

Carnegie Mellon University

CARNEGIE INSTITUTE OF TECHNOLOGY

THESIS

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF Doctor of Philosophy

TITLE Structuring Disincentives for Online Criminals

PRESENTED BY Nektarios Leontiadis

ACCEPTED BY THE DEPARTMENT OF

Engineering and Public Policy

Nicolas Christin
ADVISOR, MAJOR PROFESSOR

August 18, 2014
DATE

Douglas Sicker
DEPARTMENT HEAD

August 19, 2014
DATE

APPROVED BY THE COLLEGE COUNCIL

Vijayakumar Bhagavatula
DEAN

August 25, 2014
DATE

Structuring Disincentives for Online Criminals

Submitted in partial fulfillment of the requirements for
the degree of
Doctor of Philosophy
in
Engineering and Public Policy

Nektarios Leontiadis

B.S., Computer Science, Athens University of Economics and Business

M.S., Information Systems, Athens University of Economics and Business

Carnegie Mellon University

Pittsburgh, PA

August, 2014

UMI Number: 3691274

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI 3691274

Published by ProQuest LLC (2015). Copyright in the Dissertation held by the Author.

Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code



ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

Copyright © 2014 by Nektarios Leontiadis
All rights reserved except the rights granted by the
Creative Commons Attribution-Noncommercial 4.0 International License

Acknowledgements

I am eternally indebted to certain individuals that have not only helped me become worthy of entering this graduate program, but also guided me along this hard but life-changing journey. Words are hardly enough to express the extent of my gratitude, but, lacking a more empowering medium, they will have to suffice.

The most prominent figure in my hall of fame is unquestionably my mentor, advisor, and chair of my committee, Dr. Nicolas Christin. I am grateful to Nicolas for many reasons. From the beginning, when he hardly knew me, he had enough confidence in me to take me in as his student, and to provide me with financial support to leave Greece and join this graduate program. Ever since, he has taught me how to think in a scientifically-sound manner, and how to write high-quality technical work publishable in top research venues. His ethics, and professionalism have had an immense impact on me, equipping me with the necessary mental tools to make an impact in our society. On a personal level, he has been supportive, encouraging and a source of inspiration at the times mostly needed. Nicolas, I am extremely fortunate to having met you, and thank you for all you have done for me.

A close second is Dr. Tyler Moore, a member of my committee and co-author of the majority of the work I published while in this program. We started working together soon after I entered the PhD program, and, through our collaboration, he has been paramount to my professional and personal evolution. Tyler's skills in statistical analysis, economic modeling, data presentation, and technical writing

have been my definition of excellence and a driving force for personal advancement throughout this time. I admit having said that “I want to become Tyler when I grow up”, but with the rate with which he is advancing, I doubt I will ever be able to catch up. On the positive side, he remains a figure to look up to. Tyler, I am grateful to having met you, and for your role in my life.

I would also like to thank Professor Alfred Blumstein, and Dr. Pedro Ferreira, members of my committee, for their valuable and insightful role in completing this work. Their extensive knowledge and experience in criminology and economics have provided me with the confidence to use concepts from their respective fields, and make my contribution in science an interdisciplinary one.

In an age where the value of family and religion is increasingly underestimated for their role in providing the society with grounded individuals, I hereby attest to their invaluable role in my existence, and in my path in this life. I thank God for the family that created me, raised me, nurtured me, taught me, and equipped me for life from birth until 2009, when I left my first home, Greece; my late mother Maria, my father Nikos, and my sister Dorianna. I thank God for the family that surrounds me, inspires me, empowers me, guides me, and loves me at what has become my second home, Pittsburgh; my incredibly smart and beautiful wife Jill, and the two loving daughters she has brought into our lives, Maria and Sofia. I would also like to acknowledge the pivotal and grounding role in my journey thus far of four friends I call brothers; Yannis Mallios, Stelios Eliakis, Vagelis Kotsonis, and Dr. Thanassis Avgerinos.

In addition, I would like to thank the administrative staff at my home department, Engineering and Public Policy, and at CyLab for allowing me to focus on my research and not on the necessary but time-consuming technicalities of academic life. In this regard, special thanks goes to Vicki Finney, EPP’s graduate program administrator for being always available and resourceful.

Finally, I am grateful to the various sources of funding that supported me throughout my tenure at CMU. This research was partially supported by Carnegie Institute of Technology (CIT) Dean's Tuition Fellowship; by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office; by ICANN; by the Department of Homeland Security (DHS) Science and Technology Directorate, Cyber Security Division (DHS S&T/CSD) Broad Agency Announcement 11.02, the Government of Australia and SPAWAR Systems Center Pacific via contract number N66001-13-C-0131; and by the National Science Foundation under ITR award CCF-0424422 (TRUST) and SaTC award CNS-1223762. This dissertation represents the position of the author and not that of the aforementioned agencies.

Abstract

This thesis considers the structural characteristics of online criminal networks from a technical and an economic perspective. Through large-scale measurements, we empirically describe some salient elements of the online criminal infrastructures, and we derive economic models characterizing the associated monetization paths enabling criminal profitability. This analysis reveals the existence of structural *choke points*: components of online criminal operations being limited in number, and critical for the operations' profitability. Consequently, interventions targeting such components can reduce the opportunities and incentives to engage in online crime through an increase in criminal operational costs, and in the risk of apprehension.

We define a methodology describing the process of distilling the knowledge gained from the empirical measurements on the criminal infrastructures towards identifying and evaluating appropriate countermeasures. We argue that countermeasures, as defined in the context of situational crime prevention, can be effective for a long-term reduction in the occurrence of online crime.

“You may encounter many defeats, but you must not be defeated. In fact, it may be necessary to encounter the defeats, so you can know who you are, what you can rise from, how you can still come out of it.”

~Maya Angelou

Contents

Acknowledgements	iii
Abstract	vi
List of Tables	xv
List of Figures	xviii
List of Abbreviations	xxii
1 Introduction	1
2 Research overview	6
2.1 Thesis statement	6
2.2 Research scope	7
2.3 Research questions	9
2.4 Structure of the thesis	10
3 Background and Related Work	12
3.1 Economics and structure of online criminal markets	13
3.1.1 Abuse-based advertising on the Internet	13
3.1.2 Opportunities enabling online crime	14
3.1.3 The flow of money in online crime	15
3.1.4 Online criminal network structures	16
3.1.5 Modeling the economics of online crime	17
3.2 Legal and health aspects of online pharmacies	18

3.2.1	Regulation	18
3.2.2	Online pharmacy accreditation and reputation programs	22
3.2.3	Law enforcement operations	22
3.2.4	Health risks	24
3.3	Social and economic aspects of criminal behavior	25
3.3.1	Modeling offenders' decisions	27
4	Measuring and analyzing search-redirectation attacks in the illicit online prescription drug trade	30
4.1	Background	32
4.1.1	Search-redirectation attacks	34
4.2	Measurement methodology	36
4.2.1	Infrastructure overview	36
4.2.2	Query selection	37
4.2.3	Additional query-sample validation	38
4.2.4	Query corpus characteristics	41
4.2.5	Search-result classification	42
4.3	Empirical analysis of search results	44
4.3.1	Breakdown of search results	44
4.3.2	Variation in search position	46
4.3.3	Turnover in search results	47
4.3.4	Variation in search queries	47
4.4	Empirical analysis of search-redirectation attacks	50
4.4.1	Concentration in search-redirectation attack sources	50
4.4.2	Variation in source infection lifetimes	53
4.4.3	Characterizing the unlicensed online pharmacy network	58
4.4.4	Attack websites in blacklists	60

4.5	Towards a conversion rate estimate	62
4.6	Conclusions	66
5	Pricing and inventories at unlicensed online pharmacies	68
5.1	Background	71
5.1.1	Advertising techniques	72
5.1.2	The emergence of online black markets	74
5.2	Measurement methodology	75
5.2.1	Selecting and parsing pharmacies	75
5.2.2	Extracting inventories	79
5.2.3	Collecting supplemental data	80
5.3	Inventory analysis	82
5.3.1	Drug availability by pharmacy type	83
5.3.2	Product overlap between different types of pharmacies	86
5.3.3	Identifying drug conditions served by unlicensed pharmacies	89
5.3.4	Identifying suppliers	90
5.4	Pricing strategies	94
5.4.1	Pricing differences by seller and drug characteristics	94
5.4.2	Volume discounts as competitive advantage	96
5.4.3	How competition affects pricing	98
5.5	Conclusions	101
6	A longitudinal analysis of search-engine poisoning	103
6.1	Background	107
6.2	Data collection	108
6.2.1	Query corpus	108
6.2.2	Search result datasets	109

6.2.3	Combining the datasets	113
6.3	Search-result analysis	115
6.3.1	Overview	115
6.3.2	Search result dynamics	117
6.3.3	User intentions	123
6.4	Cleanup-campaign evolution	125
6.4.1	Cleaning up source infections	126
6.4.2	Cleaning up traffic brokers and destinations	129
6.5	Advertising network	132
6.6	Limitations	136
6.7	Conclusions	137
7	Trending-term exploitation on the web	139
7.1	Methodology	143
7.1.1	Building query corpora	144
7.1.2	Data collection	146
7.1.3	Website classification	148
7.2	Measuring trending-term abuse	153
7.2.1	Incidence of abuse	153
7.2.2	Network characteristics	156
7.2.3	MFA in Twitter	158
7.2.4	Search-term characteristics	159
7.3	Economics of trending-term exploitation	166
7.3.1	Exposed population	166
7.3.2	Revenue analysis	170
7.4	Search-engine intervention	174

7.5	Conclusion	177
8	Empirically measuring WHOIS misuse	179
8.1	Methodology	183
8.1.1	Constructing a microcosm sample	183
8.1.2	Pilot registrant survey	184
8.1.3	Experimental measurements	184
8.2	Experimental domain registrations	185
8.2.1	Registrar selection	186
8.2.2	Experimental domain name categories	187
8.2.3	Registrant identities	188
8.3	Breaking down the measured misuse	191
8.3.1	Postal address misuse	191
8.3.2	Phone number misuse	193
8.3.3	Email address misuse	194
8.3.4	Overall misuse per gTLD	197
8.4	WHOIS anti-harvesting	197
8.5	Misuse estimators	199
8.5.1	Estimators of email misuse	202
8.5.2	Estimators of phone number misuse	204
8.6	Limitations	204
8.7	Conclusion	205
9	An examination of online criminal processes to formulate and evaluate disincentives	207
9.1	Background	209
9.2	The case of illicit online prescription drug trade	211
9.2.1	Illicit advertising	212

9.2.2	Unlicensed online pharmacies	236
9.3	The case of trending term exploitation	253
9.3.1	A procedural analysis of trending term exploitation	254
9.3.2	Situational measures targeting trending term exploitation	259
9.3.3	Impact of situational measures targeting trending term exploitation	263
9.3.4	Overall Assessment	267
9.4	The case of WHOIS misuse	268
9.4.1	A procedural analysis of WHOIS misuse	269
9.4.2	Situational measures targeting WHOIS misuse	271
9.4.3	Overall assessment of situational measures targeting WHOIS misuse	272
9.5	Concluding remarks: Towards a generalizable methodology for on-line crime analysis and prevention	273
9.5.1	Commonalities in criminal infrastructures	274
9.5.2	Designing effective solutions	275
9.5.3	Limitations and future work	277
10	Summary and conclusions	279
A	Surveying registrants on their WHOIS misuse experiences	284
A.1	Methodology	285
A.1.1	Survey translations	286
A.1.2	Types of questions	286
A.2	Response and error rates	287
A.3	Analysis of responses	288
A.3.1	Characteristics of the participants	288
A.3.2	Reported WHOIS misuse	289
A.4	Discussion	292

A.4.1	Potential survey biases	293
B	Registrant survey supplemental material	294
B.1	Invitation to participate in registrant survey	294
B.2	Consent form	296
B.3	Survey questions	298
B.4	Definitions of terms	314
B.4.1	Document information	318
B.4.2	Acknowledgment of sources	318
	Bibliography	320

List of Tables

4.1	Comparing different lists of search terms to the main list used in the Chapter. All numbers are percentages.	40
4.2	Intention-based classification of the 218 queries in the drug query corpus (<i>Q</i>).	42
4.3	Classification of all search results (4–10/2010).	45
4.4	TLD breakdown of source infections.	51
4.5	Monthly search query popularity according to the Google Adwords Traffic Estimator.	62
5.1	Summary data for all four data sources.	81
5.2	Scheduled drugs, narcotics, drugs in shortage, and top drugs at <i>familymeds.com</i> , unlicensed pharmacies and Silk Road.	84
5.3	Similarities in drugs sold using different drug definitions.	86
5.4	Odds-ratios identifying the medical conditions that are over-represented or under-represented in the inventories of unlicensed pharmacies.	88
5.5	Unit price discounts for different drug categories.	95
5.6	Unit prices and percentage discounts offered by <i>familymeds.com</i> and unlicensed pharmacies for 60-pill and 90-pill orders relative to the unit price of 30-pill orders.	96
6.1	Datasets for pharmaceutical queries.	109
6.2	Search-result composition.	116
6.3	Confusion matrix for the search-redirection classification.	122
6.4	Characteristics of actively redirecting URLs.	132

6.5	Characteristics of traffic brokers.	133
6.6	Characteristics of pharmacies.	134
6.7	Connected components in the graph describing daily observed redirection chains.	134
6.8	Overlap in the criminal infrastructures.	135
7.1	Total incidence of malware and MFA in Web search and Twitter results.	154
7.2	Prevalence of malware in trending and control terms	155
7.3	Malware campaigns observed.	157
7.4	Malware and MFA incidence broken down by trending-term category.	161
7.5	Estimated number of visits to MFA and malware sites for trending terms.	167
7.6	Estimated number of visits to MFA and malware sites for trending terms.	176
8.1	Number of domains under each of the top five gTLDs	183
8.2	Breakdown of measured WHOIS-attributed misuse, broken down by gTLD and type of misuse.	197
8.3	Methods for protecting WHOIS information at 104 registrars and three registries.	200
8.4	Statistically-significant regression coefficients affecting email address misuse	203
8.5	Statistically-significant regression coefficients.	204
9.1	Costs and benefits for each of the actors involved in, or enabling illicit online advertising, before and after an intervention targeting such activity.	224
9.2	Average reduction of redirected traffic (i.e. effectiveness) per unit of complexity.	235
9.3	Average reduction of revenue from illicit online sales of prescription drugs (i.e. benefit) per unit of complexity.	252

9.4 Average reduction of traffic being subject to trending term exploitation (i.e. effectiveness) per unit of complexity. 267

List of Figures

4.1	Example of search-engine poisoning.	33
4.2	Distribution of different classes of results according to the position in the search results.	46
4.3	Change in the average domains observed each day for different classes of search results over time.	48
4.4	Search-redirection attacks appear in many queries; health resources and blog spam appear less often in popular queries.	49
4.5	Rank-order CDF of domain impact reveals high concentration in search-redirection attacks.	52
4.6	Survival analysis of search-redirection attacks shows that TLD and PageRank influence infection lifetimes.	53
4.7	Network analysis of redirection chains reveals community structure in search-redirection attacks.	57
4.8	Comparing web and email blacklists.	61
5.1	A variant of the search-redirection attack that appeared as a response to search engine intervention.	74
5.2	Example of multiple drug names, dosages, currencies and prices presented within a single page.	81
5.3	Heat map of the Jaccard distances between all pairs of pharmacies in the unlicensed pharmacy set. After reordering pharmacies, we observe a number of clusters that appear to have similarities.	91
5.4	Effect of different levels of distance threshold and different linkage criteria.	93

5.5	Cumulative distribution of pharmacies as a function of the number of clusters considered (Using average-linkage, $t = 0.31$).	94
5.6	Cumulative distribution functions of the median percentage-point price discount per pharmacy (left) and per drug (right).	97
5.7	Bar plot of the median unit price discount for drug-dosage combinations grouped in increasing number of unlicensed pharmacies selling the drug at the specified dosage.	99
6.1	Percentage of search results per category, averaged over a 7-day sliding window.	118
6.2	Percentage of unclassified search results detected as malicious based on the content by VirusTotal	121
6.3	Similar to Figure 6.1, but examining only the top 10 search result positions.	123
6.4	Percentage of search results per category, based on the type of query. Active redirections dominate results regardless of the intention of the query.	124
6.5	Survival probability for source infections.	125
6.6	Median time (in days) to cleanup source infections over time, source infections per 100 results over time, and median time (in days) to cleanup source infections by TLD.	127
6.7	Survival probability for source infections, traffic brokers and destinations over all time, and median time in days for cleanup	129
6.8	Major autonomous systems hosting traffic brokers.	131
6.9	Maximum and average degree of traffic brokers and destinations over time.	135
7.1	Ad-filled website appearing in the results for trending terms.	140
7.2	Calibration tests weigh trade-offs between comprehensiveness and efficiency for collecting trending-term results.	147
7.3	Trending-term exploitation on Twitter.	159
7.4	Exploring how popularity and ad price of trending-terms affects the prevalence of malware, and ad-laden sites.	161

7.5	Number of estimated daily victims for malware appearing in trending and control terms.	168
7.6	CDF of visits for domains used to transmit malware or ads in the search results of trending-terms.	169
7.7	MFA prevalence in the top 10 search results fell after Google announced changes to its ranking algorithm on February 24, 2011, designed to counter “low-quality” results.	174
8.1	Graphical representation of the experimental domain name combinations we register with each of the 16 registrars.	186
8.2	Targeted postal spam attributed to WHOIS misuse.	192
9.1	Components of the crime commission process in the illicit online prescription drug trade.	213
9.2	The two methods used to redirect illicitly acquired web traffic to unlicensed online pharmacies.	218
9.3	Probability density plot and cumulative distribution plot of complexity-benefit analysis for a software (CMS and web server) provider-based intervention.	231
9.4	Probability density plot and cumulative distribution plot of complexity-benefit analysis for a search engine-based intervention.	233
9.5	Probability density plot and cumulative distribution plot of complexity-benefit analysis for a registrar and Internet service provider-based intervention.	235
9.6	Probability density plot and cumulative distribution plot of complexity-benefit analysis for law enforcement based intervention.	248
9.7	Probability density plot and cumulative distribution plot of complexity-benefit analysis for registrar based intervention.	250
9.8	Probability density plot and cumulative distribution plot of complexity-benefit analysis for a US Customs and Border Protection-based intervention.	251
9.9	CDFs of the benefits of two interventions per unit of complexity to identify the stochastic dominant.	253
9.10	Components of the crime commission process in the case of trending term exploitation.	255

9.11 Plots of impact probability density and cumulative distribution functions of measures targeting trending term exploitation, when considering different actor sets. 266

9.12 Components of the crime commission process in the case of WHOIS misuse. 268

List of Abbreviations

API	Application Programming Interface
AS	Autonomous System
CBP	Customs and Border Protection, a United States federal law enforcement agency of the Department of Homeland Security.
CDF	Cumulative Distribution Function
CMS	Content Management Software
CPC	Cost per Click
CSA	Crime Script Analysis
CSIP	Center for Safe Internet Pharmacies
CTR	Click-Through Rate
DEA	Drug Enforcement Administration, a United States federal agency of the Department of Justice.
DNS	Domain Name System
FBI	Federal Bureau of Investigation, a United States federal agency of the Department of Justice.

FDA	Food and Drug Administration, a United States federal agency of the Department of Health and Human Services.
FQDN	Fully Qualified Domain Name
GNSO	Generic Names Supporting Organization
gTLD	global Top Level Domain, non-country specific TLD like .COM and .NET
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
ICANN	Internet Corporation for Assigned Names and Numbers, a non-profit organization that coordinates and regulates the use of Internet's global resources like DNS and WHOIS
IP	Internet Protocol address
ISP	Internet Service Provider
JS	JavaScript
MFA	Made-for-AdSense
NABP	National Association of Boards of Pharmacy
NDC	National Drug Code
NDF-RT	National Drug File – Reference Terminology
PDF	Probability Density Function
PPC	Pay-per-Click

RAA	Registrar Accreditation Agreement, a contract between ICANN and registrars that defines the operational responsibilities and rights of the latter.
SCP	Situational Crime Prevention
SSL	Secure Socket Layer
STD	Sexually Transmitted Disease
TF-IDF	Term Frequency – Inverse Document Frequency
TLD	Top Level Domain, for example .DE, .GR, and .COM
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
US	United States
USD	United States Dollars
USPS	United States Postal Service
VIPPS	Verified Internet Pharmacy Practice Sites

Introduction

One of the key theories integrated in the criminal justice systems is the theory of general deterrence. This theory, introduced in 1764 by Cesare Beccaria [16], focuses on how to prevent crime, instead of trying to explain criminal behavior. The general deterrence theory is based on two key assumptions; First, the punishment associated with criminal activities should forestall offenders from engaging in crime in the future. Second, the certainty of punishment should prevent others from committing criminal activities.

Accordingly, modern laws, regulations, and policies are mere implementations of this theory. These constructs are designed to introduce artificial *punishment costs* to actions that deviate from the prescribed and socially acceptable behavior. Fines and incarceration are examples of punishment costs. The punishment costs, added to the opportunity cost of engaging in an illicit activity, should introduce a significant enough loss of potential gain (compared to the gain of being an economically rational law abiding agent), thus deterring illicit actions.

The deterrent effects of punishment appear to be successful against crime for centuries in the physical world. However, the Internet has been challenging the

effectiveness of the deterrents, due to the different nature of online crime. Crime in the online domain lacks physical violence, and it is shielded by difficulties in attribution and jurisdictional complexities. Such impediments lower the perceived risk of apprehension, and increase the expected profitability of online crime, making potential offenders positively disposed to engage in online criminal activities.

In the context of this thesis, we define online crime as any activity involving the use of computers and the Internet with the intent to defraud individuals, or to trade illicit goods. Twitter and email spam that illicitly advertise male enhancement products, phishing emails looking to lure individuals into revealing their e-banking credentials, and unlicensed online pharmacies selling drugs without a prescription are examples of such illicit activities.

However different in their methods, a single common aspect motivates the actors behind such activities: *profit*. Whether the flow of money is based on volition (e.g. buying or selling prescription drugs without a prescription), or fraud (e.g. accessing another person's bank account without their permission), initiators of the illicit activity, acting as economically rational agents, are monetizing their technical skill set by bypassing the processes set by a lawful society.

Given the low expectancy of punishment, online crime is a seemingly safe way to make illicit profit. Indeed, Moore et al. have provided evidence that online criminals are economically motivated rational agents [158]. Consequently, we hypothesize that the challenges in deterring online crime lie not with the efficacy of the laws, but, rather, with their application.

Our thesis is that structural characteristics of online criminal networks can help to identify economic pressure points. These pressure points may be used as crime deterrents by making online crime less lucrative and riskier.

Legal scholars have expressed concerns in recent years on whether the existing legal frameworks in the United States (US) and abroad are providing adequate protection to the public when applied to the online domain [76, 79, 113]. More specifically, they raise two questions about the issue; First, are current laws adequate to protect society from domestic online criminal activity? Second, given the borderless nature of the Internet, where anyone can access resources and services offered in any part of the world, how can current regulations be enforced when there are issues of jurisdiction?¹We consider those concerns using the case of online pharmacies, because of its immense impact in public health.

Do laws for brick-and-mortar pharmacies regulating the distribution of controlled substances (i.e. prescription drugs) apply also to Internet-based pharmacies located in the US? The legal proceedings of *US v. Birbragher* [47] offer a short answer to this question. In this case, operators of the US-based online pharmacy *Pharmacom* were selling prescription drugs without verifying the identity and medical records of their customers, nor their possession of valid prescriptions. The operators processed more than 246,000 prescriptions, yielding revenue in excess of \$20 million United States Dollars (USD) between January of 2003 and May of 2004. They were convicted based on the Controlled Substances Act [227], which was signed into US law in 1970.

Existing US laws, therefore, are mostly adequate when dealing with traditional criminal activity transitioned into the online world, as long as it takes place within the US jurisdiction. Of course, there are cases where laws have needed amendments whenever this transition has allowed for certain loopholes.²

¹ State actors like China and Iran are capable of imposing arbitrary limitations in the online resources someone can access when within their borders. However these cases are out of the scope of this thesis.

² Examples of such amendments are the Ryan Haight Act of 2008, the Anti-Phishing Act of 2004, and the Internet False Identification Prevention Act of 2000.

However, the case of online pharmacies becomes more challenging when we factor in the globalized, Internet-based market. For example, how can the Ryan Haight Online Consumer Protection Act—a US federal statute which regulates specific aspects of online pharmacies by rendering it illegal to “deliver, distribute, or dispense a controlled substance by means of the Internet” without an authorized prescription, or “to aid and abet such activity” [223]—be enforced on an online pharmacy based outside the US jurisdiction that sells prescription drugs to US-based consumers? This is an especially significant problem, as the low prices of prescription drugs abroad [222] create incentives for US-based customers to purchase their medication over the Internet.

In this respect, Interpol-coordinated operations Pangea [110] have had considerable success in crippling activities related to illicit trafficking of drugs. These operations take place annually, last one week, and require the communication between customs, regulators, and national police forces from many countries. Indeed, in its most recent execution (2014), Operation Pangea VII had 111 participating countries. However, their limited duration is a good indicator that the required coordination effort is a prohibitive factor for running them yearlong.

Despite the noteworthy impact of operations Pangea, efforts lacking international coordination are subject to significant limitations. In the US, operations targeting illicit sales of prescription drugs from international marketplaces depend on the capability to properly identify and examine packages at the ports of entry. However, the immense number of packages that should be inspected by the US Customs and Border Protection (CBP), contrasted to the limited capabilities for inspections, allow a potentially significant amount of illicit drugs to reach US-based customers [218, 234]. Even in cases with no jurisdictional issues to prosecute offenders residing abroad, the Food and Drug Administration (FDA) depends on the foreign countries to take action against the wrongdoer [221].

While our intention is not to criticize law enforcement, the online trade of pharmaceuticals magnifies the dysfunction of the existing deterrents. In the case of unlicensed online pharmacies, it is highly uncertain whether an offender will be punished for illicitly trading prescription drugs online. The inadequacy in enforcing existing laws on the Internet questions the certainty of punishment, invalidating a fundamental assumption of general deterrence.

The Internet offers suitable opportunities to online criminals, where the reward of committing online crime is greater than the chance of being caught. While current laws are applicable against online crime, limitations in their enforcement allow for the economic incentives to break them. There is a need for a holistic approach against online crime that provides the necessary disincentives to curb such activity.

In this dissertation, we offer a methodological perspective on how to effectively impede online criminal activity. This methodology prescribes that law enforcement resources should target the critical criminal resources most expensive to acquire and maintain. We base our work predominantly on the criminal activity associated with the illicit online prescription drug trade because of its societal impact. However, our thesis is not restricted to this specific case, but rather to cases sharing similar characteristics. In addition, we show that the proposed methodology does not solve an ephemeral problem. While the advance in technical sophistication of criminal online operations is significant through the years, the characteristics of the critical criminal resources, as defined through our methodology, are found to be invariant.

Research overview

In this chapter we formally present our thesis statement followed by the scope within which it is established. In presenting the scope of the thesis, we reason on the relevance and importance of the cases of online crime we study. We then present the research questions we attempt to answer in this dissertation, and our work in providing scientifically-grounded answers.

2.1 Thesis statement

This thesis considers the structural characteristics of online criminal networks from a technical and an economic perspective. Through large-scale measurements, we empirically describe some salient elements of the online criminal infrastructures, and we derive economic models characterizing the associated monetization paths enabling criminal profitability. This analysis reveals the existence of structural *choke points*: components of online criminal operations being limited in number, and critical for the operations' profitability. Consequently, interventions targeting such components can reduce the opportunities and incentives to engage

in online crime through an increase in criminal operational costs, and in the risk of apprehension.

We define a methodology describing the process of distilling the knowledge gained from the empirical measurements on the criminal infrastructures towards identifying and evaluating appropriate countermeasures. We argue that countermeasures, as defined in the context of situational crime prevention, can be effective for a long-term reduction in the occurrence of online crime.

2.2 Research scope

The scope of this thesis covers online criminal networks exploiting public interest in products, services, and information, because of their capacity to negatively affect large portions of the population. We start focusing on the domain of online prescription drug trade, mainly due to its importance in public health. In addition, as this case is one of the most visible online criminal activities, it provides a large footprint for our analysis. However, we provide evidence that our methodology is applicable also to other forms of online criminal activity allowing miscreants to profit illicitly. In this respect, we continue with studying three additional cases of online crime, which, as we show, also affect significant portions of the online population: (i) the exploitation of trending news, (ii) the misuse of domain name registrants' personal identifying information, and (iii) the illicit online trade of a variety of counterfeit consumer goods and services utilizing similar criminal network structures. We focus on those areas for the following reasons:

- The case of prescription drugs has immense public policy implications. By enabling access to prescription drugs without a valid prescription and without proper health assessment by a medical doctor, consumers are essentially

allowed to self-medicate. This practice is a dangerous one as it can lead to severe health issues [92].

- Exploitation of trending news topics and of prescription drugs involve a similar monetization path. We use this case study to reinforce our argument that financial profit is an invariable motive for online crime. In addition, we affirm the existence of similar concentration points in the criminal network as it depends largely on a few scarce resources.
- As long as opportunity exists, online criminals do not necessarily need to employ overly elaborate technical skills. We offer a proof of concept by studying the misuse of the public directory WHOIS [53], which holds personal identifying information of domain name registrants.
- The online criminal network structure of the illicit prescription drug trade and its critical components do not manifest only in this case study. We study a variety of other commodities, like counterfeit applications, counterfeit watches, gambling and others, to reaffirm that the criminal structure and critical components is a shared resource. In addition, the same set of economic disincentives can negatively impact these illicit online markets as well.

While the specific set of criminal activities does have set particularities, they do not necessarily limit the strength or breadth of our findings. The methodology we propose for analyzing online crime suggests that, for efficient interventions, one has to look at the high-level characteristics of the underlying operations, beyond their technical implementation and realization; For example, we show in Chapter 9 that processes used to fraudulently attract potential customers, and to process payments have similar traits across criminal operations. Therefore, even though our initial empirical analysis informs the structuring of disincentives for online crim-

inals specifically engaged in these illicit activities, we are able to identify the meta-characteristics of critical criminal resources. We show that such resources are not existent only at the specific criminal operations, but are rather common in online crime. Consequently, we argue if we were to study in this thesis a different set of cases of online crime—like typosquatting or email spam—we would derive similar observations.

The empirical analysis and modeling we offer in this thesis is done in the context of the US legal framework due to the large online market which it represents. With an Internet penetration of 81% in the US in 2013 [105], we believe that our work can have a significant impact on a large portion of the population by reducing opportunities for potential offenders to engage in criminal online activity.

2.3 Research questions

We consider the various aspects of the thesis statement through five research questions. These questions also define our methodology in proving the validity of the research statement in an empirical and systematic way.

1. Are there any structural characteristics in the illicit online prescription drug trade that are a critical resource compared to other structural components? Are those critical resources the outcome of a cost limiting-process that can inform economic pressure points able to curb the trade's profitability?
2. Is the observed structure of online criminal networks an ephemeral phenomenon that would make disruption strategies we suggest in this thesis futile?
3. Do other forms of illicit online activity exhibit a similar structure, with economic pressure points as the illicit online prescription drug trade?

4. Is it the technical skills – as reflected on the complexity of online criminal structures – or the existence of suitable opportunities that enable online criminal activity?
5. Is it possible to disrupt online criminal networks by targeting critical components of their structure? What would this process involve? Would it be more efficient compared to present efforts?

2.4 Structure of the thesis

This thesis is organized as follows; We start in Chapter 3 with a review of the related work in studying online crime, which is the overarching context of this thesis. However, and considering our focus on the problem of unlicensed online pharmacies, we also offer a thorough examination of the various safety, regulatory, and law enforcement aspects. We conclude this review by presenting the criminological framework that informs our approach for effective actions targeting online crime.

In the following three chapters we empirically examine unlicensed online pharmacies from three distinct but complementary perspectives. In Chapter 4 we introduce the topic of unlicensed online pharmacies, providing insights on (i) the extent of the problem, (ii) the online criminals' methods to fraudulently promote their illicit businesses, and (iii) the structure of this online criminal network. Further on, in Chapter 5 we turn our focus on the operation of the unlicensed pharmacies. Based on our empirical measurements and analysis, we highlight the various tactics for attracting potential customers in their illicit businesses, and explain their economic viability despite the constant law enforcement efforts targeting their operation. Finally, in Chapter 6 we analyze the parallel evolution, over a period of four years, of the criminal tactics enabling the advertising and operation of illicit

online—mainly pharmaceutical—businesses, and of the frivolous measures trying to disrupt these online markets.

In Chapters 7 and 8 we take a step away from the case of unlicensed pharmacies, investigating two separate cases of online crime; One that involves the manipulation of search engines to exploit and monetize the interest towards trending news topics (Chapter 7); And one that investigates the misuse of the WHOIS directory to initiate fraudulent communication towards domain registrants (Chapter 8). While the nature of these case studies is seemingly very different compared to the unlicensed online pharmacies, our analysis brings in the foreground the underlying commonalities of online crime: it is structurally organized and enabled by the inexistence of disincentives—in other words, the availability of opportunities—to engage in online crime.

In Chapter 9 we reconsider all these cases of online crime from a procedural and a criminological perspective. Based on the preceding empirical analysis, we distill the structural characteristics of online crime by defining the related crime scripts. Moreover, we define appropriate countermeasures based on the theory of Situational Crime Prevention (SCP), and we evaluate their effectiveness, considering the complexity of their implementation. Finally, in Chapter 10 we summarize our work in the context of this thesis, providing a discussion of our contributions, and we propose future research avenues.

Background and Related Work

This thesis builds upon three different bodies of research; Web security, legal and health aspects of online pharmacies, and criminology.

Our work is directly associated with measurement studies that quantify various characteristics of online criminal activities and markets. In Section 3.1, we present related work on the economics and structure of online crime.

In Section 3.2, we use the case of the illicit online prescription drug trade to demonstrate the severe effects of online crime. We further show that, while the US legal framework pertaining to online pharmacies is mostly comprehensive, major domestic and international law enforcement operations targeting illicit online drug markets have limited effects.

Furthermore, we build upon social and economic concepts extensively studied in criminology in our effort to propose methods to discourage online criminal activity. In Section 3.3, we summarize the concepts and ideas most relevant to this work.

3.1 Economics and structure of online criminal markets

In the past decade, computer security attacks driven by fame and reputation have transformed into online crime driven by financial gain [158]. This observation has motivated measurement studies that quantify the characteristics of online criminal networks, guiding possible intervention policies. Due to the amount of related literature, we focus on work most closely related to this thesis.

3.1.1 Abuse-based advertising on the Internet

Many studies, e.g. [8, 116, 132, 249], have focused on email spam, describing the magnitude of the problem in terms of network resources being consumed, as well as some of its salient characteristics. More specifically, Anderson et al. [8] analyzed a set of 1 million spam emails in order to understand the hosting infrastructure availing the illicit content associated with spam emails.

Xie et al. in [249] were able to identify 580,466 spam emails, and reveal they sent from botnets with more than 340,000 infected hosts. Kanich et al. managed to infiltrate a botnet, and monitor its operation over a period of 26 days revealing a conversion rate¹ of 1 every 10,000 emails sent [116]. The small conversion rate indicates that email spam is a game of very large numbers, and it is not a very effective technique to advertise products.

Spamming techniques are evolving and increase their effectiveness by better targeting potential customers, as described by the flurry of spam observed in social networks [86]. In this work, the authors showed that Twitter spam has a conversion rate of 0.13%, which is 3 orders of magnitude higher than email spam. However, Lumezanu and Feamster in [143] show that online criminals often target a combination of platforms for their abuse-based advertisement. Indeed, there is

¹ Fraction of email spam that eventually result in a sale.

a significant overlap between illicit domains appearing in email spam and Twitter spam.

Spam has been increasingly supplemented by new innovative techniques that we (e.g. [128]) and others (e.g. [176,241]) have studied. Ntoulas et al. [176] measure search engine manipulation attacks caused by “keyword stuffing” at offending web sites. This technique lures search engines into believing that a certain web page is relevant to specific content, while a human observer would not make this association. Wang et al. in [241], study the prevalence of *cloaking*, a technique that allows rogue web pages to attract unintended user traffic by concealing their true nature. They found that pharmaceutical terms are more extensively targeted compared to other popular search terms.

Measurement studies of spam have also informed possible intervention policies by identifying some infrastructure weaknesses. For instance, taking down a single hosting company used by online criminal infrastructures significantly reduced the overall volume of email spam on the Internet [38]. However, the same study also highlights the unpredictable side effects of such an intervention. Infiltration of spam-generating botnets, as suggested by [183], has also been effective in designing more accurate spam filtering rules.

3.1.2 Opportunities enabling online crime

A series of papers by Moore and Clayton [156, 157, 159] investigates the economics of phishing, and reveal interesting insights on the behavior of phishers. Most importantly, and in relation to this thesis, they show that efficient communication between interested parties and speedy responses are essential characteristics in implementing deterrents against online crime. Lack of those characteristics act as opportunities for online criminals to engage in their illicit behavior.

Web hosting providers [26], Internet Service Providers (ISPs) [210], and search engines [145] prominently appear in the list of parties providing opportunities for illicit behavior. Either due to their ignorance to indicators associated with abusive activity or to their intentional inactivity, they become part of the overall problem.

3.1.3 The flow of money in online crime

A separate branch of research has focused on economic implications of online crime. While not related in content with the thesis, we were greatly inspired by employed measurement methodologies; For example, Thomas and Martin [215], Franklin et al. [73] and Zhuge et al. [255] passively monitor the advertised prices of illicit commodities exchanged in varied online environments (IRC channels and web forums). They estimate the following: the size of the markets associated with the exchange of credit card numbers, identity information, email address databases, and forged video game credentials.

Working on a topic of relevance to the focal point of this thesis, Kanich et al. in [117] examined the revenues of abuse-advertised enterprises selling counterfeit drugs and software. The work is based on ground truth data representing sales transactions from 10 such enterprises, which are often called affiliate networks in the related work. In essence, affiliate networks are businesses that use a variable set of partners (i.e. affiliates) to market their products, and, usually, drive their sales [97]. This work, Kanich et al. describe a set of inference techniques allowing for reasonable understanding of the customers' purchasing behavior, and of affiliate revenues.

3.1.4 Online criminal network structures

Evidence on the existence of affiliate criminal networks relying on illicit advertising has been informally described in [194]. However, there has recently been a heightened interest in the research community to present empirical evidence of the structural relationships among online criminals.

In a recent line of work [33, 46, 132, 135, 166] researchers have started looking at aggregate data associated with online criminal activities in order to reveal associations between online criminals.

While focusing on different online activities, Christin et al. [33] and Costin et al. [46] employ a graph-based methodology to analyze the use of a set of limited resources in criminal operations, and identify their critical structural characteristics. The limited nature of these resources is in direct association with their high operational costs (e.g. phone numbers). Similar to our work, they use the findings of their analysis to inform methods of efficient intervention.

Levchenko et al. [132] provide a thorough investigation of the different actors participating in spamming campaigns, from the spammers themselves, to the suppliers of illicit goods (e.g. luxury items, software, pharmaceutical drugs). This research, in conjunction with our work in [128], offered one of the first indications of the existence of concentration points in the structure of online criminal networks. The key difference between the two studies is that Levchenko et al. are focusing on businesses advertised by email spam, while we are looking into search engine manipulation.

The importance of a few hosts in the criminal online infrastructure is the focus of Li et al. in [135]. By analyzing the association between 4 million malicious web addresses, they reiterate on the importance of a small number of traffic brokers

in the online criminal ecosystem. These entities are similar to what we revealed in [128], but they are studied from a broader perspective of criminal activities.

Using graph analysis and algorithms for community detection, Nadji et al. process historical data from the Domain Name System (DNS) associated with malicious domains [166]. They present a set of metrics that can divulge a set of critical components in the criminal online infrastructure. This work relates to this thesis, as removing the critical components would result in significant reduction in the functionality of the criminal network.

3.1.5 Modeling the economics of online crime

Another series of papers [3,9,72,93,94] looks deeper into the economics of online crime, and less into the technical details of the online criminal activity. The authors in [93,94] highlight the importance of economic analysis in developing proper security mechanisms, and of caveats in characterizing online criminal markets.

Anderson et al. in [9] considered the different types of costs incurred because of the different flavors of online crime, namely: i) criminal revenues, ii) direct losses by the victims as a consequence of the criminal activity, iii) indirect losses in the society, and iv) costs of countermeasures. They found that the defensive mechanisms used to counter criminal online activities are financially inefficient as they are not targeting critical criminal components. In this thesis we will offer a defensible strategy to counter online crime in an efficient way.

Florêncio and Herley in [72] develop a threat model prescribing that online criminal operations need to be profitable in expectation. Reducing the criminals' expectation in profit is an integral part of this thesis.

3.2 Legal and health aspects of online pharmacies

While, as we show, the empirical work examining various aspects of online crime is rather extensive, it is rarely placed in the context of the existing regulatory framework, or juxtaposed to the ongoing law enforcement efforts. Moreover, the social implications of online criminal activity are usually either implicitly addressed, or not examined at all. Therefore, in this section we use the case of the illicit online prescription drug trade to demonstrate the severe effects of online crime. We briefly present the existing legal framework associated with online pharmacies, and we contrast it with the underwhelming law enforcement operations targeting unlicensed online pharmacies. We highlight the importance of better policing, by discussing the efforts for self-regulation by industries affected by the operation of unlicensed online pharmacies, and the health risks the operation of latter imposes.

3.2.1 Regulation

The regulatory framework in the US pertaining to drugs is laid out both on the federal and on the state level. This section focuses on the federal level, which is mostly incorporated in its state level counterparts. There are numerous scholarly articles debating on the policy aspects of the regulation (e.g. [48, 87, 207]). However, it is not in the intent of this thesis to elaborate on government regulation. Instead, we intend to offer policy recommendations based on empirically collected evidence.

In 1938, the US Congress passed a set of laws under the Federal Food, Drug, and Cosmetic Act [224] giving authority to the FDA to oversee the safety of food, drugs and cosmetics. The FDA in turn, as the primary domestic drug policy enforcer, has established cooperation with other federal agencies, namely the De-

partment of Justice, the Drug Enforcement Administration (DEA), the Federal Bureau of Investigation (FBI), the US CBP, and the Postal Inspection Service [217].

The comprehensive Drug Abuse Prevention and Control Act of 1970 [227], and especially Title II of entitled Controlled Substances Act is the core piece of legislation regulating the drug market. It defines categories of drugs named *Schedules* that characterize their potency for abuse, and dangers of misuse. In addition, it enforces certain procedures that control how drugs enter the market, how they can be sold (e.g. after a physical examination, with a prescription, etc.), and the limitations in terms of importation for personal or commercial use.

The most recent amendment in the legislative framework of drugs is the Ryan Haight Online Pharmacy Consumer Protection Act of 2008 [223]. It extends the Controlled Substances Act to regulate online pharmacies explicitly. In short, it requires that (i) online pharmacies must have an associated physical *brick-and-mortar* pharmacy that should be properly licensed in the states that it operates, (ii) online pharmacies can neither sell prescription drugs without a prescription, nor state that they do so, and (iii) issuing a prescription for the first time requires a physical in-person examination. While this law was a good step towards proper regulation of online pharmacies, the problem of international pharmacies shipping their merchandise to the US without complying with the US regulation, remains. Other legislative efforts² have tried to address different aspects of the problem of illicit online prescription drugs unsuccessfully.

Examples of those are:

- **Internet Prescription Drug Consumer Protection Act of 2000.** Focusing on the domestic online pharmacies, the law would allow the use of certain

² For example: (i) Internet Prescription Drug Consumer Protection Act of 2000, (ii) Safe Online Drug Act of 2004, (iii) The Pharmaceutical Market Access and Drug Safety Act of 2005, (iv) The Internet Drug Sales Accountability Act of 2005, (v) Safe Internet Pharmacy Act of 2007, and (vi) Safeguarding America's Pharmaceuticals Act of 2008.

judicial tools (e.g. injunction – to halt the operation of illicit online pharmacies), and it would require online pharmacies to list their address and license information. It did not pass into law as it did not address issues associated with international online pharmacies.

- **Safe Online Drug Act of 2004.** The Act would establish certification requirements for online pharmacies that would have to be renewed every two years. It also introduced liabilities for using certain clauses for advertising of online pharmacies, both on the side of the advertiser and of the advertised entity. Finally, it would enforce proper identification of drug purchases with electronic payment systems. Such transactions could then be blocked or restricted.

The bill was reintroduced in 2005 but was never enacted.

- **The Pharmaceutical Market Access and Drug Safety Act of 2005.** The high prices of prescription drugs in the domestic market compared to the prices of same drugs internationally was the focus of this Act. Acknowledging the financial difficulty, especially of senior citizens, to get a prescription, and then fill it, the Act gave support to the reimportation of cheap prescription drugs. In doing so, it also provisioned for protective measures in Internet sales similar to the Ryan Haight Act.

The bill was reintroduced in 2007, 2009, and 2011 but was never enacted.

- **The Internet Drug Sales Accountability Act of 2005.** The bill acknowledges the importance of advertising in promoting and enabling illicit online sales of prescription drugs. To this end, it makes third-party advertising networks liable for accepting illicit prescription drug advertisements, and prescribes the use of a “system” that allows for timely take-downs of offending

ads. With every violation fined up to \$1 million, it gave a good incentive for advertisers to efficiently monitor their business.

- **Safe Internet Pharmacy Act of 2007.** The act would enforce all Internet pharmacies, domestic and international, to acquire an operating license in the US. In addition, they would have to list their physical location and the states where they are allowed to market drugs. Moreover it would relinquish any liability of Internet search engines that would include illicit online pharmacies in their directories.
- **Safeguarding America's Pharmaceuticals Act of 2008.** The intent of this proposed regulation was to enable proper identification of counterfeit drugs, and to allow FDA to destroy them at the ports of entry³. With the persisting limitations in the inspection of imported packages [218], the legislation would end up being unenforceable.

