proceedings of the 2014 ACM conference on Computer and communications security*, 2014.

[17] Neha Chachra, Damon McCoy, Stefan Savage, and Geoffrey M. Voelker. Empirically Characterizing Domain Abuse and the Revenue Impact of Blacklisting. In *Proceedings of the Workshop on the Economics of Information Security (WEIS)*, pages 4:1:1–4:1:13, State College, PA, 2014.

[18] Kumar Chellapilla and David Maxwell Chickering. Improving Cloaking Detection Using Search Query Popularity and Monetizability. In *Proceedings of the SIGIR Workshop on Adversarial Information Retrieval on the Web (AIRWeb)*, Seattle, WA, August 2006.

[19] Michael Cherbonneau and Heith Copes. "Drive it like you Stole it" Auto Theft and the Illusion of Normalcy. *British Journal of Criminology*, 2006.

[20] Chia Yuan Cho, Juan Caballero, Chris Grier, Vern Paxson, and Dawn Song. Insights from the inside: A view of botnet management from infiltration. In *USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, 2010.

[21] Nicolas Christin. Traveling the silk road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22Nd International Conference on World Wide Web*, 2013.

[22] Ronald V Clarke. Situational crime prevention: Its theoretical basis and practical scope. *Crime and justice*, 1983.

[23] Ronald V Clarke. Situational crime prevention. *Crime and Justice*, 1995.

[24] Ronald V Clarke. *Situational crime prevention: Successful case studies (2nd ed.)*. Guilderland, NY: Harrow and Heston, 1997.

[25] Ronald V Clarke and Marcus Felson. Routine activity and rational choice. advances in criminological theory (Vol. 5). New Brunswic, NJ: Transaction Books, 1993.

[26] Lucian Constantin. Ransomware that demands Bitcoins is distributed by malware that steals bitcoins. `http://bit.ly/1u24YfL`, 2014.

[27] contagio. An Overview of Exploit Packs. `http://bit.ly/1xF4lru`, 2014.

[28] Derek B Cornish and Ronald V Clarke. Understanding crime displacement: An application of rational choice theory. *Criminology*, 1987.

[29] Derek B Cornish and Ronald V Clarke. Crime specialisation, crime displacement and rational choice theory. In *Criminal behavior and the justice system*. 1989.

[30] Damballa. Want to rent an 80-120k DDoS Botnet? `https://www.damballa.com/want-to-rent-an-80-120k-ddos-botnet/`, 2014.

[31] Anupam Das, Joseph Bonneau, Matthew Caesar, Nikita Borisov, and XiaoFeng Wang. The tangled web of password reuse. In *Proceedings of NDSS*, 2014.

[32] Neil Daswani and Michael Stoppelman. The anatomy of clickbot. a. In *Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets*, 2007.

[33] Vacha Dave, Saikat Guha, and Yin Zhang. Measuring and fingerprinting click-spam in ad networks. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, 2012.

[34] Vacha Dave, Saikat Guha, and Yin Zhang. Viceroi: catching click-spam in search ad networks. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 2013.

[35] David Jones. U.S. indicts hackers in biggest cyber fraud case in history. `http://reut.rs/1xnq34b`, 2013.

[36] Emiliano 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 SIGCOMM conference on Internet measurement conference*, 2014.

[37] Giancarlo De Maio, Alexandros Kapravelos, Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. Pexy: The other side of exploit kits. In *Proceedings of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2014.

[38] Department of Justice. Moneygram International Inc. Admits Anti-Money Laundering and Wire Fraud Violations, Forfeits 100 Million in Deferred Prosecution. `http://1.usa.gov/1RgMp1K`, 2012.

[39] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th Conference on USENIX Security Symposium*, 2004.

[40] Hanno Fallmann, Gilbert Wondracek, and Christian Platzer. Covertly probing underground economy marketplaces. In *Proceedings of the International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, 2010.

[41] Marcus Felson and Ronald Victor Gemuseus Clarke. *Opportunity makes the thief: Practical theory for crime prevention*, volume 98. London: Home Office, 1998.