The bill was reintroduced in 2011, and 2013. In the latest attempt, it passed the House, and it is currently being considered in the Senate [220].

Liang and Mackey in [136] have proposed a comprehensive statutory solution attending the pending problems. They discuss the risks of online drugs sales, but they also touch on issues of accountability; Most importantly, the current improper validation of the credentials of online pharmacies allow both unlicensed pharmacies and search engines to profit illicitly [60, 123, 152]. They propose the institution of low or no cost prescription medication, and strict regulation of online drug financial transactions and facilitators of the illicit activity (e.g. search engines).

³ FDA's current authority is limited to denying the delivery of imported illicit drugs

3.2.2 Online pharmacy accreditation and reputation programs

It is noteworthy that both the Internet and pharmaceutical industries have acknowledged the regulatory gap, and have made attempts for self-regulation through accreditation, verification, and reputation programs. These programs have been developed to assist consumers in making informed choices, especially when considering their ability to ship drugs across jurisdictions.

For instance, National Association of Boards of Pharmacy (NABP) is a professional association whose members are boards of pharmacies from across North America, Australia and New Zealand. Since 1999, the NABP has established the Verified Internet Pharmacy Practice Sites (VIPPS) [170] program, which provides accreditation, for a fee, to law-abiding online pharmacies. In addition, NABP provides an extensive list of “not recommended” online pharmacies, which fail to demonstrate that they abide to the law of their jurisdiction [169]. Likewise, LegitScript [124] is an online service that provides a list of law-abiding pharmacies. LegitScript is backed by the NABP, and is reportedly used by Google and Microsoft to determine whether pharmacies are legitimate or not. Many other online verification programs do exist. Their stringency varies and range from requiring valid pharmacy licenses in the US or Canada (e.g., `pharmacychecker.com`) to mere reputation forums (e.g., `pharmacyreviewer.com`). Because of the large number of online pharmacies, many pharmacies are neither accredited or licensed, nor blacklisted. For instance, `eupillz.com` an online pharmacy selling prescription drugs in 2013, did not appear at the time in any of the aforementioned databases.

3.2.3 Law enforcement operations

Considering the availability of laws in the US that regulate the operation of online pharmacies, the next logical question is what is currently being done to deter the

operation of unlicensed online pharmacies? News outlets often discuss major law enforcement operations targeting online pharmacies with the focus being on the number of storefronts that were shut down. FDA recognized – as early as 2001 – the significant complexities in investigating online pharmacies, and in enforcing current policies [91]. FDA’s efforts have focused on the shutdown of the illicit web stores, rather than on the identification of the structures that enable their operation.

Example of such operations are Cyber Chase [230] and Cyber X [231]. However, considering the extent of the problem of unlicensed online pharmacies, and the significant duration of those law enforcement operations, the outcomes are usually underwhelming, highlighting the shortcomings of current enforcement mechanisms [27]. Moreover, the unfortunate inability of US CBP to identify illegal drugs at the ports of entry [218], makes the certainty of punishment even weaker.

In the international arena, Interpol has been coordinating a series of operations to raise awareness and to identify the criminals engaging in the online prescription drug trade. Operation Pangea [110] is an annual week-long operation with a large number of participating countries – a total of 111 participated in 2014 – that enables coordinated action across many jurisdictions. Operations Mamba [107], Storm [109], and Cobra [106] are in the same spirit as Pangea, but they have regional focus⁴ and last longer.⁵

Most importantly, the effects of all those operations are short-lived with new storefronts appearing as soon as others are shut down. As we show in Section 3.3, the efforts of enforcement need to be persistent for the effects to be long-term. When applied beyond a specific threshold, proper enforcement can make criminal revenues unattractive. Our thesis aims at enabling persistent enforcement by making it more cost-efficient.

⁴ Eastern Africa, Southeast Asia, and Western Africa respectively.

⁵ On average 1 month. Storm I lasted 5 months.

3.2.4 Health risks

Beyond the legalities pertaining to the operation of online pharmacies, it is important to highlight that the operation of unlicensed pharmacies is not just a bureaucratic problem, but, most importantly, a social one. In this regard, we present two categories from the medical literature pertaining to topics of this thesis. The first category shows the risks of buying prescription drugs from unlicensed online pharmacies, and the second focuses on the socioeconomic characteristics of the customers.

The researchers in [91, 92] show that despite the convenience provided by online pharmacies (e.g. 24 hour availability), they often do not follow due diligence in issuing prescriptions, or they forfeit this requirement altogether. Moreover, by providing access to unapproved drugs, unlicensed online pharmacies put the health of their customers at risk.

Bessell et al. in [18, 19] studied the pharmacological information of prescription and over-the-counter drugs advertised at internationally-based online pharmacies. They found that the information was usually inappropriate, insufficient, or non-existent, making the use of those products unsafe.

As the health risks associated with unlicensed online pharmacies are apparent, we would expect their market penetration to be minimal. However, the high costs of health care and health insurance in the US makes them an unfortunate alternative for low income customers [136]. In addition, unlicensed online pharmacies attract customers of higher socioeconomic status, who can afford health care costs, but they are instead interested in abusing prescription drugs for recreation [139].

Possibly the most striking aspect of unlicensed online pharmacies is that they are not easily distinguishable from their legitimate counterparts. Ivanitskaya et

al. [111] found that undergraduate students, even ones enrolled in health-related studies, cannot easily identify illicit online pharmacies as such. This, in turn, indicates that Internet users of equal or lesser literacy level can easily be put at risk by illicit online pharmacies.

3.3 Social and economic aspects of criminal behavior

In the previous paragraphs we established that online crime is an important problem that can negatively impact our society, and that the legal framework, while prescribing adequate deterrents, their enforcement in the online world is rather problematic. To this end, we conclude our review of the related work, by discussing concepts from criminology associated with the understanding of criminal behavior, and with deterrents that can more effectively target online crime. We make use of those concepts to motivate the nature and structure of criminal disincentives, which we develop in this thesis. In addition, we provide the theoretical foundation of Crime Script Analysis, a framework we use in Chapter 9 to identify and then evaluate situational prevention measures, as disincentives for online crime.

Gary Becker in [17] introduced a choice model capable of explaining the mechanics of criminal behavior. Using the economic formalization of *diseconomies*, he based his analysis on the assumption that criminals are rational, economically motivated agents with a tendency to seek risk. As per his model, the cost of enforcement can be reduced by ameliorating available technologies. *Ceteris paribus*, this in turn translates into reduced occurrence of criminal activities. More importantly, given that police cannot effectively invoke its *sentinel* role in the online

domain,⁶ it is essential to advance the technologies assisting law enforcement in their online apprehension capacity.

Becker also modeled criminal activities as the *supply of offenses* (\mathcal{O}). \mathcal{O} is a function of the probability of conviction per offense p_j , and of the punishment per offense f_j . Given the risk seeking attitude of potential offenders, an increase in the probability of conviction p_j has a disproportionately greater effect in reducing \mathcal{O} than an equal increase in f_j , due to the reduction of the expected utility of a given crime.

In his recent work, Nagin [167] introduces his own choice model, attempting to further explain the process of criminal decision making. One of his main goals is to examine the hierarchy of decisions that lead to the victimization of a target. He argues that the certainty of punishment is more precisely translated as the certainty of apprehension. Based on his model, the certainty of apprehension is a more effective deterrent than the severity of punishment.

These concepts are applicable in this thesis, as we propose better technologies for increasing the possibility of online criminal apprehension. This outcome is capable of effectively reducing online crime.

The work in this domain allows us to assess the effectiveness of current enforcement methods we presented in Section 3.2.3. Their apparent lack of long-lasting effects can be theoretically predicted if they are examined as the online equivalent of “hot spot policing”. Nagin describes this approach as the targeted deployment of police forces, in physical locations where most of the crime takes place. While this method is characterized with immediate reductions in crime, Baveja et al. focusing on “hot spot” crackdowns of illicit drug markets, show the lack of long lasting effects if the enforcement is not consistent after its initial appli-

⁶ Parking a patrol car outside the premises of a hosting provider enabling illicit online criminals would clearly not have the same effect as parking the patrol car outside a convenient store prone to robberies.

cation [15]. Using Caulkins' model on the distribution and consumption of illegal drugs [28], they show that for the effect of the crackdown to persist, there is a need for continuous enforcement beyond a baseline level, otherwise the market will reinstate itself.

In this thesis we approach the effort of increasing the risk of apprehension by targeting the opportunities available to online criminals to engage in their illicit behavior. We base this approach on the fact that criminal behavior is not a mere effect of criminality or anti-social predispositions. On the contrary, crime is a result of deliberate choices by potential offenders, exploiting available opportunities to engage in crime [75].

3.3.1 Modeling offenders' decisions

Clarke discusses the major contribution of the availability of opportunities for committing crime, and the potential of Situational Crime Prevention (SCP) in reducing crime [34]. In his widely accepted perspective, crime happens when (i) there is a vulnerable target, and (ii) there is an appropriate opportunity to victimize the target. SCP prescribes that appropriately reducing the criminal opportunities would consequently reduce crime. For example, the use of safes to protect money, and of tickets in buses instead of collectors, are proven uses of the theory. However, beyond the associated opportunity-reducing prescriptions [35, 45], the theory behind SCP offers little guidance in the form of a structured method for identifying the appropriate, crime-specific preventive measures.

Clarke and Cornish's perspective of *rational choices* in the criminal decision-making process was a first step towards a systematic approach for crime prevention [36]. The authors observe criminal behavior as the "outcome of the offender's rational choices and decisions", and not as an effect of personal or societal dispositions. Given this purposeful, procedural, and rational nature of crime, the au-

thors provide an analytic framework for crime prevention, by placing the focus on the different stages of criminal events. This modeling approach is crime-specific, requiring close attention to the associated situational factors. For example, in examining cases of website compromise, the researcher would have to define separate models when the compromise is intended to manipulate search engine results as opposed to simply deface⁷ the website. In this case, the difference in the underlying motives would necessitate different—but potentially overlapping—countermeasures. Rational choice is essential in such analysis, as it highlights the goal-oriented nature of the criminal activity. Therefore, the distinct stages of a criminal process can be broken down into a series of sub-goals defining the criminal procedure. This approach incorporates the key attributes of crime which is dynamic in form (i.e. “evolutionary, adaptive, and innovatory”), and specific in content (e.g. techniques employed in a given spatio-temporal context).

Crime Script Analysis (CSA) extends the rational choice approach, using the notion of *scripts* from cognitive psychology [43]. It is a systematic framework for breaking down and examining the criminal process, and mapping situational prevention measures to every step of crime commission. In addition, crime scripts are useful in identifying the most significant steps of criminal operations (i.e. concentration points in the context of this thesis) that can be targeted with more intense or persistent measures. In a few words:

“[The] script-theoretic approach offers a way of generating, organizing and systematizing knowledge about the procedural aspects and procedural requirements of crime commission” [43]. “Crime scripts enhance understanding of crime commission, as crime can be seen as a

⁷ Website defacement is the act of gaining unauthorized write-access to the content of a website, altering its content and style, usually as an act of protest or bragging.

process rather than a single event, involving stages in which resources and locations are required and decisions are made” [30].

Crime scripts can operate at different *levels of abstraction*, and Cornish describes the following levels in a decreasing degree of abstraction [43]: (i) universal script, (ii) metascript, (iii) protoscript, (iv) script, and (v) track. This organization makes it possible to link conceptually similar crime scripts at the *track level* into more abstract categories of crime. However, as this is a bottom-up approach, it is essential not to pursue generalization too soon to avoid ignoring specific procedural details. These details are necessary in understanding the choice-structuring properties of particular crimes, and in designing the appropriate, cost-effective situational prevention measures capable of disrupting the crime scripts.

Crime scripts are not necessarily linear processes, and can be organized in scenes, each of which can be further examined as a separate script. In turn, the scenes may be organized in various combinations that represent different crime commission routes (i.e. facets) resulting in the same outcome. In such cases, a *script permutator* can reveal all possible pathways (i.e. tracks) of the crime commission process, and highlights the inherent “dynamic quality of the scripts” [43].

We argue that the concepts of SCP and CSA are highly applicable to the issue of online crime. In addition, our effort to advance the technologies available for reducing opportunities of online criminals is not only in line with the rise of evidence-based policing [201], but it is also a proven concept in the context of computer crime [246].

Measuring and analyzing search-redirection attacks in the illicit online prescription drug trade

The case of prescription drugs has immense public policy implications. By enabling access to prescription drugs without a valid prescription, and without proper health assessment by a medical doctor, consumers are essentially allowed to self-medicate. This practice is a dangerous one as it can lead to severe health issues.

In this chapter we investigate the manipulation of web search results to promote the unauthorized sale of prescription drugs [128]. We focus on a particularly pernicious variant of search-engine manipulation involving compromised web servers—which we term *search-redirection attacks*—which miscreants then use to dynamically redirect traffic to different pharmacies based upon the particular search terms issued by the consumer. We constructed a representative list of 218 drug related queries, and automatically gathered the search results on a daily basis over nine months in 2010-2011. The work presented in this chapter contributes

to the understanding of online crime and search engine manipulation—research questions 1¹ and 5²—in several ways.

First, we collected search results over a nine-month interval (April 2010 – February 2011). The data comprises daily returns from April 12, 2010–October 21, 2010, complemented by an additional 10 weeks of data from November 15th 2010–February 1st 2011. Combining both datasets, we gathered about 185,000 different Uniform Resource Identifiers (URIs)—pharmacies, benign and compromised sites—of which around 63,000 were infected. We describe our measurement infrastructure and methodology in details in Section 4.2, and discuss the search results in Section 4.3.

Second, we show that a quarter of the top 10 search results actively redirect from compromised websites to online pharmacies at any given time. We show infected websites are very slowly remedied: the median infection lasts 46 days, and 16% of all websites have remained infected throughout the study. Further, websites with high reputation (e.g., high PageRank) remain infected and appear in the search results much longer than others.

Third, we provide concrete evidence of the existence of large, connected, advertising “affiliate” networks, funneling traffic to over 90% of the unlicensed online pharmacies we encountered. Search-redirection attacks play a key role in diverting traffic to questionable retail operations at the expense of legitimate alternatives.

Fourth, we analyze whether sites involved in the pharmaceutical trade are involved in other forms of suspicious retail activities, in other security attacks (e.g., serving malware-infested pages), or in spam email campaigns. While we find occasional evidence of other nefarious activities, many of the pharmacies we inspect

¹ Are there any structural characteristics in the illicit online prescription drug trade. . .

² Is it possible to disrupt online criminal networks by targeting critical components. . .

appear to have moved away from email spam-based advertising. We discuss infection characteristics, affiliate networks, and relationship with other attacks in Section 4.4.

Fifth, we derive a rough estimate of the conversion rates achieved by search-redirecting attacks, and show they are considerably higher than those observed for spam campaigns. We present this analysis in Section 4.5.

Finally, we conclude in Section 4.6 where we also describe our initial work in tracking the fraudulent promotion of other types of goods using a similar technique of abusive advertising, in addition to a set of mitigation strategies targeting these illicit operations. However, the core of these analyses is presented in Chapters 6 and 9 respectively.

4.1 Background

Prescription drugs sold illicitly on the Internet arguably constitute the most dangerous online criminal activity. While resale of counterfeit luxury goods or software are obvious frauds, counterfeit medicines actually endanger public safety. Independent testing has indeed revealed that the drugs often include the active ingredient, but in incorrect and potentially dangerous dosages [77, 247].

In the wake of the death of a teenager, the US Congress passed in 2008 the Ryan Haight Online Pharmacy Consumer Protection Act, rendering it illegal under federal law to “deliver, distribute, or dispense a controlled substance by means of the Internet” without an authorized prescription, or “to aid and abet such activity” [223]. Yet, illicit sales have continued to thrive since the law has taken effect. In response, the White House in 2011 [96] helped form the Center for Safe Internet Pharmacies (CSIP) [51], a non-profit organization consisted of registrars, technol-

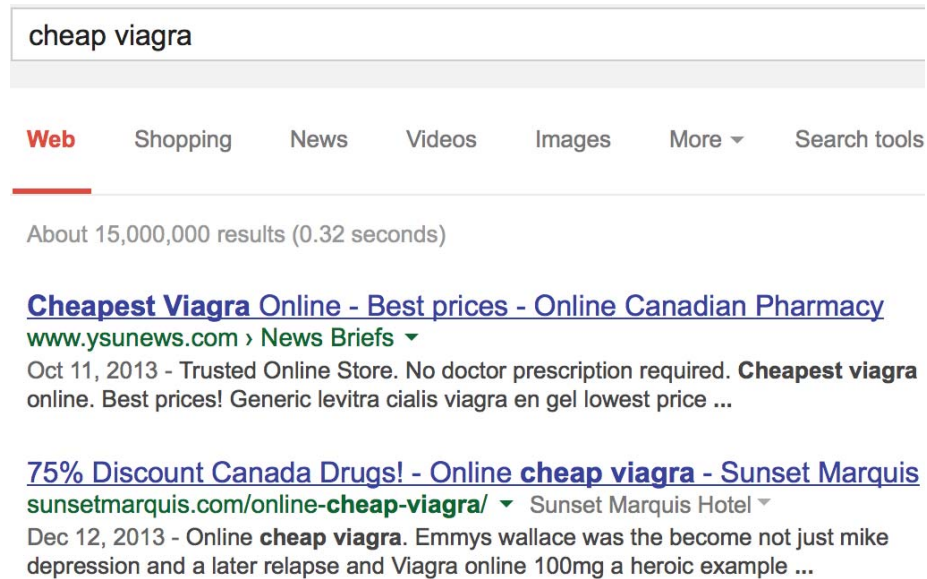


FIGURE 4.1: **Example of search-engine poisoning.** The first two results returned here are sites that have been compromised to advertise unlicensed pharmacies.

ogy companies and payment processors to counter the proliferation of unlicensed pharmacies.

Suspicious online retail operations have, for a long time, primarily resorted to email spam to advertise their products. However, the low conversion rates (realized sales over emails sent) associated with email spam [116] has led miscreants to adopt new tactics. Search-engine manipulation [242], in particular, has become widely used to advertise products. The basic idea of search-engine manipulation is to inflate the position at which a specific retailer's site appears in search results by artificially linking it from many websites. Conversion rates are believed to be much higher than for spam, since the advertised site has at least a degree of relevance to the query issued.

4.1.1 Search-redirection attacks

Figure 4.1 illustrates the attack. In response to the query “*cialis without prescription*”, the top eight results include five .edu sites, one .com site with a seemingly unrelated domain name, and two online pharmacies. At first glance, the .edu and one of the .com sites have absolutely nothing to do with the sale of prescription drugs. However, clicking on some of these links, including the top search result framed in Figure 4.1, takes the visitor not to the requested site, but to an online pharmacy store. This is an example in which the top two results obtained for the query “cheap viagra” are compromised websites. The top result is the website of a news center affiliated with a university. The site was compromised to include a pharmacy store front in a hidden directory: clicking on any of the links in that storefront sends the prospective customer to pillsforyou24.com, a known rogue Internet pharmacy [124].

The attack works as follows. The attacker first identifies high-visibility websites that are also vulnerable to code injection attacks.³ Popular targets include outdated versions of WordPress [248], phpBB [181], or any other vulnerable blogging or wiki software. The code injected on the server intercepts all incoming Hypertext Transfer Protocol (HTTP) requests to the compromised page and responds differently depending on the type of request.

Requests originating from search-engine crawlers, as identified by the *User-Agent* parameter of the HTTP request, return a mix of the compromised site’s original content plus numerous links to websites promoted by the attacker (e.g., other compromised sites, online stores). This technique, “link stuffing,” has been observed for several years [176] in non-compromised websites.

³ We defer the study of the specific exploits to future work. Our focus in this Chapter is the outcome of the attack, not the attack itself.

Requests originating from pages of search results, for queries deemed relevant to what the attacker wants to promote, are redirected to a website of the attacker's choosing. The compromised web server automatically identifies these requests based on the *Referrer* field that HTTP requests carry [68]. The *Referrer* actually contains the complete URI that triggered the request. For instance, in Figure 4.1, when clicking on any of the links, the *Referrer* field is set to `http://www.google.com/search?q=cialis+without+prescription`. Upon detecting the pharmacy-related query, the server sends an HTTP redirect with status code 302 (Found) [68], along with a *location* field containing the desired pharmacy website or intermediary. The upshot is that the end user unknowingly visits a series of websites culminating in a fake pharmacy without ever spending time at the original site appearing in the search results. A similar technique has been extensively used to distribute malware [185], while web spammers have also used the technique to hide the true nature of their sites from investigators [174].

All other requests, including typing the URI directly into a browser, return the original content of the website. Therefore, website operators cannot readily discern that their website has been compromised. As we will show in Section 4.4, as a result of this “cloaking” mechanism, some of the victim sites remain infected for a long time.

Three classes of websites are involved in search-redirection attacks. (i) **Source infections** are innocent websites that have been compromised and reprogrammed with the behavior just described; (ii) **traffic brokers** are intermediary websites that receive traffic from source infections; and (iii) retailers (here, **pharmacies**) are destination websites that receive traffic from traffic brokers.

It is not immediately obvious who the victim is in search-redirection attacks. Unlike in drive-by-downloads [185], end users issuing pharmacy searches are not necessarily victims, since they are actually often seeking to illegally procure drugs

online. In fact, here, search engines do provide results relevant to what users are looking for, regardless of the legality of the products considered. However, users may also become victims if they receive inaccurately dosed medicine or dangerous combinations that can cause physical harm or death. The operators of source infections are victims, but only marginally so, since they are not directly harmed by redirecting traffic to pharmacies. Pharmaceutical companies are victims in that they may lose out on legitimate sales. The greatest harm is a societal one, because laws designed to protect consumers are being openly flouted.

4.2 Measurement methodology

We now explain the methodology used to identify search-redirectation attacks that promote online pharmacies. We first describe the infrastructure for data collection, then how search queries are selected, and finally how the search results are classified.

4.2.1 Infrastructure overview

The measurement infrastructure comprises two distinct components: a search-engine agent that sends drug-related queries and a crawler that checks for behavior associated with search-redirectation attacks.⁴

The search-engine agent uses the Google Web Search Application Programming Interface (API) [83] to automatically retrieve the top 64 search results to selected queries. From manually inspecting some compromised websites, we found that search-redirectation attacks frequently also work on other search engines. Every 24 hours, the search-engine agent automatically sends 218 different queries for prescription drug-related terms (e.g., “*cialis without prescription*”) and stores

⁴ All results gathered by the crawler are stored in a *mySQL* database, available from <http://arima.cylab.cmu.edu/rx.sql.gz>.

all 13,952 ($= 64 \times 218$) URIs returned. We explain how we selected the corpus of 218 queries in Section 4.2.2.

The crawler module then contacts each URI collected by the search-engine agent and checks for HTTP 302 redirects mentioned in Section 4.1.1. The crawler emulates typical web-search activity by setting the *User-Agent* and *Referrer* terms appropriately in the HTTP headers. Initial tests revealed that some source infections had been programmed to block repeated requests from a single Internet Protocol (IP) address. Consequently, all crawler requests are tunneled through the Tor network [58] to circumvent the blocking.

4.2.2 Query selection

Selecting appropriate queries to feed the search-engine agent is critical for obtaining suitable quality, coverage and representativeness in the results. We began by issuing a single seed query, “*no prescription vicodin*,” chosen for the many source infections it returned at the time (March 3, 2010). We then browsed the top infected results posing as a search engine crawler. As described in Section 4.1.1, infected servers present different results to search-engine crawlers. The pages include a mixture of the site’s original content and a number of drug-related search phrases designed to make the website attractive to search engines for these queries. The inserted phrases typically linked to other websites the attacker wishes to promote, in our case other online pharmacies.

We compiled a list of promoted search phrases by visiting the linked pharmacies posing as a search-engine crawler and noting the phrases observed. Many phrases were either identical or contained only minor differences, such as spelling variations on drug names. We reduced the list to a corpus of 48 unique queries, representative of all drugs advertised in this first step.

We then repeated this process for all 48 search phrases, gathering results daily from March 3, 2010 through April 11, 2010. The 48-query search subsequently led us to 371 source infections. We again browsed each of these source infections posing as a search engine crawler, and gathered a few thousand search phrases linked from the infected websites. After again sorting through the duplicates, we got a corpus of $Q = 218$ unique search queries.

The risk of starting from a single seed is to only identify a single unrepresentative campaign. Hence, we ran a validation experiment to ensure that our selected queries had satisfactory coverage. We obtained a six-month sample of spam email (collected at a different time period, late 2009) gathered in a different context [189]. We ran SpamAssassin [12] on this spam corpus, to classify each spam as either pharmacy-related or otherwise. We then extracted all drug names encountered in the pharmacy-related spam, and observed that they defined a subset of the drug names present in our search queries. This gave us confidence that the query corpus was quite complete.

4.2.3 Additional query-sample validation

We collected two additional sets of search queries to further validate the adequate coverage of our main query corpus of 218 terms. First, we derived a query set from an exhaustive list of 9,000 prescription drugs provided by the FDA [235]. We ran a single query for each drug in the list—in the form of “no prescription [drug name]”—and collected the first 64 results per query. We executed the 9,000 queries over five days in August 2010. About 2,500 of the queries returned no search results. Of the queries that returned results, we observed redirection in at least one of the search results for 4,350 terms.

For the second list, we inspected summaries of server logs for 169 infected websites to identify drug-related search terms that redirected to pharmacies. We

obtained this information from infected web servers running Webalizer,⁵ which creates monthly reports, based on HTTP logs, of how many visitors a website receives, the most popular pages on the website, and so forth. It is not uncommon to leave these reports “world-readable” in a standard location on the server, which means that anyone can inspect their contents.

In August 2010, we checked 3,806 infected websites for Webalizer logs, finding it accessible on 169 websites. We recorded all available data, which usually included monthly reports of web activity. One of the individual sub-reports that Webalizer creates is a list of search terms that have been used to locate the site. Not all Webalizer reports list referrer terms, but we found 83 websites that did include drug names in the referrer terms for one or more months of the log reports. Since we identified the infected servers running Webalizer by inspecting results of the 218 queries from our main corpus, it is unsurprising that 98 of these terms appeared in the logs. However, the logs also contained an additional 1,179 search queries with drug terms. We use these additional search terms as an *extra queries list* to compare against the main corpus.

We collected the top 64 results for the extra queries list daily between October 20 and 31, 2010. When comparing these results to our main query corpus, we examine only the results obtained during this time period, resulting in a significantly smaller number of results than for our complete nine-month collection.

We compare our main list to the additional lists in three ways. First, we compare the classification of search results for differences in the types of results obtained. Second, we compare the distribution of Top Level Domain (TLD) and PageRank for source infections obtained for both samples. Third, we compute

⁵ Webalizer is a popular program for summarizing web server log files <http://www.mrunix.net/webalizer/>

Table 4.1: Comparing different lists of search terms to the main list used in the Chapter. All numbers are percentages.

	FDA drug list				Extra query list			
	Drug list		Main list		Extra list		Main list	
	URIs	domains	URIs	domains	URIs	domains	URIs	domains
<i>Search result classification</i>								
Source infections	24.7	4.0	43.7	22.4	35.6	14.0	49.3	27.9
Health resources	12.7	7.4	2.8	3.5	4.9	4.2	2.4	3.0
Licensed pharm.	0.5	0.1	0.03	0.07	0.1	0.1	0.02	0.05
Unlicensed pharm.	6.7	6.9	8.2	13.6	6.1	11.6	6.5	12.0
Blog/forum spam	25.4	23.7	18.6	17.8	26.3	22.7	17.8	17.7
Uncategorized	30.1	57.9	26.7	42.7	27.2	46.9	24.0	39.4
<i>Source infection TLD breakdown</i>								
.com		60.0		56.9		56.3		54.6
.org		13.8		17.0		15.4		18.0
.edu		5.6		8.9		6.2		9.3
.net		6.1		5.6		5.6		4.6
other		14.3		11.5		16.5		13.5
<i>Source infection PageRank breakdown</i>								
PR $0 \leq 3$		47.2		35.0		47.5		41.9
PR $3 \leq 6$		41.4		51.3		44.2		46.3
PR ≥ 7		11.4		13.7		8.3		11.8

the intersection between the domains obtained by both sets of queries for source infections, redirects and pharmacies.

Table 4.1 compares the FDA drugs and extra queries lists to the main list. The breakdown of search results for both samples is slightly different from what we obtained using the main queries. For instance, only 25% of the URIs in the FDA results are infections, compared to 44% for the main list during the same time period. 13% of the results in the FDA drug list point to legitimate health resources, compared to only 3% of the main sample. This is not surprising, given that the drug list often included many drugs that are not popular choices for sales by online pharmacies. Unlicensed pharmacies appear slightly less often in the drugs sample (6% vs. 8%), while blog and forum spam is more prevalent (25% to 19%).

The extra queries list follows the FDA list in some ways, e.g., more blog infections and fewer source infections than results from the corresponding main list. On the other hand, the URI breakdown in health resources is much closer (4.9% vs. 2.4%). In all samples, the number of results that point to legitimate pharmacies is very small, though admittedly biggest in the drugs sample (0.5% vs. 0.1% for the extra queries).

We next take a closer look at the characteristics of the source infections themselves. The TLD breakdown is roughly similar, with a few exceptions. `.com` is found slightly more often in the FDA drugs and extra queries results, while `.org` and `.edu` appear a bit more often in the results for the main sample. The drugs and extra queries list tend to have slightly lower PageRank than the results from the main sample, but the difference is slight.

4.2.4 Query corpus characteristics

The entire set of queries Q can itself be partitioned according to the presumed intention of the person issuing the query. For instance, in the pharmaceutical realm, queries such as “prozac side effects” appear to be seeking legitimate information—we term such queries as *Benign* queries. The set of all Benign queries is denoted by B (resp. $B(t)$ at time t). On the other hand, certain queries may denote questionable intentions. For example, somebody searching for “vicodin without a prescription” would certainly expect a number of search results to link to contraband sites. We call such queries representing potentially *illicit* intent as such, and denote them as being in a set I (resp. $I(t)$ at time t). Finally, a number of queries, e.g., “buy ativan online,” may not easily be classified as exhibiting illicit or benign intent. We refer to these queries as being in the *Gray* set, G (resp. $G(t)$ at time t).

Table 4.2: Intention-based classification of the 218 queries in the drug query corpus (Q).

Type of query	Count	%
Illicit ($ I $)	26	22%
Benign ($ B $)	75	34.4%
Gray ($ G $)	117	53.6%
Total (Q)	218	100%

Table 4.2 breaks down the query corpus Q between the illicit, benign, and gray sets I , B , and G . Overall, the queries clearly associated with illicit intentions are the minority of the total queries (22%), while the majority is placed in the gray category. This bias of the query corpus towards informative types of queries (i.e. gray and benign – 88% of total), rather than queries exhibiting illicit intent, suggests that the extent and effects of the search-redirection attack mainly affects individuals with non-illicit intentions.

4.2.5 Search-result classification

We attempt to classify all results obtained by the search-engine agent. Each query returns a mix of legitimate results (e.g., health information websites) and abusive results (e.g., spammed blog comments and forum postings advertising online pharmacies). We seek to distinguish between these different types of activity to better understand the impact of search-redirection attacks may have on legitimate pharmacies and other forms of abuse. We assign each result into one of the following categories: (i) search-redirection attacks, (ii) health resources, (iii) legitimate online pharmacies, (iv) unlicensed pharmacies, (v) blog or forum spam, and (vi) uncategorized.

We mark websites as participating in search-redirection attacks by observing an HTTP redirect to a *different* website. Legitimate websites regularly use HTTP

redirects, but it is less common to redirect to entirely different websites immediately upon arrival from a search engine. Every time the crawler encounters a redirect, it recursively follows and stores the intermediate URIs and IP addresses encountered in the database. These redirection chains are used to infer relationships between source infections and pharmacies in Section 4.4.3.

We performed two robustness checks to assess the suitability of classifying all external redirects as attacks. First, we found known drug terms in at least one redirect URI for 63% of source websites. Second, we found that 86% of redirecting websites point to the same website as 10 other redirecting websites. Finally, 93% of redirecting websites exhibit at least one of these behaviors, suggesting that the vast majority of redirecting websites are infected. In fact, we expect that most of the remaining 7% are also infected, but some attackers use unique websites for redirection. Thus, treating all external redirects as malicious appears reasonable in this study.

Health resources are websites such as `webmd.com` that describe characteristics of a drug. We used the Alexa Web Information Service API [6], which is based on the Open Directory [11] to determine each website category.

We distinguish between legitimate and unlicensed online pharmacies by using a list of registered pharmacies obtained from the non-profit organization Legitscript [124]. Legitscript, at the time, maintained a whitelist of 324 confirmed legitimate online pharmacies, which require a verified doctor's prescription and sell genuine drugs. Unlicensed pharmacies are websites which do not appear in Legitscript's whitelist, and whose domain name contains drug names or words such as "pill," "tabs," or "prescription." Legitscript's list is likely incomplete, so we may incorrectly categorize some collected legitimate pharmacies as unlicensed, because they have not been validated by Legitscript.

Finally, blog and forum spam captures the frequent occurrence where websites that allow user-generated content are abused by users posting drug advertisements. We classify these websites based only on the URI structure, since collecting and storing the pages referenced by URI is cost-prohibitive. We first check the URI subdomain and path for common terms indicating user-contributed content, such as “blog,” “viewmember” or “profile.” We also check any remaining URIs for drug terms appearing in the subdomain and path. While these might in fact be compromised websites that have been loaded with content, upon manual inspection the activity appears consistent with user-generated content abuse.

4.3 Empirical analysis of search results

We begin our measurement analysis by examining the search results collected by the crawler. The objective here is to understand how prevalent search-redirection attacks are, in both absolute terms and relative to legitimate sources and other forms of abuse.

4.3.1 *Breakdown of search results*

Table 4.3 presents a breakdown of all search results obtained during the six months of primary data collection. 137,354 distinct URIs correspond to 23,042 different domains. We observed 44,503 of these URIs to be compromised websites (*source infections*) actively redirecting to unlicensed pharmacies, 32% of the total. These corresponded to 4,652 unique infected source domains. We examine the redirection chains in more detail in Section 4.4.3.

An additional 29,406 URIs did not exhibit redirection even though they shared domains with URIs where we did observe redirection. There are several plausible explanations for why only some URIs on a domain will redirect to pharmacies.

Table 4.3: Classification of all search results (4–10/2010).

	URIs		Domains	
	#	%	#	%
Source infections	73,909	53.8	4,652	20.2
<i>Active</i>	44,503	32.4	2,907	12.6
<i>Inactive</i>	29,406	21.4	1,745	7.6
Health resources	1,817	1.3	422	1.8
Pharmacies	4,348	3.2	2,138	9.3
<i>Legitimate</i>	12	0.01	9	0.04
<i>Unlicensed</i>	4,336	3.2	2,129	9.2
Blog/forum spam	41,335	30.1	8,064	34.9
Uncategorized	15,945	11.6	7,766	33.7
Total	137,354	100.0	23,042	100.0

First, websites may continue to appear in the search results even after they have been remediated and stop redirecting to pharmacies. In Figure 4.1, the third link to appear in the search engine results has been disinfected, but the search engine is not yet aware of that. For 17% of the domains with inactive redirection links, the inactive links only appear in the search results after all the active redirects have stopped appearing.

However, for the remaining 83% of domains, the inactive links are interspersed among the URIs which actively redirect. In this case, we expect that the miscreants' search engine optimization has failed, incorrectly promoting pages on the infected website that do not redirect to pharmacies.

By comparison, very few search results led to legitimate resources. 1,817 URIs, 1.3% of the total, pointed to websites offering health resources. Even more striking, only *nine* legitimate pharmacy websites, or 0.04% of the total, appeared in the search results. By contrast, 2,129 unlicensed pharmacies appeared directly in the search results. 30% of the results pointed to legitimate websites where miscreants had posted spam advertisements to online pharmacies. In contrast to




FIGURE 4.2: Distribution of different classes of results according to the position in the search results.

the infected websites, these results require a user to click on the link to arrive at the pharmacy. It is also likely that many of these results were not intended for end users to visit; instead, they could be used to promote infected websites higher in the search results.

4.3.2 Variation in search position

Merely appearing in search results is not enough to ensure success for miscreants perpetrating search-redirect attacks. Appearing towards the top of the search results is also essential [114]. To that end, we collected data for an additional 10 weeks from November 15th 2010 to February 1st 2011 where we recorded the position of each URI in the search results.

Figure 4.2 presents the findings. Around one third of the time, search-redirect attacks appeared in the first position of the search results. 17% of

the results were actively redirecting at the time they were observed in the first position. Blog and forum spam appeared in the top spot in 30% of results, while unlicensed pharmacies accounted for 22% and legitimate health resources just 5%.

The distribution of results remains fairly consistent across all 64 positions. Active search-redirectation attacks increase their proportion slightly as the rankings fall, rising to 26% in positions 6–10. The share of unlicensed pharmacies falls considerably after the first position, from 22% to 14% for positions 2–10. Overall, it is striking how consistently all types of manipulation have crowded out legitimate health resources across all search positions.

4.3.3 Turnover in search results

Web search results can be very dynamic, even without an adversary trying to manipulate the outcome. We count the number of unique domains we observe in each day's sample for the categories outlined in Section 4.2. Figure 4.3 shows the average daily count for two-week periods from May 2010 to February 2011, covering both sample periods. The number of unlicensed pharmacies and health resources remains fairly constant over time, whereas the number of blogs and forums with pharmaceutical postings fell by almost half between May and February. Notably, the number of source infections steadily increased from 580 per day in early May to 895 by late January, a 50% increase in daily activity.

4.3.4 Variation in search queries

As part of its AdWords program, Google offers a free service called Traffic Estimator to check the estimated number of global monthly searches for any phrase.⁶ We fetched the results for the 218 search terms we regularly checked; in total,

⁶ <https://adwords.google.com/select/TrafficEstimatorSandbox>




FIGURE 4.3: Change in the average domains observed each day for different classes of search results over time.

over 2.4 million searches each month are made using these terms. This gives us a good first approximation of the relative popularity of web searches for finding drugs through online pharmacies. Some terms are searched for very frequently (as much as 246,000 times per month), while other terms are only searched for very occasionally.

We now explore whether the quality of search results vary according to the query's popularity. We might expect that less-popular search terms are easier to manipulate, but also that there could be more competition to manipulate the results of popular queries.

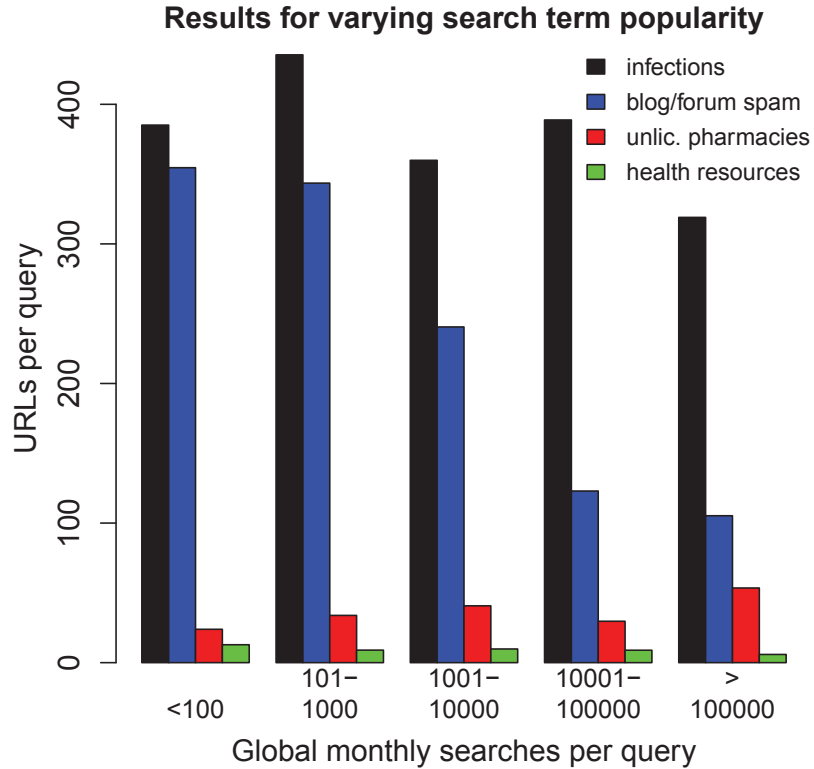


FIGURE 4.4: Search-redirection attacks appear in many queries; health resources and blog spam appear less often in popular queries.

Figure 4.4 plots the average number of unique URIs observed per query for each category. For unpopular searches, with less than 100 global monthly searches, search-redirection attacks and blog spam appear with similar frequency. However, as the popularity of the search term increases, search-redirection attacks continue to appear in the search results with roughly the same regularity, while the blog and forum spam drops considerably (from 355 URIs per query to 105).

While occurring on a smaller scale, the trends of unlicensed pharmacies and legitimate health resources are also noteworthy. Health resources become increasingly crowded out by illicit websites as queries become more popular. For unpopular queries (< 100 global monthly searches), 13 health URIs appear. But for queries with more than 100,000 results, the number of results falls by more

than half to 6. For unlicensed pharmacies, the trends are opposite. On less popular terms, the pharmacies appear less often (24 times on average). For the most popular terms, by contrast, 54 URIs point directly to unlicensed pharmacies. Taken together, these results suggest that the more sophisticated miscreants do a good job of targeting their websites to high-impact results.

4.4 Empirical analysis of search-redirection attacks

We now focus our attention on the structure and dynamics of search-redirection attacks themselves. We present evidence that certain types of websites are disproportionately targeted for compromise, that a few such websites appear most prominently in the search results, and that the chains of redirections from source infections to pharmacies betray a few clusters of concentrated criminality.

4.4.1 Concentration in search-redirection attack sources

We identified 7,298 source websites from both data sets that had been infected to take part in search-redirection attacks—4,652 websites in the primary 6-month data set and 3,686 in the 10-week follow-up study (1,130 sites are present in both datasets). We now define a measure of the relative impact of these infected websites in order to better understand how they are used by attackers.

$$\mathcal{I}(\text{domain}) = \sum_{q \in \text{queries}} \sum_{d \in \text{days}} u_{qd} * 0.5^{\frac{r_{qd}-1}{10}}$$

where

u_{qd} : 1 if domain in results of query q on

day d & actively redirects to pharmacy

u_{qd} : 0 otherwise

r_{qd} : domain's position (1..64) in search results

Table 4.4: TLD breakdown of source infections.

	.com	.org	.edu	.net	other
% global Internet	45%	4%	< 3%	6%	42%
% infected sources	55%	16%	6%	6%	17%
% inf. source impact	30%	24%	35%	2%	10%

The goal of the impact measure \mathcal{I} is to distill the many observations of an infected domain into a comparable scalar value. Essentially, we add up the number of times a domain appears, while compensating for the relative ranking of the search results. Intuitively, when a domain appears as the top result it is much more likely to be utilized than if it appeared on page four of the results. The heuristic we use normalizes the top result to 1, and discounts the weighting by half as the position drops by 10. This corresponds to regarding results appearing on page one as twice as valuable as those on page two, which are twice as valuable as those on page three, and so on.

Some infected domains appeared in the search results much more frequently and in more prominent positions than others. The domain with the greatest impact—`unm.edu`—accounted for 2% of the total impact of all infected domains. Figure 4.5 plots using a logarithmic x -axis the ordered distribution of the impact measure \mathcal{I} for source domains. The top 1% of source domains account for 32% of all impact, while the top 10% account for 81% of impact. This indicates that a small, concentrated number of infected websites account for most of the most visible redirections to online pharmacies.

We also examined how the prevalence and impact of source infections varied according to TLD. The top row in Table 4.4 shows the relative prevalence of different TLDs on the Internet [237]. The second row shows the occurrence of infections by TLD. The most affected TLD, with 55% of infected results, is `.com`, followed by `.org` (16%), `.edu` (6%) and `.net` (6%). These four TLDs account for





FIGURE 4.5: Rank-order CDF of domain impact reveals high concentration in search-redirection attacks.

83% of all infections, with the remaining 17% spread across 159 TLDs. We also observed 25 infected `.gov` websites and 22 governmental websites from other countries.

One striking conclusion from comparing these figures is how more ‘reputable’ domains, such as `.com` (55% of infections vs. 45% of registrations), `.org` (16% vs. 4%) and `.edu` (6% vs. < 3%), are infected than others. This is in contrast to other research, which has identified country-specific TLDs as sources of greater risk [146].

Furthermore, some TLDs are used more frequently in search-redirection attacks than others. While `.edu` domains constitute only 6% of source infections, they account for 35% of aggregate impact through redirections to pharmacy websites. Domains in `.com`, by contrast, account for more than half of all source domains but 30% of all impact. We next explore how infection durations vary across domains, in part with respect to TLD.



Cox-proportional hazard model

$$h(t) = \exp(\alpha + \text{PageRank}x_1 + \text{TLD}x_2)$$

	coef.	exp(coef.)	Std. Err.)	Significance
PageRank	-0.085	0.92	0.0098	$p < 0.001$
.edu	-0.26	0.77	0.086	$p < 0.001$
.net	0.08	1.1	0.084	
.org	0.055	1.0	0.054	
other TLDs	0.34	1.4	0.053	$p < 0.001$

log-rank test: $Q=158, p < 0.001$

FIGURE 4.6: Survival analysis of search-redirection attacks shows that TLD and PageRank influence infection lifetimes.

4.4.2 Variation in source infection lifetimes

One natural question when measuring the dynamics of attack and defense is how long infections persist. We define the “lifetime” of a source infection as the number of days between the first and last appearance of the domain in the search results while the domain is actively redirecting to pharmacies. Lifetime is a standard metric in the empirical security literature, even if the precise definitions vary by the attacks under study. For example, Moore and Clayton [155] observed that phishing

websites have a median lifetime of 20 hours, while Nazario and Holz [171] found that domains used in fast-flux botnets have a mean lifetime of 18.5 days.

Calculating the lifetime of infected websites is not entirely straightforward, however. First, because we are tracking only the results of 218 search terms, we count as “death” whenever an infected website disappears from the results or stops redirecting, even if it remains infected. This is because we consider the harm to be minimized if the search engine detects manipulation and suppresses the infected results algorithmically. However, to the extent that our search sample is incomplete, we may be overly conservative in claiming a website is no longer infected when it has only disappeared from our results.

The second subtlety in measuring lifetimes is that many websites remain infected at the end of our study, making it impossible to observe when these infections are remediated. Fortunately, this is a standard problem in statistics and can be solved using survival analysis. Websites that remain infected and in the search results at the end of our study are said to be *right-censored*. 1,368 of the 4,652 infected domains (29%) are right-censored.

The survival function $S(t)$ measures the probability that the infection’s lifetime is greater than time t . The survival function is similar to a complementary cumulative distribution function, except that the probabilities must be estimated by taking censored data points into account. We use the standard Kaplan-Meier estimator [119] to calculate the survival function for infection lifetimes, as indicated by the solid black line in the graphs of Figure 4.6. The median lifetime of infected websites is 47 days; this can be seen in the graph by observing where $S(t) = 0.5$. Also noteworthy is that at the maximum time $t = 192$, $S(t) = 0.160$. Empirical survival estimators such as Kaplan-Meier do not extrapolate the survival distribution beyond the longest observed lifetime, which is 192 days in our sample. What we can discern from the data, nonetheless, is that 16% of infected domains were in

the search results throughout the sample period, from April to October. Thus, we know that a significant minority of websites have remained infected for at least six months. Given how hard it is for webmasters to detect compromise, we expect that many of these long-lived infections have actually persisted far longer.

We next examine the characteristics of infected websites that could lead to longer or shorter lifetimes. One possible source of variation to consider is the TLD. Figure 4.6 (left) also includes survival function estimates for each of the four major TLDs, plus all others. Survival functions to the right of the primary black survival graph (e.g., `.edu`) have consistently longer lifetimes, while plots to the left (e.g., other and `.net`) have consistently shorter lifetimes. Infections on `.com` and `.org` appear slightly longer than average, but fall within the 95% confidence interval of the overall survival function.

The median infection duration of `.edu` websites is 113 days, with 33% of `.edu` domains remaining infected throughout the 192-day sample period. By contrast, the less popular TLDs taken together have a median lifetime of just 28 days.

Another factor beyond TLD is also likely at play: the relative reputation of domains. Web domains with higher PageRank are naturally more likely to appear at the top of search results, and so are more likely to persist in the results. Indeed, we observe this in Figure 4.6 (center). Infected websites with PageRank 7 or higher have a median lifetime of 153 days, compared to just 17 days for infections on websites with PageRank 0.


One might expect that `.edu` domains would tend to have higher PageRanks, and so it is natural to wonder whether these graphs indicate the same effect, or two distinct effects. To disentangle the effects of different website characteristics on lifetime, we use a Cox proportional hazard model [50] of the form:

$$h(t) = \exp(\alpha + \text{PageRank}x_1 + \text{TLD}x_2)$$


Note that the dependent variable included in the Cox model is the hazard function $h(t)$. The hazard function $h(t)$ expresses the instantaneous risk of death at time t . Cox proportional hazard models are used on survival data in preference to standard regression models, but the aim is the same as for regression: to measure the effect of different independent factors (in our case, TLD and PageRank) on a dependent variable (in our case, infection lifetime). PageRank is included as a numerical variable valued from 0 to 9, while TLD is encoded as a five-part categorical variable using deviation coding. (Deviation coding is used to measure each categories' deviation in lifetime from the overall mean value, rather than deviations across categories.) The results are presented in the table in Figure 4.6. PageRank is significantly correlated with lifetimes—lower PageRank matches shorter lifetimes while higher PageRank is associated with longer lifetimes. Separately, `.edu` domains are correlated with longer lifetimes and other TLDs to shorter lifetimes.

Coefficients in Cox models cannot be interpreted quite as easily as in standard linear regression; exponents (column 3 in the table) offer the clearest interpretation. $\exp(\text{PageRank}) = 0.92$ indicates that each one-point increase in the site's PageRank decreases the hazard rate by 8%. Decreases in the hazard leads to longer lifetimes. Meanwhile, $\exp(\text{.edu}) = 0.77$ indicates that the presence of a `.edu` domain, holding the PageRank constant, decreases the hazard rate by 23%. In contrast, the presence of any TLD besides `.com`, `.edu`, `.net` and `.org` increases the hazard rate by 40%.


Therefore, we can conclude from the model that *both* PageRank and TLD matter. Even lower-ranked university websites and high-rank non-university websites are being effectively targeted by attackers redirected traffic to pharmacy websites.



(a) **Structure of the giant component G_0 that links 96% of infected domains.** Links between vertices are based on observed traffic redirection chains. Vertices are colored according to their community.



(b) **CDF of nodes in the giant component belonging to different communities.** The largest 7 (out of 73) communities comprise over half the nodes.



(c) **Scatter plot of in- and out-degree of nodes in the giant component.** (Log-log scale, where 0 is technically represented as 0.1.)

FIGURE 4.7: Network analysis of redirection chains reveals community structure in search-redirection attacks.

4.4.3 Characterizing the unlicensed online pharmacy network

We now extend consideration beyond the websites directly appearing in search results to the intermediate and destination websites where traffic is driven in search-redirection attacks. We use the data to identify connections between a priori unrelated online pharmacies.

We construct a directed graph $G = (V, E)$ as follows. We gather all URIs in our database that are part of a redirection chain (source infection, traffic broker, unlicensed online pharmacy) and assign each second-level domain to a node $v \in V$. We then create edges between nodes whenever domains redirect to each other. Suppose for instance that `http://www.example.com/blog` is infected and redirects to `http://1337.attacker.test` which in turns redirects to `http://www32.cheaprx4u.test`. We then create three nodes $v_1 = \text{example.com}$, $v_2 = \text{attacker.test}$ and $v_3 = \text{cheaprx4u.test}$, and two edges, $v_1 \rightarrow v_2$ and $v_2 \rightarrow v_3$. Now, if `http://hax0r.attacker.test` is also present in the database, and redirects to `http://www.otherrx.test`, we create a node $v_4 = \text{otherrx.test}$ and establish an edge $v_2 \rightarrow v_4$.

In the graph G so built, online pharmacies are usually leaf nodes with a positive in-degree and out-degree zero.⁷ Compromised websites feeding traffic to pharmacies are generally represented as sources, with an in-degree of zero and a positive out-degree. Traffic brokers, which act as intermediaries between compromised websites and online pharmacies have positive in- and out-degrees.

The resulting graph G for our entire database consists of 34 connected subgraphs containing more than two nodes. The largest connected component G_0 contains 96% of all infected domains, 90% of the redirection domains and 92% of the pharmacy domains collected throughout the six-month collection period.

⁷ Manually checking the data, we find a few pharmacies have an out-degree of 1, and redirect to other pharmacies.

In other words, we have evidence that most unlicensed pharmacies are connected by redirection chains. While this does not necessarily indicate that a single criminal organization is behind the entire online pharmacy network, this does tell us that most unlicensed pharmacies in our measurements are obtaining traffic from a large interconnected network of advertising affiliates. Undercover investigations have confirmed the existence of such affiliate networks and provided anecdotal evidence on their operations [194], but they have not precisely quantified their influence. These affiliate networks consist of a loosely organized set of independent advertising entities that feed traffic to their customers (e.g., online retailers) in exchange for a commission on any resulting sales.

Communities and affiliated campaigns.. To uncover affiliate networks, we locate *communities* within G_0 , i.e., sets of vertices closely interconnected with each other and only loosely connected to the rest of the graph. Here, each community represents a set of domains in close relationship with each other, possibly part of the same business operation, or in the same manipulation campaigns. Several algorithms have recently been proposed for community detection, e.g., [178, 187, 191]. We use the spin-glass model proposed by Reichardt and Bornhold [191] (with $q = 500$, $\gamma = 1$) because its stochastic nature allows it to complete quickly even on large graphs like ours, and because it works on directed graphs.

In Figure 4.7(a), we plot a visual representation of G_0 . Different colors denote different communities. The community detection algorithm identifies a total of 73 distinct communities. Most larger communities can be observed in the dense clusters of nodes in the center of the figure, and it appears that less than a dozen of communities play a significant role. More precisely, we plot in Figure 4.7(b) the Cumulative Distribution Function (CDF) of nodes in G_0 as a function of the

number of communities considered. The graph shows that the seven largest communities account for more than half of the nodes in the graph, and that about two thirds of the nodes belong to one of the top twelve communities. In other words, a relatively small number of loosely interconnected, possibly distinct, operations is responsible for most attacks.

Manual inspection confirms these insights. For instance, the third largest community (400 nodes) consists of compromised hosts primarily sending traffic to a single redirector, which itself redirects to a single unlicensed pharmacy (`securetabs.net`).

Figure 4.7(c) is a scatter-plot of the in- and out-degree of each node in G_0 . A vast majority of nodes are source infections (null in-degree, high out-degree, i.e., points along the y -axis) or unlicensed pharmacies (low out-degree, high in-degree, i.e., along the x -axis). Traffic brokers, with non-zero in- and out- degrees are comparatively rare. We identify 314 traffic brokers in G_0 , out of which only 127 have both an in- and an out-degree greater than two. 103 of these 127 traffic brokers (80%) are *cut vertices* for G_0 . That is, removing any of these 103 traffic brokers would partition G_0 .

4.4.4 Attack websites in blacklists

The websites we have identified here have either been compromised (in the case of source infections) or have taken advantage of compromised servers (in the case of traffic brokers and pharmacies). Given such insalubrious circumstances, we wondered if any of the third party blacklists dedicated to identifying Internet wickedness might also have noticed these same websites. To that end, we consulted three different sources: Google's Safe Browsing API, which identifies web-based malware; the `zen.spamhaus.org` blacklist, which identifies email spam

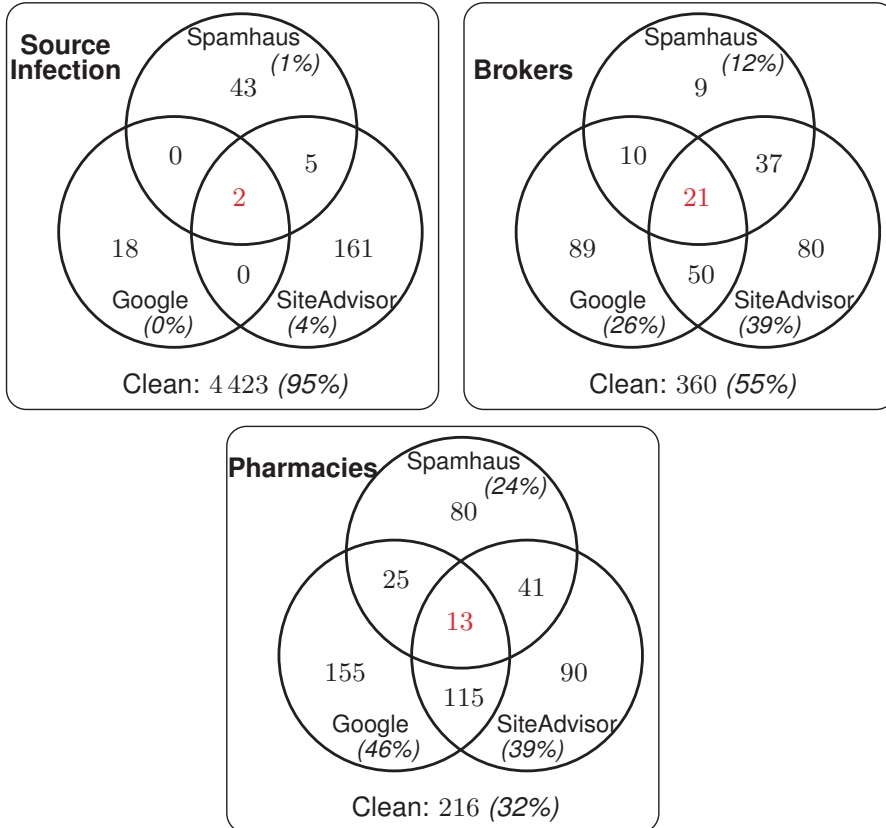


FIGURE 4.8: Comparing web and email blacklists.

senders; and McAfee SiteAdvisor, which tests websites for “spyware, spam and scams”.

Figure 4.8 plots sets of Venn diagrams of the three blacklists for each class of attack domain. Several trends are apparent from inspecting the diagrams. First, source infections are not widely reported by any of the blacklists (95% do not appear on a single blacklist), but around half of the redirects are found on at least one blacklist and over two thirds of unlicensed pharmacy websites show up on at least one blacklist. Surprisingly, 12% of traffic brokers appear on the email spam blacklist, as well as 24% of unlicensed pharmacies. We speculate that this could be caused by affiliates advertising pharmacy domains in email spam, but

Table 4.5: Monthly search query popularity according to the Google Adwords Traffic Estimator.

	Mean	Median	% Searches > 0	Total
Main	14,388	1600	73%	2,374,085
FDA drugs	74	0	6%	323,104
Extra queries	46,380	1,300	59%	32,652,121
Total	6,771	0	20%	35,343,610

it could also be that the pharmacies directly send email spam advertisements or use botnets for both hosting and spamming.

The level of coverage of Google and SiteAdvisor are comparable, which is somewhat surprising given SiteAdvisor’s relatively broader remit to flag scams, not only malware. Google’s more comprehensive coverage of pharmacy websites in particular suggests that some pharmacies may also engage in distributing malware. We conclude by noting that the majority of websites affected by the traffic redirection scam are not identified by any of these blacklists. This in turn suggests that relatively little pressure is currently being applied to the miscreants carrying out the attacks.

4.5 Towards a conversion rate estimate

While it is difficult to measure precisely as an outsider, we nonetheless would like to provide a ballpark figure for how lucrative web search is to the illicit on-line prescription drug trade. Here we measure two aspects of the demand side: search-query popularity and sales traffic.

For the first category, we once again turn to the Google Traffic Estimator to better understand how many people use online pharmacies advertised through search-redirection attacks. Table 4.5 lists the results for each of the three search

query corpora described in Sections 4.2.2 and 4.2.3. The main and extra queries attract the most visitors, with a median of 1,600 monthly searches for the main sample and 1,300 for the extra queries. Several highly popular terms appeared in the results: “viagra” and “pharmacy” each attract 6 million monthly searches, while “cialis” and “phentermine” appear in around 3 million each. By contrast, only 6% of the search queries in the FDA sample registered with the Google tool. The FDA query list includes around 6,500 terms, which dwarfs the size of the other lists. Since over 90% of the FDA queries are estimated to have no monthly searches, the overall median popularity is also zero.

While these search terms do not cover all possible queries, taken together they do represent a useful lower bound on the global monthly searches for drugs. To translate the aggregate search count into visits to pharmacies facilitated by search-redirection attacks, we assume that the share of visits websites receive is proportional to the number of URIs that turn up in the search results. Given that 38% of the search results we found pointed to infected websites, we might expect that the monthly share of visits to these sites facilitated by Google searches to be around 13 million. Google reportedly has a 64.4% market share in search [64]. Consequently we expect that the traffic arriving from other search engines to be $\frac{1-0.644}{0.644} * 13 \text{ million} = 7 \text{ million}$.

We manually visited 150 unlicensed pharmacy websites identified in our study and added drugs to shopping carts to observe the beginning of the payment process. We found that 94 of these websites in fact pointed to one of 21 different payment processing websites. These websites typically had valid Secure Socket Layer (SSL) certificates signed by trusted authorities, which helps explain why multiple pharmacy storefronts may want to share the same payment processing website.

The fact that these websites are only used for payment processing means that if we could measure the traffic to these websites, then we could roughly approximate how many people actually purchase drugs from these pharmacies. Fortunately for us, these websites receive enough traffic to be monitored by services such as Alexa. We tallied Alexa’s estimated daily visits for each of these websites; in total, they receive 855,000 monthly visits.

We next checked whether these payment websites also offered payment processing other than just for pharmacy websites. To check this, we fetched 1,000 backlinks for each of the sites from Yahoo Site Explorer [252]. Collectively, 1,561 domains linked in to the payment websites. From URI naming and manual inspection, we determined that at least 1,181 of the backlink domains, or 75%, are online pharmacies. This suggests that the primary purpose of these websites is to process payments for online pharmacies.

Taken together, we can use all the information discussed above to provide a lower bound on the sales conversion rate of pharmacy web search traffic:

$$\text{Conversion} \approx \frac{0.75 \times 855,000}{20,000,000} = 3.2\% .$$

To ensure that the estimate is a lower bound for the true conversion rate, whenever there is uncertainty over the correct figures, we select smaller estimates for factors in the numerator and larger estimates for factors in the denominator. For example, it is possible that the estimate of visits to payment sites is too small, since pharmacies could use more than the 21 websites we identified to process payments. A more accurate estimate here would strictly increase the conversion rate. Similarly, 20 million visits to search-redirection websites may be an overestimate, if, for instance, more popular search queries suffer from fewer search-redirection attacks. Reducing this estimate would increase the conversion rate since the figure is in the denominator.

There is likely one slight overestimate present in the numerator. It is not certain that every single visitor to a payment processing site eventually concluded the transaction. However, because these sites are *only* used to process payments, we can legitimately assume that most visitors ended up purchasing products. Even with a conservative assumption that only 1 in 10 visitors to the payment processing site actually complete a transaction, the lower bound on the conversion rates we would obtain (in the order of 0.3%) far exceeds the conversion rates observed for email spam [116] or social-network spam [86].

While email spam has attracted more attention, our research suggests that more unlicensed pharmacy purchases are facilitated by search-redirection attacks than by email spam. One study estimated that the entire Storm botnet—which accounted for between 20-30% of email spam at its peak [59, 179]—attracted around 2,100 sales per month [116]. The payment processing websites tied to search-redirection attacks collectively process many hundreds of thousands of monthly sales. Even allowing for the possibility that these websites may also process payments for pharmacies advertised through email spam, the bulk of sales are likely dominated by referrals from web search. In this regard, the work of Kanich et al. [117], reporting an estimated amount of monthly sales for online pharmacies advertised mainly through email spam in the order of 82,000—i.e. one order of magnitude lower than our estimation—supports this claim. This is not surprising, given that most people find it more natural to turn to their search engine of choice than to their spam folder when shopping online. However, disqualifying the 24% of the spam-advertised pharmacies identified in our measurements from inclusion in the conversion rate analysis would have allowed for a more robust estimate of this rate. Consequently, we state this as a limitation of our analysis.

4.6 Conclusions

Given the enormous value of web search, it is no surprise that miscreants have taken aim at manipulating its results. We have presented evidence of systematic compromise of high-ranking websites that have been reprogrammed to dynamically redirect to online pharmacies. These search-redirection attacks are present in one third of the search results we collected in 2010. The infections persist for months, and a static analysis of the redirection chains shows that 96% of the infected hosts are connected through redirections. In addition, a few collections of traffic brokers are critical to the connection between source infections and pharmacies. We have also observed that legitimate businesses are nearly absent from the search results, having been completely drawn out of the search results by blog and forum spam and compromised websites. In Chapter 6 we revisit and validate these observations from a longitudinal perspective, providing also better insights on the temporal characteristics of the redirection chains.

Even though counterfeit drugs are the most pressing issue to deal with due to their inherent danger, other purveyors of black-market goods, such as counterfeit software, or luxury goods replicas, might also hire affiliates that manipulate search results with infected websites for advertising purposes. We ran a brief (12 days) pilot experiment to assess how search-redirection attacks applied to counterfeit software in October 2010. After collecting results from 466 queries, created using input from Google Adwords Keyword Tool, we gathered 328 infected source domains, 72 redirect domains and 140 domains selling counterfeit software. Using the same clustering techniques described earlier in the chapter, we discovered two connected components dominating the network, each in its own way: one component was responsible for 44% of the identified infections, and the other was responsible for 30% of the software-selling sites. We also observed a small but

substantial (12.5%) overlap in the set of redirection domains with those used for online pharmacies. Some redirection domains thus provide generic traffic redirection services for different types of illicit trade. However, the small overlap is also a sign of fragmentation among the different fraudulent trading activities. In Chapter 6 we examine in greater detail our findings based on longitudinal measurements of all retail operations benefiting from search-redirection attacks, in order to better understand the economic relationships between advertisers and resellers. In Chapter 6 we examine this overlap in a more systematic way.

Systematic monitoring of web search results will likely become more important due to the value miscreants have already identified in manipulating outcomes. Indeed, this work has shown that understanding the structure of the attackers' networks gives defenders a strong advantage when devising countermeasures. Indeed, the measurements we gathered lead us to consider three complementary mitigation strategies to reduce the impact of search-redirection attacks. One can target the infected sources, advocate search-engine intervention, or try to disrupt the affiliate networks. In Chapter 9 we provide an in-depth examination and evaluation of these countermeasures.

5

Pricing and inventories at unlicensed online pharmacies

Normally, online pharmacies need to meet a number of licensing requirements before they can operate legally in a large number of countries. Because these requirements can be quite stringent, many entrepreneurs decide to forgo them and operate unlicensed online pharmacies instead. Consequently, such pharmacies face several hurdles designed to stymie their success. First, unlicensed pharmacies encounter considerable scrutiny when advertising online, triggering many operators to employ questionable techniques (e.g., spam, search-engine poisoning), which are likely a lot less effective at bringing customers than legitimate advertising channels (e.g., Google AdWords). Second, the payment processors they rely on to complete transactions may be pressured into cutting off service [147]. Third, unlicensed pharmacies face stiff competition from established pharmacy stores, and even from online black markets [32]. However, unlicensed online pharmacies have managed not only to survive, but even to generate considerable revenue,

with reported annual revenues between \$12.8M and \$67.7M (USD) for some of the largest pharmaceutical networks [148].

In this chapter we attempt (i) to understand the economic reasons for their success, while facing stiff competition from both legal and illegal alternatives, and (ii) to identify characteristics of their supply chains that could be used to disrupt illicit sales. Different from most related work though and from the previous chapter, our focus here is on inventories and prices, rather than on advertising techniques, payment systems, or affiliate network structure [129]. The goal is to learn more about the incentives consumers face when purchasing from unlicensed pharmacies; in economic terms, we analyze the supply to understand the demand better.

We attempt to address this goal through a systematic study of pharmacy inventories. We conjecture that unlicensed online pharmacies either provide inventory that is not available or that is restricted at licensed online pharmacies—e.g., certain types of scheduled drugs, which would require a prescription—or offer considerable price differentials—e.g., they are much cheaper for certain products. We test these hypotheses through a series of controlled measurements.

To that end, we collect and analyze six months worth of inventories and prices at 265 unlicensed online pharmacies identified from a corpus of pharmacies that advertise through search-engine poisoning—a concept we discussed in detail in Section 4.1.1. We compare these inventories and prices with those of another group of 265 pharmacies characterized as “not recommended” by the NABP, but which may not necessarily resort to spam or search-engine poisoning. We also compare unlicensed pharmacy inventories with the inventory of a licensed pharmacy (*familymeds.com*), and with goods that can be found on Silk Road [202], a notorious online black market with a focus on narcotics [32].

This work is, on the one hand, related to measurement studies that focus on specific aspects of online markets to gain better insights over some observed

behavior. In this category, we can relate to Scott et al. [198] that manually collect inventory information from a number of Internet-based stores, and analyze anomalies in the pricing of diamonds; and to the work of Lin et al. [137] that study the effect of reputation systems in online auction markets, by collecting auction information pertaining to specific categories of items.

On the other hand, the work we present in this chapter is closer in spirit to the study of illicit online markets. In particular, it builds on our analysis of search-redirection attacks (Chapter 4), where we identify concentration effects among participants of the underground market. Similarly, McCoy et al. [148] provide ground-truth data on the transactions information of three major pharmaceutical affiliate programs, totaling about 170 Million/year. Our work also builds on the study of the Silk Road marketplace [32], which shows overall revenue to be in the range of 15 Million/year.

Our findings center on two broad areas of investigation: (i) examining drug inventories, and (ii) inferring pricing strategies at unauthorized pharmacies. With respect to inventories, we find that narcotic and schedule drugs are rare: 0.6% of the 486 scheduled ingredients are sold by *familymeds.com*, compared to 6% in unlicensed pharmacies and 9% at Silk Road. Drugs treating chronic medical conditions such as cardiac and psychiatric disorders are found disproportionately often at unlicensed pharmacies, while cancer medications are under-represented. Finally, we can cluster unlicensed pharmacies by the similarity of their inventories, finding that half of those inspected belong to one of eight clusters likely sharing supply chains.

With respect to pricing, we present evidence that pharmacy operators strategically price drugs to entice budget-conscious customers. First, unlicensed pharmacies are simply cheaper than *familymeds.com* overall—median \$2.14 (56%) per unit cheaper. But the discounts vary considerably: fake generics (where the

pharmacy claims to offer a non-existent generic version of a branded drug) are \$1.54 cheaper than other prices at unlicensed pharmacies. While most legitimate pharmacies do not offer volume discounts, unlicensed pharmacies do: the median discount for a 90-day supply is 17% off the price of a 30-day supply. While pharmacies can influence some discounts, they must also react to market competition. We find that the more pharmacies sell a given drug, the deeper the discount they offer.

The work we present in this chapter informs various aspects of research questions 1¹ and 5² by (i) enhancing our understanding of the financial incentives and opportunities allowing the operation of illicit online pharmacies, (ii) characterizing the structure of the illicit prescription drug supply network, and (iii) by outlining financially-motivated disincentives for the criminal actors engaged in the illicit online prescription drug trade.

The rest of this chapter is organized as follows. We start by contrasting the different types of online pharmacies and discuss some of the advertising techniques employed by unlicensed pharmacies in Section 5.1. We describe our data collection methodology in Section 5.2. We analyze inventories in Section 5.3, and examine pricing strategies in Section 5.4, before drawing conclusions in Section 5.5.

5.1 Background

Categorizing online pharmacies as either “legitimate” or “illicit” oversimplifies the diversity of the market. A first reasonable distinction one can make is between *licensed* and *unlicensed* pharmacies. Licensed pharmacies are either online front-ends to brick-and-mortar stores with a valid pharmacy license, or online phar-

¹ Are there any structural characteristics in the illicit online prescription drug trade. . .

² Is it possible to disrupt online criminal networks by targeting critical components. . .

macies that obtain prescriptions only through third-party pharmacies with verified licenses. Licensing requirements themselves vary from country to country, or even from state to state in the case of the US. Thus, an online pharmacy may have a perfectly valid license in Barbados, but would not necessarily be licensed to sell drugs in the US. To this end, as discussed in Section 3.2.2, accreditation and verification programs have been developed to help assess the legitimacy of online pharmacies, and assist consumers in making informed decisions.

In this section, we first discuss advertising techniques, which may themselves be a good indicator of whether an online pharmacy is engaging in questionable business or not. We then briefly introduce emerging online black markets, which are at the far end of the legitimacy spectrum.

5.1.1 Advertising techniques

An indicator of the potential legitimacy of an online pharmacy is the type of advertising techniques it employs. Licensed, accredited pharmacies can purchase Google AdWords for instance, while unlicensed pharmacies have been barred from doing so since 2003 [177]. Thus, some unlicensed pharmacies resort to illicit advertising techniques. Of those, email spam [116] is perhaps the best known, but blog and forum spam, as well as search-engine poisoning have established their prominence [130]. Because this latter form of advertising involves active compromise of unsuspecting Internet hosts (and doing so is criminal offense in many countries), we can almost assuredly categorize the pharmacies resorting to search-engine poisoning as illicit.

Variants of search-redirection attacks. In Chapter 4 we presented in detail a method employed by online pharmacies to fraudulently advertise, by compromising vulnerable websites, and manipulating search engines. While this methodol-

ogy is still largely in use, in 2011 we identified two additional variants of search redirection attacks.

These new variants appeared as a response to search-engine interventions that concealed the HTTP referrer field when users click on search results. For instance, in secure HTTP (i.e. HTTPS) searches, that are the default in modern versions of the popular web browsers (e.g. Firefox), the referrer field only shows that a given visitor is coming from Google, but not the specific terms used in the query. This defeats the attack outlined in Section 4.1.1. In response, attackers started placing simple pharmacy storefronts *within* the compromised domain, and display them if they notice the traffic is coming from Google (as opposed to coming from a different page in the same domain, or visiting directly the location of the storefront), regardless of the type of query being made. These storefronts typically consist of a few pictures with links; clicking on any of these links redirects the visitor to an online pharmacy.

The second variant is slightly more complex, and is outlined in Figure 5.1. Upon connection to a compromised site (step 1), the visiting client receives a cookie (step 2), and is simultaneously redirected to a key generator site (step 3) which simply passes back a response key to the client (step 4), and redirects the client back to the compromised servers (step 5). The visiting client produces both the cookie received earlier and the response key, which triggers the compromised server to display a pharmacy storefront as in the previous variant. Clicking on any link takes the client to an actual pharmacy store (step 6). From a user standpoint, there is no difference between this attack and the previously described attack; from the attacker's standpoint, however, the use of cookies make this type of attack significantly more difficult to detect by automated crawlers, which tend not to keep any state. Indeed if the cookie is not produced, an empty page, rather than a pharmacy storefront, is displayed.




FIGURE 5.1: A variant of the search-redirectation attack that appeared as a response to search engine intervention.

In Chapter 6 we offer additional details on the evolution of criminal advertising tactics, such as the ones we describe here, considering the context and the temporal characteristics of deployed countermeasures.

5.1.2 The emergence of online black markets

Unlicensed pharmacies are in addition facing a novel form of concurrence: online black markets. Thanks to significant usability efforts in the past couple of years, Tor [58] is now usable by computer novices, who can simply download the “Tor browser” and access the Internet anonymously. Tor also supports hidden services, which are essentially web servers whose IP address is concealed. Coupled with the recent emergence of Bitcoin [168], a peer-to-peer distributed currency without any central governing authority, a number of hidden services have emerged selling

contraband or illicit items [20, 202, 212]. Different from unlicensed pharmacies, these black markets do not make any claims of legitimacy: Users know that they are purchasing contraband items.

Perhaps the best known of those black markets is Silk Road, which primarily focuses on narcotics and prescription drugs, and has an estimated total yearly revenue of approximately \$15 million [32]. The Silk Road market operators do not themselves sell any goods, but instead provide an anonymous online forum for sellers and buyers to engage in transactions. As such, it is not a “pharmacy” per se so much as a middleman bringing together vendors of pharmaceutical goods (among others) with prospective customers. Of course, on Silk Road and other black markets, no prescription or verification of any kind is required to make a purchase.

5.2 Measurement methodology

We next discuss how we collected inventory and pricing data, which we make publicly available for reproducibility purposes.³ We first discuss how we selected pharmacy sites, before explaining how we extracted inventories from each pharmacy.

5.2.1 *Selecting and parsing pharmacies*

We gathered data from four groups of pharmacies: 265 pharmacies that have been advertising using one of the variants of the search-redirection attacks discussed in Section 5.1, an additional set of 265 pharmacies that are listed as “not recommended” by the NABP, 708 distinct vendors on Silk Road, and the licensed pharmacy *familymeds.com*.

³ See <https://arima.cylab.cmu.edu/rx/>.

Search-redirection advertised pharmacies. We identified pharmacies advertising through search-redirection attacks by adapting the crawler we defined in Section 4.2 to analyze results from both Bing and Google. That crawler simply follows chains of HTTP 302 redirects found in response to a corpus of 218 drug-related queries until it reaches a final site, which it labels as an online pharmacy. As discussed in Chapter 4, this simple heuristic is surprisingly accurate at identifying online pharmacies. We enhanced the crawler to reach pharmacies advertised using the novel attacks described in Section 5.1.1.

We then scraped all the candidate pharmaceutical sites our crawler identified. Around April 3rd, 2012, we attempted to scrape all the candidate pharmaceutical sites our crawler had identified until then; many of these domains had been taken offline, which is not overly surprising given the relatively short life span of online pharmacies. Then, between April 3rd, 2012 and October 16th, 2012, we scraped all candidate pharmaceutical sites at the time our crawler detected them.

We used `wget` to scrape the content of the candidate pharmacy domains. We used random delays between different web page accesses in the same domain to avoid detection. As we have previously observed (Section 4.2), operators actively monitor visitor connections and respond to abnormal activity. To overcome the risk of being banned, we anonymized our traffic using Tor [58], changing Tor circuits every 15 minutes to evade IP blacklisting. Traffic anonymization came at the price of longer latencies, however. Depending on the size of each pharmaceutical domain, the scraping process took from 4 to 12 hours to complete. As a result, we decided to scrape each pharmacy only once.

After removing, from our set of 583 candidate pharmacies, false positives (non-pharmaceutical sites), parked domains, and pharmacies for which we could not easily retrieve inventories, we obtained complete inventories for a total of 265 online pharmacies that advertise through variants of search-redirection attacks.

By a slight abuse of terminology, we will refer to this set of pharmacies as the *unlicensed pharmacy set*.

NABP's "not recommended" pharmacies. We complement the unlicensed pharmacy set by a random sample of pharmacies labeled as "not recommended" by the NABP. There are 9,679 such such pharmacies. The details of how the NABP has assembled this list are unclear, but only 60 domains from the unlicensed pharmacy set are among the 9,679 "not recommended" pharmacies. This shows that the NABP is applying a set of criteria very different from ours to identify illicit pharmacies. Therefore a sample drawn from these 9,679 pharmacies can be useful to determine whether pharmacies that use search engine manipulation as their advertising vector exhibit different behavior compared to other illicit pharmacies.

Out of the 9,679 pharmacies in the list, after excluding pharmacies in the unlicensed pharmacy set we draw a random sample of 265 domain names. We scrape these pharmacies to acquire their inventory as described above. Scraping took place between October 30th, 2012 and November 4th, 2012. We will denote this set of pharmacies as the *blacklisted pharmacy set*.

Familymeds.com. Finding licensed online pharmacies for which we can collect inventory information was a surprisingly difficult task. The vast majority of popular online pharmacies we examined require active membership to grant access to their inventories and pricing information. Since becoming a member often requires producing valid private information of a sensitive nature, such as health and/or prescription insurance contract numbers, we opted not to register for any of these domains. Instead, we chose *familymeds.com*, a VIPPS [170] accredited pharmacy based in Connecticut as our source of legitimate prices and invento-

ries. *familymeds.com*'s inventories and prices are freely available to anybody who browses their site.

There are certainly additional licensed online pharmacies that we could consider to enrich this dataset. However, it is important to realize two characteristics of licensed pharmacies. First, inventories should considerably overlap from one legitimate pharmacy to the next, since drug names and active ingredients are fixed. Second, drug prices vary significantly (50% on average) between different pharmacies, even within small communities [205]. Price variations appear to be more associated with the consumer behavior and less with the pharmacy itself. For example, prices of frequently prescribed drugs (e.g. drugs that treat chronic conditions) tend to vary less than one-time prescriptions (e.g. antibiotics). Consequently, including prices from additional licensed pharmacies would only introduce additional noise in the data, without much added benefit for our analysis. While *familymeds.com* provides an interesting datapoint to which we can compare illicit pharmacy prices, studying price differentiation between legitimate pharmacies is indeed outside the scope of this Chapter.

Silk Road. Finally, we use data from Silk Road, an online anonymous black market. As part of a related study [32], we obtained the entire inventory of 24,385 items available on Silk Road between February 3, 2012 and July 24, 2012. Then, we matched each item against a comprehensive list of drug names provided by the FDA [235]. Excluding items for which no match was found this narrowed down the list to 5,511 items, which were offered by 708 different vendors. After additional inspection, we discarded a number of items that either were completely irrelevant, or did not have all the information we need (e.g., missing dosage or number of units sold), which further reduced this list to 4,208 unique items. Even though vendors are different entities, we will consider Silk Road as a unique “pharmacy” in the rest

of this Chapter, because we conjecture that minor differences from one vendor to the next are small in comparison to the differences between Silk Road-like black markets and online pharmacies.

5.2.2 *Extracting inventories*

Once we have the webpages of interest, we need to extract inventories from these pages. We wrote a generic Hypertext Markup Language (HTML) parser to accomplish this task. More importantly, building inventories requires us to identify what constitutes a “drug,” and associating it the right data. Defining the notion of drug is not as simple as it sounds: is a drug defined by its brand name, or by its active ingredient? Should we include dosage in the definition, considering that, at different doses, a medication might shift from over-the-counter to prescription only? For instance, ibuprofen is available over-the-counter at less than 200 mg, and requires a prescription for higher dosages.

To build a drug price index, economists have previously discussed possible sets of features that put together, could adequately describe and track a drug [4]. Following their lead, we decided to collect as much information as possible regarding a given “drug.” Specifically, we gather the following 5-tuples for each medication: (1) *drug name* (e.g., “Viagra”), (2) *active ingredient(s)* (e.g., Sildenafil), (3) *dosage* (e.g. 10mg, 10mcg, 10%), (4) Number and type of units (e.g. 10 tablets, 1 bottle, 2 vials), which we will collectively refer to as *unit*, (5) and the *type of drug* (i.e., generic vs. brand). We then associate each of these tuples with a *price* (e.g., 10.83) and a *currency* (e.g, USD, GBP, Euro).

For each pharmacy page we scraped, we identified the main drug advertised using a list of known prescription drugs [235], computing the Term Frequency – Inverse Document Frequency (TF-IDF) score [193] and picking the drug name

with the highest score. In a few cases, we were able to determine the name of the main drug simply by looking at the HTML file name.

We used the same method to determine the type of drug—brand or generic. However, these terms are often used ambiguously, or even deceptively by unlicensed pharmacies. For example, many online pharmacies advertise “generic Viagra.” However, a generic can only be produced and traded when the associate intellectual property rights have expired, or in jurisdictions where the intellectual property rights do not apply. In the case of Viagra the relevant patent is still in effect, which means that “generic Viagra” does not legally exist in most countries.⁴ Whether this means the product sold is counterfeit medication, or simply mislabeled, is unclear without making a purchase and analyzing the drug.

Using the displayed drug names, we identify the active ingredients by querying the RxNorm database of normalized names for clinical drugs [140]. Illicit pharmacies often sell drugs that are either not licensed in the US (e.g. Silagra, Kamagra) or are simply counterfeit combinations of existing drugs (e.g. Super Hard ON). Such drugs do not have any associated ingredient in the RxNorm database, and we exclude the 119,701 such tuples from our our analysis.

We then collect pricing information for each tuple collected. Figure 5.2 shows a typical example of how pricing information associated with a drug is presented. In this figure, our parser would produce three separate inventory entries. For instance, the entry corresponding to the first row in the figure would be “Viagra, Sildenafil, 200mg, 20 pills, brand, USD 150.”

5.2.3 *Collecting supplemental data*

We complement our inventory entries by gathering supplemental drug attributes from several different sources.

⁴ India is an exception in this case, as its patent laws permit the production of “generic” Viagra [5].

The screenshot shows the Canadian Health & Care Mall website. At the top, there are navigation links for ALL PRODUCTS, ABOUT US, HOW TO ORDER, and TESTIMONIALS. Below the header is a search bar with the placeholder text "Enter product name" and a "Your Cart" button. On the left side, there is a currency selector with options for USD, GBP, CAD, EUR, AUD, and CHF. Below that is a "MEN'S HEALTH" section with a list of products including Viagra, Cialis, and Levitra, each with a red star icon. The main content area displays three separate entries for Viagra pills: 200mg, 150mg, and 130mg. Each entry includes a description, a price with a 10% discount, and a "Testir" price. The 200mg entry shows a price of \$166.67 (discounted from \$185.19) and a "Testir" price of \$150.00. The 150mg entry shows a price of \$155.33 (discounted from \$172.59) and a "Testir" price of \$139.80. The 130mg entry shows a price of \$144.44 (discounted from \$160.48) and a "Testir" price of \$127.00. Each entry also includes a "FREE BONUS PILLS" offer.

FIGURE 5.2: Example of multiple drug names, dosages, currencies and prices presented within a single page (rxcaresign.com). From this page, our parser produces three separate inventory entries.

Table 5.1: **Summary data for all four data sources.** In the case of Silk Road, we show the number of different “vendors” rather than a number of pharmacies.

Data Source	# Pharmacies	# Drug names			# Ingredients	# Records	Inventory size (median / mean)	Diseases targeted
		Scheduled	Narcotics	All				
unlicensed pharmacies	265	42	9	1,000	557	1,022,635	157 / 170	652
<i>familymeds.com</i>	1	4	0	657	500	7,277	697	616
Silk Road	708	69	12	237	183	4,208	272	335
blacklisted pharmacies	265	51	7	1,283	774	417,467	64 / 107	726
Total	532	90	15	1,611	939	1,451,587		755

Schedule drugs and narcotics. We collect information related to the Schedule and Narcotic status of each drug [229]. The Schedule classification was established as part of the Controlled Substances Act in 1970 [227] and includes five ordered classes of drugs. Drugs are assigned to any of the schedules based on their potential for abuse and addiction. Schedule I drugs (e.g., marijuana), have the highest potential for abuse and are not deemed to have any acceptable medical use in the US, while Schedule V drugs (e.g., Robitussin) have the lowest potential for abuse compared to the other schedules.

Diseases treated. We use the National Drug File – Reference Terminology (NDF-RT) [23, 138] to collect the associations between active ingredients and the diseases they treat or prevent.

WebMD drug classification. We supplement the NDF-RT information with data collected from WebMD [244]. WebMD groups drugs into 100 categories of medical conditions that the drugs are designed to treat, such as “Acne” or “Headache”. We extracted the drug names associated with each condition. We also used WebMD to get an idea of the drug popularity, by extracting the 180 drug names classified by WebMD as “top drugs”, which were selected “according to the number of searches submitted on WebMD for each individual drug”.

FDA drug shortage list. The FDA tracks when drugs are in currently in short supply [233]. We gathered the list of 110 drug ingredients listed as in shortage, in order to check their availability at unlicensed pharmacies, *familymeds.com* and Silk Road. The information in [233] is relatively unstructured, but it provides the National Drug Code (NDC) identifiers of the drugs, which are directly associated with specific combinations of drug names and dosage. We used the RxTerms database [74] to decode the collected NDCs into information compatible with our drug data.

We combine the inventory information given by the 5-tuples discussed earlier, with this supplemental information, and create separate *records* in our database for each drug so observed.

5.3 Inventory analysis

We next present an analysis of the inventory data we gathered. In this section, we focus on item availability, rather than prices. We start with an overview of the data we have, before discussing the granularity which we will use to define “drugs.”

We then compare the availability of different drug classes (schedule drugs, for instance) across pharmacies, and we specify main types of medical conditions targeted by the unlicensed pharmacy set. Last, we perform clustering analysis on the available inventories, to identify a common pattern in the suppliers of online pharmacies.

5.3.1 Drug availability by pharmacy type

Table 5.1 presents a breakdown of the collected data from the unlicensed pharmacy set, the blacklisted pharmacy set, *familymeds.com* and Silk Road. We collected a total of 1,451,587 distinct (drug name, active ingredient, dosage, unit) records. These records contain 1,611 different drug names.

Drug availability. Both unlicensed pharmacies and blacklisted pharmacies exhibit the largest number of different drugs being sold, but the total number of different actual active ingredients are similar to those available on *familymeds.com*. A possible explanation is that, compared to licensed pharmacies, unlicensed pharmacies try to offer a wide variety of drug names to attract a wider range of customers. In addition, unlicensed pharmacies also target markets outside the United States, where same active ingredients often carry different market names. For instance, generic variants of Tylenol (acetaminophen) in the United States are sold as “paracetamol” in the United Kingdom. This seems to be confirmed by the fact that there are between 4.4 and 4.7 different drug names listed per disease/condition treated in the unlicensed and blacklisted pharmacies, compared to 3.4 different drug names associated with a given condition in *familymeds.com*. The corresponding number in Silk Road is 2.7.

Scheduled drugs. 90 of the 1,611 drug names we found are listed under Schedules I to V, including 15 drugs categorized as narcotics. The licensed pharmacy

Table 5.2: Scheduled drugs, narcotics, drugs in shortage, and top drugs at *familymeds.com*, unlicensed pharmacies and Silk Road.

Category	total #	Unlic. pharm.		Unlic. pharm.		Drug ingredients		Sig. diff.?	Silk Road	Sig. diff.?	#Unlic. Pharm.	
		(all)	(%)	(median)	(%)	<i>familymeds.com</i>	(%)					
Drugs in shortage	150	75	(50%)	8	(5.3%)	32	(21.3%)	✓	21	(14%)	✓	265
Top WebMD Drugs	283	255	(90.1%)	57	(20.1%)	146	(51.6%)	✓	93	(32.9%)	✓	265
Narcotics	166	10	(6.0%)	2	(1.2%)	0	(0%)		11	(6.6%)		8
Schedule (all)	486	33	(6.8%)	1	(0.2%)	3	(0.6%)		44	(9.1%)		63
Schedule I	132	0	(0%)	0	(0%)	0	(0%)		0	(0%)		0
Schedule II	93	10	(10.8%)	2	(2.2%)	0	(0%)		15	(16.1%)		8
Schedule III	116	9	(7.8%)	1	(0.9%)	1	(0.9%)		7	(6.0%)		46
Schedule IV	135	14	(10.4%)	2	(1.5%)	2	(1.5%)		21	(15.6%)		28

familymeds.com does not sell any narcotics and only four scheduled drugs. Both blacklisted pharmacies and unlicensed pharmacies, on the other hand, appear to sell more scheduled drugs and narcotics, and both sets appear relatively similar to each other. Silk Road tops the list in both scheduled drugs and narcotics.