[42] Dennis Fisher. Microsoft settles with no-ip over malware takedown. `http://bit.ly/1swgV8Y`, 2014.

[43] Dinei Florêncio and Cormac Herley. Phishing and money mules. In *Information Forensics and Security (WIFS), 2010 IEEE International Workshop on*, 2010.

[44] Dinei Florêncio and Cormac Herley. Is everything we know about password stealing wrong? *Security & Privacy, IEEE*, 2012.

[45] Jason Franklin, Vern Paxson, Adrian Perrig, and Stefan Savage. An inquiry into the nature and causes of the wealth of internet miscreants. In *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, VA, October 2007.

[46] Laura Frieder and Jonathan Zittrain. Spam works: Evidence from stock touts and corresponding market activity. *Hastings Comm. & Ent. LJ*, 2007.

[47] Hongyu Gao, Jun Hu, Christo Wilson, Zhichun Li, Yan Chen, and Ben Y Zhao. Detecting and characterizing social spam campaigns. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010.

[48] Vaibhav Garg, Chris Kanich, and L. Jean Camp. Analysis of ecrime in crowd-sourced labor markets: Mechanical turk vs. freelancer. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, 2012.

[49] Chris Grier, Lucas Ballard, Juan Caballero, Neha Chachra, Christian J. Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Damon McCoy, Antonio Nappa, Andreas Pitsillidis, Niels Provos, M. Zubair 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*, pages 821–832, 2012.

[50] Chris Grier, Kurt Thomas, Vern Paxson, and Michael Zhang. @ spam: the underground on 140 characters or less. In *Proceedings of the 17th ACM conference on Computer and communications security*, 2010.

[51] Doug Gross. Facebook cracking down on fake 'Likes'. http://www.cnn.com/2012/09/27/tech/social-media/facebook-fake-likes/index.html, 2012.

[52] H. Haddadi. Fighting Online Click-fraud Using Bluff Ads. *SIGCOMM Computer Communication Review*, 40, April 2010.

[53] Michael Hanke and Florian Hauser. On the effects of stock spam e-mails. *Journal of Financial Markets*, 2008.

[54] Shuang Hao, Matthew Thomas, Vern Paxson, Nick Feamster, Christian Kreibich, Chris Grier, and Scott Hollenbeck. Understanding the domain registration behavior of spammers. In *Proceedings of the Internet Measurement Conference (IMC)*, 2013.

[55] Cormac Herley and Dinei Florêncio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Economics of Information Security and Privacy*, 2010.

[56] Chase Hoffberger. Youtube strips universal and sony of 2 billion fake views. http://bit.ly/10MpDse, 2012.

[57] Thomas J Holt. Exploring the social organisation and structure of stolen data markets. *Global Crime*, 2013.

[58] Thomas J Holt and Eric Lampke. Exploring stolen data markets online: products and market forces. *Criminal Justice Studies*, 2010.

[59] Thorsten Holz, Christian Gorecki, Konrad Rieck, and Felix C Freiling. Measuring and detecting fast-flux service networks. In *Proceedings of the ISOC Network and Distributed Systems Symposium (NDSS)*, 2008.

[60] Ting-Kai Huang, Md Sazzadur Rahman, Harsha V Madhyastha, Michalis Faloutsos, and Bruno Ribeiro. An analysis of socware cascades in online social networks. In *Proceedings of the 22nd international conference on World Wide Web*, 2013.

[61] Jelena Isacenkova, Olivier Thonnard, Andrei Costin, Davide Balzarotti, and Aurelien Francillon. Inside the scam jungle: A closer look at 419 scam email operations. In *Security and Privacy Workshops (SPW), 2013 IEEE*, 2013.

[62] Bruce A Jacobs. Crack dealers' apprehension avoidance techniques: A case of restrictive deterrence. *Justice Quarterly*, 1996.

[63] John P John, Alexander Moshchuk, Steven D Gribble, and Arvind Krishnamurthy. Studying spamming botnets using botlab. In *NSDI*, 2009.

[64] John P. John, Fang Yu, Yinglian Xie, Arvind Krishnamurthy, and Martin Abadi. deSEO: Combating Search-Result Poisoning. In *Proceedings of the USENIX Security Symposium*, San Francisco, CA, August 2011.

[65] Andrew Johnson. U.S. Authorities Investigating Western Union. http://www.wsj.com/articles/SB10001424052702304610404579403544220536418, 2014.

[66] Jackie Jones and Damon McCoy. The check is in the mail: Monetization of craigslist buyer scams. In *IEEE eCrime Research Summit*, 2014.

[67] Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M Voelker, Vern Paxson, and Stefan Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM conference on Computer and communications security*, 2008.

[68] Chris Kanich, Nicholas Weaver, Damon McCoy, Tristan Halvorson, Christian Kreibich, Kirill Levchenko, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage. Show Me the Money: Characterizing Spam-advertised Revenue. In *Proceedings of the USENIX Security Symposium*, pages 219–234, 2011.

[69] M. Karami, S. Ghaemi, and D. Mccoy. Folex: An analysis of an herbal and counterfeit luxury goods affiliate program. In *eCrime Researchers Summit (eCRS), 2013*, 2013.

[70] Mohammad Karami and Damon McCoy. Understanding the emerging threat of ddos-as-a-service. In *Presented as part of the 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats*, Berkeley, CA, 2013. USENIX.

[71] Kevin McLaughlin. TJX Confirms Largest Credit-Card Breach Ever. http://bit.ly/1GJhNAY, 2007.

[72] Eduard Kovacs. Cryptolocker 2.0 appears to be the work of copycats. http://bit.ly/1pJ7tnt.

[73] Brian Krebs. $72m scareware ring used conficker worm. http://bit.ly/1uWJZx9, Jun 2011.

[74] Brian Krebs. Fake antivirus industry down, but not out. http://bit.ly/1ssjaJN, Aug 2011.

[75] Brian Krebs. Shady Reshipping Centers Exposed, Part I. http://bit.ly/1uTNjZ3, 2011.

[76] Brian Krebs. Fool Me Once... http://krebsonsecurity.com/2013/04/fool-me-once/, 2013.

[77] Brian Krebs. Meet Paunch: The Accused Author of the BlackHole Exploit Kit. http://bit.ly/1u8rXV3, 2013.

[78] Brian Krebs. Pavel Vrublevsky Sentenced to 2.5 Years. http://bit.ly/1ukWsJ5, 2013.

[79] Christian Kreibich, Chris Kanich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, and Stefan Savage. On the spam campaign trail. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, San Franciso, CA, April 2008.

[80] N. Kshetri. The Economics of Click Fraud. *Security Privacy, IEEE*, 8(3):45 –53, May-June 2010.

[81] Do kyum Kim, Marti Motoyama, Geoffrey M. Voelker, and Lawrence K. Saul. Topic Modeling of Freelance Job Postings to Monitor Web Service Abuse. In *Proceedings of the ACM Workshop on Artificial Intelligence and Security (AISEC)*, pages 11–20, Chicago, IL, October 2011.

[82] McAfee Lab. Mcafee threats report: First quarter 2013, 2013.

[83] McAfee Labs. Pws-zbot. https://kc.mcafee.com/resources/sites/MCAFEE/content/live/PRODUCT_DOCUMENTATION/23000/PD23030/en_US/McAfee_Labs_Threat_Advisory_PWS-ZBot.pdf, June 2014.

[84] Nektarios Leontiadis, Tyler Moore, and Nicolas Christin. Measuring and analyzing search-redirection attacks in the illicit online prescription drug trade. In *USENIX Security Symposium*, 2011.

[85] Nektarios Leontiadis, Tyler Moore, and Nicolas Christin. A nearly four-year longitudinal study of search-engine poisoning. In *Proceedings of the ACM Conference on Computer and Communications Security (Scottsdale, AZ*, 2014.

[86] Kirill Levchenko, Andreas Pitsillidis, Neha Chachra, Brandon Enright, Márk Félegyházi, Chris Grier, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Nicholas Weaver, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of the IEEE Symposium and Security and Privacy*, pages 431–446, Oakland, CA, May 2011.