Comparison of drug availability by ingredient type. Beyond the absolute numbers of drugs for sale at different types of pharmacies, we are also interested in studying how comprehensive the inventories are. For instance, while it is useful to know that scheduled drugs are offered under 15 different names, it would also be nice to know how many scheduled drugs *cannot* be found at unlicensed pharmacies, and whether the proportion offered is greater or less than at licensed pharmacies. To answer these questions, we compare the ingredients observed to comprehensive listings of drug *ingredients*, since drugs may be marketed under many names that cannot easily be enumerated completely.

Table 5.2 reports on the prevalence of different categories of drug ingredients on unlicensed pharmacies, *familymeds.com* and Silk Road. In addition to the schedule and narcotics categories mentioned above, the table also reports on the availability of popular drugs and those currently in shortage. For example, 75 of the 150 drug ingredients currently in shortage are for sale at one or more unlicensed pharmacies. While this is much higher than the 32 shortage ingre-

dients for sale at *familymeds.com*, it would be wrong to conclude that there is better availability at unlicensed pharmacies than at licensed ones, since we are comparing the inventories of 265 pharmacies to just one. A fairer comparison is between *familymeds.com* and the median number of shortage ingredients offered by unlicensed pharmacies (8). Using a χ^2 test, we conclude that this difference in proportions is statistically significant ($p < .0001$).

By contrast, when comparing inventories on the Silk Road to unlicensed pharmacies, it is better to compare the complete inventory for unlicensed pharmacies since both rely on many sellers. In the case of shortages, unlicensed pharmacies offer much greater coverage than do sellers on Silk Road (14% vs. 50%).

Notably absent from our datasets are Schedule I drugs. We could not find any evidence of such drugs being sold at unlicensed pharmacies. While they are, on the other hand, frequently sold on Silk Road [32], we purposefully excluded them from our data collection, since they are not in the FDA list we used [235]—this list focuses on drugs with therapeutic effects.

The main takeaway is that, the more illicit the market, the more controlled substances are available. We note that the differences between pharmacy types were statistically significant (using a χ^2 test) for schedule drugs, which is not unexpected given how uncommon schedule drugs and narcotics are in unlicensed pharmacies and *familymeds.com*. This result is not overly surprising: Guidelines that regulate the distribution of scheduled drugs (and therefore narcotics, which are a subset of the scheduled drugs) make electronic ordering from licensed online pharmacies difficult. For instance, codeine, which is listed under Schedule II, requires a written prescription that the pharmacy needs to verify before dispensing the drug. Processing such orders is cumbersome for law-abiding online pharmacies, and they may limit their inventories of controlled substances. At the other end of the

Table 5.3: **Similarities in drugs sold using different drug definitions.** While pharmacies may sell the same drugs, it is somewhat less common to sell the same drug and dosage, rarer still to sell the same drug, dosage and number of pills.

Drug name <i>example tuple</i>	Dosage	Units	Unlicensed pharmacies		
			# matches	# pharmacies	% records
⊆ familymeds.com					
<i>Viagra</i>			391	260	54.6%
<i>Viagra</i>	<i>100mg</i>		318	247	25.6%
<i>Viagra</i>	<i>100mg</i>	<i>30 pills</i>	299	243	15.4%
⊆ Silk Road					
<i>Viagra</i>			164	261	32.1%
<i>Viagra</i>	<i>100mg</i>		138	257	11.1%
<i>Viagra</i>	<i>100mg</i>	<i>30 pills</i>	62	250	1.2%
⊆ familymeds.com ⊆ Silk Road					
<i>Viagra</i>			85	256	25.2%
<i>Viagra</i>	<i>100mg</i>		69	245	7.4%
<i>Viagra</i>	<i>100mg</i>	<i>30 pills</i>	26	234	0.6%

spectrum, an anonymous online black market like Silk Road thrives on offering controlled substances to anybody willing and able to pay for them [32].

5.3.2 Product overlap between different types of pharmacies

We next investigate the extent to which products offered by different types of pharmacies overlap. Recall from Section 5.2, that a drug is fully described by five features: active ingredient(s), name, dosage, units, and whether the drug is a brand or a generic. As such the definition of “overlap” in inventory is actually dependent on the level of granularity we choose to define what a “drug” is. Table 5.3 shows the effect of choosing a specific level of granularity to look for matches across pharmacies. The left-hand side of the table displays the set of pharmacies and drug features we are using to look for matches within the unlicensed pharmacy set. The numbers on the right hand-side of the table indicate the number of matches, number of pharmacies that contain a match, and overall fraction of

records for which a match is found. For example, the first row describe matches when we only use the drug name for comparison, ignoring other attributes. An example would be to simply search matches for “Viagra,” ignoring differences in dosages and units. Matching by drug names only, we find that there are 391 drugs sold both by *familymeds.com* and unlicensed pharmacies; we are able to find a match in 260 of the unlicensed pharmacies. These matches correspond to 54.6% of the 1,022,635 records (drug/price combination) we collected from the unlicensed pharmacy set.

Obviously, the more features we use to identify matching drugs, the fewer records we have available to draw conclusions from. On the other hand, these finer records are of better quality, since we know that we are comparing similar items. A particularly interesting result in Table 5.3 is that, regardless of the level of granularity considered, inventories in unlicensed pharmacies and *familymeds.com* are considerably different.

This shows that one of the ways unlicensed pharmacies compete with legitimate pharmacies is by offering different items. The fact that a large number of unlicensed pharmacies actually appear in the matches indicates that unlicensed pharmacies collectively offer a larger inventory than we can find at *familymeds.com*. This finding is confirmed by what we observe when looking at Silk Road. Silk Road, as described above, has a much richer inventory in controlled substances than both unlicensed pharmacies and *familymeds.com*. In other words, a key lesson from Table 5.3 is that, rather than purely competing on substitutes with legitimate pharmacies, unlicensed pharmacies and black market vendors are providing complementary inventories.

Table 5.4: Odds-ratios identifying the medical conditions that are over-represented or under-represented in the inventories of unlicensed pharmacies.

Condition	odds ratio	95% CI	p value	Meta-condition
<i>Conditions with more drugs sold by unlicensed pharmacies</i>				
Bipolar Disorder	6.0	(3.4,11.6)	0.0000	Psychiatric
Congestive Heart Failure	4.6	(2.9,7.6)	0.0000	Cardiac
Heart Attack	4.3	(2.8,6.8)	0.0000	Cardiac
Stroke Prevention	7.7	(2.8,27.7)	0.0000	Cardiac
Sinus Infection	10.5	(2.8,73.9)	0.0002	Allergies
Syphilis	7.3	(2.6,26.2)	0.0001	STD
Chlamydia	5.1	(2.5,11.1)	0.0000	STD
High Blood Pressure	3.4	(2.4,4.7)	0.0000	Cardiac
Bronchitis	4.9	(2.4,10.7)	0.0000	
Depression	4.0	(2.4,7.1)	0.0000	Psychiatric
Cold Sores	13.2	(2.4,332.7)	0.0015	
Acid Reflux	5.1	(2.3,12.4)	0.0000	
Strep Throat	5.3	(2.3,13.7)	0.0000	
Tonsillitis	5.5	(2.3,15.5)	0.0001	
Gonorrhea	4.2	(2.2,8.5)	0.0000	STD
Anxiety	3.5	(2.2,5.9)	0.0000	Psychiatric
Ear Infection	3.4	(2.0,5.8)	0.0000	
Diabetes	2.7	(1.9,4.0)	0.0000	
Asthma	2.6	(1.8,3.7)	0.0000	
COPD	2.9	(1.6,5.4)	0.0004	
Dementia	2.9	(1.6,5.4)	0.0007	Psychiatric
Lyme Disease	3.7	(1.5,9.9)	0.0039	
Fibromyalgia	3.5	(1.5,8.8)	0.0039	
Bursitis	2.9	(1.3,6.4)	0.0065	
Staph Infection	2.0	(1.3,3.1)	0.0012	
Gout	4.1	(1.3,15.6)	0.0155	
Hives	2.3	(1.3,4.3)	0.0053	
Chest Pain	2.2	(1.3,3.6)	0.0038	Cardiac
Ulcer	3.7	(1.3,12.1)	0.0157	
Tendonitis	2.6	(1.3,5.7)	0.0108	
Pneumonia	1.7	(1.2,2.6)	0.0054	
Stomach Flu	3.3	(1.1,11.0)	0.0312	
High cholesterol	2.1	(1.1,3.8)	0.0227	Cardiac
Arthritis	1.6	(1.1,2.2)	0.0136	
Edema	2.3	(1.1,5.1)	0.0324	
Bladder Infection	2.1	(1.1,4.1)	0.0315	STD
<i>Conditions with fewer drugs sold by unlicensed pharmacies</i>				
Psoriasis	0.66	(0.43,0.98)	0.0408	
Leukemia	0.60	(0.38,0.92)	0.0179	Cancer
Lymphoma	0.54	(0.36,0.79)	0.0012	Cancer
Anemia	0.38	(0.23,0.60)	0.0000	
Endometriosis	0.34	(0.16,0.65)	0.0006	
Lung Cancer	0.31	(0.11,0.68)	0.0026	Cancer
Constipation	0.17	(0.01,0.88)	0.0316	

5.3.3 Identifying drug conditions served by unlicensed pharmacies

In epidemiology, it is common to observe a disease and only afterwards identify risk factors that promoted transmission. *Case-control* studies are suited to this task [196], and we can use this method to identify which medical conditions are at greater “risk” of being served by unlicensed pharmacies. Lee first employed the case-control method to cybercrime [122], identifying which academic departments were targeted most by spear-phishing emails laced with malware. We use the data mapping drug ingredients to 100 medical conditions from WebMD to construct risk factors. We then check how many of these ingredients are offered at unlicensed pharmacies. For each category we calculate the following probabilities:

	Case (in unlicensed pharmacies)	Control (not in unlicensed pharmacies)
Drug in condition	p_{11}	p_{10}
Drug not in condition	p_{01}	p_{00}

We can then compute an odds ratio for each category:

$$\text{odds ratio} = \frac{p_{11} * p_{00}}{p_{10} * p_{01}}$$

95% confidence intervals for the odds ratio are calculated using the mid-p method. Any risk factor with lower 95% confidence bound greater than 1 is positively correlated with drugs appearing in unlicensed pharmacies. Similarly, any risk factor with upper 95% confidence bound less than 1 is negatively correlated with drugs appearing in unlicensed pharmacies.

Table 5.4 lists the 36 conditions positively correlated with appearing in unlicensed pharmacies along with 7 negatively-correlated conditions. The remaining 57 conditions are not included in the table due to space constraints. We can see from the table that cardiac conditions, Sexually Transmitted Diseases (STDs)

and psychiatric conditions are among the meta-categories with multiple conditions positively associated with drug ingredients offered by unlicensed pharmacies. It makes sense that cardiac drugs would be featured prominently by unlicensed pharmacies, given their widespread use as ongoing maintenance medication and considerable expense. STDs and psychiatric disorders are also often chronic conditions, which require ongoing drug treatment and consequently, recurring expenses that many consumers would opt to reduce. Furthermore, some psychiatric drugs may be abused for recreational purposes, e.g., Xanax.

By contrast, three of the seven conditions negatively associated with unlicensed pharmacies are forms of cancer. Cancer medications are frequently administered by hospitals, and so consumers are less likely to fill prescriptions directly. Furthermore, many people might be willing to try an online pharmacy to treat chronic conditions such as diabetes and cardiac medication, but they would balk at doing so for drugs to treat cancer.

In sum, we have found evidence that unlicensed pharmacies do not simply offer a random selection of drugs in their inventories. Instead, they choose to sell drugs favoring chronic conditions such as cardiac and psychiatric disorders, while selling fewer drugs to treat cancer.

5.3.4 Identifying suppliers

We next turn to looking at similarities in inventories *among* unlicensed pharmacies. As has been described in previous work [128, 132, 147, 148], unlicensed pharmacies often operate as parts of affiliate networks. That is, affiliates essentially set up storefronts, and are in charge of finding ways of bringing traffic to them. On the other hand, once a sale is completed, they are not actually involved in the shipping and delivery of the drugs. This task is handled by the affiliate network operators, who collect most of the sales revenues [148]. Hence, we expect to see




FIGURE 5.3: Heat map of the Jaccard distances between all pairs of pharmacies in the unlicensed pharmacy set. After reordering pharmacies, we observe a number of clusters that appear to have similarities.

striking similarities in inventories offered by various members of the same affiliate network. In fact, prior work by Levchenko et al. [132] observed similarities in web pages from identical affiliate programs. Here, we focus on inventories to further determine whether or not different networks might have common suppliers.

As in related work on malware classification [112] or webpage classification [132], we use the Jaccard distance to determine how (dis)similar two pharmacy inventories are. If A is the inventory of pharmacy \mathcal{A} and B the inventory of pharmacy \mathcal{B} , the Jaccard distance J_δ between A and B is given by:

$$J_\delta(A, B) = 1 - J(A, B) = 1 - \frac{|A \cap B|}{|A \cup B|} \quad (5.1)$$

If two pharmacies share the same exact inventories their Jaccard distance will be equal to 0, and if their inventories have nothing in common, then their distance is 1.

We plot a heat map of the Jaccard distances between all pharmacy pairs in Figure 5.3. After reordering columns to pool together Jaccard distances that are close to each other, clusters of similar inventories appear quite clearly in the figure.

We can define two pharmacies as belonging to the same cluster if their Jaccard distance is below a threshold t . To recursively merge clusters we consider three alternatives:

—*Single linkage*, where the distance between two clusters of pharmacies X and Y is defined as the distance of the two most similar members of the clusters. That is, the distance between two clusters is

$$\min\{J_\delta(x, y) : x \in X, y \in Y\} ,$$

where x and y correspond to inventories of pharmacies in each cluster, respectively.

—*Complete linkage*, where the distance between two clusters of pharmacies X and Y is defined as the distance of the two most dissimilar members of the clusters. That is,

$$\max\{J_\delta(x, y) : x \in X, y \in Y\} .$$

—*Average linkage* [204], where the distance between two clusters of pharmacies X and Y is defined as the average distance between all pairs of members in both clusters:

$$\frac{1}{|X| \cdot |Y|} \sum_{x \in X} \sum_{y \in Y} J_\delta(x, y) .$$

Figure 5.4 shows how many clusters are identified as a function of the distance threshold. The left plot corresponds to the unlicensed pharmacy set, while the right plot corresponds to the blacklisted pharmacy set. The lines correspond to the different linkage criteria. A good threshold value is empirically defined as a value

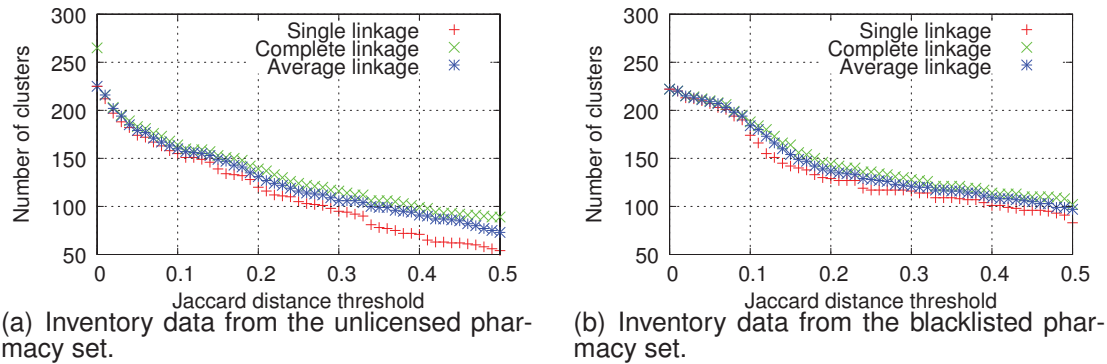


FIGURE 5.4: Effect of different levels of distance threshold and different linkage criteria.

for which the number of clusters remains constant even if we slightly increase the threshold. Using average linkage, we find that $t = 0.31$ is a good choice for the threshold. This value is incidentally very close to the value ($t = 0.35$) used by Levchenko et al. in their related analysis [132]. More interestingly, we find that $t = 0.31$ is an appropriate choice for both the unlicensed pharmacy set and the blacklisted pharmacy set.

In Figure 5.5 we plot the cumulative distribution of the pharmacies as a function of the number of clusters considered. Clusters are ranked by decreasing size. While we observe 82 singletons, the key finding here is that, for unlicensed pharmacies, half of the pharmacies belong to one of eight clusters. Presumably, these map to the larger pharmaceutical affiliates. We obtain similar results for blacklisted pharmacies (101 singletons, 9 clusters corresponding to 50% of all pharmacies), which is another piece of evidence that the unlicensed pharmacy set and the blacklisted pharmacy set have roughly similar properties.

In short, we do observe fairly large concentrations in similar inventories. This confirms that unlicensed pharmacies operate with a relatively small set of suppliers. From an intervention standpoint, this is good news: if the few factories

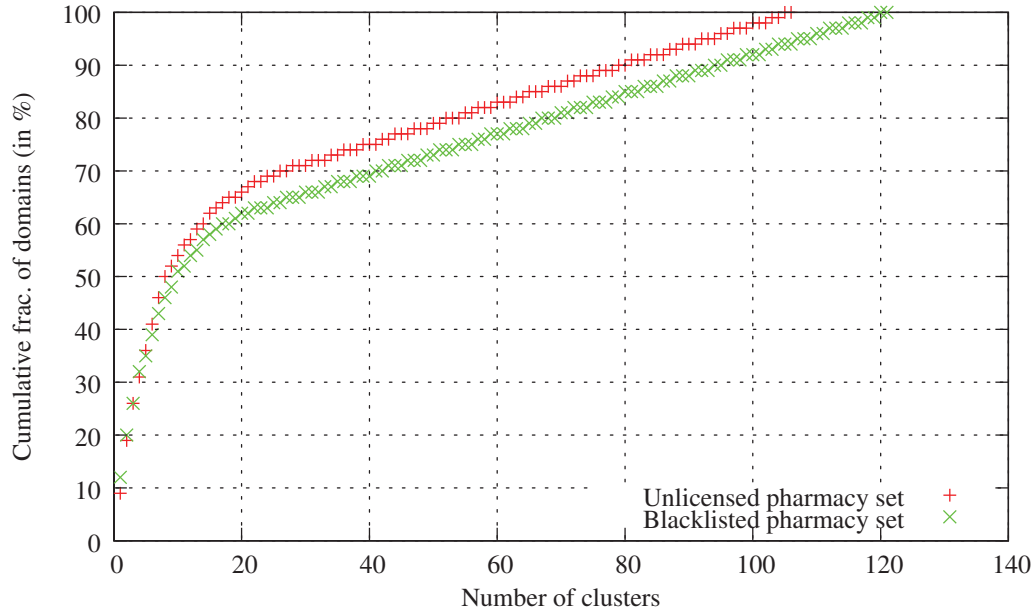


FIGURE 5.5: Cumulative distribution of pharmacies as a function of the number of clusters considered (Using average-linkage, $t = 0.31$).

supplying these drugs can be subject to more stringent controls, potential harm will be greatly reduced.

5.4 Pricing strategies

We now turn our attention to the product prices offered by the sets of pharmacies we are studying. We measure the price variation from one type of pharmacy to another, and we look at factors that might be affecting it.

5.4.1 Pricing differences by seller and drug characteristics

Table 5.5 summarizes several price differences we examined. In the first test, we confirmed that prices are considerably cheaper at illicit pharmacies than at *familymeds.com*. For this test, we compare the prices for drugs that are available at *familymeds.com* and an unlicensed pharmacy when a direct comparison is possible—that is, when both pharmacies sell the drug name at the same dosage

Table 5.5: Unit price discounts for different drug categories.

Difference (median)	95% C.I.	Sig.?	# Records
<i>familymeds.com- unlicensed pharmacy price</i>			
\$2.14	(\$2.12, \$2.17)	✓	171,098
<i>Fake generic (illicit price discount)</i>			
\$3.14	(\$3.09, \$3.18)	✓	41,669
<i>Drug in shortage (illicit price discount)</i>			
\$0.72	(\$0.59, \$0.85)	✓	3,966
<i>Popular drugs (illicit price discount)</i>			
\$0.36	(\$0.32, \$0.41)	✓	95,308
<i>Silk Road - unlicensed pharmacy price</i>			
-\$0.46	(-\$0.54, -\$0.37)	✓	3,821

and in the same number of units (e.g., a 10-pack of Lipitor 10mg pills costs \$8.99 at pills4everyone.com, compared to \$41.90 at *familymeds.com*). We normalize all prices to the per-unit price (e.g., the aforementioned pills cost \$0.89 each at pills4everyone.com, a discount of \$3.30 compared to the \$4.19 at *familymeds.com*).

Overall, the median difference in per-unit prices between *familymeds.com* and unlicensed pharmacies is \$2.14. This difference is statistically significant at greater than the 0.01% level according to the Mann-Whitney U-test, while the 95% confidence interval is (\$2.12, \$2.17). Thus, we can safely conclude that unlicensed pharmacies are a lot cheaper than at least one legitimate alternative.

We are also interested in whether any other characteristics of the drugs on sale might influence the magnitude of the pricing advantage. To that end, we next study differences in the size of discount offered by unlicensed pharmacies relative to *familymeds.com*. One common deceptive tactic employed by unlicensed pharmacies is to offer “generic” versions of drugs where no such generic exists (e.g., because the patent is still in effect). We found around 42,000 such “fake generic” discrepancies in our dataset. The median per-unit price discount for fake

Table 5.6: Unit prices and percentage discounts offered by *familymeds.com* and unlicensed pharmacies for 60-pill and 90-pill orders relative to the unit price of 30-pill orders.

	30 pills		60 pills discount				90 pills discount				
	unit price	unit price	\$	95% CI	Sig. diff.?	%	unit price	\$	95% CI	Sig. diff.?	%
<i>familymeds.com</i>	\$3.86	\$3.86	\$0.00	(\$0.00,\$0.00)		0%	\$3.86	\$0.00	(\$0.00,\$0.00)		0%
unlic. pharm.	\$1.77	\$1.60	\$0.16	(\$0.15,\$0.18)	✓	10.0%	\$1.48	\$0.27	(\$0.25,\$0.29)	✓	16.9%

generics is \$3.14, compared to a \$1.70 discount for other drugs not mislabeled as generic. The Mann-Whitney U-test estimates a median difference of \$1.54 in the discount for fake generics. This suggests that deceiving customers with promises of branded generics can be financially enticing. We also find smaller, yet still statistically significant price discounts for drugs in shortage and those identified in WebMD as “top drugs”.

How do prices compare between unlicensed pharmacies and drugs sold on Silk Road? While Silk Road has become notorious for selling narcotics even though other unlicensed pharmacies do not, sellers on Silk Road also offer non-narcotics for sale, many of which can also be bought from unlicensed pharmacies. Overall, drugs found on Silk Road are \$0.46 cheaper per unit than their unlicensed counterparts. This is somewhat surprising, given that privacy-concerned customers drawn to Silk Road might have been expected to be willing to pay a premium for purchasing anonymity.

5.4.2 Volume discounts as competitive advantage

Another way for unlicensed pharmacies to entice prospective customers is to offer discounts when buying at higher volumes. We examined the prices of drugs offered in both *familymeds.com* and unlicensed pharmacies at the same dosage and number of units. Of the 171,098 matching tuples, 156,136 (91%) offered 30, 60, or 90 pills. These drugs were offered by 221 unlicensed pharmacies, 83% of the total. We therefore focus our analysis on only these drugs.




FIGURE 5.6: Cumulative distribution functions of the median percentage-point price discount per pharmacy (left) and per drug (right).

Note that a single unlicensed pharmacy can sell same combination of drug, dosage and units at several prices. This happens for two reasons. First, the drug may be sold in different currencies. Second, the pharmacy may sell multiple variants of the same drug, e.g., Super Viagra, at different prices. To simplify comparison, for every pharmacy and drug, dosage and unit combination, we compute the median of all per-unit prices. For example, `rx-pharm-shop.com` sells a 30-pack of Viagra 10mg in USD, GBP and EUR in different varieties, totaling 9 different prices. Its median per-unit price is \$2.74, falling to \$2.44 for 60-pack prices and \$2.06 for 90-pack prices. In total, we observe median per-unit 30, 60, and 90-day prices for 20,124 distinct drug-dosage-pharmacy combinations.

We check for discounts in the per-unit prices of 60- and 90-unit supplies relative to the per-unit price of a 30-unit supply. Table 5.6 presents our findings. First, we never observed a per-unit discount on drugs from *familymeds.com*. By contrast, the 221 unlicensed pharmacies offered a median discount of 10%-pts. for 60-day supplies, rising to a 16.8% pt. discount for 90-day supplies. The unit price on unlicensed pharmacies falls from \$1.77 on 30-day supplies to \$1.60 for 60-

day supplies and \$1.48 for 90-day supplies. These discounts are found to be statistically significant according to the Mann-Whitney U-test.

But how do the discounts vary by pharmacy? The left plot in Figure 5.6 presents a CDF of the median percentage point discount offered by each unlicensed pharmacy. We can see that over 80% of pharmacies offer a discount of at least 7%. While volume discounts are the norm, around 15% of pharmacies actually charge more per-unit for larger volumes, which is surprising since a consumer could simply buy multiple 30-unit supplies instead. The discounts are consistently greater for 90-unit supplies than for 60-unit supplies. Finally, a few pharmacies offer very deep discounts at higher volume—around 5% of pharmacies offer median discounts exceeding 15% for 60-day supplies and 25% for 90-day supplies.

We also observe substantial variation in discounting according to the drug sold, as shown in Figure 5.6 (right). The median discount for drugs is 11.3% for 30-day supplies and 18.8% for 90-day supplies. However, the 10% most deeply-discounted drugs save at least 20% for 60-day supplies and 29% for 90-day supplies. We conclude that unlicensed pharmacies can use volume discounting as a way to attract prospective customers, particularly as the tactic may not be used widely by legitimate pharmacies.

5.4.3 How competition affects pricing

We have already seen that unlicensed pharmacies adjust prices strategically in order to attract customers, ranging from discounting volume sales to offering fake generics. They must also react to competition from other unlicensed pharmacies. Some common drugs are sold by nearly all the pharmacies we studied, while other more obscure drugs are sold by just a few. Microeconomic theory predicts that competition drives prices down; we now examine whether prices set by unlicensed

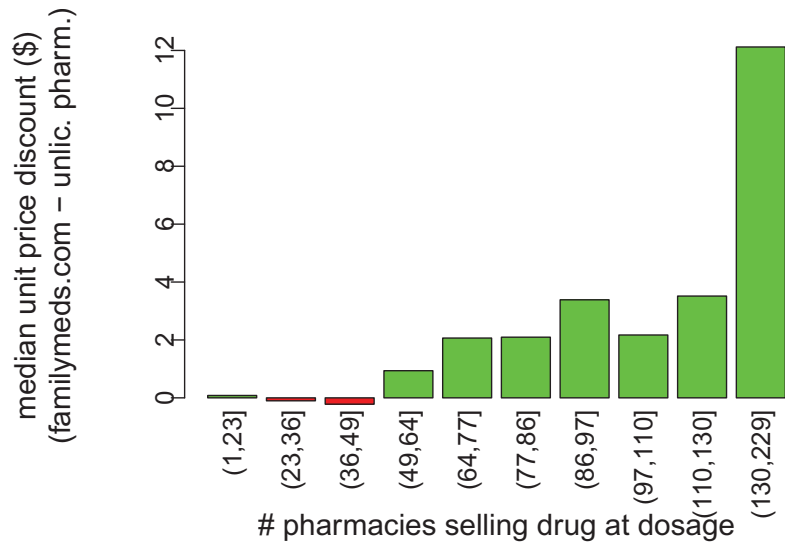


FIGURE 5.7: Bar plot of the median unit price discount for drug-dosage combinations grouped in increasing number of unlicensed pharmacies selling the drug at the specified dosage.

pharmacies do in fact fall when competition among sellers is high and rise when competition is low.

To answer this question, we examine all prices for combinations of drugs and dosages. We normalize each price by the number of units sold at the drug's dosage and then compare the median normalized prices offered at *familymeds.com* and unlicensed pharmacies. We compute the price difference between *familymeds.com* and each unlicensed pharmacy selling a drug-dose combination. We then compute the median of this difference across all pharmacies selling that drug-dose combination. For example, the following 7 pharmacies sell Mirapex 1mg at these prices per pill:

Pharmacy	unit price (unlic. pharm.)	unit price (<i>familymeds.com</i>)	discount
yourhealthylife.cc	\$4.37	\$4.57	\$0.20
drugs-medshop.com	\$1.89	\$4.57	\$2.67
pharmaluxe.com	\$4.14	\$4.57	\$0.43
24medstore.com	\$4.00	\$4.57	\$0.56
online-canadian-drugshop.com	\$4.37	\$4.57	\$0.20
safetymedsonline.com	\$4.37	\$4.57	\$0.20
7-rx.com	\$1.86	\$4.57	\$2.71
Median	-	-	\$0.43

We then check whether the number of unlicensed pharmacies influences the median discount offered. We group the drug-dosage combinations into deciles according to how many pharmacies sell them. Figure 5.7 plots the median discount offered for each decile. We can see that less popular drugs offer a very small discount, and sometimes even charge slightly more than *familymeds.com* does. However, as more pharmacies sell the drug, competition drives pharmacies to sell at a higher discount relative to the price charged by *familymeds.com*. For example, the median discount for drugs sold by 87–97 pharmacies is \$3.39. The discount rises to \$12.12 for the 10% most popular drugs.

Discounts in the blacklisted pharmacies. We performed a similar analysis to check for a significant difference between the discounts in the two sets of unlicensed pharmacies. The result of Mann-Whitney U-test showed that the difference in the observed discounts between the unlicensed pharmacy set and the blacklisted pharmacy set are statistically insignificant. In other words, the observation of discounts for volume purchases is not limited to the main set of unlicensed pharmacy set, and is not caused by any measurement bias. On the contrary, the discounting phenomenon is characteristic of all unlicensed pharmacies.

5.5 Conclusions

Unlicensed pharmacies circumvent legal requirements put in place to protect consumers from physical harm. But they operate in the context of a broader ecosystem where consumers can choose among licensed pharmacies, unlicensed pharmacies and anonymous contraband marketplaces. Consequently, unauthorized pharmacies must offer a compelling reason for consumers to do business with them instead of more legitimate alternatives. One approach is for unlicensed operators to fake legitimacy through clever website design and deception. The web suffers from asymmetric information—it can be very hard for the average consumer to distinguish good websites from bad. Licensed pharmacies combat this with certification schemes such as VIPPS and LegitScript. But the findings of this Chapter suggest that seals cannot do the job on their own.

Unauthorized pharmacies are already competing hard by offering deep inventories and discount prices. Inventories at unlicensed pharmacies can rival those at licensed pharmacies, and can be more extensive for certain classes of drugs (e.g., schedule drugs). We have shown evidence of sophistication in how prices are set by unlicensed pharmacies. While cheaper across the board than at a reference licensed pharmacy, unlicensed pharmacies also employ deceptive tactics such as fake generics to attract customers, in addition to more straightforward volume discounts.

So what interventions are available to counter online pharmacies more effectively? One option is to devote more resources to blacklisting unlicensed pharmacies. Unfortunately, blacklists only offer a partial solution, since online criminals have shown resilience in changing web domains rapidly in many contexts. One promising option we found is to cluster pharmacies by their inventories in order to identify a smaller number of suppliers. Shutting down pharmacy websites is

futile, since the cost to criminals of setting up new sites is too low compared to the cost of take-down. Disrupting supply chains, on the other hand, could be much more cost-effective. In Chapter 9, we explore this solution in a systematic way, evaluating and comparing its effectiveness in the context of a comprehensive set of situational prevention measures.

6

A longitudinal analysis of search-engine poisoning

The previous two chapters offer empirical insights on the online criminals networks and components supporting an end-to-end monetization of the fraudulent trade of prescription drugs. In this chapter, we set off to display the long-term effects of online criminal activity, when enforcement is simply misplaced [130]. We investigate the evolution of search-engine poisoning using data on over 5 million search results collected over nearly 4 years. We build on our prior work investigating search-redirection attacks, where criminals compromise high-ranking websites and direct search traffic to the websites of paying customers, such as unlicensed pharmacies who lack access to traditional search-based advertisements. In addition, we overcome several obstacles to longitudinal studies by amalgamating different resources and adapting our measurement infrastructure to changes brought by adaptations by both legitimate operators and attackers. Our goal is to empirically characterize how strategies for carrying out and combating search poisoning have evolved over a relatively long time period.

Such search-engine result poisoning has been getting increased attention from the research community since 2011, when we published [128] our work in Chapter 4. Researchers have attempted to measure and describe specific campaigns [115, 142, 240], infection techniques [22, 132], or even economic properties [148]. For instance, Levchenko et al. [132] focus primarily on email spam, but also provide some insights on “SEO” (search-engine optimization) by people involved in the online trade of questionable products. A follow up work by the same group [148] analyzed the finances of several large pharmaceutical “affiliate networks” and provided evidence that search-result poisoning accounted for a non-trivial part of the traffic brought to these pharmacies. Most of the aforementioned studies tend to either describe phenomena observed on relatively short time-spans (e.g., volume of orders at online pharmacies measured over a period of a few weeks [117]), or to describe longer-term activities of specific actors (e.g., specific pharmaceutical affiliate networks [148], or a specific search-engine optimization botnet [240]).

While originally the compromised sites participating in search-redirectation attacks did little more than simply send HTTP 302 redirects (Section 4.1.1), they have evolved toward more complex and evasive forms of redirection, apparently in response to deployed defenses from search engines. For instance, in Chapter 5 we describe how a more modern search-redirectation variant uses cookies to store state, in order to look innocuous to web crawlers while still actively redirecting users behind a “real” browser. We also explain that attackers increasingly host “store fronts” under hidden directories in the compromised webserver as shown in Figure 4.1 (second result). Borgolte et al. [22] describe more recent advances in redirecting techniques, in particular JavaScript (JS) injections that are particularly hard for crawlers to detect. Li et al. [134] describe techniques to detect these JS

injections, and show that JS injections often are used to support a peer-to-peer network of compromised hosts distributing malware.

Coming from a different angle, a recent paper by Wang et al. [239] explores the effect of interventions against search-poisoning campaigns targeting luxury goods, both by search-engine providers who demote poisoned results and by brand-protection companies enforcing intellectual property law by seizing fraudulent domains.

Different from the previous work, our empirical analysis in this chapter is the first to look at data on such a large scale and over a long time period. This in particular allows us to observe trends in how attackers and defenders have been adapting to each others' strategies over the years. In addition, it provides us with interesting insights on the criminal ecosystem that facilitates abuse. We combine multiple data sources to gain insights into the long-term evolution of search-engine poisoning. With a primary focus on how unlicensed pharmacies are advertised, we analyze close to four years (April 2010-September 2013) of search-result poisoning campaigns. We investigate how the composition of search results themselves have changed. For instance, we find that search-redirection attacks have steadily grown to take over a larger share of results (rising from around 30% in late 2010 to a peak of nearly 60% in late 2012), despite efforts by search engines and browsers to combat their effectiveness. We also study the efforts of hosts to remedy search-redirection attacks. We find that the median time to clean up source infections has fallen from around 30 days in 2010 to around 15 days by late 2013, yet the number of distinct infections has increased considerably over the same period. Finally, we show that the concentration of traffic to the most successful brokers has persisted over time. Further, these brokers have been mostly hosted on a few autonomous systems, which indicates a possible intervention strategy. We do not focus on a specific campaign or affiliate network, but

instead analyze measurements taken from the user's standpoint. In particular, we study what somebody querying Google for certain types of products would see. While we focus here on Google due to their dominance in the US web search market [40], previous work (e.g., [22]) showed other search engines (e.g., Yandex) are not immune to search-result poisoning.

The analysis we present here has three primary objectives. First, we describe the relationship between attackers' actions and defensive interventions. We are notably interested in identifying the temporal characteristics of attackers' reactions to defensive changes in search-engine algorithms. At the same time we describe the long term structural characteristics of online criminal networks primarily in the illicit online prescription drug market, and in the illicit online market of: (i) counterfeit applications and antivirus software, (ii) books (iii) gambling, and (iv) counterfeit watches. Second, we aim to determine whether, over a long enough interval, we can observe changes in attitudes among the victims. For instance, are compromised sites getting cleaned up faster in 2013 than they were in 2010 (Section 4.4.2)? Have defenders been trying to target critical components of the infrastructure search-result poisoning relies on? In this regard we present evidence of the persistence of the specific criminal networks over the years, regardless of the domestic and international efforts against illicit online pharmacies. Third, we want to better understand the long-term evolution of the thriving search-poisoning ecosystem, notably in terms of consolidation or diversification of the players. All these objectives are essential in addressing research question 2,¹ among others.

¹ Is the observed structure of online criminal networks an ephemeral phenomenon. . .

6.1 Background

Conceptually, there are three distinct components to a successful search-redirection attack (Section 4.1.1): (i) *Source infections* are sites that have been compromised to participate in a search-redirection campaign. Their owners frequently do not suspect a compromise has taken place. These source infections are the sites that appear in search-engine results to queries for illicit products. Source infections redirect to an optional intermediate set of (ii) *traffic brokers*. The (set of) traffic broker(s) ultimately redirects traffic to a (iii) *destination*, typically an illicit business, e.g., an unlicensed pharmacy when entering pharmaceutical search terms or a distributor of counterfeit software when entering software-related terms.

Among source infections, we can distinguish between results that *actively redirect* at the time t of the measurement; *inactive redirects*, i.e., sites that used to be redirecting at some point prior to t but are not redirecting anymore—possibly because they have been cleaned up, but have not yet disappeared from the Google search results; and *future redirects* that appear in Google search results at time t without redirecting yet, but that will eventually redirect at a time $t' > t$. Presumably those are sites that have been compromised and already participate in link-farming [88], but have not yet been configured to redirect.

As described above, the technology behind search redirections has evolved over time. In this chapter, active redirects include fully automated redirections by HTTP 302, as well as “embedded storefronts,” which result on HTTP 302 redirects when a link is clicked on. Other types of redirections, such as JS-based redirects, or HTML “Refresh” meta-tags, could also be considered as active redirects, but we will treat these separately.

6.2 Data collection

Besides the time-consuming nature of such an endeavor, collecting nearly four years' worth of data is in itself a complex process. Software and APIs used to acquire the data change over time, attackers' techniques evolve, and new defensive countermeasures are frequently deployed. In other words, the target of the measurements itself changes over time. Thus, we must rely on several distinct sources of data for our analysis. Because of the heterogeneous nature of these datasets, not all the data available can be used for all the analyses we conduct here. We first characterize the queries used to produce these different datasets, then the contents of the datasets, and finally our methodology to combine the datasets.

6.2.1 Query corpus

The corpus of queries we use has a considerable influence on the results we obtain. Owing to the prevalence of the trade of pharmaceutical products among search-engine poisoning activities, we use a primary set of queries Q related to drugs. We complement this first set with queries related to other types of goods and services routinely sold through abusive means: luxury counterfeit watches, software, gambling, and books. We refer to this second query set as Q' .

Drug-related queries. For our set of drug-related queries, we use the set Q of 218 queries we defined in Section 4.2.2. There are two reasons for that choice. First, using an identical query set allows us to produce directly comparable results, and expand this relatively short-term initial analysis. Second, we have showed that this relatively small set of queries provides adequate coverage of the entire online prescription drug trade.

Other queries. We construct an additional query corpus Q' composed of an extra 600 search terms. We create and track Q' to provide evidence that search-

Table 6.1: **Datasets for pharmaceutical queries.** Dataset 1 only contains search results and no ranking information. Dataset 2 contains search results and overall rankings, but no individual rankings per query. Dataset 3 contains everything we need, but only for a strict time-varying subset of all queries.

Dataset	1	2	3	4
Period covered	T_1 4/12/2010–11/15/2010	T_2 11/15/2010–10/08/2011	T_3 10/08/2011–9/16/2013	
Queries used	Q	Q	$Q(t) \subseteq Q$	$Q'(t) \subseteq Q'$
Search results/query	64	64	16 to 32	
Ranking info?	No	Aggregate only	Yes	
Mapping queries-results	No	Partial	Yes	
Total size of result corpus	260,824	3,609,675	1,530,099	2,244,723
Unique URLs in results	150,955	189,023	122,382	122,567
Unique domains in results	25,182	36,557	30,881	24,339
Total size of redir. corpus	50,821	929,809	522,017	111,361
Unique redir. URLs	50,784	71,935	62,288	27,973
Unique redir. domains	5,546	8,738	11,157	3,974

poisoning is not strictly tied to pharmaceutical terms, and to study whether or not miscreants share parts of their infrastructure to advertise different products and services. Q' consists of six categories: antivirus, software (in general), pirated software, e-books, online gambling, and luxury items (specifically, watches). We choose these topics based on the amount of email spam we have received in spam traps we are running. For each category, we use Google’s Keyword Planner to select the 100 most queried keyword suggestions associated with the category name. Except for pirated software queries, we manually filter out queries that do not denote benign or gray intent.

6.2.2 Search result datasets

We use data collected on a daily basis between April 12, 2010, and September 16, 2013. Each dataset has its own particularities, summarized in Table 6.1, which we discuss next.

Dataset 1 (4/12/2010-11/15/2010). This first dataset represents data collected daily between April 12, 2010 and November 15, 2010 (time interval T_1), and was initially used for the analysis in Chapter 4. The data contains daily search results for the pharmaceutical query corpus Q , without preserving any ranking information, beyond noting that only the top-64 results—at most—are collected. Likewise, the redirection corpus contains all the sites visited (including “redirection chains”) at a given time t , but those are not mapped to specific queries. In other words, if two queries q_1 and q_2 produce results $\{u, v, w\}$, we do not know which of q_1 or q_2 yielded each of u, v, w , nor how u, v and w ranked among all search results. Redirections in this first corpus are only gathered by following HTTP 302 redirects.

Dataset 2 (11/15/2010-10/09/2011). The second dataset spans from November 15, 2010 through October 8, 2011, and was used partially in the analyses presented in Chapters 4 and 5. Different from Dataset 1, this dataset contains information about the search rankings for the pharmaceutical query corpus. Here again, only the top 64 results per query are collected. We furthermore have the mappings between a given query and the results it produces, but, regrettably, not the full mapping between a given query, its results, and the ranking of the results. Going back to our previous example, for two queries q_1 and q_2 , we know that q_1 yielded (u, v) and q_2 yielded (v, w) , and we know the ranks at which each result appeared overall, but we do not know if v appeared as the top result in response to q_1 or q_2 . Here too, redirections are gathered by following HTTP 302 redirects.

Dataset 3 (10/13/2011-9/16/2013). The third dataset was collected specifically for the analysis we perform in this chapter, and we make it publicly available for reproducibility purposes.² With this dataset have the complete mapping between a

² See <https://arima.cylab.cmu.edu/rx/>.

query, the results it produces and their associated rankings, as well as the possible redirection chains that follow from clicking on each result.

Our collection infrastructure is markedly different from that used for Datasets 1 and 2. Datasets 1 and 2 were assembled by having a graphical web browser run the queries against Google's search engine. Here, we use an automated (command-line) script, increasing the level of automation in collecting search results.

Because attackers are known to perform *cloaking*, that is, to make malicious results look benign when suspecting a visit from an automated agent as opposed to a customer, we periodically spot-checked the results our automated infrastructure collection gathered with what a full-fledged graphical browser would obtain. In addition, we ran all of our queries over the Tor network [58], changing Tor circuits frequently. This had two effects: we obtained geographical diversity in the results since queries were apparently issued by hosts in various countries; and we escaped IP-based detection (and potential identification), which is frequently used as a decision to cloak results. We were worried that, because Tor exit IP addresses are well-known, they could be subject to cloaking as well. Spot-checking the results we obtained by comparing results from Tor exits as opposed to non-Tor exits did not yield any significant indication this was the case. In short, if unlicensed pharmacy operators are aware of the existence of Tor, they seem to tolerate people connecting over the Tor network, perhaps because some of their intended customers desire anonymity.

Regrettably, on November 30th 2011 the Google API introduced certain restrictions, reducing both the number of queries we could run on a daily basis, and the number of search results we could collect per query.³ These restrictions came

³ Recent research, e.g., [22], uses the Yandex search engine instead of Google search in an apparent effort to overcome some of the limitations of the Google API. For the sake of comparability

one year after Google announced the deprecation of the Search API, giving it a phasing out period of three years.⁴

The upshot is that we could only run a random strict subset of Q on a daily basis. The size and composition of the query set varies over time, but, on average, consists of 64 queries. Likewise, instead of collecting $N = 64$ results per query, we were limited to between $N = 16$ and $N = 32$.

We refer as T_3 the collection interval over which we collected this dataset. During the collection of this third dataset, on April 9, 2012, we updated our collection infrastructure. Instead of simply considering redirections characterized by HTTP 302 messages, our crawler became able to detect more advanced (cookie-based) redirection techniques, as described in Section 6.1. We did not observe “Refresh” META tag redirections. We also realized that we can never be sure that we are able to detect all forms of attacks, as attackers always deploy new attack variants. To address this limitation, we elected to capture the first 200 lines of raw HTML content present at each source infection, using *both* a user-agent string denoting a search-engine spider and a user-agent string denoting a regular browser. The data so captured can then be analyzed after the fact to determine if there was cloaking, and to attempt to reverse-engineer types of attacks that were unknown at data collection time. For instance, while our crawler was not able to detect JavaScript-redirections at data collection time, we were ultimately able to analyze how prevalent they were in our data corpus.