[87] Liana Baker and Jim Finkle. Sony PlayStation suffers massive data breach. http://reut.rs/1cOiytt, 2011.

[88] John Lowman. Street prostitution control some canadian reflections on the finsbury park experience. *British Journal of Criminology*, 1992.

[89] Long Lu, Roberto Perdisci, and Wenke Lee. SURF: Detecting and Measuring Search Poisoning. In *Proceedings of the ACM Conference on Computer and Communications Security*, Chicago, IL, October 2011.

[90] Lucian Constantin. Target point-of-sale terminals were infected with malware. http://bit.ly/1pJVsy8, 2014.

[91] Cristian Lumezanu, Nick Feamster, and Hans Klein. # bias: Measuring the tweeting behavior of propagandists. In *Sixth International AAAI Conference on Weblogs and Social Media*, 2012.

[92] Lydia DePillis. Click farms are the new sweatshops. http://wapo.st/1bKLYVg, 2014.

[93] William R Marczak, John Scott-Railton, Morgan Marquisboire, and Vern Paxson. When governments hack opponents: A look at actors and technology. In *Proceedings of the 23rd USENIX Security Symposium*, 2014.

[94] John Markoff. Before the gunfire, cyberattacks. http://www.nytimes.com/2008/08/13/technology/13cyber.html, 2013.

[95] Damon McCoy, Hitesh Dharmdasani, Christian Kreibich, Geoffrey M. Voelker, and Stefan Savage. Priceless: The Role of Payments in Abuse-advertised Goods. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 845–856, 2012.

[96] Damon McCoy, Andreas Pitsillidis, Grant Jordan, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey M. Voelker, Stefan Savage, and Kirill Levchenko. PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs. In *Proceedings of the USENIX Security Symposium*, pages 1–16, Bellevue, WA, August 2012.

[97] Damon McCoy, Andreas Pitsillidis, Grant Jordan, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey M Voelker, Stefan Savage, and Kirill Levchenko. Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs. In *Proceedings of the 21st USENIX conference on Security symposium*, 2012.

[98] Brian S McWilliams. *Spam Kings: The real story behind the high-rolling hucksters pushing porn, pills, and%*@)# enlargements*. " O'Reilly Media, Inc.", 2004.

[99] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the ACM Internet Measurement Conference*, pages 127–140, Barcelona, Spain, October 2013.

[100] A. Metwally, D. Agrawal, A. El Abbadi, and Qi Zheng. On Hit Inflation Techniques and Detection in Streams of Web Advertising Networks. In *Proceedings of Distributed Computing Systems*, 2007.

[101] Brad Miller, Paul Pearce, Chris Grier, Christian Kreibich, and Vern Paxson. What's clicking what? techniques and innovations of today's clickbots. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, 2011.

[102] T. Moore and N. Christin. Beware the middleman: Empirical analysis of Bitcoin-exchange risk. In *Proceedings of IFCA Financial Cryptography'13*, 2013.

[103] Tyler Moore and Richard Clayton. Examining the impact of website take-down on phishing. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*, 2007.

[104] Tyler Moore and Richard Clayton. Evil searching: Compromise and recompromise of internet hosts for phishing. In *Financial Cryptography and Data Security*, 2009.

[105] Tyler Moore, Nektarios Leontiadis, and Nicolas Christin. Fashion crimes: trending-term exploitation on the web. In *Proceedings of the 18th ACM conference on Computer and communications security*, 2011.

[106] Marti Motoyama, Kirill Levchenko, Chris Kanich, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. Re:

CAPTCHAs – understanding CAPTCHA-solving from an economic context. In *Proceedings of the USENIX Security Symposium*, pages 435–452, Washington, D.C., August 2010.

[107] Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage, and Geoffrey M. Voelker. Dirty Jobs: The Role of Freelance Labor in Web Service Abuse. In *Proceedings of the USENIX Security Symposium*, pages 203–218.

[108] Marti Motoyama, Damon McCoy, Kirill Levchenko, Stefan Savage, and Geoffrey M. Voelker. An Analysis of Underground Forums. In *Proceedings of the ACM Internet Measurement Conference*, pages 71–80, Berlin, CA, November 2011.

[109] Yacin Nadji, Manos Antonakakis, Roberto Perdisci, David Dagon, and Wenke Lee. Beheading hydras: performing effective botnet takedowns. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, 2013.

[110] Graeme R Newman and Ronald V Clarke. *Superhighway Robbery: Preventing E-commerce Crime*. Cullompton, United Kingdom: Willan, 2013.

[111] Nicole Perlroth. Tally of Cyber Extortion Attacks on Tech Companies Grows. http://nyti.ms/11fmqT7, 2014.

[112] Nick Nikiforakis, Federico Maggi, Gianluca Stringhini, M Zubair Rafique, Wouter Joosen, Christopher Kruegel, Frank Piessens, Giovanni Vigna, and Stefano Zanero. Stranger danger: exploring the ecosystem of ad-based url shortening services. In *Proceedings of the 23rd international conference on World wide web*, 2014.

[113] Yuan Niu, Hao Chen, Francis Hsu, Yi-Min Wang, and Ming Ma. A quantitative study of forum spamming using context-based analysis. In *NDSS*, 2007.

[114] Youngsam Park, Jackie Jones, Damon McCoy, Elaine Shi, and Markus Jakobsson. Scambaiter: Understanding targeted nigerian scams on craigslist. In *Proceedings of the ISOC Network and Distributed Systems Symposium (NDSS)*, 2014.

[115] Paul Pearce, Vacha Dave, Chris Grier, Kirill Levchenko, Saikat Guha, Damon McCoy, Vern Paxson, Stefan Savage, and Geoffrey M. Voelker. Characterizing large-scale click fraud in zeroaccess. In *Proceedings of the ACM Conference on Computer and Communications Security*, Scottsdale, Arizona, November 2014.

[116] Ken Pease. The kirkholt project: Preventing burglary on a british public housing estate. *Security Journal*, 1991.

[117] Kimberly Kiefer Peretti. Data breaches: what the underground world of carding reveals. *Santa Clara Computer & High Tech. LJ*, 2009.

[118] Peter Svensson. ATM Fraud Allowed Thieves To Steal 45 Million In Hours. http://huff.to/1pK7nMp, 2013.

[119] Kevin Poulsen. Cybercrime Supersite 'DarkMarket' Was FBI Sting, Documents Confirm. http://www.wired.com/2008/10/darkmarket-post/, 2008.

[120] Kevin Poulsen. The Secret Service Agent Who Collared Cybercrooks by Selling Them Fake IDs. http://wrd.cm/1xz4X1Q, 2013.

[121] Emil Protalinski. Facebook: 8.7 percent are fake users. http://www.cnet.com/news/facebook-8-7-percent-are-fake-users/, 2012.

[122] Niels Provos, Panayiotis Mavrommatis, Moheeb Abu Rajab, and Fabian Monrose. All your iframes point to us. In *17th USENIX Security Symposium*, 2008.

[123] Moheeb Abu Rajab, Lucas Ballard, Panayiotis Mavrommatis, Niels Provos, and Xin Zhao. The nocebo effect on the web: an analysis of fake anti-virus distribution. In *USENIX workshop on large-scale exploits and emergent threats (LEET)*, 2010.

[124] Jacob Ratkiewicz, Michael Conover, Mark Meiss, Bruno Gonçalves, Alessandro Flammini, and Filippo Menczer. Detecting and tracking political abuse in social media. In *ICWSM*, 2011.

[125] Reed Albergotti and Jeffrey Sparshott. U.S. Says Firm Laundered Billions. http://on.wsj.com/1GQr1eE, 2013.

[126] Riva Richmond. Web Gang Operating in the Open. http://nyti.ms/1qH7pzi, 2012.

[127] Roni Robbins. Twitter says over 13 million accounts may be bots and fakes. http://bit.ly/1kBx4M8, 2014.

[128] Robert Lemos. Point-of-sale malware has now infected over 1,000 companies in US. http://bit.ly/1nCVjW1, 2014.

[129] Eugene Rodionov and Aleksandr Matrosov. The evolution of tdl: Conquering x64. *ESET, June*, 2011.