Dataset 4 (10/31/2011-9/16/2013). This dataset has the same properties as Dataset 3, but uses the query set Q' . As with Dataset 3, the number of actual

with Datasets 1 and 2, and also because it appears that search-redirection attacks primarily target the Google search engine, we continued to use the Google API.

⁴ <http://googlecode.blogspot.com/2010/11/introducing-google-apis-console-and-our.html>

queries $Q'(t)$ issued every day is a varying subset of Q' . On average, 64 queries per day are issued for each category (gambling, watches,...).

Finally, given the long term nature of measurements, there are periods with incomplete or no daily measurements. These measurements gaps are attributed to glitches with the measurement equipment (e.g. power or network outage), or upgrades to the measurement infrastructure. Out of the 1,254 days in the measurement period, we have complete measurements for 1,004 days.

6.2.3 Combining the datasets

Since, in Datasets 3 and 4, all mappings between queries, results, and rankings are recorded, as well as more complete redirection information, we can carry out more in-depth analysis than with the first two datasets. On the other hand, the reduced number of queries used and results collected per query makes it slightly more complicated to combine Dataset 3 with Datasets 1 and 2. (Dataset 4 concerns a different set of queries, and as such does not need to be combined with the other datasets.)

It also means that we cannot necessarily claim to have the same desirable coverage properties, as described in Section 4.2.2. However, we can attempt to combine all datasets to obtain results over the entire collection interval; this essentially consists of sampling some of the queries and some of the results in Datasets 1 and 2 to match the statistical properties of Dataset 3.

Sampling queries. In Datasets 1 and 2, for all t , the whole set Q of queries is issued. In Dataset 3, a different random subset $Q(t) \subsetneq Q$ of all queries is used every day. Within that subset, the proportion of illicit $I(t)$ and benign $B(t)$ queries follows the Beta distribution with parameters $(\alpha = 22.49, \beta = 194.29)$. The proportion of gray queries $G(t)$ follows the normal distribution with parameters $(\mu = 0.57,$

$\sigma^2 = 0.03$). Because these results are slightly different from the proportions in Q (see Table 4.2), we also need to sample from Q in the first two datasets to be able to perform meaningful comparisons when looking at the entire measurement interval. Unfortunately, as there is no association between individual queries and results in Dataset 1, we may only be able to use Datasets 2 and 3 when looking at metrics for which the specific types of queries used has importance. Given the known expected probabilities of $I(t)$, $B(t)$, and $G(t)$ in Dataset 3, we create samples of queries for each day in T_2 that follow the same distributions. In turn, we consider only the daily results in Dataset 2 associated with each daily query sample.

Sampling results. Dataset 3 (and 4) is often limited to $N = 32$ results, while Datasets 1 and 2 contain the top-64 results for each query. Arguably, from a user standpoint, the difference is minimal: Given that the probability of clicking on a link decreases exponentially with its position in the search results [114], results in position 33 and below are unlikely to have much of an impact. Unfortunately, Dataset 1 does not contain any ranking information; as such we cannot use it for direct comparisons with Dataset 3 in terms of search-result trends. We can, however, use Dataset 1 when we are only concerned about measuring how long certain hosts appear in the measurements (e.g., for survival analysis).

Dataset 2, on the other hand, contains some ranking information. From the above discussion, for each result we obtained, we know what was its ranking at the time; there may however be uncertainty as to which query produced that result when results occur in response to more than one query. We include each result u with a probability $p(u)$ corresponding to the number of times u appears at a rank below 32 divided by the total number of times u appears in the whole dataset. That is, (i) results that never appear in the top-32 results are always excluded ($p = 0$),

(ii) results that always appear in the top-32 results are always included ($p = 1$), and (iii) results appearing both in and out of the top-32 results are included with a probability dependent of how often they are in the top 32.

Combining query and result sampling, we use approximately 14.7% of the search results in Dataset 2. Another 12.3% appear both in ranks 1–32 and above 32 and are probabilistically included.

6.3 Search-result analysis

We now turn to analyzing the datasets we have, and first look at the evolution of search results over intervals T_2 and T_3 (November 2010 through September 2013), corresponding to Datasets 2 and 3.⁵ We start with an analysis of the whole interval, before looking into the dynamics of the search results.

6.3.1 Overview

We focus here on pharmaceutical goods, where we identify several different categories of search results issued in response to queries containing drug names. For the sake of comparison, we use some of the definitions provided in Section 4.2.5, extending this taxonomy whenever required.

Licensed pharmacies, are those online pharmacies having been verified by Legitscript [124].

Health resources, associated with (usually benign) websites, and providing information about drugs. We use information from the Open Directory Project [11] to make that determination.

Unlicensed pharmacies, characterized as such by Legitscript and directly appearing in the organic search results.

⁵ Recall that the information available from Dataset 1 is too coarse to be useful in this section.

Table 6.2: **Search-result composition.** Results collected between November 2010 and September 2013.

Result category	% of results	Range (%)	# of results
Active search-redirectation	38.8	[8.7, 61.7]	621,623
Unclassified	18.8	[6.3, 35.4]	300,427
Unlicensed pharmacies	16.9	[12.1, 30.1]	271,045
Health resources	7.7	[4.2, 14.5]	123,883
Blog & forum spam	7.1	[3.0, 16.4]	113,250
Content injection (compromised)	4.7	[1.9, 10.0]	74,556
Future search-redirectation	4.1	[0.0, 6.7]	65,548
Inactive search-redirectation	1.8	[0.0, 10.6]	28,976
Licensed pharmacies	0.2	[0.0, 0.9]	2,779
Total			1,602,087

Content injection (blog and forum spam), which point to discussion websites with drug-related spam posts. We identify such sites through URL parameter names they commonly use—containing terms such as “blog,” or “forum” for instance.

Search-redirectations, as defined in Section 6.1. Domains in this category have generally nothing to do with prescription drugs and are merely used as a feed to online pharmacies.

Content injection (compromised), which represent websites other than blogs and forums, in which an attacker injected drug-related content, but never exhibit signs of search-redirectation. For this category, we consider the characteristics of URLs that are search-redirecting with embedded storefronts; The Fully Qualified Domain Names (FQDNs) contain no drug- or pharmacy-related keywords, while the trailing paths do. We then apply this heuristic to the set of results not placed in any of the previous five categories.

Finally, we mark as **unclassified** sites that do not fit into any of the above categories.

Table 6.2 shows the breakdown of results in each category over the roughly three years that T_2 and T_3 span. We combine Datasets 2 and 3 by sampling Dataset 2 as described in Section 6.2. In the end, we examine 1,602,087 search results over the entire interval. Out of those, more than 38% are active redirections; on any given day between 8.7% and 61.7% of the obtained results actively redirect. Inactive and future redirects represent another 5.9% altogether, while blog and forum spam, and compromised sites, taken together, account for another 11.8%. Shortly stated, the vast majority of results are illicit or abusive. Particularly telling is the fact that legitimate pharmacies only consist of 0.2% of the entire results!

The fairly large proportion of “unclassified” results (18.8% of all results) led us to further examine them. Unclassified results may be (i) benign websites with information about drugs, (ii) malicious websites (compromised or redirections) that we failed to identify as such, or (iii) results only marginally related to the search query. We need to obtain the contents of these sites rather than their mere URL to make this determination. By using the Internet Archive Wayback Machine [213], we attempted to access the content of all 45,213 unclassified results collected in 2013. We managed to find matches archived roughly at the time of our own crawls for 41,547 of them. 14,993 (33.1%) of the examined unclassified results did in fact contain drug-related terms, which is an indication that a non-negligible number of unclassified results may actually present some different form of illicit behavior.

6.3.2 *Search result dynamics*

In Figure 6.1, we examine how search results, which appear to be dominated by malicious links, dynamically evolve over time. The graph shows, as a function of time, the proportion of results belonging to each category, averaged over a 7-day sliding window. Vertical lines denote events of interest that occur during data




FIGURE 6.1: **Percentage of search results per category**, averaged over a 7-day sliding window. Minor categories are excluded. The vertical lines correspond to documented changes in search-engine behavior ($G1$, $G2$, $G3$), browser behavior ($B1$, $B2$, $B3$) and in our own collection infrastructure ($C1$, $C2$).

collection. In particular, $C1$ corresponds to the switch from Dataset 2 to Dataset 3, and $C2$ corresponds to an update in our crawler to detect more advanced types of search-redirections. From late 2010 through late 2012 active redirects have not only been dominating the search results, but they have also been steadily growing to a peak of nearly 60%. Meanwhile, unclassified results are decreasing overall, unlicensed pharmacies remain stable around 15–20%, and licensed pharmacies constantly hover near zero. Spam contents seems to marginally decrease until late 2012 as well.

Then, in early 2013 we notice a change in trends: active redirections seem to finally decrease somewhat steadily, while, on the other hand content injection (both spam and compromised websites), as well as unclassified results enjoy a bit of a resurgence. Even more interestingly, we also observe that unlicensed pharmacies mirror very closely the trend of active redirections in 2013. Whenever redirections become more frequent, direct links to unlicensed pharmacies become rarer, and vice versa. This suggests that attackers use direct links to pharmacies as a kind of alternative to search redirections.

Search-engine interventions. The lines marked $G1$, $G2$ and $G3$ correspond to documented changes in search-engine behavior. We examine their impact on the search results using the Mann-Whitney non-parametric U-test of significance, and data we collected within 30 days before and after each event.

On February 23, 2011 ($G1$) Google deployed an improved ranking algorithm to demote low quality search results [203]. This apparently caused a statistically-significant drop in redirecting results by 2.3% ($p = 0.003$), and by 2.7% for spam websites ($p < 0.001$). However, the improvement was only transient: Starting in May 2011 we observe a sharp increase until August 2011 in the proportion of results that are actively redirecting. Specifically, the median difference in the proportions of redirecting results collected in April and in June of 2011 shows an increase by 15.5% ($p < 0.001$). Apparently, after being initially impacted, attackers managed to find countermeasures to defeat Google’s improved ranking algorithm.

Between October 2011 ($G2$, [118]) and March 2012 ($G3$, [164]), Google updated its service again to gradually remove information from the HTTP Referrer field about the query that produced the result. In theory, this should have reduced active redirects, which originally relied primarily on the Referrer information to determine how to handle incoming traffic. In practice, the effect was non-existent, as redirects continued increasing in the time interval $G2$ – $G3$. Indeed, comparing the proportion of results identified as redirecting within 30 days before $G2$ and 30 days after $G3$, we find a statistically-significant median increase by 9.9% ($p < 0.001$). Here again, attackers seem to have been able to adapt to a countermeasure from the search engine. Furthermore, since Google announced the change well in advance of its implementation in order to accommodate the many legitimate websites affected by the change, those perpetrating poisoning attacks also had plenty of time to adapt before being adversely impacted.

Browser evolution. A series of major changes to Internet browsers occurred in the second half of 2012 and beginning of 2013. On July 17, 2012 (*B1*) Firefox 14 was released. This was the first major browser (roughly 25% of reported market share at the time according to StatCounter) to use secure HTTP (HTTPS) search by default, which only lists the previous domain (but no URL parameters) in the Referrer. On September 19, 2012, Safari followed suit (*B2*); and on January 13, 2013, Google Chrome, the browser with the dominant market share also switched to HTTPS search (*B3*). At that point, the majority of desktop browsers were using HTTPS search by default. Perhaps coincidentally, we started observing a stagnation and eventual decrease in the number of active redirections. While we emphasize we cannot affirm causality, a plausible explanation is that traditional, simple Referrer-based redirection techniques, by early 2013, stopped working for a large proportion of the population, which led to alternative techniques being used (e.g., cookie-based redirections). We periodically still see some large spikes (e.g., in early Summer 2013), perhaps attributable to short-lived campaigns. We conversely observe an increase in “direct advertising” of unlicensed pharmacies.

Undetected infections. An alternative explanation for the plateauing and decrease of search-redirections observed since early 2013 might be that attackers’ tactics have evolved, and are not captured by our crawlers anymore. To determine whether that is the case, we take a closer look at the “unclassified” category. Recall, that from April 2012 (*C2*) through the end of our measurement interval, we record the first 200 lines of HTML code of each source infection, posing both as a search-engine spider, and as a regular browser. When we observe a difference in the HTML returned between the two treatments, we infer there might have been cloaking.




FIGURE 6.2: Percentage of unclassified search results detected as malicious based on the content by VirusTotal (May 2012–August 2013).

In February 2014, we submitted to VirusTotal [238] the 213,705 unique samples we had collected (based on their SHA1 hash) for examination. The idea was that evidence of malicious injections in webpages (e.g., JS redirects as used by RedKit [135] or other variants described earlier [22]) would likely be detected by at least some malware URL blacklists.

Figure 6.2 presents the proportion of unclassified results detected as malicious by VirusTotal. Typically, the malicious websites contain trojans (e.g. JS/Redirector.GR), backdoors (e.g. PHP/WebShell.J, C99), and exploits (e.g. HTML/IframeRef.AS). Overall, 19.5% of unclassified results appear as malicious. We see that websites with malicious content are relatively infrequent when search-redirection attack is experiencing its peak towards the end of 2012 (Figure 6.1). However, in 2013 we observe an increase of malicious websites among unclassified results. This may be an indication that miscreants are increasingly using other forms of manipulation our crawler did not detect, like JavaScript-based compromises. However, returning to Figure 6.1, this potential increase in infections does not compensate for the decrease observed in redirections overall. At most one third of all unclassified results (up to 7% of all results in 2013) are compro-

Table 6.3: Confusion matrix for the search-redirection classification.

		Predicted as compromised	
		Yes	No
Actually compromised	Yes	96.4%	3.6%
	No	0%	100%

mised in this way, whereas the active redirections have themselves dropped by roughly 20 percentage points.

Despite the decrease observed in 2013, claiming success in solving the search-redirection problem would be a stretch. Indeed, redirections still constitute the largest proportion of results for the query set we used.

Overall detection rate. Considering the proportion of undetected infections (3.6% of total) that we retrospectively identify from the category of unclassified results, in Table 6.3 we show that the overall true positive rate is estimated as 94.6% on average. In addition, following a manual clean-up of detected infections that should not have been classified as such, the false-positive rate is estimated to be close to 0%. We note that this manual clean-up process involved the inspection of the domain names placed in the search-redirecting category, while we would reasonably expect them to not be classified as such. For example, various `.gov` websites of federal agencies that provide information about unlicensed pharmacies (e.g. `fda.gov`).

Top 10 search result positions. In Figure 6.3 we present the evolution of search result considering only the top 10 result positions—contrary to Figure 6.1 that characterizes the top 32 positions. The reason for paying attention to this subset of results is their importance in terms of generating traffic. Indeed, Joachims et al. [114] have shown that 98.8% of users click on results appearing in the first 10 positions.




FIGURE 6.3: Similar to Figure 6.1, but examining only the top 10 search result positions.

While the previous observations are still valid at a high-level for the top 10 results, we point out a few important differences. First, the actively redirecting results have about 10% less daily occurrence, but the resulting vacancies are occupied by a different type of malicious results: organic results pointing directly to unlicensed pharmacies. Second, *G2* has apparently a stronger impact in the top 10 results, where we see that organic results pointing to unlicensed pharmacies briefly become the majority. The drop in redirecting results is explained by their dependence (at least in those early variants of the search-redirection attack) to the HTTP referrer.

6.3.3 User intentions

Our measurements appear to point at a large amount of malicious search results overall. A natural question is then whether or not users are actively looking for questionable results. If that is the case, it would then be hard to fault search engines for actually providing the users with what they are seeking.

To answer this question, we assess the impact of user intentions on search results by plotting, in Figure 6.4, the proportion of results we get for illicit, gray, and benign queries over time. The key take-away is that regardless of the type

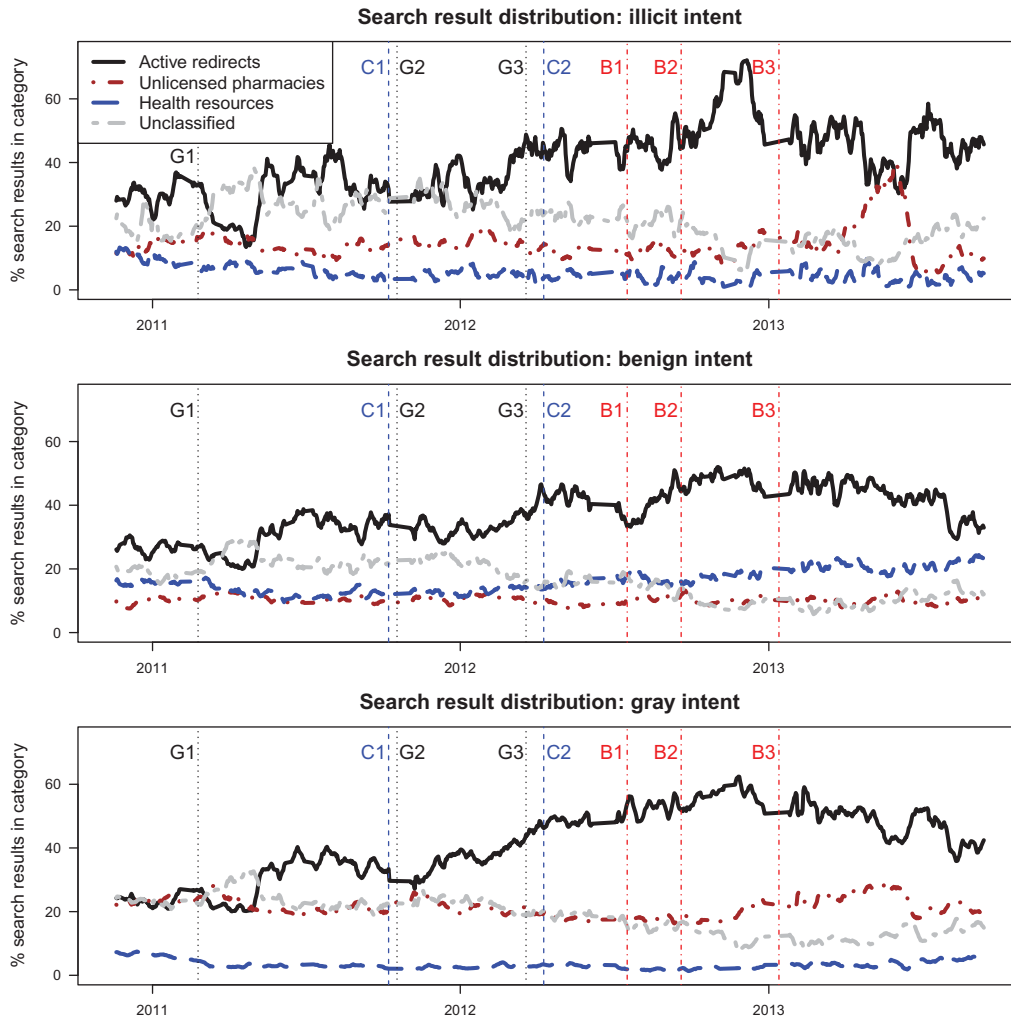


FIGURE 6.4: **Percentage of search results per category, based on the type of query.** Active redirections dominate results regardless of the intention of the query.

of query, active redirects dominate results. Unlicensed pharmacies also appear significantly not only in the results for illicit queries, but also for gray queries. We therefore reject the notion that active redirects only appear in search engines because users are seeking access to unlicensed pharmacies. Rather, unlicensed pharmacies appear to be successfully poisoning search results regardless of the queries' intent.




FIGURE 6.5: **Survival probability for source infections.** We use the entire measurement interval T_1, T_2, T_3 to compute this metric.

6.4 Cleanup-campaign evolution

Thus far we have examined how the proportion of search results with search-redirect attacks has changed over time. This helps in understanding the overall attack impact, and gives us a sense of the progress defenders (such as search engines) have made in combating this method of abuse. We now study much more explicitly how the interplay between those perpetrating search-redirect attacks and those working to stop them has evolved.

Several conditions must simultaneously hold for a search-redirect attack to be successful. First, the source infection must appear in the search results for popular queries. Second, the infection must remain on the website appearing in the results. Third, any intermediate traffic brokers must remain operational. Fourth, the destination website must stay online. Defenders may disrupt any one of these components to counter search-redirect attacks. In this section we examine how effective defenders have been in combating each component of the attack infrastructure. We first study the persistence of source infections over time, before investigating traffic brokers and destinations.

6.4.1 *Cleaning up source infections*

A key measure of defense is the time source infections persist in the search results and continue redirecting traffic elsewhere. We calculate the survival time of a source infection as the number of days a FQDN is first and last observed to be actively redirecting to different domains while appearing in the search results.⁶ Thus, source infections can be “cleaned” in two ways: either the responsible webmaster removes the infection that triggers the redirection or the website gets demoted from the search results because the search engine detects foul play.

Figure 6.5 shows the survival probability of the 26,673 source infections observed throughout the entire time period. Any measure of infection lifetimes involves “censored” data points, that is, infections that have not been remedied by the end of the observation period. In our dataset, 1,178 source infections were still actively redirecting at the end of data collection and are therefore censored. Survival analysis can deal with such incompleteness in the data by building an estimated probability distribution that takes censored data points into account. Figure 6.5 plots the survival probability as calculated using the Kaplan-Meier estimator [119].

We can see from the figure that many infections are short-lived. One-third last five days or less, while the median survival time for infections is 19 days. Nonetheless, it is noteworthy that some infections persist for a very long time. 17% of infections last at least six months, while 8% survive for more than one year. 459 websites, 1.7% of the total, remain infected for at least two years! Hence, while most infections are remedied in a timely fashion, a minority persist for much longer.

⁶ We treat different Uniform Resource Locator (URL) on the same FQDN as coming from a single infection. The reason we consider different FQDNs sharing the same second-level domain name as distinct infections is that frequently differing FQDNs represent distinct servers (e.g., `bronx.mit.edu` and `strategic.mit.edu` both appear in our sample). There is one exception to this policy. Whenever we observe multiple FQDNs cleaned up on the same day, we treat them as a single infection.

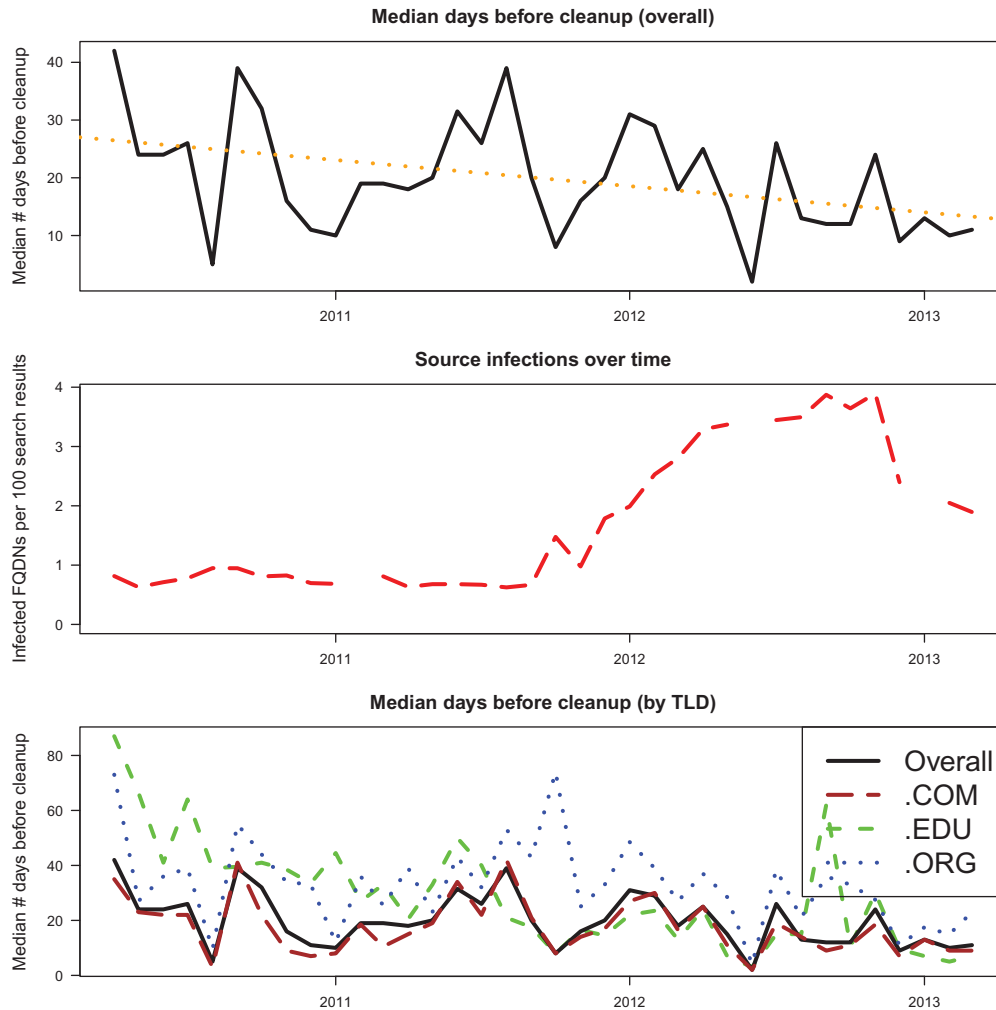


FIGURE 6.6: (Top) **Median time (in days) to cleanup source infections over time.** (Middle) **Source infections per 100 results over time.** (Bottom) **Median time (in days) to cleanup source infections by TLD.**

We next investigate how the time required to cleanup source infections has changed over time. We computed a survival function for each month from April 2010 to March 2013. We included all source infections that were first identified in that month. To make comparisons consistent across months, we censored any observed survival time greater than 180 days.⁷

⁷ This censoring also explains why we do not report anything for the final six months of the study.

Figure 6.6 (top) reports the median survival time (in days) for each monthly period. We can immediately see that the median time is highly volatile, ranging from 42 days in April 2010 to 2 days in June 2012. However, the overall trend is down, as indicated by the best-fit orange dotted line. Judging by the trend line, it appears that the median time to clean up source infections has fallen by around 10 days in three years.

While this is a welcome trend, we wondered what impact, if any, expedited cleanup times could have on the attacker's strategy. In particular, shorter-lived source infections could lead attackers to simply compromise more websites than before. Figure 6.6 (middle) plots the number of source infections per 100 search results observed each month.⁸ Here we observe a strongly positive effect. While the number of infected FQDNs hovered around 1 per 100 search results in 2010 and early 2011, observed infections increased substantially beginning in late 2011. This rose to nearly 4 infections per 100 search results by late 2012, before falling somewhat. Hence, it does appear that any crackdown in cleaning up source infections has been matched by an uptick in new infections, which helps to explain the increase in the percentage of search results that redirect as shown in Section 6.3.

Finally, Figure 6.6 (bottom) examines how cleanup times have changed for source infections on different TLDs. In Section 4.4.2 we found that `.edu` websites remained infected for much longer than others, and that `.org` and `.com` were cleaned more quickly. The figure shows that `.com` websites (denoted by the long dashed brown line) still in fact closely follow the overall trends in cleanup times. Notably, however, `.edu` websites (indicated by the dashed green line) went from considerably above-average survival times in 2010 to following the average by

⁸ The missing points in Figure 6.6 (middle) are from months when there are temporary 50% or greater drops in gathered search results.

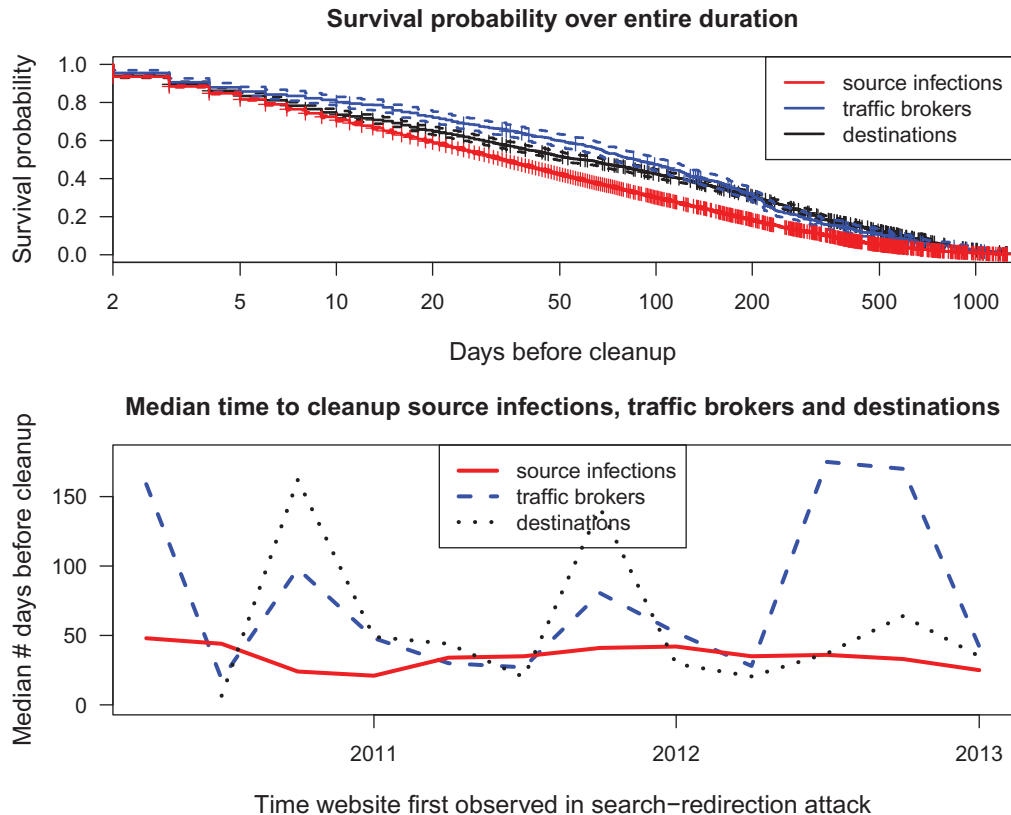


FIGURE 6.7: (Top) **Survival probability for source infections, traffic brokers and destinations over all time.** (Bottom) **Median time in days (survival time) to cleanup source infections, traffic brokers and destinations.**

mid-2011. In their place, however, `.org` websites began to lag behind starting in mid-2011. The timing suggests that attackers may have even shifted to targeting `.org` websites once `.edu` websites started to be cleaned up.

6.4.2 *Cleaning up traffic brokers and destinations*

Source infections are not the only hosts that can be targeted by defenders when combating search-redirectation attacks. Traffic brokers and destinations can also be shut down. We now compare the survival times of these to source infections.

Figure 6.7 (top) plots the survival time for source infections, traffic brokers and destinations. For traffic brokers and destinations, we report the second-level do-

main survival time, since subdomains often change to match drug names (e.g., `zolofit.example.com`).⁹ We also report the survival time for websites appearing for at least two days, since this removes a substantial number of false positives.

The graph shows that source infections are removed fastest, followed by destinations and traffic brokers. For example, 43% of sources are removed within three weeks, compared to 29% of traffic brokers and 36% of destinations. The median survival time for source infections is 34 days, compared to 59 days for destinations and 86 days for traffic brokers. So while the median traffic broker performs worst, the story changes slightly in the tail of the distribution: the 20% longest-lived source infections survive at least 6 months, compared to 9 months for traffic brokers and 11 months for destinations.

Figure 6.7 (bottom) tracks how the median survival time changes over time for source infections, traffic brokers and destinations. The median times are calculated quarterly, rather than monthly as in Figure 6.6 (bottom), due to the smaller number of traffic brokers and destinations compared to sources. We see once again the slow but steady improvement in reduced survival times for source infections. However, we see much greater vacillation for the survival times of traffic brokers and destinations. For some quarters the median time is around 5 months, whereas in others it follows more closely the survival times of sources. Notably, the survival times of traffic brokers and destinations are positively correlated.

We conclude from this analysis that traffic brokers and destinations have not received the same levels of pressure from defenders as source infections have. This is reflected in the longer survival times, as well as in the smaller number of domains ultimately used.

⁹ We removed 7 traffic brokers and 5 destinations from consideration here because they are known URL shortening services.

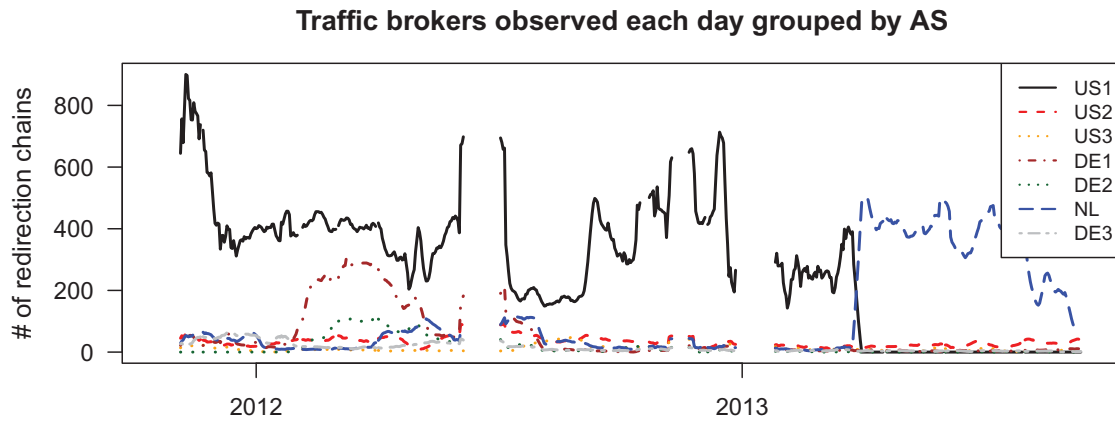


FIGURE 6.8: **Major autonomous systems hosting traffic brokers.** The plot shows the number of redirection chains using brokers from these ASs. In early 2013, US1 stopped hosting traffic brokers, which seemingly moved to NL.

Where are traffic brokers hosted?

The previous set of findings led us to look up the Autonomous System (AS) each traffic broker belongs to. It turns out that only 7 ASs (3 in the US, 3 in Germany, 1 in the Netherlands) support more than 10 traffic brokers every day. We plot on Figure 6.8, the number of redirection chains supported by brokers belonging to these 7 ASs as a function of time. None of these autonomous system provides “bulletproof hosting.” In fact, US1 is a known cloud-service provider. Some time in 2013, US1 seemingly decided to shutdown these brokers that had been using their service for more than a year. Some of them consequently shifted to NL, but what is most striking in this plot is the high concentration in traffic brokers over a few autonomous systems, especially since mid-2012. Coordinated take-downs among these ASs could be a very promising avenue for intervention.

Table 6.4: **Characteristics of actively redirecting URLs.** Averages over T_3 of the mean values obtained for each quantity, for each day (“averages of daily means”). The data does not sum to 100% because some infections “switch” categories over the measurement interval.

Active redirections	URLs		Per redirecting URL	
	#	%	Brokers (FQDN)	Pharmacies (FQDN)
With dedicated broker	428.8	42.8%	1	1
With shared broker	193.2	14.8%	1	2.4
Without broker	286.5	25.1%	-	1

6.5 Advertising network

We next turn to a deeper discussion of the redirection chains involved in search-redirection attacks. Redirection chains can indeed yield valuable insights about the “advertising network” used by criminals to peddle their products. We study traffic brokers and destinations in this section. We only focus on interval T_3 , since from Table 6.1, neither Datasets 1 nor 2 contain enough information to be able to extract the insights we discuss here. In the remainder of this section, we always look at traffic brokers and pharmacies at the FQDN level.

Source infections to traffic brokers. We start by looking at the connections between source infections and traffic brokers. On average, over 95% of the source infections a given day actually work; that is, less than 5% fail to take the visitor to a questionable site, instead landing on a parking page.

In Table 6.4 we describe the characteristics of actively redirecting URLs. About a quarter (25.1%) of these source infections send traffic directly to a pharmacy without any intermediate traffic broker. Another 42.8% use dedicated brokers that only get traffic from a single infection. More interestingly, on average about 14.8% of source infections send traffic to a broker shared with other source infections. Such brokers on average send traffic to 2.4 different pharmacies.

Table 6.5: **Characteristics of traffic brokers.** The data is given in averages of daily means over T_3 .

Traffic brokers	Daily average (FQDNs)		Per broker	
	#	%	Infections	Pharmacies
Redirecting to a single pharmacy	23.1	61.1%	18.9 URLs	1 URL
Redirecting to many pharmacies	14.4	33.8%	11.8 URLs	2.8 URLs
Redirecting to other brokers	3.8	5.2%	-	-

Traffic broker characteristics. (Table 6.5). Unsurprisingly, in light of what we saw above, 61.1% of brokers drive traffic to a single pharmacy, receiving traffic from 18.9 infected URLs on average. 33.8% of brokers redirect to multiple pharmacies, and receive on average traffic from 11.8 URLs. Finally only 5.2% of traffic brokers send traffic to other traffic brokers.

Pharmacies. (Table 6.6) We see that 56% of pharmacies do not rely on any broker and get their traffic, on average, from 4.6 infected URLs. 17.8% of all pharmacies get traffic from a dedicated broker, which feeds them traffic coming from about 24.2 distinct infected URLs. Slightly less than a third of all pharmacies use a shared traffic broker, which—interestingly enough—forward traffic from only 5.2 infected URLs. In other words, dedicated traffic brokers appear to be driving considerably more traffic than “co-hosted” solutions using shared traffic brokers. This in turn seems to give further credence to the belief that “advertising networks” (e.g., pharmaceutical affiliates) are highly heterogeneous, with actors ranging from powerful “dedicated” brokers to others operating on a shoe-string budget. The proportion of pharmacies directly linked to infections, without a traffic broker, is high—and can be explained by the difficulties search-redirection attacks experienced in 2013, and evidenced in Figure 6.1.

Network characteristics. Table 6.7 provides an overview of the graphs consisting of all redirection chains on any given day. We observe a very strong network heterogeneity, with large connected components that appear to dominate the graph.

Table 6.6: **Characteristics of pharmacies.** Data given as averages of daily means (over T_3).

Pharmacies	Daily average (FQDNs)		Per pharmacy	
	#	%	Infections	Traffic brokers
Without traffic broker	59.0	55.9%	4.6 URLs	-
With dedicated traffic broker	17.8	18.1%	24.2 URLs	1.3 URLs
With shared traffic broker	32.0	28.4%	5.4 URLs	2.2 URLs

Table 6.7: Connected components in the graph describing daily observed redirection chains.

Graph characteristics	#	Daily average		Range
			%	
Number of nodes	1055.4		100	[228, 2309]
Redirecting results	908 (URLs)		86.0	[193, 1, 927]
Traffic brokers	41.3 (FQDNs)		3.9	[9, 238]
Pharmacies	106.1 (FQDNs)		10.1	[26, 181]
Connected components	82.6		-	[25, 129]
Smallest connected component				
Number of nodes	2 nodes		5.7 (combined)	[2, 2]
2-node components	30.0		35.9	[9, 56]
Largest connected component				
Size of largest connected component	390 nodes		39.1	[72, 1, 091]
Redirecting results	379.6 (URLs)		38.1	[66, 1067]
Traffic brokers	5.8 (FQDNs)		0.6	[0, 16]
Pharmacies	4.6 (FQDNs)		0.4	[1, 31]

In other words, the illicit advertising business is dominated by a few large players. The same observation was reported by McCoy et al. [148], and in Section 4.4.3.

It is worth examining whether this concentration in advertisers changes over time. Figure 6.9 provides some elements of answer. We plot, as a function of time the maximum (top) and average (bottom) degree of traffic brokers and destinations. The degree is defined here as the sum of the number of links going in (in-degree) and out (out-degree) of a given “node” (traffic broker or destination). Each datapoint represents a 7-day moving average. The vertical lines correspond to the events introduced in Section 6.3. The size of the largest traffic brokers




FIGURE 6.9: Maximum and average degree of traffic brokers and destinations over time.

Table 6.8: **Overlap in the criminal infrastructures.** The fourth column is computed as the Jaccard index (Equation 5.1) between the two sets.

Type and granularity of node	Drugs	Other markets combined	Shared #	Jaccard index (%)
Source infection FQDNs	14,770	3,975	167	0.9
Traffic broker Domains	382	202	34	6.2
Traffic broker FQDNs	735	297	33	3.3
Destination domains	2,232	1,388	120	3.4
Destination FQDNs	2,249	1,388	119	3.4

varies drastically over time—the spikes observed in late 2012 seem to have been caused by particularly virulent campaigns (where a few brokers received a large amount of traffic from many infected sites) that took time to be fended off by search engines. Since early 2013, the size of the largest brokers has decreased a fair bit, reflecting the trend that search-redirection might be less popular than it was in 2012.

Shared infrastructure. We complete our analysis of the redirection network by looking at the traffic brokers used for different (non-pharmaceutical) types of

trades, and the extent to which they overlap with the pharmaceutical trade. Table 6.8 gives an overview of these results over the time interval 10/31/2011–09/16/2013. Over a long enough time interval, there is modest overlap between the various types of products. Source infections are rarely used for multiple campaigns; traffic broker domains tend to show a bit more overlap, presumably due to the fact that miscreants take advantage of lax verification policies at certain hosting providers. At the FQDN level, though, both destinations (i.e., shops) and brokers show little evidence of overlap, which is surprising given the known fact that certain botnets operate over multiple markets. Even in such cases, the different business domains appear to be kept separate.

6.6 Limitations

In addition to the numerous difficulties one faces when dealing with such long-range datasets, this study presents two major limitations. First, we have only looked at Google results. We justify this by the market share dominance of Google, at least in the US [40], but we also point out that the related work (e.g. Chapter 7, [22]) has shown other search engines are not immune to search-poisoning. Second, we have mostly looked at search results based on their presence or not in the result corpora. What is more important, however, is their *position* in the results. While top links are frequently clicked on, it has been shown that links past the 10th result have close to zero probability of being used [114]. Weighing the results we obtained by click probability would probably yield a better insight into which operations are profitable. However, we have shown in Section 4.3.2 (Figure 4.2) that the *type* of results (e.g., search-redirection attacks vs. health resources) is fairly consistent regardless of position.

6.7 Conclusions

Search engines are invaluable tools that deliver enormous value to consumers by referring them to the most relevant resources quickly and effortlessly. Search-engine poisoning threatens to undermine this value proposition, and could conceivably lead users to reduce their online activities [9].

We have presented the results of a long-term, large-scale empirical investigation into search-engine poisoning. Building up on our work in Chapter 4, this longitudinal analysis has enabled us to draw several new and important insights. First, despite the best efforts of search engines to demote low-quality content and browsers to protect the privacy of search queries, miscreants have readily adapted. In fact, the share of results taken over by search-redirectation attacks doubled from late 2010 to late 2012, before falling slightly. Second, efforts to clean up the compromised websites that initiate the redirections have improved: the persistence of source infections has steadily fallen from one month to two weeks. But here too, the attackers have adapted, notably by simply compromising more websites. Third, we continue to observe extensive concentration in the funneling of traffic from source infections to destinations via a small number of central brokers.

A key takeaway from this investigation is that uncoordinated interventions by individual stakeholders – a search engine ranking algorithm tweak here, a push by some hosting providers to clean up infected servers there – is not sufficient to disrupt persistent poisoning attempts. Instead, focusing on key points of concentration and in cooperation across stakeholders is required to effect dramatic change. For instance, coordinated traffic broker take-downs at the AS level, held in conjunction with the demotion or removal of poisoned search results at the search engine level (e.g., using proactive identification techniques [206]) could impact the economics of search-engine poisoning significantly, and, hopefully, durably. In

Chapter 9, we take a comprehensive look at all possible intervention approaches, and evaluate their effectiveness from an economic and a criminological perspective.

Trending-term exploitation on the web

Exploitation of trending news topics and of prescription drugs involve a similar monetization path. In both cases the money is in the traffic, not on the specific commodity being sold. We use this case study to reinforce our argument that financial profit is an invariable motive for online crime. In addition, we affirm the existence of similar concentration points in the criminal network as it depends largely on a few scarce resources.

Blogs and other websites pick up a news story only about 2.5 hours on average after it has been reported by traditional media [131]. This leads to an almost continuous supply of new “trending” topics, which are then amplified across the Internet, before fading away relatively quickly. However narrow, these first moments after a story breaks present a window of opportunity for attackers to infiltrate web and social network search results in response. The motivation for doing so is primarily financial. Websites that rank high in response to a search for a trending-term are likely to receive considerable amounts of traffic, regardless of their quality. Web traffic can in turn be monetized in a number of ways, as shown

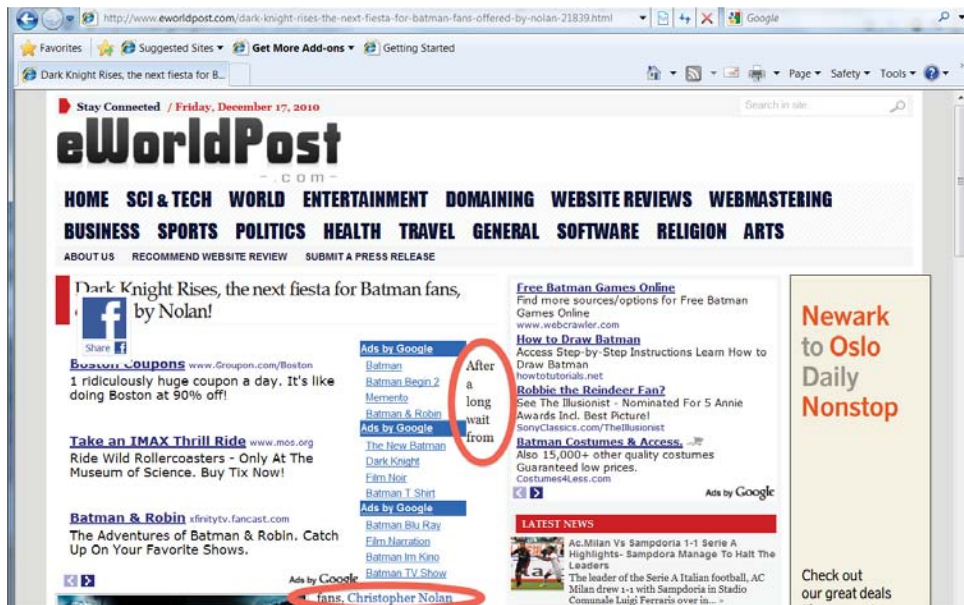


FIGURE 7.1: Ad-filled website appearing in the results for trending-terms (only 8 words from the article, circled, appear on screen).

in related work [33, 86, 115, 128]. In short, manipulation of web or social network search engine results can be a profitable enterprise for its perpetrators.