[130] Romain Dillet. Feds Seize Another 2.1 Million From Mt. Gox, Adding Up To 5 Million. http://tcrn.ch/1EDAWlR, 2013.

[131] Michael S Scott and Kelly Dedel. Street prostitution: Problem oriented policing guide series. US Department of Justice, Office of Community Oriented Policing Services, 2006.

[132] Per Stangeland. Other targets or other locations? An analysis of opportunity structures. *British journal of criminology*, 1998.

[133] Brett Stone-Gross, Ryan Abman, Richard Kemmerer, Christopher Kruegel, Doug Steigerwald, and Giovanni Vigna. The Underground Economy of Fake Antivirus Software. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, Washington, DC, June 2011.

[134] Brett Stone-Gross, Marco Cova, Lorenzo Cavallaro, Bob Gilbert, Martin Szydlowski, Richard Kemmerer, Christopher Kruegel, and Giovanni Vigna. Your botnet is my botnet: analysis of a botnet takeover. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 635–647. ACM, 2009.

[135] Brett Stone-Gross, Ryan Stevens, Richard Kemmerer, Christopher Kruegel, Giovanni Vigna, and Apostolis Zarras. Understanding Fraudulent Activities in Online Ad Exchanges. In *Proceedings of the Internet Measurement Conference (IMC)*, Berlin, Germany, November 2011.

[136] Gianluca Stringhini, Manuel Egele, Christopher Kruegel, and Giovanni Vigna. Poultry Markets: On the Underground Economy of Twitter Followers. In *Proceedings of the Workshop on Online Social Networks (WOSN)*, Helsinki, Finland, August 2012.

[137] Gianluca Stringhini, Christopher Kruegel, and Giovanni Vigna. Detecting Spammers on Social Networks. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*, Austin, TX, December 2010.

[138] J. Surowiecki. Why did criminals trust liberty reserve. *The New Yorker, May 31*, 2013.

[139] Symantec. State of spam. http://goo.gl/QhCgVI, 2009.

[140] Samaneh Tajalizadehkhoob, Hadi Asghari, Carlos Gañán, and Michel van Eeten. Why them? extracting intelligence about target selection from zeus financial malware. 2014.

[141] Kurt Thomas, Chris Grier, and Vern Paxson. Adapting social spam infrastructure for political censorship. In *Proceedings of the 5th USENIX conference on Large-Scale Exploits and Emergent Threats*. USENIX Association, 2012.

[142] Kurt Thomas, Chris Grier, Dawn Song, and Vern Paxson. Suspended accounts in retrospect: an analysis of twitter spam. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, 2011.

[143] Kurt Thomas, Dmytro Iatskiv, Elie Bursztein, Tadek Pietraszek, Chris Grier, and Damon McCoy. Dialing back abuse on phone verified accounts. In *Proceedings of the 21st Annual Conference on Computer and Communications Security*, 2014.

[144] Kurt Thomas, Frank Li, Chris Grier, and Vern Paxson. Consequences of connectivity: Characterizing account hijacking on twitter. In *Proceedings of the 21st Annual Conference on Computer and Communications Security*, 2014.

[145] 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*, 2013.

[146] Rob Thomas and Jerry Martin. The underground economy: Priceless. In *USENIX ;login:*, 2006.

[147] Robert Tillman and Michael Indergaard. *Pump and dump: The rancid rules of the new economy*. Rutgers University Press, 2008.

[148] Tim Wilson. Source Code For SpyEye Trojan Published; More Exploits On The Horizon, Researcher Says. http://ubm.io/1wck0w9, 2011.

[149] United States Attorney's Office. Manhattan U.S. Attorney Charges Seven Individuals for Engineering Sophisticated Internet Fraud Scheme That Infected Millions of Computers Worldwide and Manipulated Internet Advertising Business. http://1.usa.gov/1pgFXgW, 2011.

[150] United States Attorney's Office. Nine Individuals Indicted in One of the Largest International Penny Stock Frauds and Advance Fee Schemes in History. http://1.usa.gov/1evqjDM, 2013.

[151] United States Attorney's Office. Cyber Criminal Pleads Guilty To Developing And Distributing Notorious SpyEye Malware. http://1.usa.gov/1crV7XT, 2014.

[152] David Wang, Matthew Der, Mohammad Karami, Lawrence Saul, Damon McCoy, Stefan Savage, and Geoffrey M. Voelker. Search + seizure: The effectiveness of interventions on seo campaigns. In *Proceedings of the ACM Internet Measurement Conference*, Vancouver, BC, Canada, November 2014.

[153] David Wang, Stefan Savage, and Geoffrey M. Voelker. Cloak and Dagger: Dynamics of Web Search Cloaking. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 477–490, Chicago, IL, October 2011.

[154] David Wang, Stefan Savage, and Geoffrey M. Voelker. Juice: A longitudinal study of an seo campaign. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, pages 7:4:1–7:4:17, San Diego, CA, February 2013.

[155] Gang Wang, Tristan Konolige, Christo Wilson, Xiao Wang, Haitao Zheng, and Ben Y Zhao. You are how you click: Clickstream analysis for sybil detection. In *USENIX Security*, 2013.

[156] Gang Wang, Christo Wilson, Xiaohan Zhao, Yibo Zhu, Manish Mohanlal, Haitao Zheng, and Ben Y Zhao. Serf and turf: crowdturfing for fun and profit. In *Proceedings of the 21st international conference on World Wide Web*, 2012.

[157] Yi-Min Wang and Ming Ma. Detecting Stealth Web Pages That Use Click-Through Cloaking. Technical report, Microsoft Research, December 2006.

[158] Yi-Min Wang, Ming Ma, Yuan Niu, and Hao Chen. Spam Double-Funnel: Connecting Web Spammers with Advertisers. In *Proceedings of the International World Wide Web Conference (WWW)*, Banff, Alberta, May 2007.

[159] Mark Ward. Cryptolocker victims to get files back for free. http://www.bbc.com/news/technology-28661463.

[160] David Weisburd and Lorraine Green. Policing drug hot spots: The jersey city drug market analysis experiment. *Justice Quarterly*, 1995.

[161] Thomas Whiteside. *Computer capers: tales of electronic thievery, embezzlement and fraid*. New American Library, 1979.

[162] Alma Whitten and J Doug Tygar. Why johnny can't encrypt: A usability evaluation of pgp 5.0. In *Usenix Security*, 1999.

[163] Baoning Wu and Brian D. Davison. Cloaking and Redirection: A Preliminary Study. In *Proceedings of the SIGIR Workshop on Adversarial Information Retrieval on the Web*, Chiba, Japan, May 2005.

[164] Baoning Wu and Brian D. Davison. Detecting Semantic Cloaking on the Web. In *Proceedings of the International World Wide Web Conference (WWW)*, Edinburgh, United Kingdom, May 2006.

[165] Yinglian Xie, Fang Yu, Kannan Achan, Rina Panigrahy, Geoff Hulten, and Ivan Osipkov. Spamming botnets: signatures and characteristics. *ACM SIGCOMM Computer Communication Review*, 2008.

[166] Sandeep Yadav, Ashwath Kumar Krishna Reddy, AL Reddy, and Supranamaya Ranjan. Detecting algorithmically generated malicious domain names. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010.

[167] Michael Yip, Nigel Shadbolt, and Craig Webber. Structural analysis of online criminal social networks. In *Intelligence and Security Informatics (ISI), 2012 IEEE International Conference on*, 2012.

[168] Michael Yip, Nigel Shadbolt, and Craig Webber. Why forums?: an empirical analysis into the facilitating factors of carding forums. In *Proceedings of the 5th Annual ACM Web Science Conference*, 2013.

[169] Michael Yip, Craig Webber, and Nigel Shadbolt. Trust among cybercriminals? carding forums, uncertainty and implications for policing. *Policing and Society*, 2013.

[170] Jianwei Zhuge, Thorsten Holz, Chengyu Song, Jinpeng Guo, Xinhui Han, and Wei Zou. Studying malicious websites and the underground economy on the chinese web. In *Proceedings of the Workshop on Economics of Information Security (WEIS)*, 2009.