We study the abuse of “trending” search terms, which miscreants exploit to link to malware-distributing or ad-filled web sites. In particular, the sole goal of many sites designed in response to trending-terms is to produce revenue through the advertisements that they display in their pages, without providing any original content or services. Figure 7.1 presents a screenshot for `eworldpost.com`, which has appeared in response to 549 trending-terms between July 2010 and March 2011. The actual article (circled) is hard to find, when compared to the amount of screen real estate dedicated to ads. Such sites are often referred to as Made-for-AdSense (MFA) after the name of the Google advertising platform they are often targeting. Whether such activity is deemed to be criminal or merely a nuisance remains an open question, and largely depends on the tactics used to prop the sites up in the search-engine rankings. Some other sites devised to

respond to trending-terms have more overtly sinister motives. For instance, a number of malicious sites serve malware in hopes of infecting visitors' machines [185], or peddle fake anti-virus software [49, 188, 209].

In this chapter we detail our analysis on a large-scale measurement of trending-term exploitation on the web. Based on our collection of over 60 million web search and Twitter results associated with trending-terms gathered over nine months, we characterize how trending-terms are used to perform web search-engine manipulation and social-network spam. We further devise heuristics to identify ad-filled sites, and we report on the prevalence of malware and ad-filled sites in trending-term search results.

We uncover collusion across offending domains using network analysis, and through regression analysis, we conclude that both malware and ad-filled sites thrive on less popular and less profitable trending-terms. We build an economic model informed by our measurements, and find that ad-filled sites and malware distribution may be economic substitutes.

Additionally, we measure the success in blocking such content. Both MFA and malware-hosting sites are enough of a scourge to trigger response from search engine operators. Google modified its search algorithm in February 2011 in part to combat MFA sites [203], and has long been offering the Google Safe Browsing API to block malware-distribution sites. Trending-term exploitation makes both MFA and malware sites even more dynamic than they used to be, thereby complicating the defenders' task. Because our measurement interval spanned February 2011, when Google announced changes to its ranking algorithm to root out low-quality sites, we assess the impact of search-engine intervention on the profits that miscreants can achieve. An important feature of our work is that we bring an outsider's perspective. Instead of relying on proprietary data tied to a specific search engine, we use comparative measurements of publicly observable data

across different web search engines (Google, Yahoo!/Bing) and social network (Twitter) posts.

Our work here inscribes itself in the body of literature on understanding the underground online economy. Some of the early econometric work in that domain revolves around quantities bartered in underground forums [73], and on email spam campaigns [116, 132]. Grier et al. [86] extend this literature to Twitter spam. Along the same lines, Moore and Clayton have published a series of papers characterizing phishing campaigns [155, 157, 159].

A number of papers have also started to investigate web-based scams. Christin et al. [33] study a specific web-based social engineering scam (“one click fraud”). Provos et al. describe in details how so called “drive-by-downloads” are used to automatically install malware [185, 186]. Cova et al. [49] and Stone-Gross et al. [209] focus on fake anti-virus malware, and provide estimates of the amount of money they generate. Stone-Gross et al. calculate, through recovery of the miscreants’ transactions logs, that fake antivirus campaigns gross between \$3.8 and \$48.4 million a year. Affiliates funneling traffic to miscreants get between \$50,000 and \$1.8 million in over two months. These totals are markedly higher than what we obtain, but they consider all possible sources of malware (botnets, search engine manipulation, drive-by-downloads) whereas we only look at the much smaller subset of search engine manipulation based on trending-term exploitation.

Our approach differs from the related work in that we focus on a specific phenomenon—trending-term exploitation—by investigating how it is carried out (e.g., search-engine manipulation, Twitter spam), as well as its purpose: malware distribution and monetization through advertisements. Our analysis thus sheds light on a specific technique used by miscreants that search-engine operators are battling to fend off.

Our specific contributions are as follows. We (i) provide a methodology to automate classification of websites as MFA, (ii) show salient differences between tactics used by MFA site operators and malware peddlers, (iii) construct an economic model to characterize the trade-offs between advertising and malware as monetization vectors, quantifying the potential profit to the perpetrators, and (iv) examine the impact of possible intervention strategies. The last two contributions are especially important in the context of this thesis, as they extend our understanding of concepts related to the research questions 3¹ and 5.²

The rest of this chapter is organized as follows. We introduce our measurement and classification methodology in Section 7.1. We analyze the measurements collected in Section 7.2 to characterize trending-term exploitation on the web. Notably, we uncover collusion across offending domains using network analysis, and we use regression analysis to conclude that both malware and MFA sites thrive on less popular and profitable trending-terms. We then use these findings to build an economic model of attacker revenue in Section 7.3, and examine the effect of search-engine intervention in Section 7.4, before drawing brief conclusions in Section 7.5.

7.1 Methodology

We start by describing our methodology for data collection and website classification. At a high level, we need to issue a number of queries on various search engines for current trending-terms, follow the links obtained in response to these queries, and classify the websites we eventually reach as malicious or benign. Within the collection of malicious sites so obtained, we have to further distinguish between malware-hosting sites and ad-laden sites. Moreover, we need to com-

¹ Do other forms of illicit online activity exhibit a similar structure. . .

² Is it possible to disrupt online criminal networks by targeting critical components. . .

pare the results obtained with those collected from “ordinary,” rather than trending, terms.

The data collection hinges on a number of design choices that we discuss and motivate here. Specifically, we must determine how to build the corpus of trending-terms to use in queries (“*trending set*”); identify a set of control queries (“*control set*”) against which we can compare responses to queries based on trending-terms; decide on how frequently, and for how long, we issue each set of queries; and find mechanisms to classify sites as benign, malware-distributing, and MFA.

7.1.1 *Building query corpora*

Building a corpus of trending-terms is not in itself a challenging exercise. Google, through Google Hot Trends [82], provides a list of 20 current “hot searches,” which we determined, through pilot experiments, to be updated hourly. Likewise, Twitter avails a list of 10 trending topics [216], and Yahoo! gives a “buzz log” [251] containing the 20 most popular searches over the past 24 hours.

These different lists sometimes have very little overlap. For instance, combining the 20 Yahoo! Buzz logs, 20 Google Hot Trends, and 10 Twitter Trending Topics, it is not uncommon to find more than 40 distinct trending-terms over short time intervals. This would seem to make the case for aggregating all sources to build our query corpus. However, all search APIs limit the rate at which queries can be issued. We thus face a trade-off between the time granularity of our measurements and the size of our query corpus.

Trending set. Fortunately, we can capture most of the interesting patterns we seek to characterize by solely focusing on Google Hot Trends. Indeed, a measurement study conducted by John et al. [115] shows that over 95% of the terms used in search engine manipulation belong to the Google Hot Trends. However,

because Twitter abuse may not necessarily follow the typical search engine manipulation patterns, we use both Google Hot Trends and the Twitter current trending topics in our Twitter measurements.

Hot trends, by definition, are constantly changing. We update our trending-term corpus every hour by simply adding the current Google Hot Trends to it. Determining when a term has “cooled” and should be removed from the query corpus is slightly less straightforward. We could simply remove terms from our query corpus as soon as they disappear from the list of Google Hot Trends. However, unless all miscreants stop poisoning search results with a given term as soon as this term has “cooled”, we would likely miss a number of attempts to manipulate search engine results. Furthermore, Hot Trends are selected based upon their rate of growth in query popularity. Terms that have fallen out of the list in most cases still enjoy a sustained period of popularity before falling.

We ran a pilot experiment collecting Google and Twitter search results on 20 hot terms for up to four days. As Figure 7.2(a) shows, 95% of all unique Google search results and 81% of Twitter results are collected within three days. Thus, we settled on searching for trending-terms while they remain in the rankings, plus up to three days after they drop out of the rankings.


Control set. It is necessary to compare results from the trending set to a control set of consistently popular search terms, to identify which phenomena are unique to the trending nature of the terms as opposed to their overall popularity. We build a control list of the most popular search terms in 2010 according to Google Insights for Search [80]. Google lists the top 20 most popular search terms for 27 categories. These reduce to 495 unique search terms, which we use as a control set.

7.1.2 *Data collection*

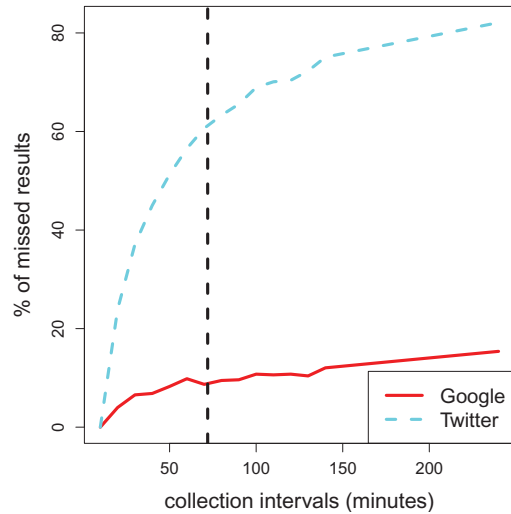
For each term in our trending and control sets, we run automated searches on Google and Yahoo! between July 24, 2010 and April 24, 2011. We investigate MFA results throughout that period, and study the timeliness of malware identification between January 26 and April 24, 2011. We study Twitter results gathered between March 10 and April 18, 2011.

We use the Google Web Search API [83] to pull the top 32 search results for each term from the Google search engine, and the Yahoo! BOSS API to fetch its top 100 Yahoo! results for each term. Since the summer of 2010 Yahoo! and Bing search results are identical [151]. Consequently, while in the chapter we refer to Yahoo! results, they should also be interpreted as those appearing on Bing. Likewise, we use the Twitter Atom API to retrieve the top 16 tweets for each term in Google's Hot Trends list and Twitter's Current Trends list. We resolve and record URLs linked from tweets, as well as the authors of these tweets linking to other sites.

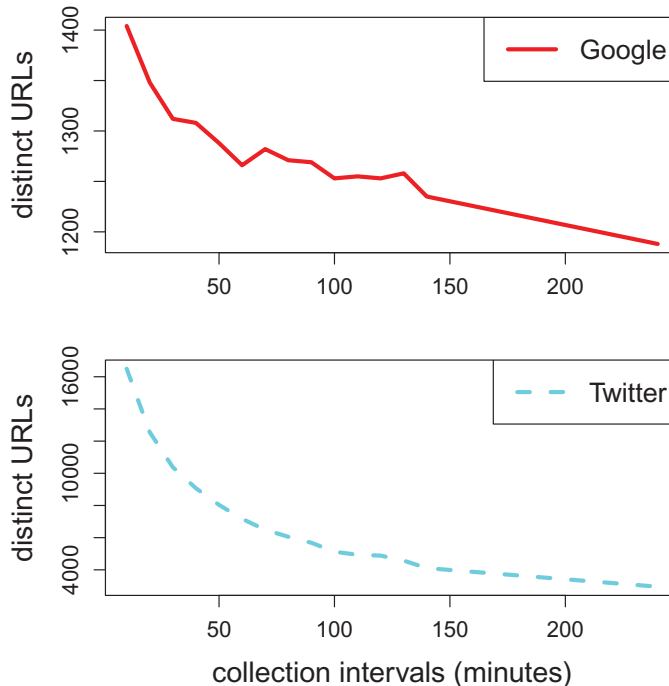
Because all these APIs limit the number of queries that can be run, we had to limit the frequency with which we ran the search queries. To better understand the trade-offs between search frequency and comprehensiveness of coverage, we selected 20 terms from a single trending list and ran searches using the Google API every 10 minutes for one week. We then compared the results we could obtain using the high-frequency sampling to what we found when sampling less often. The results are presented in Figures 7.2(c) and 7.2(b). Sampling once every 20 minutes, rather than every 10 minutes, caused 4% of the Google search results to be missed. Slower intervals caused more sites to be missed, but only slightly: 85% of the search results found when reissuing the query every 10 minutes could also be retrieved by sampling only once every 4 hours. So, even for trending topics,



(a) New search results as a function of time in Twitter and Google. More than 80% of Google results appear within 3 days, while Twitter continuously produces new results.



(b) Number of distinct URLs that we failed to collect using different collection intervals. The measurement lasted for two weeks using a fixed set of terms that was trending at the beginning of the experiment.



(c) Number of distinct URLs collected using different collection intervals. The measurement lasted for two weeks using a fixed set of terms that was trending at the beginning of the experiment.

FIGURE 7.2: Calibration tests weigh trade-offs between comprehensiveness and efficiency for collecting trending-term results.

searching for the hot terms once every four hours provides adequate coverage of Google results. For consistency, we used the same interval on Twitter despite the higher miss rate. Twitter indeed continues producing new results over a longer time interval, primarily due to the “Retweet” function which allows users to simply repost existing contents.

7.1.3 Website classification

We next discuss how we classified websites as benign, malware-distributing, or ad-filled. We define a website as a set of pages hosted on the same second-level DNS domain. That is, `this.example.com` and `that.example.com` belong to the same website.³ While we realize that different websites may be hosted on the same second-level domains, they are ultimately operated or endorsed by the same entity—the owner of the domain. Hence, in a slight abuse of terminology we will equivalently use “website” and “domain” in the rest of this discussion.

Malware-distributing sites. We pass all search results to Google’s Safe Browsing API, which indicates whether a URL is currently infected with malware by checking it against a blacklist. Because the search results deal with timely topics, we are only interested in finding which URLs are infected near the time when the trending topic is reported. However, there may be delays in the blacklist updates, so we keep checking the results against the blacklist for 14 days after the term is no longer hot.

When a URL appears in the results and is only later added to the blacklist, we assume that the URL was already malicious but not yet detected as such. It is, of course, also possible that the reason the URL was not in the blacklist is that the site had not yet been infected. In the case of trending terms, however, a

³ So do `this.example.co.uk` and `that.example.co.uk`, as `co.uk` is considered a TLD; as are a few others (e.g., `ac.jp`) for which we maintain an exhaustive list.

site appearing in results indicates a likely compromise, since the attacker's modus operandi is to populate compromised web servers with content that reflects trending results [115].

The possibility of later compromise further justifies our decision to stop checking the search results against the blacklist after two weeks have passed. While it is certainly possible that some malware takes more than two weeks to be detected, the potential for prematurely flagging a site as compromised also grows with time. Indeed, in a study of spam on Twitter [86], the majority of tweets flagged by the Google Safe Browsing API as malicious were not added to the blacklist until around a month had passed. We suspect that many of the domains marked as malicious were in fact only compromised much later. Consequently, our decision to only flag malware detected within two weeks is a conservative one that minimize false positives while slightly increasing false negatives.

Dealing with long-delayed reports of malware poses an additional issue for terms from the control set, because these search results are more stable over time. Sometimes a URL appears in the results of a term for years. If that website becomes infected, then it would clearly be incorrect to claim that the website was infected but undetected the entire time. In fact, most malware appearing in the results for the control set are for websites that have only recently "pushed" their way into the top search results after having been infected. For these sites, delays in detection do represent harm.

We thus exclude from our analysis of malware in the control set URLs that appeared in the results between December 20-31, 2010, when we began collecting results for the control set. To eliminate the potential for edge effects, our analysis of malware does not begin until January 26, 2011. As in the trending set, we also only flag results as malware when they are detected within 14 days.

Finally, we note that sometimes malware is undetected by the Safe Browsing API on the top-level URL, but that URLs loaded externally by the website are blocked. Consequently, our analysis provides an upper bound on malware success.

MFA sites. Automated identification of MFA sites is a daunting task. There are no clear rules for absolutely positive identification, and even human inspection suggests a certain degree of subjectivity in the classification. We discuss here a set of heuristics we use in determining whether a site is MFA or not.

While 182,741 different domains appeared in the top 32 Google and Yahoo! search results for trending-terms over 9 months, only 6,558 (3.6%) appeared in the search results for at least 20 different trending-terms. Because the goal of MFA sites is to appear high in the search results for as many terms as possible, we investigate further which of these 6,558 websites are in fact legitimate sources of information, and which are low-quality, ad-laden sites. To that effect we selected a statistically significant (95% confidence interval) random sample of 363 websites for manual inspection. From this sample, we identified five broad categories of websites indicative of MFA sites. All MFA sites appear to include a mechanism for automatically updating the topics they cover; differences emerge in how the resulting content is presented.

1. Sites which reuse snippets created by search engines and provide direct links to external sites with original content (e.g., <http://newsblogged.com/tornado-news-latest-real-time>).
2. Sites in blog-style format, containing a short paragraph of content that is likely copied from other sources and only slightly tweaked—usually by a machine algorithm, rather than a human editor (e.g., <http://toptodaynews.com/water-for-elephants-review>).

3. Sites that automatically update to new products for sale pointing to stores through paid advertisements (for instance, <http://tgibblackfriday.com/Online-Deals/-261-up-Europe-On-Sale-Each-Way-R-T-required--deal>).
4. Sites aggregating content by loading external websites into a frame so that they keep the user on the website along with their own overlaid ads (e.g., <http://baltimore-county-news.newslib.com/>).
5. Sites containing shoddily, but seemingly manually written content based on popular topics informed by trending-terms (e.g., <http://snarkfood.com/mel-gibsons-mistress-says-hes-not-racist/310962/>).

Based on manual inspection of our random sample of 363 sites, we decided to classify websites in any of the first four categories as MFA, while rejecting sites in the fifth category. (Including those would have driven up the false positive rate to unacceptable levels.) This results in 44 of the 363 websites being tagged as MFA.

Subsequently, we used a supervised machine-learning algorithm (Bayesian Network [180] constructed using the K2 algorithm [42]) to automatically categorize the remaining 6,195 candidate websites.

The set of measures used to describe each page is a combination of structural and behavioral characteristics: (i) the number of internal links, i.e. links to the same domain as the web page under examination; (ii) the number of external links, i.e. links directed to external domains; and (iii) the existence of advertisements in the web page. We calculate these three quantities for each of the 6,558 domains by parsing the front page of the domain and a set of five additional web pages within the same domain, randomly chosen among the direct links existing in the front page.

We experimented with many more features in the classifier (e.g., time since the website was registered, private WHOIS registration, number of trending-terms where a website appears in the search results, presence of JavaScript, etc). As manual inspection confirmed, this did not improve classification accuracy beyond the three features described here. MFA sites exhibit large numbers of external links but few internal links, because unlike external links to ads, internal links do not (directly) generate revenue.

We determine whether a website has advertisements by looking for known advertising domains in the collected HTML. Because these domains often appear in JS, we use regular expressions to search throughout the page. We use manually-collected lists of known advertising domains used by Google and Yahoo!, complemented by the “Easy List” maintained by AdBlock Plus [1] (Jan. 12, 2011).

We used a subset of the 363 sample domains as a training set for the machine learning algorithm. We did not use the entire set because it is overcrowded with non-MFA domains (87% non-MFA vs. 13% MFA), which would lead to over-training the model towards non-MFA websites. By using fewer non-MFA websites in the training set (80% vs. 20%), we kept our model biased towards non-MFA websites, thereby maintaining the assumption of innocence while remaining able to identify obvious MFA instances.

We assessed the quality of our predictive model by performing 10 rounds of cross-validation [120], yielding a 87.3% rate of successful classifications. In the end, the algorithm classified 838 websites—0.46% of all collected domains—that appear in the trending set results as MFAs. The relatively small number of positive identifications allows for manual inspection to root out false positives. We find that 120 of the websites—consistent with the predicted 87.3% success rate—are likely

false positives. We remove these websites from consideration when conducting the subsequent quantitative analysis of MFA behavior.

7.2 Measuring trending-term abuse

7.2.1 *Incidence of abuse*

We now discuss the prevalence of malware and MFA in the trending search results. There are many plausible ways to summarize tens of millions of search results for tens of thousands of trending-terms gathered over several months. We consider four categories: terms affected, search results, URLs, and domains.

Table 7.1 presents totals for each of these categories. For web search, we observed malware in the search results of 1,232 of the 6,946 terms in the trending set. Running queries six times a day over three months yielded 9.8 million search results. Only 7,889 of these results were infected with malware—0.08% of the total. These results corresponded to 607,156 unique URLs, only 1,905 of which were infected with malware. Finally, 495 of the 108,815 domains were infected.

How does this compare to popular search terms? As a percentage, more control terms were infected with malware, but that is due to their persistent popularity. Around the same number of search results were infected, but the control set included nearly twice as many overall results—because there were around 300 trending-terms “hot” at any one time compared to the 495 terms always checked in the control set. 1,905 URLs were infected in the trending set, compared to only 302 in the control set.

The prevalence of malware on Twitter is markedly lower: only 2.4% of terms in the trending set were found to have malware, compared to 18% for search, and only 101 URLs on 13 distinct domains were found infected. While the number of infections observed is very small (0.03%), it is consistent with the proportion of

Table 7.1: Total incidence of malware and MFA in Web search and Twitter results.

	Terms			Results			URLs			Domains		
	Total	Inf.	%	Total	Inf.	%	Total	Inf.	%	Total	Inf.	%
Malware												
<i>Web Search</i>												
Trending set	6,946	1,232	18	9.8M	7,889	.08	607K	1,905	.30	109K	495	.50
Control set	495	123	25	16.8M	7,332	.04	231K	302	.13	86K	123	.14
<i>Twitter</i>												
Trending set	1,950	46	2.4	466K	137	.03	355K	101	.03	43K	13	.03
Control set	495	53	11	1M	139	.01	825K	129	.02	98K	101	.02
Twitter trnd.	1,176	20	1.7	180K	24	.01	139K	21	.02	26K	9	.03
MFA sites												
<i>Web Search</i>												
Trending set	19,792	15,181	76.7	32.3M	954K	3.0	1.35M	83,920	6.2	183K	629	.34
<i>Twitter</i>												
Trending set	1,950	1,833	94	466K	32,152	6.9	355K	32,130	9.0	43K	141	.3
Twitter trnd.	1,176	1,012	86	179K	12,145	6.6	139K	12,144	8.7	26K	42	.2

malicious URLs observed by Grier et al. [86] on a significantly larger dataset of 25 million unique URLs. The control and Twitter-trending sets also reveal similarly low levels of infection.

Grier et al. observed a much higher proportion of “spammy” behavior on Twitter. Likewise, we observe substantial promotion of MFA websites on Twitter: 94% of trending-terms contained tweets with MFA domains. While most terms are targeted, only a small number of domains are promoted—141 in the trending set and 42 in the Twitter’s trending set. Web search is also targeted substantially by MFA sites. 77% of terms in the trending set included one or more of the 629 MFA domains in at least one result.

From the figures in Table 7.1 alone, it would appear that malware on trending-terms is largely under control, while MFA sites are relatively rampant. However, aggregating figures across a large period of time can obscure the potential harm of malware distributed via trending-terms. Table 7.2 presents the malware infection rate at a single *point in time*: counting the number of terms and search results that are infected with malware for each of the trending-terms within a 3-day window of rising. For example, on average, 12.8 trending-terms are infected with malware

Table 7.2: Prevalence of malware in trending and control terms, presented as the average prevalence of malware at every point in time when searches are issued.

	Terms		Results		Domains		URLs	
	#	%	#	#	%	#	%	
<i>Trending terms—web search (point in time)</i>								
detected	12.8	4.4	14.8	13.8	0.089	8.7	0.146	
top 10	2.9	1.0	3.2	3.1	0.020	2.4	0.040	
undetected	6.2	2.1	7.6	6.7	0.0	3.718	0.061	
top 10	1.2	0.4	1.5	1.4	0.009	0.9	0.015	
<i>Control terms—web search (point in time)</i>								
detected	9.5	1.9	14.1	11.5	0.043	8.9	0.067	
top 10	3.1	0.6	3.9	3.7	0.014	3.1	0.023	
undetected	1.0	0.2	1.0	1.0	0.0	0.856	0.006	
top 10	0.1	0.0	0.1	0.1	0.000	0.1	0.001	

that has already been flagged by the Safe Browsing API, which corresponds to 4.4% of recently hot terms at any given moment. A further 6.2 trending terms are infected but not yet detected by the blacklist. On average, 1.2 terms include a top 10 result that distributes malware and has not yet been detected by the Safe Browsing API. Viewed in this manner, the threat from web-based malware appears more worrisome.

But is the threat worse for trending-terms? 9.5 control terms include detected malware at a given point in time, with one term infected but not yet detected. Hence, popular terms are still targeted for malware, but less frequently than trending-terms and with less success. Finally, the false negative rate for the trending set is much higher than for the control set: 34% (7.6 results undetected compared to 14.8 detected) vs. 7% (1 undetected result compared to 14.1 detected).

7.2.2 Network characteristics

We next turn to characterizing how sites preying on trending-terms are connected to each other. To prop up their rankings in Google, one would expect a group of sites operated by a same entity to link to each other—essentially building a “link farm” [88]. Thus, we conjecture that looking at the network structure of both MFA and malware-serving sites may yield some insight on both the actors behind these attacks, and the way campaigns are orchestrated.

MFA domains. We build a directed graph G_{MFA} where each node corresponds to one of the 629 domains we identified as MFA, and each of the 3,221 (directed) edges corresponds to an HTML link between two domains. We construct the graph by fetching 1,000 backlinks for each of the sites from Yahoo! Site Explorer [252]. Extracting the strongly connected components from G_{MFA} yields family of sites that link to each other. We find 407 distinct strongly connected components, most (392) only contain singletons. More interestingly, 193 sites—30.7% of all MFA sites—form a strongly connected component. These nodes have on average a degree (in- and out-links) of 12.83, and an average path length between two nodes of 3.92, indicating a quite tightly connected network. It thus appears that a significant portion of all MFA domains may be operated by the same entity—or at the very least, by a unique group of affiliates all linking to each other. Further inspecting where these sites are hosted indicates that 130 of the 193 sites belong to one of only seven distinct ASs; here, sites within a same AS are usually hosted by the same provider, which confirms the presence of a fairly large, collusive, MFA operation.

Malware-serving sites. Examining the network characteristics of malware-distributing sites serves a slightly different purpose. Here, sites connected to each other are unlikely to be operated by the same entity, but are likely to have

Table 7.3: Malware campaigns observed.

Campaign ID	# Domains	Duration	Distinct ASes
949	590	>1 year	>200
5100	36	>8 months	1
5101	25	>8 months	1
5041	11	4 days	2
5053	10	2 days	1
4979	9	11 days	2
4988	9	8 days	2

been *compromised* by the same group or as part of the same campaign. This is consistent with the behavior observed by John et al. [115], who found that miscreants add links between malicious websites to elevate PageRank. As with MFA sites, we build a directed graph G_{mal} where each node corresponds to one of the 6,133 domains we identified as malware-serving based on a longer collection of trending-terms gathered from April 6, 2010 to April 27, 2011. Each (directed) edge corresponds to an HTML link between two malware-serving domains. G_{mal} contains 6,133 nodes and 18,864 edges, and 5,125 distinct strongly connected components, only 216 of which contain more than one node.

Table 7.3 lists the largest strongly connected components (“campaigns”) in G_{mal} . For each of the nodes in these campaigns, we look up the time at which they were first listed as infected. By comparing the first and last nodes to be infected within a given campaign, we can infer the campaign’s duration. We also look up the number of distinct ASs in each campaign.

We observe divergent campaign behaviors, each characterized by markedly different attacker tactics. The largest campaign (949) was still ongoing at the time of this analysis (mid-2011): nodes are compromised at a relatively constant rate, and are hosted on various ASs. This indicates a long-term, sustained effort. This

campaign affects at least 9.6% of all the malware-infested sites we observed. Campaigns 5100 and 5101 are likely part of the same effort: all nodes share the same set of servers, and seem compromised by the same exploit. Interestingly, this campaign went unabated for at least 8 months (until Dec. 2010). Finally, the other four notable campaigns we observed target small sets of servers, that are compromised almost simultaneously, and all immediately link to each other.

Our definition of a campaign is extremely conservative: we are only looking for strongly connected components in the graph we have built. It is thus likely that many of the singletons we observed are in fact part of larger campaigns. Further detection of such campaigns would require more complex clustering analysis. For instance, one could try to use the feature set of the classification algorithm as a coordinate system, and cluster nodes with nearby coordinates. However, it is unclear that this specific coordinate system would provide definitive evidence of collusion.

7.2.3 MFA in Twitter

We turn our attention now to the use of MFA links in Twitter posts. We are interested in measuring the amount of unique MFA-related URLs each malicious user posts, and the popularity of the MFA websites among them.

Figure 7.3(a) shows that 95% of the authors who post MFA URLs link to 5 domains or less—this amounts to about 20,000 posts. However, the remaining 5% is responsible for about 55,000 posts, and links to 870 domains. The control set gives similar numbers.

In other words, a small number of authors are responsible for wide promotional campaigns of MFA websites. The vast majority of authors post a small number of MFA links, and it is unclear whether they are actually malicious or not.

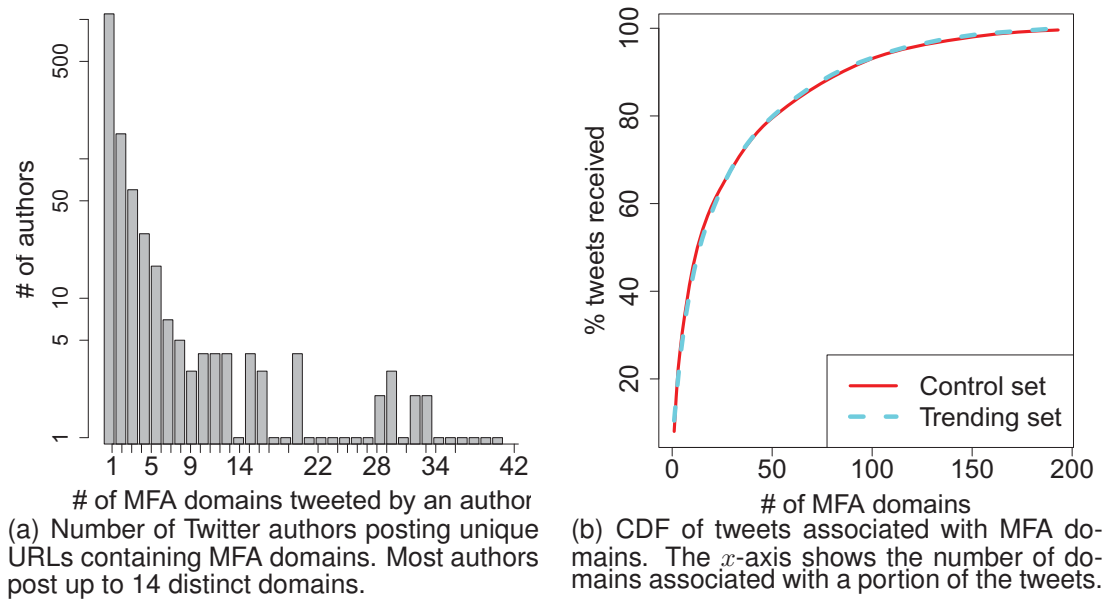


FIGURE 7.3: Trending-term exploitation on Twitter.

Similarly, the number of MFA domains that receive the majority of related tweets is small as Figure 7.3(b) shows. 50% of the MFA infected tweets direct users to 14 MFA domains, with the remaining 50% distributed across 180 MFA domains.

7.2.4 Search-term characteristics

We now examine how characteristics of the trending-terms themselves influence the prevalence of malware and MFA sites in their search results. We focus on the importance of the term's category, popularity in searches, and expected advertising revenue.

Measuring term category, popularity and ad prices. We combine results from several Google tools in order to learn more about the characteristics of each of the trending-terms. First, we classify the trending terms into categories using Google Insights for Search, which assigns arbitrary search terms to the most likely

category, out of the same 27 categories used for constructing the control sample described in Section 7.1.1 above.

Second, Google offers a free service called Traffic Estimator that estimates for any phrase the number of global monthly searches averaged over the past year [81]. For trending-terms, averaging over the course of a year significantly underestimates the search traffic when a term is peaking in popularity. Fortunately, Google also offers a measure of the relative popularity of terms through Google Trends [82], provided at the granularity of one week. The relative measure is normalized against the average number of searches for the past year, precisely the figure returned by the Traffic Estimator. We obtain the *peak-popularity estimate* $\text{Pop}(s)$ for a term s by multiplying the relative estimate for the week when the term peaked by the absolute long-run popularity estimate.

The Google Traffic Estimator also indicates the advertising value of trending terms, by providing estimates of the anticipated Cost per Click (CPC) for keywords. We collect the CPC for all trending and control terms. Many trending-terms are only briefly popular and return the minimum CPC estimate of 0.05 USD. We use the CPC to approximate the relative revenue that might be obtained for search results on each term. The CPC is a natural proxy for the prospective advertising value of user traffic because websites that show ads are likely to present ads similar to the referring term.

Empirical analysis. Table 7.4 breaks down the relative prevalence of trending-terms, and their abuse, by category. Over half of the terms fall into three categories—Entertainment, Sports and Local. These categories feature topics that change frequently and briefly rise from prior obscurity. 18% of trending-terms include malware in their results, while 38% feature MFA websites in at least 1% of the top 10 results.




FIGURE 7.4: Exploring how popularity and ad price of trending-terms affects the prevalence of malware (left) and ad-laden sites (right).

Table 7.4: Malware and MFA incidence broken down by trending-term category.

Category name	%	CPC	Malware		MFA	coef.
			% terms	% terms	% top 10	
Arts & Humanities	2.7	\$0.44	20.1	40.6	6.8	
Automotive	1.3	\$0.67	16.0	29.2	5.2	-0.0062
Beauty & Personal Care	0.8	\$0.76	19.6	32.5	6.9	
Business	0.4	\$0.87	7.4	32.9	6.9	
Computers & Electronics	2.4	\$0.61	14.5	31.7	5.9	
Entertainment	30.6	\$0.34	18.6	41.0	6.4	-0.0043
Finance & Insurance	1.4	\$1.26	20.2	30.4	5.6	
Food & Drink	2.9	\$0.43	17.1	49.5	7.9	+0.0105
Games	2.3	\$0.32	13.4	30.0	5.6	-0.0073
Health	2.5	\$0.85	14.1	27.6	5.9	-0.0046
Home & Garden	0.5	\$0.76	7.1	29.7	7.2	
Industries	1.6	\$0.50	26.1	38.6	6.6	-0.0072
Internet	0.7	\$0.49	7.7	43.7	6.0	
Lifestyles	4.5	\$0.33	25.4	45.8	6.5	
Local	11.0	\$0.51	21.8	39.2	6.9	-0.0027
News & Current Events	3.6	\$0.39	19.7	45.0	7.0	
Photo & Video	0.2	\$0.59	0.0	21.9	6.4	
Real Estate	0.2	\$1.02	6.2	34.2	6.5	
Recreation	1.0	\$0.43	13.7	43.5	6.5	
Reference	1.4	\$0.43	14.5	55.4	8.7	+0.0203
Science	1.4	\$0.40	16.0	44.9	9.1	+0.0095
Shopping	3.2	\$0.56	11.6	43.7	8.8	+0.0106
Social Networks	0.5	\$0.19	27.8	59.1	6.4	
Society	5.1	\$0.62	15.2	33.7	5.6	-0.0085
Sports	15.4	\$0.38	20.7	44.9	6.9	-0.0044
Telecommunications	0.8	\$0.91	10.9	36.4	4.6	
Travel	1.7	\$0.88	10.1	29.3	6.4	
Average (category)	3.7	\$0.59	18.4	38.3	6.6	

We observe some variation in malware and MFA incidence across categories. However, perhaps the most striking result from examining the table is that all categories are targeted, irrespective of the category's propensity to "trendiness." Miscreants do not seem to be specializing yet by focusing on particular keyword categories.

If we instead look at popularity and ad prices, substantial differences emerge. Figure 7.4 shows how the incidence of malware and MFA varies according to the peak popularity and ad price of the trending-term. The left-most graph shows how malware varies according to the term's peak popularity. The least popular terms—less than 1,000 searches per day at their peak—attract the most malware in their top results. 38% of such terms include malware, while 9% of these terms include malware that is not initially detected. As terms increase in peak popularity, fewer are afflicted by malware: only 6.2% of terms with peak popularity greater than 100,000 daily searches include malware in their results, and only 2% of terms include malware that is not immediately detected. A similar pattern follows for malware incidence according to the term's ad price. 30% of terms with ad prices under 10 cents per click had malware in their results, compared to 8.8% of terms with ad prices greater than \$1 per click.

A greater fraction of terms overall include MFA websites in their results than malware—37% vs. 19%. Consequently, all proportions are larger in the two graphs on the right side of Figure 7.4. 60% of terms with peak popularity of less than 1,000 daily searches include MFA sites in their results. This proportion drops steadily until only 17.4% of terms attracting over 100,000 daily visits include MFA in the top 10 search results. A similar reduction can be seen for varying ad prices in the right-most figure. The two right figures show the percentage of all terms that have MFA, followed by the percentage of top 10 results that are MFA, for only those terms that have MFA terms present. Here we can see that the percentages

remain relatively steady irrespective of term popularity and price. For unpopular terms, 10% of their results point to MFA, dropping modestly to 8% for the most popular terms. The drop is more significant for ad prices—from 10% to 6%. Consequently, while the success in appearing in results diminishes with popularity and rising ad prices, when a term does have MFA, a similar proportion of its results are polluted.

Of course, ad prices and term popularity are correlated—more popular search terms tend to attract higher ad prices, and vice versa. Consequently, we use linear regression to disentangle the effect both have on the prevalence of abuse.

Because the dependent variable is binary in the case of malware—either the term has malware present or does not—we use a logit model for the regression of the following form:

$$\text{logit}(p_{\text{HasMalware}}) = \beta + \text{AdPrice}x_1 + \log_2(\text{Popularity})x_2$$

We also ran a logit regression with the term’s categories, but none of the category values were statistically significant. Thus, we have settled on this simpler model. The results of the regression reveal that a term’s ad price and search popularity are both negatively correlated with the presence of malware in a term’s search results, and the relationship is statistically significant:

	coef.	odds	Std. Err.	Significance
AdPrice	-0.509	.601	0.091	$p < 0.001$
$\log_2(\text{Popularity})$	-0.117	0.889	0.012	$p < 0.001$

These coefficients mean that a \$1 increase in the ad price corresponds to a 40% decrease in the odds of having malware in the term’s results. Likewise, when the popularity of a term doubles, the odds of having malware in the term’s results decreases by 11%.

We also devised a linear regression using the fraction of a term’s top 10 results classified as MFA as the dependent variable:

$$\text{FracTop10MFA} = \beta + \text{AdPrice}x_1 + \log_2(\text{Popularity})x_2 + \text{Category}x_3 .$$

The Category variable is encoded as a 27-part categorical variable using deviation coding. This coding scheme is used to measure each categories’ deviation from the overall mean value, rather than deviations across categories.

For this regression, the term’s ad price and search popularity are both statistically significant and negatively correlated with the fraction of a trending-term’s top 10 results classified as MFA:

	coef.	Std. Err.	Significance
AdPrice	-0.0091	0.091	$p < 0.001$
$\log_2(\text{Popularity})$	-0.004	0.012	$p < 0.001$
Coefficients for category variables in Tab. 7.4, R^2 : 0.1373			

A \$1 increase in the ad price corresponds to a 0.9% decrease in the MFA rate, while a doubling in the popularity of a search term matches a 0.4 percentage point decrease. This may not seem much, but recall that, on average, 6.6% of a term’s top 10 results link to MFA sites. A 0.9% decrease in MFA prevalence represents a 13.2% decrease from the average rate.

Each of the coefficients listed in the right-most column in Table 7.4 are statistically significant—all have p values less than 0.001, except Local, Health, and Automotive, where $p < 0.05$. For instance, Food & Drink terms correspond to a 1 percentage point increase in the rate of MFA domains in their top 10 results, while Reference terms suffer a 2% higher MFA rate.

Implications of analysis. The results just presented demonstrate that, for both malware and MFA sites, miscreants are struggling to successfully target the more

lucrative terms. An optimistic interpretation is that defenders manage to relegate the abuse to the more obscure terms that have less overall impact. A more pessimistic interpretation is that miscreants are having success in the tail of hot terms, which are more difficult to eradicate.

It is not very surprising that malware tends to be located in the results of terms that demand lower ad prices, given that higher ad prices do not benefit malware distribution. However, it is quite unexpected that the prevalence of MFA terms is negatively correlated with a term's ad price, since those promoting MFA sites would much prefer to appear in the search results of more expensive terms. One reason why malware and MFA appears less frequently on pages with higher ad prices could be that there is stronger legitimate competition in these results than for results fetching lower ad prices.

Furthermore, there is a potential incentive conflict for search engines to eradicate ad-laden sites, when many of the pages run advertisements for the ad platforms maintained by the search engines. It is therefore encouraging that the evidence suggests that search engines do a better job at expelling MFA sites from the results of terms that attract higher ad prices.

Finally, the data helps to answer an important question: are malware and ad abuse websites competitors, or do they serve different parts of the market? The evidence suggests that, in terms of being a technique to monetize search traffic, malware and MFA behave more like substitutes, rather than complements. Both approaches thrive on the same types of terms, low-volume terms where ads are less attractive. Consequently, a purely profit-motivated attacker not fearful of arrest might choose between the two approaches, depending on which method generates more revenue.

7.3 Economics of trending-term exploitation

We next examine the revenues possible for both malware and ads, by first characterizing the volumes of population affected, before deriving actual expected revenues.

7.3.1 Exposed population

We first estimate the number of visits malware and MFA sites attract from trending-term searches. The cumulative number of visits over an interval t to a website w for a search term s is given by

$$V(w, s, t) = C(\text{Rank}(w, s)) \cdot \text{Pop}(s) \cdot \frac{4}{30 \times 24} \times t,$$

where $\text{Pop}(s)$ is the monthly peak popularity of the term, as defined in Section 7.2.4. $\text{Rank}(w, s)$ is the position in search results website w occupies in response to a query for s , and $C(r)$ defines a click probability function for search rank $1 \leq r \leq 10$ following the empirical distribution observed by Joachims et al. [114]. They found that 43% of users clicked on the first result, 17% on the second result, and 98.9% of users only clicked on results in the first page. We ignore results in ranks above 10 (i.e., $C(r) = 0$ for $r > 10$).

$\text{Pop}(s)$ is measured at a monthly rate, so we normalize the visits to the four-hour interval between each search. We also weigh Google and Yahoo! search results differently. Google has reportedly an 64.4% market share in search, while Yahoo! and Bing have a combined market share of 30% [64]. Since our estimates are based on what Google observes, we anticipate that Yahoo! and Bing attract $\frac{30\%}{64.4\%} = 46.5\%$ of the searches that Google does.

The results are given in Table 7.5. MFA sites attract 39 million visits over nine months, or 4.3 million visits per month. For the malware results, we compare the

Table 7.5: Estimated number of visits to MFA and malware sites for trending terms.

	Total	# Visitors	
		Period	Monthly Rate
MFA	39,274,200	275 days	4,284,458
Malware (trending set)			
detected	454,198	88 days	154,840
Bing, Yahoo!	189,511	88 days	64,606
undetected	143,662	88 days	48,975
Malware (control set)			
detected	12,825,332	88 days	4,372,272
Bing, Yahoo!	6,352,378	88 days	2,165,583
undetected	83615	88 days	28505

estimated visits for both control and trending terms. While more users see malware in the results of control terms than trending-terms—about 4.4 million versus about 200,000 per month over three months—over 99% of the visits from control terms are blocked by the Safe Browsing API. By contrast, 24% of the visits triggered from the results of trending-terms are not blocked by the Safe Browsing API. In aggregate, trending-terms expose around 49,000 victims per month to undetected malware, compared to about 28,000 for control terms.

The table also lists the number of Bing and Yahoo! users that encounter malware detected by Google’s Safe Browsing API. We cannot say for certain whether or not these users will be exposed to malware. If they attempt to visit the malicious site using the Chrome or Firefox browser then they would be protected, since Google’s Safe Browsing API is integrated into those browsers. Internet Explorer users would be protected only if the sites appear in IE’s internal blacklist. Unfortunately, we could not verify this since the blacklist is not made publicly accessible.




FIGURE 7.5: Number of estimated daily victims for malware appearing in trending and control terms.

The sums presented in the table mask several peculiarities of the data. First, for malware, the number of visitors exposed is highly variable. Figure 7.5 plots the number of daily victims over time. Most days the number of victims exposed is very small, often zero. Because terms in the control set are always very popular, successful attacks cause large spikes, but tend to be rare. On the other hand, trending-terms exhibit frequent spikes, but many of the spikes are small. This is because many trending-terms are in fact not very popular, even at their peak. A big spike, as happened around March 5, results from the conjunction of three factors: (1) the attacker must get their result towards the top of the search results; (2) the result cannot be immediately spotted and flagged; and (3) the trending-term has to be popular enough to draw in many victims. Consequently, there is a downside to the constantly replenishing pool of trending-terms for the attacker—they are often not popular enough for the attacker to do much damage. This is further exacerbated by the finding from the last section—more popular terms are

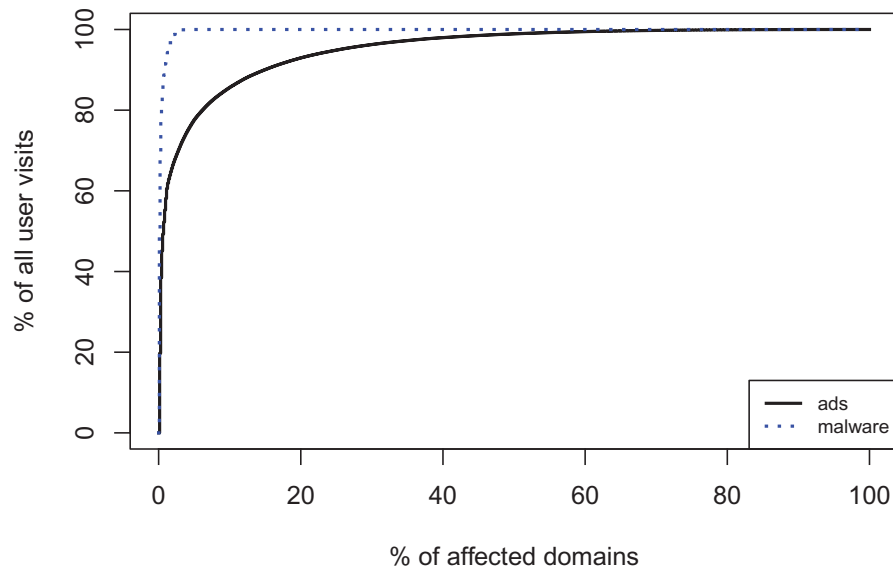


FIGURE 7.6: CDF of visits for domains used to transmit malware or ads in the search results of trending-terms.

less likely to be manipulated. At the same time, the figure demonstrates that even the odd success can reel in many victims.

Figure 7.6 plots the CDF of user visits compared to the affected domains. The graph indicates high concentration—most of the traffic is drawn to a small number of domains. The concentration of visitors is particularly extreme for malware, which makes sense given the spikes observed in Figure 7.5. The concentration in MFA sites shows that a few websites profit handsomely from trending-terms, and that many more are less successful. This is consistent with our earlier finding that there are only a few large connected clusters of MFA sites linking to each other. One consequence of this concentration is that we can approximate the revenue to the biggest players simply by considering aggregate figures.

7.3.2 Revenue analysis

We next compare revenues miscreants generate from MFA sites and from malware-hosting sites.

MFA revenue. Essentially, the aggregate revenue for MFA websites is a sum of the revenues generated by all MFA sites w obtained in response to all the search terms s considered. Each website generates a revenue equal to the number of website visitors times the advertising revenue that can be obtained from these visitors:

$$R_{\text{MFA}}(t) = \sum_{w \in \text{MFA}(s)} \sum_s V(w, s, t) \cdot (p_{\text{PPC}} \cdot p_{\text{clk}} \cdot r_{\text{PPC}} + p_{\text{banner}} \cdot r_{\text{banner}} + p_{\text{aff}} \cdot p'_{\text{clk}} \cdot r_{\text{aff}}) \quad (7.1)$$

There are three broad classes of online advertising in use on MFA domains—Pay-per-Click (PPC) (e.g., Google AdSense), banners (e.g., Yahoo! Right Media) and affiliate marketing (e.g., Commission Junction). Banner advertisements are paid r_{banner} by the visit, PPC only pays r_{PPC} when the user clicks on an ad (which happens with probability p_{clk}), and affiliate marketing pays r_{aff} whenever a visitor clicks the ad and then buys something (which happens with probability p'_{clk}). By inspecting our corpus of MFA sites, we discover that 83% include PPC ads, 66% use banner ads, and 16% include affiliate ads. 50% of sites use two types of advertising, and 7% use all three. We include each type of advertisement in the revenue calculation with probability $p_{\text{ad type}}$, and we assign the probability according to the percentage of MFA site visits that include each class of ad. For the MFA websites we have identified, $p_{\text{PPC}} = 0.94$, $p_{\text{banner}} = 0.53$, and $p_{\text{aff}} = 0.33$.

To calculate the earning potential of each ad type, we piece together rough measures gathered from outside sources. Estimating the Click-Through Rate (CTR) p_{clk} is difficult, as click-through rates vary greatly, and ad platforms such

as Google keep very tight-lipped on average click-through rates. One Google employee reported that an average CTR is “in the neighborhood of 2%” [197]. We anticipate that the CTR for MFA sites is substantially higher than 2%, since sites have multiple ads aggressively displayed and little original content. Nonetheless, we assign $p_{\text{clk}} = 0.02$.

To measure per-click ad revenue r_{PPC} , we turn to the CPC estimates Google provides for advertising keywords. We expect that more persistent search terms are likely to appear as keywords for ads, even on websites about trending-terms. Hence, we assume that advertising revenue for trending-terms matches the CPC for most popular keywords in the corresponding category. We assign the expected advertising revenue to the mean of ad prices for the 20 most popular search terms weighted by the amount each category is represented in the results from the trending set (see Table 7.4, column 1). This yields $r_{\text{PPC}} = \$0.97$.

Calculating banner advertising revenue is a bit easier, since no clicks are required to earn money. Public estimates of average revenue are hard to come by, but the ad network Adify issued a press release stating that its median cost per 1,000 impressions in Q2 2010 was \$5.29 [2], so we assign $r_{\text{banner}} = \$0.00529$.

For affiliate marketing, we assume that $p'_{\text{clk}} = p_{\text{clk}} = 0.02$, the same as for PPC ads. To estimate the revenue r_{aff} that can be earned, we turn to Commission Junction (CJ), one of the largest affiliate marketing networks that matches over 2,500 advertisers with affiliates. CJ provides an estimate of expected earnings from advertisers per 100 clicks; we collected this estimate for all advertisers on Commission Junction in December 2009, and found it to be \$26.49. Consequently, we estimate that $r_{\text{aff}} = \$0.265$.

Putting it all together, we estimate the monthly revenue to MFA sites to be:

$$\begin{aligned}
 R_{\text{MFA}}(1 \text{ month}) &= 4,284,458 \times (0.94 \times 0.02 \times \$0.97 \\
 &\quad + 0.53 \times \$0.00529 + 0.33 \times 0.02 \times \$0.265) \\
 &= \$97,637 .
 \end{aligned}$$

So, MFA sites gross roughly \$100,000 per month from trending-term exploitation. There are, however, costs that are not factored into the above derivation, which makes it an upper bound. For instance, Google generally imposes a 32% fee on advertising revenues [154]. Furthermore, servers have to be hosted and maintained. As an example, most sites in the largest cluster in Section 7.2.2 are hosted by the same service provider, which charges \$140/server/month. That cluster contains 193 nodes hosted on 155 unique servers, which, ignoring economies of scale, would come up to \$21,700/month in maintenance. Nevertheless, it is worth noting that these costs can be amortized over other businesses—it is unlikely that such servers are only set up for the purpose of trending-term exploitation.

Malware revenue. Attackers have experimented with several different business models to monetize drive-by-downloads, from adware to credential-stealing trojans [186]. However, researchers have observed that attackers exploiting trending-terms have tended to rely on fake antivirus software [49, 188, 209]. We therefore define the revenue due to malware in trending results as:

$$R_{\text{mal}}(t) = \sum_{w \in \text{mal}(s)} \sum_s V(w, s, t) \cdot p_{\text{exp}} \cdot p_{\text{pay}} * r_{\text{AV}} \quad (7.2)$$

where we multiply the number of visits times the likelihood of exposure, the probability of a victim paying for the software, and the amount paid. For these figures, we turn to the analysis of Stone-Gross et al. [209], who acquired a copy of

back-end databases detailing the revenues and expenses of three large fake antivirus programs, each of which were advertised by compromising trending search results. They found that 2.16% of all users exposed to fake antivirus ultimately paid for a “license,” at an average cost of \$58. We can use these figures directly in our model for the revenues due to malware, setting $p_{\text{pay}} = 0.0216$ and $r_{\text{AV}} = \$58$.

Unlike most drive-by-downloads, fake antivirus software does not need to exploit a vulnerability in the client visiting the infected search result in order for a user to be exposed. Instead, the server will use a server-side warning designed to appear as though it is on the client’s machine, and then prompt a user to install software [188]. Because of this, every user that visits a link distributing fake antivirus is exposed, and so we assign $p_{\text{exp}} = 1$. These parameters yield a monthly revenue from malware of:

$$R_{\text{mal}}(1 \text{ month}) = 48,975 \times 1 \times 0.0216 \times 58 \approx \$61,356 \quad (7.3)$$

Thus, malware sites (e.g., fake antivirus sites) generate roughly \$60,000/month just from trending-term exploitation.

Here too, there are costs associated with deploying these sites, but server maintenance is a lot cheaper than in the case of MFA sites, given that most machines hosting malware have been compromised rather than purchased. Bots go for less than a dollar [31, and references therein], while a compromised server—presumably with high quality network access—goes at most for \$25 according to Franklin et al. [73]. Note that we do not adjust the returns on malware for the risk of being caught because the likelihood of being arrested for cyber-criminal activity is currently negligible in many jurisdictions where cyber-criminals operate.

One conclusion of this analysis is that malware and MFA hosting have quite different revenue models, but yield surprisingly similar amounts of money to their per-




FIGURE 7.7: MFA prevalence in the top 10 search results fell after Google announced changes to its ranking algorithm on February 24, 2011, designed to counter “low-quality” results.

petrators. This lends further support to the hypothesis that they could be treated as substitutes.

7.4 Search-engine intervention

On February 24, 2011, following a series of high-profile reports of manipulation of its search engine (e.g., [199, 200]), Google announced changes to its search ranking algorithm designed to eradicate “low-quality” results [203]. Google defined low-quality sites as those which are “low-value add for users, copy content from other websites or sites that are just not very useful.” The MFA sites examined in this chapter certainly appear to match that definition. Because we were already collecting search results on the trending set, we can measure the effectiveness of the intervention in eradicating abuse targeting trending-terms.

Figure 7.7 plots over time the average percentage of top 10 search results marked as MFA for terms in the trending set. From July to February, 3.1% of Google's top 10 results (solid line) for trending-terms pointed to MFA sites, compared to 2.0% for Yahoo!'s top 10 results (dotted line). The vertical dashed line marks February 24, 2011, the day of Google's announcement. The proportion of MFA sites quickly fell, stabilizing a month later at a rate of 0.47% for Google. Curiously, Yahoo!'s share of top 10 MFA results also fell, to an average of 0.56%.

Landing in the top results tells only part of the story. The underlying popularity of the trending-terms is also important. We compute the estimated site visits to MFA sites, which is more directly tied to revenue. Table 7.6 compares the number of visits referred to by Google and Yahoo! search results before and after the intervention. Between July 24, 2010 and February 24, 2011, MFA sites attracted 4.67 million monthly visits on average. Between March 10 and April 24, 2011, the monthly rate fell 31% to 3.2 million.

However, the changes differed greatly across search engines. Referrals from Google search results fell by 47%, while on Yahoo! and Bing the visits increased by 11%. The table also distinguishes between whether the MFA site uses Google ads or another provider. 81% of MFA sites show Google ads, which is not surprising given Google's dominance in PPC advertising. It is an open question whether Google might treat MFA sites hosting its own ads differently than sites with other ads. Striking them from the search results reduces Google's own advertising revenue. However, it is in Google's interest to provide high-quality search results, the amount of foregone revenue is small, and is likely to be partly replaced by other search results. Our figures support the latter rationale. Sites with Google ads fell by 1.2 million visits, or 41%. Visits to sites not using AdSense fell by 91%, but, in absolute terms, the reduction was smaller than for sites with Google ads. By contrast, Yahoo! results with Google ads rose by 18%.

Table 7.6: Estimated number of visits to MFA and malware sites for trending terms.

	Monthly MFA visits		
	Pre-intervention	Post-intervention	% change
Google search	3,364,402	1,788,480	-47%
<i>Google ads</i>	2,989,821	1,763,709	-41%
<i>Other ads</i>	374,556	24,770	-93%
Yahoo!/Bing search	1,302,314	1,448,058	+11%
<i>Google ads</i>	1,204,928	1,424,323	+18%
<i>Other ads</i>	95,363	23,734	-75%
Total	4,666,716	3,236,538	-31%

Using the pre- and post-intervention MFA visit rates into the revenue equations developed in Section 7.3.2, the average monthly take for MFA sites has fallen from \$106,000 to \$74,000. If this reduction holds over time, what are the implications for miscreants? First, they may decide to devote more effort to manipulating Yahoo! and Bing, despite their lower market penetration, since the MFA revenues are growing more equitable in absolute terms. Second, malware becomes more attractive as an alternative source of revenue, so one unintended consequence of the intervention to improve search quality could be to foster more overtly criminal activities harming consumers. Third, revenue models based on advertising require volume, and external efforts that reduce traffic levels can cause significant pain to the miscreant. By contrast, malware offers substantially more expected revenue per visitor, and is therefore likely to be much more difficult to eradicate.

Given the striking change in MFA prevalence following Google’s intervention, it is worth checking whether this intervention alters the significance of the empirical conclusions reached in Section 7.2.4. We included a dummy variable into the MFA regression reflecting whether Google’s intervention had yet occurred, and

found that this inclusion does not alter the significance of the dependent variables presented in Section 7.2.4.

Finally, we contrast the success of Google’s intervention in reducing the profitability of trending-term exploitation with the inability of the same intervention to affect search-redirection attacks, as discussed in Chapter 6. Specifically, in the previous chapter we showed that this intervention was unable to limit the long-term prevalence of compromised websites redirecting traffic to unlicensed online pharmacies (Section 6.3.2). However, here we present evidence of a reduction of web traffic landing at MFA websites by 31%. We attribute these contradictory findings to the limited duration of the measurements we analyze in this chapter. Indeed, in Section 6.3.2 we show that the number of search-redirecting results dropped immediately after the change in the ranking algorithm, but, within a few weeks, such results appeared much more prominently in the search results. Therefore, we conjecture that if our measurements on trending-term manipulation lasted longer, we would be able to observe a similar trend. Online crime is dynamic in nature, and online criminals adapt to and circumvent deployed countermeasures to their benefit. This observation highlights the importance of longitudinal measurements in the study and understanding of online crime.

7.5 Conclusion

In this chapter we have presented our large-scale investigation into the abuse of “trending” terms, focusing on the two primary methods of monetization: malware and ads. We have found that the dynamic nature of the trends creates a narrow opportunity that is being effectively exploited on web search engines, and social-media platforms. We have presented statistical evidence that the less popular and less financially lucrative terms are exploited most effectively. In addition, we

found that the spoils of abuse are highly concentrated among a few players. We have developed an empirically grounded model of the earnings potential of both malware and ads, finding that each attracts aggregate revenues on the order of \$100,000 per month. Finally, we have found that Google's intervention to combat low-quality sites has likely reduced revenues from trend exploitation by more than 30%.

There is a connection in our economic modeling to the battle over how to profit from typosquatting [160]. In both cases, Internet "bottom feeders" seek to siphon off a fraction of legitimate traffic at large scale. Several years ago, typosquatting was used in phishing attacks and to distribute malware. Today, however, typosquatting is almost exclusively monetized through PPC and affiliate marketing ads [160], attracting hundreds of millions of dollars in advertising revenue to domain squatters via ad platforms.

The open question is whether a significant crackdown on, say, fake antivirus sales, will simply shift the economics in favor of low-quality advertising. However, while ad platforms might tolerate placing ads on typosquatted websites, advertising that lowers the quality of search results directly threatens the ad platform's core business of web search. Consequently, we are more optimistic that search engines might be willing to crack down on all abuses of trending terms, as we have found in our initial data analysis. However, we acknowledge that this optimism is constricted by the limited duration of the measurements we analyzed in this chapter, which may overestimate the success of search engine interventions. To this end, in Chapter 9 we explore the effectiveness of various intervention strategies towards a long-term reduction of this illicit activity.

Empirically measuring WHOIS misuse

In the previous chapters we examined a set of cases of online crime with rather complex characteristics in terms of the underlying criminal networks supporting their operation and monetization. However, one of our main arguments in this thesis is that online crime—similar to traditional crime [35, 36]—is enabled by the availability of opportunities to victimize a vulnerable target, rather than by the technical sophistication of criminal operations. We argue that the degree of sophistication impacts only the level of commitment and expertise required to characterize the criminal infrastructures. In turn, this derived understanding needs to be “translated” into a set of available opportunities which should be targeted with appropriate countermeasures. In this chapter, we examine WHOIS misuse, a rather simple case of online crime, in an effort to show that, as long as opportunity exists, online criminals do not need to employ overly elaborate technical skills. The characteristics of WHOIS misuse, as we show, are appropriately simple, stripped off the technical sophistication that characterizes the previous cases of online crime, allowing us to focus mainly on the enabling opportunities.

WHOIS is an online directory that primarily allows anyone to map domain names to the registrants' contact information [53]. Based on their operational agreement with the Internet Corporation for Assigned Names and Numbers (ICANN) [101], all global Top Level Domain (gTLD) *registrars*¹ are required to collect this information during domain registration, and subsequently publish it into the WHOIS directory; how it is published depends on the specific *registry*² used. While the original purpose of WHOIS was to provide the necessary information to contact a registrant for legitimate purposes—e.g. abuse notifications, or other operational reasons—there has been increasing anecdotal evidence of misuse of the data made publicly available through the WHOIS service. For instance, some registrants³ have reported that third-parties used their publicly available WHOIS information to register domains similar to the reporting registrants', using contact details identical to the legitimate registrants'. The domains registered with the fraudulently acquired registrant information were subsequently used to impersonate the owners of the original domains.

Elliot in [63] provides an extensive overview of issues related to WHOIS. Researchers use WHOIS to study the characteristics of various online criminal activities, like click fraud [33, 55] and botnets [253], and have been able to gain key insights on malicious web infrastructures [128, 135]. From an operational perspective, the FBI has noted the importance of WHOIS in identifying criminals, but the presence of significant inaccuracies hinder such efforts [225]. Moreover, online criminals often use privacy or proxy registration services to register malicious domains, complicating further their identification through WHOIS [39].

¹ Registrars are entities that process individual domain name registration requests

² Registries are entities responsible for maintaining an authoritative list of domain names registered in each gTLD

³ <http://www.eweek.com/c/a/Security/Whois-Abuse-Still-Out-of-Control>

ICANN has acknowledged the issue of inaccurate information in WHOIS [245], and has funded research towards measuring the extent of the problem [175]. ICANN's Generic Names Supporting Organization (GNSO), which is responsible for developing WHOIS-related policies, identified in [104] the possibility of misuse of WHOIS for phishing and identity theft, among others. Nevertheless, ICANN has been criticized [63, 243] for its inability to enforce related policies.

This sad state of affairs brings into question whether the existence of the WHOIS service is even needed in its current form. One suggestion is to promote the use of a structured channel for WHOIS information exchange, capable of authenticated access, using already available web technologies [98, 173, 211]. An alternate avenue is to completely abandon WHOIS, in favor of a new Registration Data Service. This service would allow access to verified WHOIS-like information only to a set of authenticated users, and for a specific set of *permissible purposes* [65].

The work we present in this Chapter attempts to illuminate this policy discussion by empirically characterizing the extent to which WHOIS misuse occurs, and which factors are statistically correlated with WHOIS misuse incidents [127]. In addition, we provide a quantitative and qualitative assessment of the types of WHOIS misuse experienced by domain name registrants, the magnitude of these misuse cases and defense measures—i.e. anti-harvesting mechanisms—that may impact misuse. A separate three-month measurement study from ICANN's Security and Stability Advisory Committee (SSAC) [103] examined the potential of misuse of email addresses posted exclusively in WHOIS. The authors registered a set of domain names composed as random strings, and monitored the electronic mailboxes appearing in the domains' WHOIS records for spam emails, finding WHOIS to be a contributing factor to received spam. We generalize this work with a much more comprehensive study using 400 domains across the five largest global top

level domains .com, .net, .org, .info, and .biz) which, in aggregate, are home to more than 127 million domains [100]. In addition, we not only look at email spam but also at other forms of misuse (e.g., of phone numbers or postal addresses).

The initial motivation of this research was to respond to the decision of ICANN's GNSO to pursue WHOIS studies [78] to scientifically determine if there is substantial WHOIS misuse warranting further action from ICANN. However, in the context of this thesis, it provides a proof-of-concept for the fact that as long as opportunities exist, online criminals do not need to employ overly elaborate technical skills to profit illicitly, addressing research question 4.⁴

We validate the hypothesis that public access to WHOIS leads to a measurable degree of misuse, identify the major types of misuse, and, through regression analysis, discover factors that have a statistically-significant impact on the occurrence of misuse. Most importantly, we prove that a mere reduction in the availability of opportunities to engage in WHOIS misuse through the implementation of appropriate anti-harvesting measures, can thwart this fraudulent activity.

The remainder of this chapter is organized as follows. We discuss our methodology in Sections 8.1 and 8.2. We present a breakdown of the measured misuse in Section 8.3, and the deployed WHOIS anti-harvesting countermeasures in Section 8.4. We perform a regression analysis of the characteristics affecting the misuse in Section 8.5, note the limitations of our work in Section 8.6, and conclude in Section 8.7.

⁴ Is it the technical skills or the existence of opportunities enabling. . .

Table 8.1: Number of domains under each of the five global Top Level Domains within scope in March 2011 [100].

gTLD	.com	.net	.org	.info	.biz	Total
# of domains	95,185,529	14,078,829	9,021,350	7,486,088	2,127,857	127,694,306
Proportion in population	75.54%	11.03%	7.06%	5.86%	1.67%	100%

8.1 Methodology

To whittle down the number of possible design parameters for our measurement experiment, we first conducted a pilot survey of domain registrants to collect experiences of WHOIS misuse. We then used the results from this survey to design our measurement experiment.

8.1.1 *Constructing a microcosm sample*

In November of 2011 we received from ICANN, per our request, a sample set of 6,000 domains, collected randomly from gTLD zone files with equal probability of selection. Of those 6,000 domains, 83 were not within the five gTLDs we study, and were discarded. Additionally, ICANN provided the WHOIS records associated with 98.7% (5,921) of the domains, obtained over a period of 18 hours on the day following the generation of the domain sample.

Out of these nearly 6,000 domains, we created a proportional probability microcosm of 2,905 domains representative of the population of 127 million domains, using the proportions in Table 8.1. In deciding the size of the microcosm we use as a baseline the 2,400 domains used in previous work [175], and factor in the evolution in domain population from 2009 to 2011.

Finally, we randomly sampled the domain microcosm to building a representative sample of $D = 1,619$ domains from 89 countries—country information is available through WHOIS.

8.1.2 Pilot registrant survey

We use the domains' WHOIS information to identify and survey the 1,619 registrants associated with domains in D , about their experiences on WHOIS misuse. Further details on the survey questions, methodology, and sample demographics are available in Appendix A.

Despite providing incentives for response (participation in a random drawing to be eligible for prizes such as iPads or iPods) we only collected a total of 57 responses, representing 3.4% of contacted registrants. As a result, this survey could only be used to understand some general trends, but the data was too coarse to obtain detailed insights.

With the actual margin of error at 12.7%, 43.9% of registrants claim to have experienced some type of WHOIS misuse, indicating that the public availability of WHOIS data leads to a measurable degree of misuse. The registrants reported that email, postal, and phone spam were the major effects of misuse, with other types of misuse (e.g. identity theft) occurring at insignificant rates.

These observations are based on limited, self-reported data, and respondents may incorrectly attribute misuse to WHOIS. Nevertheless, the pilot survey tells us that accurately measuring WHOIS misuse requires to primarily look at the potential for spam, not limited to email spam, but also including phone and postal spam.

8.1.3 Experimental measurements

We create a set of 400 domain names and register them at 16 registrars (25 domains per registrar) across the five gTLDs, with artificial registrant identities. Each artificial identity consists of (i) a full name (i.e. first and last name), (ii) an email address, (iii) a postal address, and (iv) a phone number.

All registrants' contact details are created solely for the purpose of this experiment, ensuring that they are only published in WHOIS. Through this approach, we eliminated confounding variables. From the moment we register each experimental domain, and the artificial identity details become public through WHOIS, we monitor all channels of communication associated with every registrant. We then classify all types of communication and measure the extent of illicit or harmful activity attributed to WHOIS misuse targeting these registrants.

Given the wide variety of registrars and the use of unique artificial identities, the registration process did not lend itself to automation and was primarily manual. We registered the experimental domains starting in the last week of June 2012, and completed the registrations within four weeks. We then monitored all incoming communications over a period of six months, until the last week of January 2013. All experimental domains were registered using commercial services offered by the 16 registrars; we did not use free solutions like DynDNS.

8.2 Experimental domain registrations

We associated the WHOIS records of each of the 400 domains with a unique registrant identity. Whenever the registration process required the inclusion of an organization as part of the registrant information, we used the name of the domain's registrant. In addition, within each domain, we used the registrant's identity (i.e. name, postal/email address, and phone numbers) for all types of WHOIS contacts (i.e., registrant, technical, billing, and administrative contacts).

Figure 8.1 provides a graphical breakdown of the group of 25 domains we register per registrar. Every group contains five subgroups of domains, one for each of the five gTLDs. Finally, each subgroup contains a set of five domains, one for each type of domain name, as discussed later.

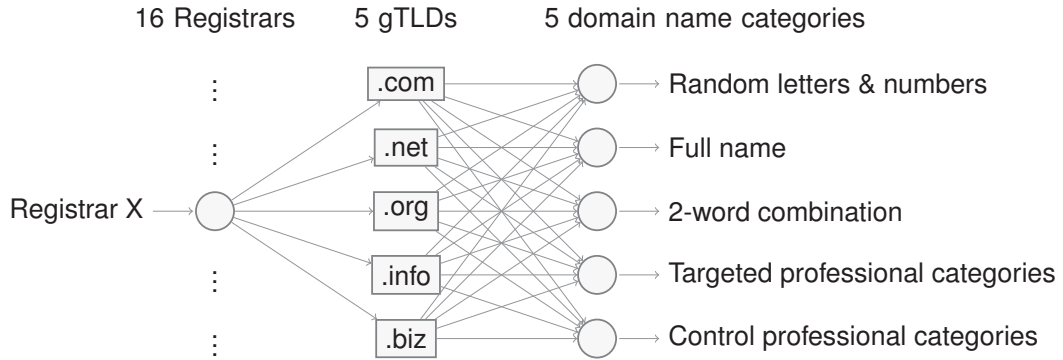


FIGURE 8.1: Graphical representation of the experimental domain name combinations we register with each of the 16 registrars.

8.2.1 Registrar selection

We selected the sixteen registrars used in our measurement study as follows. Using the WHOIS information of the 1,619 domains in D , we first identify the set R of 107 registrars used by domains in D . Some registrars only allow domain registration through “affiliates.” In these cases we attempt to identify the affiliates used by domains in D , by examining the name server information in the WHOIS records.

We then sort the registrars (or affiliates, as the case may be) based on their popularity in the registrant sample. More formally, if $D_r \subset D$ is the set of domains in the registrant sample associated with registrar r , we define r 's popularity as $S_r = |D_r|$. We sort the 107 registrars in descending order of S_r , and then select the 16 most popular registrars as the set of our *experimental registrars* that allow:

- The registration of domain names in all five gTLDs. This restriction allows us to perform comparative analysis of WHOIS misuse across the experimental registrars, and gTLDs.
- Individuals to register domains. Registrars providing domain registration services only to legal entities (e.g. companies) are excluded from consideration.

- The purchase of a single domain name, without requiring purchasing of other services for that domain (e.g. hosting).
- The purchase of domains without requiring any proof of identity. Given our intention to use artificial registrant identities, a failure to hide our identity could compromise the validity of our findings.

8.2.2 Experimental domain name categories

We study the relationship between the category of a domain name, and WHOIS misuse. Specifically, we examine the following set of name categories:

1. Completely random domain names, composed by 5 to 20 random letters and numbers (e.g. unvdazzihevqnky1das7.biz).
2. Synthetic domain names, representing person full names (e.g. randall-bilbo.com).
3. Synthetic domain names composed by two randomly selected words from the English vocabulary (e.g. neatlimbed.net).
4. Synthetic Domain names intended to look like businesses within specific professional categories (e.g. hiphotels.biz).

To construct the last category, we identify professional categories usually targeted in cases of spear-phishing and spam, by consulting two sources. We primarily use the “Phishing Activity Trend” report, periodically published by the Anti-Phishing Working Group (APWG) [10]. We identify the professional categories mostly targeted by spam and phishing in the second quarter of 2010 with percentages of more than 4% in total. These categories are: (i) Financial services, (ii) payment services, (iii) gaming, (iv) auctions, and (v) social networking. We

complement this list with the following professional categories appearing in the subject and sender portions of spam emails we had previously received: (i) medical services, (ii) medical equipment, (iii) hotels, (iv) traveling, and (v) delivery and shipping services.

In addition, we define a control set of professional categories that are not known to be explicitly targeted. We use the control set to measure the potential statistical significance of misuse associated with any of the previous categories. The three categories in the control set are : (i) technology, (ii) education, and (iii) weapons.

8.2.3 Registrant identities

We create a set of 400 unique artificial registrant identities, one for each of the experimental domains. Our ultimate goal is to be able to associate every instance of misuse with a single domain, or a small set of domains.

A WHOIS record created during domain registration contains the following publicly available pieces of registrant information: (i) full name, (ii) postal address, (iii) phone number, and (iv) email address. In this section we provide the design details of each portion of the artificial registrant identities.

Registrant name. The registrant's full name (i.e. first name-last name) serves as the unique association between an experimental domain and an artificial registrant identity. Therefore we need to ensure that every full name associated with each of the 400 experimental domains is unique within this context.

We create the set of 400 unique full names, indistinguishable from names of real persons, by assembling common first names (male and female) and last names with Latin characters.

Email address. We create a unique email address for each experimental domain in the form *contact@example.com*. We use this email address in the domain’s WHOIS records, and we therefore call it *public email address*.

However, any email sent to a recipient other than *contact* (e.g. *foo@example.com*), is still collected for later analysis under a *catchall* account. We refer to these as *unpublished email addresses*, as we do not publish them anywhere, including WHOIS.

Mail exchange (MX) records are a type of DNS record pointing to the email server(s) responsible for handling incoming emails for a given domain name [153]. The MX records for our experimental domains all point to a single IP address functioning as a proxy server. The proxy server, in turn, aggregates and forwards all incoming requests to an email server under our control. The use of a proxy allows us to conceal where the “real” email server is located (i.e., at our university); our email server functions as a spam trap—i.e., any potential spam mitigation at the network- or host-level is explicitly disabled.

Postal address. We examined the possibility of using a postal mail-forwarding service to register residential addresses around the world. Unfortunately, and, given the scale of this experiment, we were unable to identify a reasonably-priced and legal solution.

In most countries—the US included—such services often require proof of identification prior to opening a mailbox,⁵ and limit the number of recipients that can receive mail at one mailbox. Moreover, we were hesitant to trust mail-forwarding services from privately owned service providers,⁶ because the entities providing such services may themselves misuse the postal addresses, contaminating our mea-

⁵ For example United States Postal Service (USPS) form 1583: *Application for Delivery of Mail Through Agent* in the US.

⁶ Also known as “virtual office” services.

surements. For example, merely requesting a quote from one service provider, resulted in our emails being placed on marketing mailing lists without our explicit consent.

We eventually decided to use three Post Office (PO) boxes within the US; and, randomly assigned to each registrant identity one of these addresses. Traditionally, the address of a PO box with number *123* is of the following format: *PO Box 123, City, Zip code*. However, we utilize the *street addressing* service offered by USPS to camouflage our PO boxes as residential addresses. Street addressing enables the use of the post office's street address to reach a PO box located at the specific post office. Through this service, the PO box located at a post office with address *456 Imaginary avenue*, is addressable at *456 Imaginary avenue #123, City, Zip code*.

In addition, PO boxes are typically bound to the name of the person who registered them. However, each experimental domain is associated with a unique registrant name, even when sharing the same postal address, different than the owner of the PO box. We evaluated possible implications of this design in receiving postal mail to a PO box addressee not listed as the PO box owner. We originally acquired five PO boxes across two different US states, and sent one letter addressed to a random name to each of these PO boxes. We successfully received letters at three of the PO boxes indicating that mail addressed to any of the artificial registrant names would be delivered successfully. The test failed at the other two PO boxes—we got back our original test letters marked as undeliverable—making them unsuitable for the study.

Phone number. Maintaining individual phone numbers for each of the 400 domains over a period of six months would be prohibitively expensive. Instead, we group the 400 domains into 80 sets of domains having the same gTLD and reg-

istrar, and we assign one phone number per such group. For example all .com domains registered with GoDaddy share the same phone number.

We acquire 80 US-based phone numbers using Skype Manager⁷ with area codes matching the physical locations of the three PO boxes. We further assign phone numbers to registrant identities with area codes matching their associated PO box locations.

8.3 Breaking down the measured misuse

In this section we present a breakdown of the empirical data revealing WHOIS-attributed misuse. The types of misuse we identify fall within three categories: (1) *postal address misuse*, measured as postal spam, (2) *phone number misuse*, measured as voice mail spam, and (3) *email address misuse*, measured as email spam.

8.3.1 *Postal address misuse*

We monitor the contents of the three PO boxes biweekly, and categorize the collected mail either as *generic spam* or *targeted spam*. Generic spam is mail not associated with WHOIS misuse, while targeted spam can be directly attributed to the domain registration activity of the artificial registrant identities.

When postal mail does not explicitly mention the name of the recipient, we do not associate it with WHOIS misuse, and we classify it as generic spam. Common examples in this category are mail addressed to the “PO Box holder”, or to an addressee not in the list of monitored identities.

In total, we collected 34 pieces of generic spam, with two out of the three PO boxes receiving the first kind of generic spam frequently. Additionally, we collected

⁷ <http://www.skype.com/en/features/skype-manager/>

Please make checks payable to [redacted]
 Please write your reference number on the front of your check
 Enclose check in the addressed envelope provided
 DO NOT SEND CASH

WEBSITE ADDRESS LISTING INCLUDES:
 Domain name submission with up to 8 keywords / phrases to 20 established search engines
 Initial search engine position and ranking report sent to you via e-mail
 Monthly search engine position and ranking reports sent to you via e-mail
 Complete details are located online at [redacted]

PAYMENT INFORMATION:
 To ensure listing by October 19, 2012, please remit payment on or before October 12, 2012.
 All listings are final

Current payment details

	Amount	Total
Annual Listing (October 19, 2012 to October 19, 2013)	\$85.00	\$85.00
Total		\$85.00

(a) Advertisement of search engine optimization services.

No special equipment
 Print your cards
 Instantly print official
 mail class, for anyth
 Confirmation™ with
 postage discount

Mail and shipping
 Drop your letters at
 mail carrier, or sche
 our toll-free Custom

Print postage in 3 simple steps...

1. Insert the enclosed CD.
2. Create an account (No commitment. No upfront costs. Cancel anytime.)
3. Enter promo code found above your mailing address on this mailer.

ACT NOW! \$100 BONUS OFFER EXPIRES SOON

PEEL & SAVE ▶



CLICK
Instantly buy and calculate exact postage.

PRINT
Print postage on envelopes, labels or plain paper.

POSTAGE
Affix postage and mail anywhere in the world.

(b) Advertisement of postal and shipping services.

FIGURE 8.2: Targeted postal spam attributed to WHOIS misuse.

four instances of the second type of generic spam, received at a single PO box. A reasonable explanation for the latter is that previous owners of the PO box still had mail sent to that location.

Postal mail is placed in the targeted spam category when it is addressed to the name and postal address of one the of the artificial registrant identities. We observed targeted spam at a much lower scale compared to the generic spam, with a total of four instances.

Two instances of targeted postal spam, were sent to two different PO locations, but were identical in terms of (i) their sender, (ii) the advertised services, (iii) the date of collection from the PO boxes, and (iv) the posting date. The purpose of the letters, as shown in Figure 8.2(a), was to sell domain advertising services. This advertising scheme works with the registrant issuing a one-time payment for \$85 USD, in exchange for the submission of the registrant's domain to search engines in combination with search engine optimization (SEO) on the domains. The two experimental domains subjected to this postal misuse were registered using the same registrar, but under different registrant identities, and gTLDs.

The purpose of the third piece of targeted postal spam (Figure 8.2(b)) was to enroll the recipient in a membership program that provides postal and shipping services. Finally, the fourth piece of postal mail spam was received very close to the end of the experiment and offered a free product in exchange for signing up on a website.

Overall, the volume of targeted WHOIS postal spam is very low (10%), compared to the portion classified as generic spam (90%). However, this is possibly due to the small geographical diversity of the PO boxes.

8.3.2 Phone number misuse

We collected 674 voicemails throughout the experiment. We define the following five types of content indicative of their association—or lack thereof—to WHOIS misuse, and manually classify each voicemail into one of these five categories:

WHOIS-attributed spam Unsolicited calls offering web-related services (e.g. website advertising), or mentioning an experimental domain name or artificial registrant name.

Possible spam Unsolicited phone calls advertising services that cannot be associated with WHOIS misuse, given the previous criteria. (e.g. credit card enrollment based on random number calling)

Interactive spam Special case of possible spam with a fixed recorded message saying “press one to accept”.

Blank Voice mails having no content, or with incomprehensible content.

Not spam Accidental calls, usually associated with misdialing, or with a caller having wrong contact information (e.g. confirmation for dental appointment)

Two of these categories require further explanation. First, in the case of *possible spam*, we cannot tell if the caller harvested the number from WHOIS, or if it was obtained in some other way (e.g., exhaustive dialing of known families of phone numbers). We therefore take the conservative approach of placing such calls in a category separate from WHOIS-attributed spam. Second, calls marked as *interactive spam* did not contain enough content to allow for proper characterization of the messages. However, the large number of these calls—received several times a day, starting in the second month of the experiment—suggests a malicious intent.

Of the 674 voicemails, we classify 5.8% as WHOIS-attributed spam, 4.2% as possible spam, 38% as interactive spam, and 15% as not spam. Finally, we classify 36.9% of voicemails as blank due to their lack of intelligible content.

Of the 39 pieces of WHOIS-attributed spam, 77% (30) originated from a single company promoting website advertising services. This caller placed two phone calls in each of the numbers, one as an initial contact and one as a follow up. These calls targeted `.biz` domains registered with 5 registrars, `.com` domains registered with 4 registrars, and `.info` domains registered with 6 registrars. In total, the specific company contacted the registrants of domains registered with 11 out of the 16 registrars.

The remaining spam calls targeted `.biz` domains registered with 4 registrars, `.com` domains registered with 4 registrars, and `.info`, `.net`, and `.org` domains associated with 1 registrar each. In one case we observed a particularly elaborate attempt to acquire some of the registrant's personally identifiable information.

8.3.3 *Email address misuse*

We classify incoming email either as solicited or spam, using the definition of spam in [214]. In short, an email is classified as spam if (i) it is unsolicited, and

(ii) the recipient has not provided an explicit consent to receive such email. For this experiment, this means that all incoming email is treated as spam, except when it originates from the associated registrars (e.g., for billing).

The contract between registrar and registrant, established upon domain registration, usually permits registrars to contact registrants for various reasons (e.g. account related, promotions, etc.). We identify such email by examining the headers of the emails received at the public addresses, and comparing the domain part of the sender's email address to the registrar's domain.

However, under the Registrar Accreditation Agreement (RAA) [101], ICANN-accredited registrars are prohibited from allowing the use of registrant information for marketing, or otherwise unsolicited purposes. Nevertheless, we acknowledge the possibility that some registrars may share registrant information with third parties that may initiate such unsolicited communication. We do not distinguish between registrars that engage in such practices and those that do not, and we classify all communications originating from a party other than the registrar as spam.

Throughout the experiment, published email addresses received 7,609 unsolicited emails out of which 7,221 (95%) are classified as spam. Of the 400 experimental domains, 95% received unsolicited emails in their published addresses with 71% of those receiving spam email. Interestingly, 80% of spam emails targeted the 25 domains of a single registrar.

In an effort to explain this outlier, we reviewed the terms of domain registration for all 16 registrars. We discovered that four registrars (including the registrar that appears as an outlier) mention in their registrant agreements the possibility of use of WHOIS data for marketing purposes. Since this is only a hypothesis, we do not factor it into the regression analysis we propose later. It is, however, a plausible explanation for the outlier.

We classified all 1,872 emails received at the unpublished addresses as spam, targeting 15% of the experimental domains. Since the unpublished addresses are not shared in any way, all emails received are unsolicited, and therefore counted as spam, including some that may have been the result of the spammers attempting some known account guessing techniques.

Two domains received a disproportionate amount of spam in their unpublished mailboxes. We ascribed this to the possibility that (i) these domains had been previously registered, and (ii) the previous domain owners are the targets of the observed spam activity. Historical WHOIS records confirm that both domains had been previously registered—12 years prior, and 5 years prior, respectively—which lends further credence to our hypothesis.

We examine the difference in proportions of email spam between published and unpublished addresses. Using the χ^2 test, we find that the difference is statistically significant considering the gTLD ($p < 0.05$), and the registrar ($p < 0.001$), but not the domain name category ($p > 0.05$).

Attempted malware delivery

We use VirusTotal [238] to detect malicious software received as email file attachments during the first 4 months of the experiment. In total, we analyze 496 emails containing attachments. Only 2% of emails with attachments (10 in total) targeted published email addresses, and they were all innocuous. The 15.6% of emails (76 in total) containing malware, targeted exclusively unpublished addresses, and VirusTotal classified them within 12 well-known malware families. As none of the infected attachments targeted any published email address, we do not observe any WHOIS-attributed malware delivery.

Table 8.2: **Breakdown of measured WHOIS-attributed misuse, broken down by gTLD and type of misuse.** Per the experimental design (Section 8.2), each gTLD group contains 80 domains.

Type of misuse	gTLD of affected experimental domains					Total
	.com	.net	.org	.info	.biz	
Postal address misuse	1 domain	1 domain	1 domain	1 domain	–	4 domains
Phone number misuse	5.0%	1.3%	1.3%	7.5%	10.0%	5.0%
Email address misuse	60.0%	65.0%	56.3%	77.5%	93.8%	70.5%

8.3.4 Overall misuse per gTLD

In Table 8.2 we present the portion of domains affected by all three types of WHOIS misuse, broken down by gTLD and type of misuse. We find that the most prominent type of misuse is the one affecting the registrants’ email addresses, followed by phone and postal misuse. Due to the small number of occurrences of postal misuse, we present the absolute value of affected domains. For both phone and email misuse, we present the misuse as the portion of affected domains, out of the 80 experimental domains per gTLD. Clearly, email misuse is common; phone misuse is also not negligible (especially for `.biz` domains).

The stated design limitations, especially the limited number of postal addresses we use, potentially affect the rates of misuse we measure. We nevertheless find that misuse of registrant information is measurable, and causally associated with the unrestricted availability of the data through WHOIS. We acknowledge though that this causal link is only valid based on the assumption that all ICANN-accredited registrars comply with the relevant RAA provisions (e.g., no resale of the registrant data for marketing purposes), as discussed in Section 8.3.3.

8.4 WHOIS anti-harvesting

WHOIS “anti-harvesting” techniques are a proposed solution deployed at certain registrars to prevent automatic collection of WHOIS information. We next present

a set of measurements characterizing WHOIS anti-harvesting implemented at the 16 registrars and the three thick WHOIS registries.⁸ Later on we use this information to examine the correlation between measures protecting WHOIS, and the occurrence of misuse.

More specifically, we test the rate-limiting availability on port 43, which is the well-known network port used for the reception of WHOIS queries, by issuing sets of 1,000 WHOIS requests per registrar and registry, and analyzing the responses. Each set of 1,000 requests repeatedly queries information for a single domain from the set of 400 experimental domains. We use different domain names across request sets. We select domains from the `.com` and `.net` pool when testing the registrars' defenses, and from the appropriate gTLD pool when testing thick WHOIS gTLD registries.

In addition, we examine the defenses of the remaining 89 registrars in the registrar sample. In this case we query domains found in the registrant sample instead of experimental domains. In three occasions, all domains associated with three out of the 89 registrars had expired at the time we ran this experiment. Therefore, we exclude these registrars from this analysis.

The analysis of WHOIS responses reveals the following methods of data protection:

Method 1: Limit number of requests, then block further requests.

Method 2: Limit number of requests, then provide only registrant name and offer instructions to access complete the WHOIS record through a web form.

⁸ *Thick WHOIS* registries maintain a central database of all WHOIS information associated with registered domain names, and they respond directly to WHOIS queries with all available WHOIS information. From the five gTLDs under consideration, the three registries maintaining the `.biz`, `.info`, and `.org` zones are thick registries.

Method 3: Delay WHOIS responses, using a variable delay period of a few seconds.

Method 4: No defense.

In Table 8.3 we present in aggregate form the distribution of registrars and registries using each one of the four defense methods. We find that one of the three registries does not use any protection mechanism, while the remaining two take a strict rate-limiting approach. For instance, one registry employs relatively strict measures by allowing only four queries through port 43 before applying a temporary blacklist.

Only 41.6% of the experimental registrars employ rate-limiting, allowing, on average, 83 queries, before blocking additional requests. Just two registrars in this group provide information (as part of the WHOIS response message) on the duration of the block, which, in both cases, was 30 minutes. The remaining registrars either use a less strict approach (Method 2, 18.8%), or no protection at all (Method 4, 37.5%)

One registrar would not provide responses in a timely manner (method 3), causing our testing script to identify the behavior as a temporary blacklisting. It is unclear if this is an intended behavior to prevent automated queries, or if it was just a temporary glitch with the registrar.

The remaining 89 registrars (not in the experimental set) follow more or less the same pattern as our experimental set. The majority does not use any protection mechanism, and a relatively large minority uses Method 1.

8.5 Misuse estimators

We finally examine the correlation of a set of parameters (i.e. estimators) with the measured phone and email misuse, attributed to WHOIS. These estimators are

Table 8.3: Methods for protecting WHOIS information at 104 registrars and three registries.

Tested entities	Total #	Type of WHOIS harvesting defense			
		Method 1	Method 2	Method 3	Method 4
Thick WHOIS registries	3	2 (66.6%)	–	–	1 (33.3%)
Experimental registrars	16	7 (43.7%)	2 (12.5%)	1 (6.3%)	6 (37.5%)
Remaining registrars	89	37 (41.6%)	1 (1.1%)	3 (3.4%)	48 (53.9%)

descriptive of the experimental domain names, and of the respective registrars and (thick) WHOIS registries. We do not examine postal address misuse, as the number of observed incidents in this case is very small and unlikely to yield any statistically-significant findings.

More specifically, we consider the following estimators:

- β_1 : Domain gTLD.
- β_2 : Price paid for domain name acquisition.
- β_3 : Registrar used for domain registration.
- β_4 : Existence of WHOIS anti-harvesting measures at the registrar level for `.com` and `.net` domains (thin WHOIS gTLDs), and at the registry level for `.org`, `.info`, and `.biz` domains (thick WHOIS gTLDs).
- β_5 : Domain name category.

We disentangle the effect of these estimators on the prevalence of WHOIS misuse through regression analysis. We use logistic regression [99], which is a generalized linear model [172] extending linear regression. This approach allows for the response variable to be modeled through a binomial distribution given that we examine WHOIS misuse as a binary response (i.e. either the domain is a victim of misuse or not).

In addition, using a generalized linear model instead of the ordinary linear regression allows for more relaxed assumptions on the requirement for normally distributed errors. In this analysis, we use the iteratively reweighted least squares [57] method to fit the independent variables into maximum likelihood estimates of the logistic regression parameters.

Our multivariate logistic regression model takes the following form:

$$\text{logit}(p_{\text{DomainEmailMisuse}}) = \beta_0 + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_4 + \beta_5x_5 \quad (8.1)$$

$$\text{logit}(p_{\text{DomainPhoneMisuse}}) = \beta_0 + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_4 \quad (8.2)$$

Equation 8.2 does not consider β_5 as an estimator, since the experimental design does not permit the association between measured misuse and the composition of the domain name.

We considered the use of multinomial logistic regression (MLR) for the analysis of phone number misuse, given the five classes of voicemails we collected. Such regression models require a large sample size (i.e. observations of misuse in this case) to calculate statistically-significant correlations [254]. However, in the context of our experiment, the occurrence of voicemail misuse is too small to analyze with MLR.

Therefore, we reverted to using a basic logistic regression by transforming the multiple-response dependent variable into a dichotomous one. We did this by conservatively transforming observations of possible spam into observations of not spam. In addition, we did not consider the categories of interactive spam and blank, as they do not present meaningful outcomes.

All estimators, except β_2 , represent categorical variables, and they are coded as such. Specifically, we code estimators β_1 , β_3 , and β_5 as 5-part, 16-part, and 5-part categorical variables respectively, using deviation coding. Deviation cod-

ing allows us to measure the statistical significance of the categorical variables' deviation from the overall mean, instead of deviations across categories.

We code WHOIS anti-harvesting (β_4) as a dichotomous categorical variable denoting the protection of domains by any anti-harvesting technique. While the 16 registrars, and 3 thick WHOIS registries employ a variety of such techniques (Section 8.4), the binary coding enables easier statistical interpretation.

8.5.1 Estimators of email misuse

In Table 8.4 we report the statistically-significant regression coefficients, and associated odds characterizing email misuse. Overall, we find that some gTLDs, the domain price, WHOIS anti-harvesting, and domain names representing person names are good estimators of email misuse.

Domain gTLD. The email misuse measured through the experimental domain names is correlated with all gTLDs but `.info`. Specifically, the misuse at `.biz` domains is 21 times higher than the overall mean, while domains registered under the `.com`, `.net`, and `.org` gTLDs experience less misuse.

Domain price. The coefficient for β_2 means that each \$1 increase in the price of an experimental domain corresponds to a 15% decrease in the odds of the registrants experiencing misuse of their email addresses. In other words, the more expensive the registered domain is, the lesser email address misuse the registrant experiences.

The reported correlation does not represent a correlation between domain prices and differentiation in the registrars' services. Even though we did not systematically record the add-on services the 16 registrars offer, we did not observe any considerable differentiation of services based on the domain price. Most im-

Table 8.4: Statistically-significant regression coefficients affecting email address misuse (Equation 8.1).

Estimator	coefficient	odds	Std. Err.	Significance
Domain gTLD (β_1)				
.com	-1.214	0.296	0.327	$p < 0.001$
.net	-0.829	0.436	0.324	$p = 0.01$
.org	-1.131	0.322	0.318	$p < 0.001$
.biz	3.049	21.094	0.566	$p < 0.001$
Domain price (β_2)				
	-0.166	0.846	1.376	$p < 0.001$
Lack of WHOIS anti-harvesting (β_4)				
	0.846	2.332	0.356	$p = 0.01$
Domain name composition (β_5)				
Person name	-0.638	0.528	0.308	$p = 0.04$

portantly, we did not use any such service for any of the experimental domains we registered, even when such services were offered free of charge.

What this correlation may suggest is that higher domain prices may be associated with other protective mechanisms, like the use of blacklists to prevent known harvesters from unauthorized bulk access to WHOIS. However, such mechanisms are transparent to an outside observer, so we may only hypothesize on their existence and their effectiveness.

WHOIS anti-harvesting. The analysis shows that the existence of WHOIS anti-harvesting protection is statistically-significant in predicting the potential of email misuse. The possibility of experiencing email misuse without the existence of any anti-harvesting measure is 2.3 times higher than when such protection is in place.

Domain name category. We identify the category of domains denoting person names (e.g. randall-bilbo.com) as having negative correlation to misuse. In this case, the possibility of experiencing email address misuse is slightly lower than the overall mean.

This appears to be an important result. However, we point out that all the domain names in this category contain a hyphen (i.e. -), contrary to all other

Table 8.5: Statistically-significant regression coefficients in Equation 8.2.

Estimator	coefficient	odds	Std. Err.	Significance
Domain gTLD (β_1)				
.info	1.634	5.124	0.554	$p = 0.003$
.org	-2.235	0.106	0.902	$p = 0.01$
.biz	2.000	7.393	0.661	$p = 0.002$

categories. Therefore, it is unclear whether the reported correlation is due to the domain name category itself, or due to the different name structure.

8.5.2 Estimators of phone number misuse

The gTLD is the only variable with statistical significance in Equation 8.2. Table 8.5 presents the 3 gTLDs with a significant correlation to the measured WHOIS-attributed phone number misuse. Domains under the .biz and .info gTLDs correlate with 7.4 and 5.1 times higher misuse compared to the overall mean, respectively. On the other hand, .org domains correlate with lower misuse, being close to the mean.

There is no verifiable explanation as to why gTLD is the sole statistically-significant characteristic affecting this type of misuse. A possible conjecture is that domains usually registered under the .biz and .info gTLDs have features that make them better targets.

8.6 Limitations

Specific characteristics of the experimental design (e.g., budgetary constraints) result in some limitations in the extent or type of insights we are able to provide.

In particular, we were not able to use postal addresses outside the United States, due to mail regulations requiring proof of residency, in most countries. In addition, “virtual office” solutions are prohibitively expensive at the scale of our

experiment, and, as discussed earlier, could introduce potential confounding factors. Therefore, we were not able to gain major insights on how different regions, and countries other than the US are affected by WHOIS-attributed postal address misuse.

Similarly, we were not able to assign a unique phone number to each of the 400 artificial registrant identities. Instead, every phone number was reused by five (very similar) experimental domains. This design limits our ability to associate an incoming voice call with a single domain name, especially if the caller does not identify a domain name or a registrant name in the call. Nevertheless, we were able to associate every spam call with a specific [registrar, gTLD] pair.

8.7 Conclusion

We examined and validated through a set of experimental measurements the hypothesis that public access to WHOIS leads to a measurable degree of misuse in the context of five largest global Top Level Domains. We identified email spam, phone spam, and postal spam as the key types of WHOIS misuse. In addition, through our controlled measurements, we found that the occurrence of WHOIS misuse can be empirically predicted taking into account the cost of domain name acquisition, the domains' gTLDs, and whether registrars and registries employ WHOIS anti-harvesting mechanisms.

The last point is particularly important, as it evidences that anti-harvesting is, to date, an effective deterrent with a straightforward implementation. This can be explained by the economic incentives of the attacker: considering the type of misuse we observed, the value of WHOIS records appears rather marginal. As such, raising the bar for collecting this data ever so slightly might make it unprofitable to the attacker, which could in turn lead to a considerable decrease in

the misuse, at relatively low cost to registrars, registries, and registrants. In our thesis statement we argue that choke points exist because of certain economic incentives. Indeed, the fact that very simple techniques can be used for abuse in certain contexts—when there is no protection whatsoever as we described in this chapter—is another factor that creates very low costs for abuse.

An examination of online criminal processes to formulate and evaluate disincentives

In this chapter we address the second part of our thesis statement, by using the empirically-grounded findings outlined in the previous five chapters to structure appropriate disincentives for online criminals. In Chapters 4 to 8 we characterized the components of the criminal infrastructures of various cases of online crime, and the associated monetization paths. We now take a structured approach, informed by our empirical analyses, to examine the procedural aspects of those cases, and understand the processes enabling their operation and profitability. We structure these findings as a set of crime scripts, and we map them to Situational Crime Prevention (SCP) measures capable of disrupting the criminal operations.

We define a two-staged methodological approach. First, we use Crime Script Analysis (CSA) [43] to structure the empirically-derived knowledge on the online criminal infrastructures. This streamlined understanding reveals the motivating properties of the criminal networks, critical for their operation. Further on, we propose and empirically evaluate appropriate countermeasures based on SCP,

capable of affecting the profitability and risk associated with engaging in such illicit activities. To this end, we consider the SCP measures prescribed by Clarke and Cornish [35, 45], adapting them to the specific characteristics of online crime. While the various individual components of our methodology—i.e. empirical measurements of online crime, crime scripts, and SCP measures—have been widely used in the past, the novelty of our approach is in combining them into a coherent and solution-oriented method against crime in the digital domain.

In addition, we evaluate the expected impact of the suggested situational measures considering two key notions: the effectiveness and complexity of situational measures. The first aspect represents the expected reduction of illicit activity that follows a given intervention. This estimation is largely informed by the empirical-based insights we provide in earlier chapters. The notion of complexity, on the other hand represents the difficulty of enforcement and of a sustainable intervention. It is estimated as a function of the size of the homogeneous groups of actors that are capable of undertaking a given set of interventions. Further on, following sensitivity analyses on the potential values of effectiveness and complexity, we characterize the impact distributions of measures. Finally, we consider characteristics of the impact distributions (e.g. mean and median) to rank potential interventions, towards identifying better “choke points”. Whenever the use of such descriptive measures are not capable of providing meaningful insights, we examine the Probability Density Functions (PDFs) of impacts to characterize their comparative stochastic dominance [89].

The rest of this chapter is organized as follows. We start in Section 9.1 discussing the theoretical framework supporting crime script analysis, and the related work that uses CSA to study criminal cases. Then, the following three sections revisit the online crime case studies we have empirically examined in this thesis, using CSA to structure their processes, to suggest appropriate situational pre-

vention measures, and to evaluate the effectiveness of measures. Specifically, in Section 9.2 we examine the case of illicit online prescription drug trade, in Section 9.3 we focus on the case of trending term exploitation, and then in Section 9.4 we turn our attention to WHOIS misuse. We conclude in Section 9.5 with an attempt to generalize the methodological aspects of effective online crime analysis, combining empirical measurements with situational crime prevention.

9.1 Background

SCP associates crime commission with the existence of two principal components: (i) a vulnerable target, and (ii) an opportunity to victimize the target. While it is not always possible or feasible to remove the vulnerable target (e.g. a vulnerable website that can be used to fraudulently funnel traffic to unlicensed online pharmacies), it is usually possible to affect the existence opportunities to victimize this target in various ways [35, 45]. In a manner complementary to SCP, Cornish has shown the equal importance of taking a methodical approach for identifying and mapping the appropriate opportunity-reducing prescriptions to the various stages of online criminal activity, through CSA.

Analysis of crime with CSA. There is a number of studies that use CSA to understand criminal cases, and inform efficient situational countermeasures in mostly the physical [30, 121, 133, 163, 195], but also in the digital domain [246].

At a high level, Levi and Maguire [133], and Sanova [195] show the importance of using situational measures to fight organized crime through crime scripts. Morselli and Roy examine two stolen-vehicle exportation operations through crime script analysis [163]. While these operations take place in the physical world, the relevance to our work is in terms to the importance of brokers that enable such

criminal operations. They reveal that removal of key brokers would result in a significant disruption to the underground market.

Willison [246] examines a case of insider threat in computer related crime, where a city employee accessed the city's financial systems to create fraudulent invoices. He defines the crime script explaining the various actions that allowed the criminal to be successful in defrauding the city, and, based on this script, he suggests situational measures to prevent future occurrences of the specific crime. Chiu et al. take a look at illicit drug manufacturing labs using data from transcripts of 30 Australian courts [30]. The authors use the information from the transcripts to build a crime script characterizing (i) the manufacturing and storage locations, (ii) the resources used (i.e. chemicals, and equipment), and (iii) the actions and interactions among the various actors. Finally, they identify measures for effective intervention at every step of the crime commission process, organized by location, target, and offender involvement, as prescribed by the problem analysis triangle¹ [37].

Displacement effects. A common question in research that examines crime reduction techniques through situational prevention measures is what happens to the net amount of criminal activity deflected through such measures—i.e. the *displacement effects* [44, 66]. Indeed, there are various types of crime displacement that may occur after an intervention; For example, criminals can alter (i) the location, (ii) the temporal characteristics, (iii) the individual targets, and (iv) their techniques in committing their crime, or even (v) switch to a completely different criminal activity altogether [66].

Hesseling, in his examination of displacement effects identified in 55 published articles [95], found that there is little to no evidence of such effects when crimi-

¹ Also known as crime triangle.

nal activity is targeted through situational prevention measures. Moreover, when displacement does occur, the new levels of observed criminality are lower than before the implementation of situational measures—i.e. incomplete displacement—resulting in a net benefit. Hesseling also reported that the two main empirical approaches to measure displacement effects are based either on ethnographic studies on the rational decisions of offenders, or on quantitative measurements of the criminal activity after the implementation of such measures.

In evaluating the impact of SCP measures in this chapter, we assess the potential of displacement, whenever possible. However, due to the lack of available empirical data that would make a quantitative analysis possible, our assessment is rather qualitative.

9.2 The case of illicit online prescription drug trade

In this section we focus on the case of the illicit online prescription drug trade, using the findings from Chapters 4, 5, and 6. We use crime scripts to structure the crime commission process, and we evaluate the impact of various situational prevention measures on the criminal profitability and risk of apprehension.

We identify two key components that enable this illicit trade: (i) the illicit advertising, that is responsible for driving potential customers (i.e. web traffic) to the unlicensed online pharmacies, and (ii) the unlicensed pharmacy, which is the process responsible for monetizing the received web traffic. In the context of CSA, the two processes are termed *scenes*, and we list their key sub-processes (termed *script actions*) in Figure 9.1. We note that while the two processes function independently, they should be considered as complementary to each other; The output of the illicit advertising is used as input for the pharmacy operation, and we indicate this “communication” with a dotted arrow in Figure 9.1.

Their complementary nature is evident when considering the multitude of uses for the hijacked web traffic. For example, the same traffic can be directed to other illicit online markets (Chapter 6), and even to websites that can potentially infect their visitors with malware (Chapter 7). Similarly, unlicensed online pharmacies can attract potential customers through means other than traffic hijacking, like email spam [116, 184], and organic search results (Section 6.3).

In the rest of this Section we will delve into the details of each scene separately, suggesting appropriate preventive measures, and evaluate their impact on the criminal operations. In this regard we define a novel metric we term complexity-effectiveness, which assesses countermeasures considering their implementation complexity, and their effect per unit of complexity.

9.2.1 Illicit advertising

We start by providing the crime script detailing the process of illicit online advertising. As mentioned earlier, the specific crime script is applicable also in criminal operations distinctly different from the online prescription drug trade, like fake watches and counterfeit software (Section 6.5). However, the present analysis is informed by the specific case study, and therefore we often refer to its association to the unlicensed online pharmacies. Nevertheless, this association is circumstantial, and can be easily adapted to examine the effectiveness of countermeasures in other cases of online crime.

The procedural components of illicit advertising

Illicit advertising, in this context, represents the various methods used by unlicensed online criminals to direct potential customers to the online pharmacies. We have empirically examined in depth a set of such illicit methods that we clas-

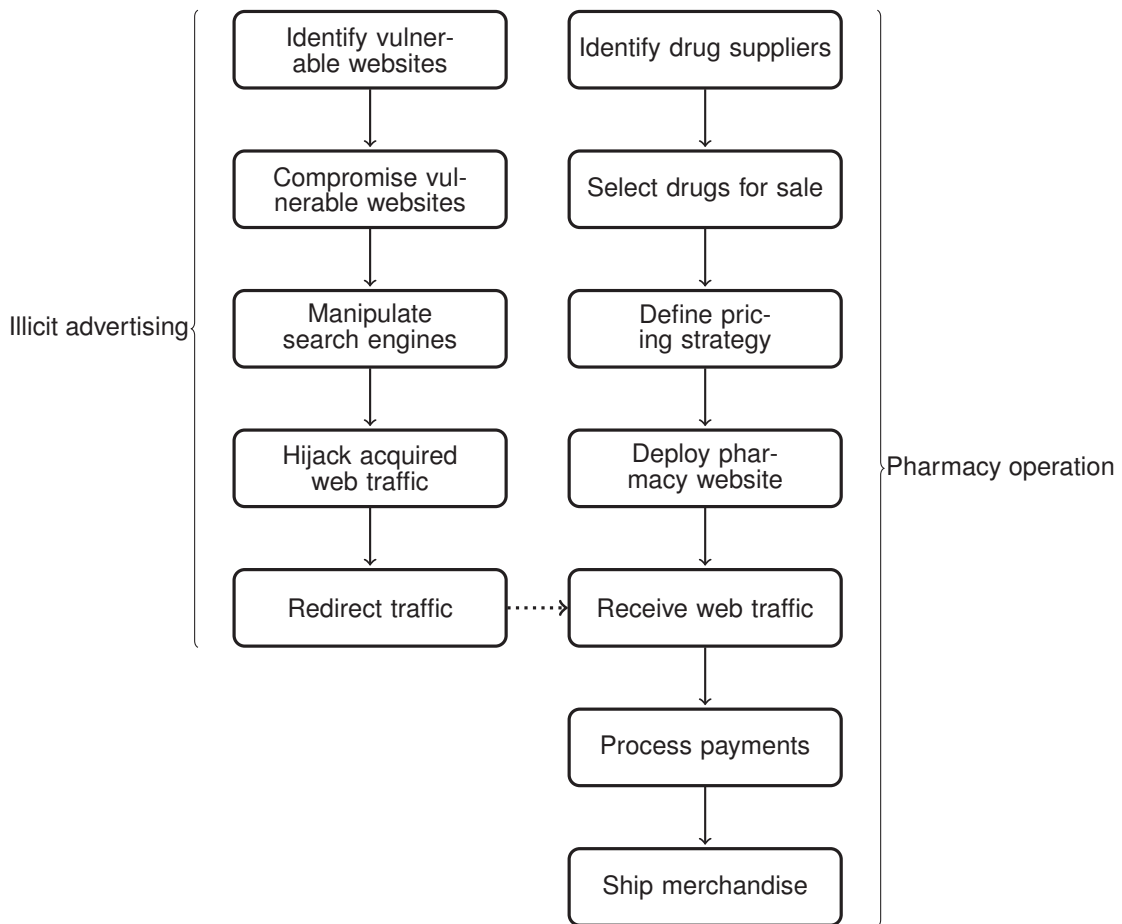


FIGURE 9.1: Components of the crime commission process in the illicit online prescription drug trade.

sify as search-redirectation attacks, and the present analysis allows for a detailed identification of the criminal procedures.

The search-redirectation attack works in four steps (Section 4.1.1). Initially, the criminals identify vulnerable websites, and they compromise them by injecting malicious code altering the functionality of those websites. In essence, the compromised websites perform two main actions controlled by their perpetrators: (i) they manipulate search engines into associating the compromised websites with drug-related terms, even if these terms are completely irrelevant to the original content

of those websites, and (ii) they redirect web traffic originating from search engine results to online pharmacies, often through one or more *traffic brokers*.

We now examine each of the four steps of the criminal process, identifying the commonly employed criminal methods.

Identifying vulnerable websites. Online criminals mainly employ scanners and search engines to identify vulnerable websites or hosting providers [134,141,157]. Through both methods, attackers look for specific characteristics of the hosting operating systems, web servers, and web content that are exploitable, allowing them to gain unauthorized access.

The motivation behind the use of these techniques is the reduction of criminal operational costs. They are automated, and capable of identifying a large portion of potential victims at low marginal cost. Florêncio and Herley [72] have discussed the validity of this threat model from an economic perspective, showing that online criminal operations need to be effective at a large scale. However, while the authors associate the reduction of expected criminal gains with the *sum-of-efforts* of defenders, this argument is not applicable in this case, due to the well-known vulnerable state of these websites. If the argument was applicable, reducing the number of vulnerable websites would actually increase the risk of victimization [25].

The process of identifying vulnerable websites is precise in nature, as it reveals the websites that are known to lack the required defenses [157]. We may therefore assume that the majority of vulnerable websites identified with the aforementioned techniques are eventually compromised.

Compromising vulnerable websites. Rather than examining the technical aspects of the attacks, we focus instead on the observational factors that positively correlate with compromised websites. We support this approach through the ob-

servation that the methods of compromise become more sophisticated as they adapt to deployed countermeasures, but the characteristics of compromised websites (i.e. the targets of criminals) are similar across time (Section 6.2.2).

Vasek and Moore [236] have examined the risk factors that correlate with a website being vulnerable to compromise and used for search-redirection. The authors found that (i) running a Content Management Software (CMS) system, (ii) using popular CMSs,² (iii) using often exploitable CMSs, (iv) using outdated versions of CMSs, and (v) the website being hosted on a specific set of server types are factors positively correlated with search-redirection infected websites.³ In the same context, Soska and Christin have demonstrated a highly automated method to predict if a website will be compromised within a one year horizon, based on an adaptive set of extracted features with a recall rate of 66% [206].

In addition, we consider the popularity of compromised websites, as it is represented through their ranking (i.e. position) in the search results (Section 4.3.2). High popularity positively correlates with the amount of traffic landing at the website [114], and, therefore, can result to greater amounts of redirected traffic. As the compromised websites inherit the popularity of the infected domains, we may reasonably assume that online criminals have the incentive to specifically target vulnerable websites with high ranking, like educational websites under the .EDU top level domain.

Once these requirements for compromise have been met, the miscreants use tools available online (like Metasploit [190]) to deploy their attack, taking control of the vulnerable websites, and injecting their malicious code. Within the scope

² The popularity of a CMS is equivalent of its market share.

³ We also note the argument that hiding the version information of the CMS being used can reduce the potential for compromise [54]. However, this argument not only lacks empirical support [236], but it also interferes with the maintenance efforts of web admins [157].

of search-redirection attack, this malicious code manages to manipulate search engines, and hijack the web traffic directed to the compromised websites.

Manipulating search engine results. One of the two key “responsibilities” of a compromised website is to manipulate the search engine crawlers into associating the legitimate-but-compromised website with drug-related queries. Examining the methods for accomplishing this goal since 2010, we have identified two such prevalent techniques: *cloaking* and *pharmacy storefront injection* (Section 5.1.1).

Cloaking is the act of serving substantially different web content, depending on the characteristics of the requestor. In the case of search-redirection attack, a compromised website can detect the presence of a search engine crawler,⁴ and provide a version of the compromised website that is filled with drug related terms, and links to other compromised websites (an act termed as *link farming* [88]). However, when the request is initiated from a non-crawler entity (i.e. normal web traffic), the compromised server either (i) presents the original content of the compromised website to avoid detection, or (ii) redirects the traffic to a different web location under the control of the attackers. The exact behavior is dependent on the variant of injected malware, and is often triggered using the information in the referrer field of the HTTP request.

A relatively new variant of the search-redirection attack injects a pharmacy storefront on an attacker-defined location within the compromised web server. In this case, the web server presents the illicit content regardless of the referrer information. This approach reduces the risks of and increases the benefits to the attackers in the following two ways; First, it does not involve cloaking, a tactic that is usually against the terms of use of search engines [85, 150, 250]. Therefore,

⁴ A search engine crawler is an automated process that crawls the web, and retrieves the content of website. This content is then associated with search queries, based on relevance criteria like TF-IDF [193].

the chances of being the focus of a search engine intervention are lower than with the previous method. Second, it overcomes a deployed countermeasure that involves hiding the referrer information when a request originates from a search result page. This piece of information has been the cornerstone for previous attack variants, and withholding it nullifies the effects of the attack. However, we have shown that by injecting a pharmacy storefront, online criminals effectively overcome the deployed countermeasures (Section 5.1.1).

Traffic hijacking. The second “responsibility” of the compromised websites is to redirect incoming traffic originating from search results. This function is essentially responsible for directing the illicitly acquired web traffic to the online pharmacies. On the technical level, this is accomplished either through web server directives, or through injected JS and HTML code.

Through the first method, the web server issues a HTML 302 redirection, when the web traffic meets certain requirements based on the attack variant. Such requirements are an appropriate referrer value (implicit redirection), or a click on an embedded storefront (explicit redirection). Detecting this compromise requires auditing the web server configuration files, and the outbound links. The second method accomplishes the same objective, but through the injection of malicious JS libraries, which in turn generate the appropriate HTML redirection code [134]. In this case, the attacker manipulates certain broadly-used JS libraries, and detection is more complicated.

Traffic redirection. We have identified two criminal methodologies to redirect traffic: (i) using one or more traffic brokers that act as intermediate redirectors before reaching one or possibly more unlicensed pharmacies, and (ii) without traffic brokers, redirecting traffic directly from compromised websites to unlicensed pharmacies. In Figure 9.2 we graphically present the possible methodological

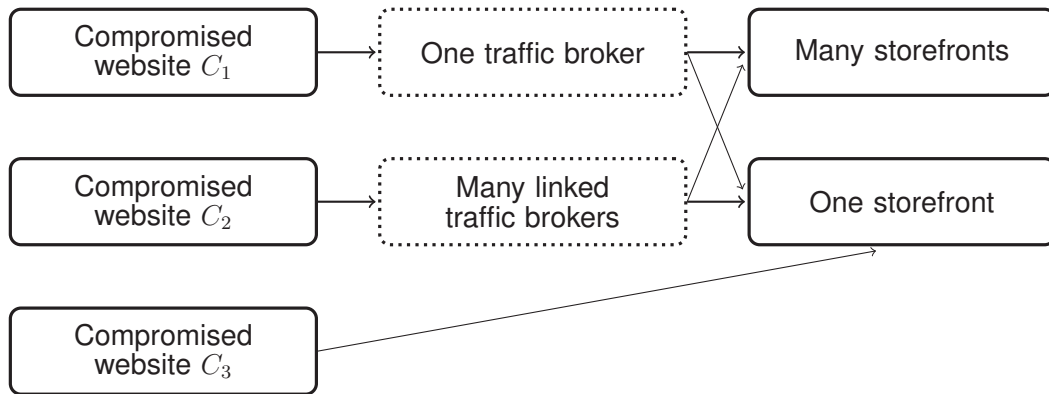


FIGURE 9.2: The two methods used to redirect illicitly acquired web traffic to unlicensed online pharmacies.

combinations. At this point we reiterate the fact that the brokers are not used exclusively to funnel traffic to unlicensed pharmacies, but they are rather an important resource for other types of shady online markets.

We have empirically measured that on a daily basis, the vast majority of compromises make use of one or more brokers to redirect traffic to one or more pharmacies ($|C_1| + |C_2| = 74.9\%$ on average—Section 6.5). *Dedicated brokers* that redirect traffic to a single pharmacy (per broker) are 61.1% of the total, and are linked to an average of 18.9 compromised URLs. On the other hand, *shared brokers* being 33.8% of the total, redirect traffic to 2.8 pharmacies (per broker), and are linked to an average of 11.8 infected URLs.

These figures show the importance of traffic brokers. First, both types enable the dynamic management of the pool of compromised web sites, by making it possible to redirect to an alternative pharmacy location, when the one previously used is taken down. Secondly, shared brokers can distribute the hijacked traffic to a large set of potential destinations, by allowing the dynamic redirection of traffic to a different pharmacy location at any point in time. The latter type of brokers is especially important for the online criminals, as each of the brokers is responsible for 33.04 possible infection-to-pharmacy daily average combinations.

Situational measures targeting illicit advertising

We examine situational measures capable of affecting the criminal opportunities for engaging in illicit online advertising. This examination is performed from two distinct perspectives; Before, and after the occurrence of website compromise that facilitates the illicit operation. We make this distinction as the situational measures affect distinctly different opportunities at each stage. In addition, we consider measures targeting the infrastructures of traffic brokers.

Measures Applicable *Before* Website Compromise. The situational measures in this category are specifically designed to prevent the compromise of vulnerable websites.

– **Utilize webmasters for website hardening.** The vulnerable websites are the main driving force of this type of illicit advertising (Section 9.2.1). Therefore, providing proper incentives or education to website owners in keeping their web space secure would effectively reduce the target availability. This would consequently increase the efforts required by the online criminals to succeed in their illicit goals.

Considering the expected lack of interest of webmasters in implementing security countermeasures [93], such incentives would need to highlight the mandatory nature of taking action in this direction—e.g. by imposing fines. However, the enforcement of such penalties on a global scale is dubious, a fact that we consider in the evaluation of such measures later in the chapter.

– **CMS and web server hardening.** We have shown that certain aspects of CMSs are the enablers of website compromises. Incentives for adequate penetration testing [14, 70], and inclusion of self-updating mechanisms that fix identified vulnerabilities could reduce the number of compromised web-

sites. In addition, Vulnerability Reward Programs (VRP) are a cost-effective method for fixing software problems, especially when they are appropriately structured to provide rewards proportional to the severity of identified problems [69]. In essence, VRPs provide incentives for independent researchers to discover and submit vulnerabilities to the respective software vendors in exchange for monetary rewards, instead of selling this information to the black market.

– **Utilize search engines to increase the effort and risks of compromise.**

Search engines are a key facilitator of this criminal operation, and can be utilized in a number of ways:

- **Deflect offenders.** The use of search engines from offenders to identify vulnerable websites can be thwarted through the active identification and blocking of queries capable of revealing possible target websites from the search engines.
- **Conceal vulnerable websites.** Using the same methods as the offenders for identifying vulnerable websites (i.e. queries), search engines can completely remove such websites their indexes, or decrease their ranking while they remain vulnerable. In terms of the latter type of action, Edwards et al. suggest that search engines can prevent the spread of hosted infections by demoting—or “depreferencing”—compromised websites [62]. While their analysis covers websites that are—potentially—already compromised, considering the predictive power of the methodology suggested by Soska and Christin [206], we argue that the same approach could be effective for vulnerable websites with high potential for compromise.

- **Extend guardianship for high-value targets.** Given the popularity of websites with specific characteristics (e.g. under certain gTLDs discussed in Section 9.2.1), search engines could take routine precautions to identify vulnerabilities and attempts for compromise at these locations.
- **Reduce anonymity for suspect queries.** Target-revealing queries could be permitted only for authenticated (i.e. signed-in) users, while blocked for mischievous purposes.

Measures Applicable *After* Website Compromise. Once a website has been compromised resulting into search engine manipulation, the effort of the following situational measures shifts towards reducing the rewards to the offenders.

- **Utilize search engines to conceal victimized targets.** Search engines can reduce the benefits of compromise, by first detecting and then removing or depreferencing compromised websites. Based on the attack variant, the following two heuristics have been proven adequate to detect compromise: (i) cloaking, and (ii) injected storefront detection. The second heuristic can be implemented either through link analysis, as we have demonstrated in [129], or by identifying unexpected content, considering the historical profile of the investigated websites.
- **Utilize webmasters to identify compromise.** Webmasters should have the proper incentives (e.g. accountability), and receive proper education and assistance to regularly maintain and monitor their online property for indicators of compromise. This would be a distributed effort towards effectively stopping traffic redirection to malicious destinations. However, as discussed earlier, such measures are inapplicable in expectation.

Measures affecting the illicit advertising infrastructure In Section 9.2.1 we discussed that more than half of the compromised websites, victims of the search-redirection attack, are linked to traffic brokers. In essence, we want to identify measures that can disconnect the traffic brokers from the rest of the criminal infrastructure.⁵ The *internet service providers* and the *domain registrars*, being the “place managers” that facilitate the operation of brokers—by providing them with IP addresses and domain names—meet this operational requirement. An intervention at this level would result in an increase in (i) the operational risk (by increasing the possibility of punishment), (ii) the efforts of criminals (by making it harder to find a “friendly” hosting provider), and (iii) would reduce the associated rewards by forcing offenders to use more expensive (i.e. “bulletproof”) hosting providers.

It is important to note that before making a request to the service providers to discontinue the services and resources of brokers, there is a need for empirically-based investigative work for the proper identification of the traffic brokers. Nevertheless, there are well-defined methodologies capable of meeting this requirement outlined in previous chapters (e.g. by targeting the few ASs that support the operation of traffic brokers—Section 6.4.2).

Impact of situational measures

While the proposed measures properly target the various components of the illicit advertising crime script, we do not suggest that they share the same degree of applicability or effectiveness. Therefore we attempt to assess their effectiveness using, whenever possible, available empirical data. In this assessment we also consider possible displacement effects which could occur whenever a counter-

⁵ Obviously, whenever traffic brokers are not used for traffic redirection, such measures are irrelevant, and we should place our attention to the appropriate points of the criminal operations instead (e.g. through search engine intervention).

measure forces online criminals to change the parameters of their illicit activity, effectively circumventing the countermeasure.

As the cost of a preventive measure burdens the actors that are expected to implement it, we consider the following groups of actors separately: (i) webmasters, (ii) software providers (or vendors), (iii) search engines, and (iv) registrars and Internet service providers (collectively referred to as service providers). However, before engaging in this analysis, it is necessary to examine these actors from an economic perspective.

Welfare economics and externalities Before starting our evaluation of the effectiveness of SCP measures, we need to examine the economic incentives that are essential in motivating the involved actors to take action. Indeed, while society as a whole suffers—at least financially—from criminal activity [17, 158], an economically rational entity is expected to act upon a situation \mathcal{A} , only if the resulting status \mathcal{B} will provide higher levels of utility [93]. Therefore, in the following paragraphs we examine the degree to which the aforementioned illicit activity is burdening the actors capable of implementing the countermeasures, in an effort to assess their willingness to engage in such action.

To this end, it is important to realize that the described illicit activity does not impose any direct cost (i.e. financial loss) to the actors capable of implementing the suggested countermeasures. Also, they receive no direct (i.e. private in economic terms) financial benefits by implementing any of the prescribed actions, making the allocation of their effort in this direction, an inefficient one. In essence, our measures suggest that actors should use resources to transition to state $\mathcal{B} = \{\text{reduced crime rate}\}$, even though state $\mathcal{A} = \{\text{established crime rate}\}$ provides an equal—or better in some cases—amount of utility. Therefore, the costs associated with the implementation effort constitute a negative external-

Table 9.1: Costs and benefits for each of the actors involved in, or enabling illicit online advertising, before and after an intervention targeting such activity.

Actors	\mathcal{A} : Current status		\mathcal{B} : Post-countermeasure	
	Costs	Benefits	Costs	Benefits
Attackers	Direct: Operational	Direct: Profit	Direct: Operational (↑)	Direct: Profit (↓)
Webmasters	none	none	Indirect: Education, impl.	none
Software vendors	Indirect: Reputation	none	Indirect: VRP, impl.	Indirect: Reputation
Search engines	Indirect: Reputation	none	Indirect: Implementation	Indirect: Reputation
Service providers	Indirect: Reputation	none	Indirect: Loss of revenue	Indirect: Reputation

ity [24], and, from a public policy perspective, we may not expect any rational agent to undertake the cost of action [52]. Table 9.1 offers a vivid outline of this situation.

The described plane can be viewed as the digital equivalent of the *tragedy of the commons* [90]. This is an economic theory used to describe a situation where a common public resource (e.g. a grazing field) is utilized by private entities (e.g. herders) in a way that maximizes their private benefits without considering the social costs of their activities (e.g. grazing field being depleted). In the case we are examining, the common resource is the Internet, with the aforementioned actors being the private entities participating in the illicit advertising activity (willingly or not).

This discussion, however, has yet to identify the actors undertaking the direct costs of illicit advertising. These actors would naturally be expected to undertake (at least partially) the costs of intervention to minimize their financial loss. Based on Figure 9.1 though, the web traffic (i.e. the traded commodity in this context) is bought by the unlicensed online pharmacies, which have no incentive to take measures curbing the availability of the commodity they are purchasing.

Therefore, an economically meaningful approach would have to deal with internalizing those negative externalities. There are two key approaches in doing so: (i) privatizing the share resource, which would effectively make this problem *someone's* problem with an economic incentive to fix it, and (ii) imposing taxes equal to

the negative externalities. Such taxes, often termed *Pigovian taxes*,⁶ would then be used to compensate the actors implementing the countermeasures. However, both options are inapplicable in this case; The Internet is a globally distributed resource, and, as such (i) cannot be privatized, and (ii) collection and proper allocation of Pigovian taxes from all governments is an unfeasible expectation.

Consequently, in examining the impact of the proposed situational measures, we will be considering the nature and significance of indirect costs in motivating an actor to take action.

Estimating impact through a complexity-effectiveness perspective We evaluate the impact of the situational measures through a form of cost-effectiveness analysis, whenever this is reasonable from an economic perspective. In the next paragraphs we structure the methodological aspects of this evaluation, which is reused throughout the chapter.

In essence, the goal of all situational measures is to reduce the output of the illicit activity, which, in this case, is the traffic directed to the online pharmacies. Therefore, rather than using the financial benefit of the situational measures as one measure of effectiveness—which essentially requires a rather arbitrary price tag per unit of redirected traffic—we elect to examine their effectiveness through the estimated reduction in traffic redirections. Using a policy's effectiveness instead of its monetary benefit is common whenever (i) it is hard to estimate effectively the economic benefits associated with a specific action, and, most importantly, (ii) when the existence of externalities requires that the researcher considers the *social* benefits of a policy instead of the *private* benefits. Thus, if U_{Δ} is the fraction of removed redirecting results achieved through a specific set of

⁶ The pigovian taxes take their name from Arthur Pigou, the economist that introduced their concept [182].

measures, we define a measure's effectiveness E as the achieved reduction of generated traffic:

$$E = U_{\Delta} \tag{9.1}$$

As the base case of redirected traffic we use our estimation from Section 4.5, placing the number at 20 million visitors per month. We make this estimation considering the median popularity for a fixed set of queries (i.e. the estimated number of monthly searches—median 1,600 per month), and the measured proportion of redirecting search results (38% of total). In addition, our measurements in Chapter 6 place the daily average of redirecting search results (using the same methodology) at 908 URLs, or about $S = 27,240$ on a monthly basis (Section 6.5). Therefore, we estimate that the marginal traffic generated by each compromised website per month to be $T_{\Delta} = \frac{20,000,000}{27,240} = 734$ visits.

It is important to note in the analysis that follows, we evaluate the effectiveness of measures using relative values—i.e. percentages (%)—instead of absolute values, like the one we devised in the previous paragraph. The reason for this approach is that different measurement methodologies can yield largely different absolute values. For example, in Section 4.5 we place the number of traffic landing at unlicensed online pharmacies at $0.75 \times 855,000 = 641,250$ visitors. However, Kanich et al. [117], using a different approach—and measuring a slightly different illicit activity—report an estimated monthly traffic landing at unlicensed pharmacies at one order of magnitude lower than our estimate—82,000. Therefore, we use our estimation of the absolute amount of visitors only for demonstration purposes, and our evaluation is not limited by any absolute traffic estimates.

For example, if a set of situational measures was capable of reducing the number of redirecting websites by 10%, then the total reduction in redirected traffic would be $S \times U_{\Delta} \times T_{\Delta} = 27,240 \times 0.1 \times 734 = 1,999,416$ visitors. Of course, not all

search queries generate the same traffic, and neither do all compromised results as they are placed at different rankings. However, we argue that this measure is an estimator that provides the desired level of simplicity and accuracy for this analysis.

Estimating the *cost* of each intervention though for use in a cost-effectiveness valuation is a much more difficult task. The various actors we identified not only have significant geographical diversity, but they also vary in terms of available resources and technical expertise for implementing countermeasures. Consequently, we use a non-monetary estimator of the cost which is the *complexity* of implementing a measure with a given a homogeneous group of actors. In this regard, we represent the complexity as a function of the number of actors in a group capable of implementing a countermeasure. Therefore, with A being a set of actors (e.g. search engines), we define the complexity of a countermeasure as:

$$C_A = |A| \quad (9.2)$$

This measure dictates that the complexity of an intervention is directly proportionate to the number of actors required to implement it. Thus, the problem of estimating the cost of an intervention is now reduced to an estimation of the number of entities in the four actor groups.

With this evaluation framework in mind, we define the impact I_A of a countermeasure implemented by actors A , as the ratio of the resulting reduced traffic E per unit of complexity C_A . Using the previous example, if the associated complexity was $C_A = 5,000$ then the solution's marginal effectiveness, as estimated though Equation 9.3, would be $I_A = 0.002\%$ or 399.88 visitors. In plain terms this means that each actor in A is capable of reducing the redirected traffic by about 400 visitors.

$$I_A = \frac{E}{C_A} \quad (9.3)$$

Obviously, when comparing different options through their impact measure, we would prefer the one that maximizes Equation 9.3.

There are two aspects of our definition for complexity that affect the impact measure. First is the matter quality of the implementation. For example, if a specific measure is implemented by x different actors (i.e. $C_A = x$), a logical question is if all x implementations are equally targeting the illicit activity. In this regard, in our analysis we make the assumption that all implementations are, indeed, equivalent based on the prescribed actions, and they can successfully affect the criminal activity. Therefore, our assessment of the impact reflects the upper bound of effectiveness per unit of complexity; if one or more actors provide a limited or malfunctioning implementation of a countermeasure, its impact would be lower than the one we estimate. However, this discussion also highlights the benefit of using the number of actors as a proxy for a measure's implementation cost; the more actors involved in an implementation, the more probable is to face problematic implementations.

Second, our definition for complexity and impact does not allow for an evaluation of combinations of situational measures. We support this approach by restating the goal of the present analysis. Our intention is not to provide the fine-grained details on the outcome of each countermeasure—mainly due to the many assumptions we make along the way—but rather, to identify the characteristics of measures that have the potential to require the least effort and to be most effective. We achieve this by looking into the effects of measures grouped by homogeneous sets of actors, while arguing that the alternative (i.e. examining compound measures at various degrees of engagement by different actor groups) would be difficult to practically interpret.

Impact per actor group We now proceed by considering the four groups of actors, and estimating their potential impact against illicit advertising. Per our earlier discussion, while in absolute terms all actors are capable of implementing specific sets of countermeasures, when considering the context of economic incentives (Table 9.1) we show that not all are expected to do so.

The perspective of webmasters. Herley in [93], examined the economic decision-making process—through a cost-benefit analysis—of Internet users receiving security-related advice on how to protect themselves by choosing better passwords. He found that even after receiving such advice, they often rationally choose to take no action, because they perceive the implementation costs as a negative externality.

This finding is directly relevant to the situation we are examining here. The measures that webmasters are expected to implement (Section 9.2.1) are forms of security advice. In addition, the cases examined by Herley involve users that can experience direct losses from an illicit activity, and, still, take no action. However, in the present case, we do not expect webmasters to experience any direct financial loss. Therefore, the negative externalities could possibly have an even stronger effect, thus making the implementation of countermeasures ineffective in expectation.

The perspective of software vendors. Through casual investigation on the number of available CMSs and web servers, we can estimate this number to be between 200 to 500, uniformly distributed, when studying the intervention complexity. Managing to convince all the software vendors to implement the situational measures can be, in itself, a daunting task. Assuming though that this could be achieved at some degree, we may not reasonably assume that all security problems can be identified, and fixed in an automated and instantaneous way.

Therefore, we estimate the effectiveness of the related countermeasures using the adoption rate of new WordPress (WP) versions as soon as they become available.

We use WP, as it is one of the most popular CMSs,⁷ and one of the most targeted for search-redirectation attacks [236]. In 2011 it was estimated that 15% of WP websites would switch to the (at the time) upcoming newer version containing security enhancements.⁸ We use this value as the low value of U_{Δ} since the proposed situational measures suggest an automated method of security patch deployment. Moreover, as we have no way of estimating an upper bound of effectiveness, we examine this variable through a uniform distribution with parameters $min = 0.35$, and $max = 1$.

We perform sensitivity analysis through a Monte Carlo simulation,⁹ to examine the potential impact of intervention measures requiring the participation of software vendors. In Figure 9.3 we present the plots of the PDF and the Cumulative Distribution Function (CDF) of the derived impact distribution. At the 50th percentile, each actor is capable of reducing the monthly redirected traffic by 0.16% (14,370 visitors), with the average being at 0.17% (15,260 visitors).

Of course, as the degree of effectiveness increases, we expect that the remaining “unprotected” websites will be targeted more intensely by the online criminals. However, considering that each compromised website can be under the control of a single actor—or group of actors—at a given point in time [157], we do not expect that the variation of impact changes as the effectiveness rate grows.

⁷ <http://www.forbes.com/sites/jjcolao/2012/09/05/the-internets-mother-tongue/>

⁸ <http://www.dev4press.com/2011/blog/slow-adoption-rate-of-new-wordpress-versions/>

⁹ For all sensitivity analyses in this thesis, we perform Monte Carlo simulations with 1,000 iterations, using the distributions for effectiveness and complexity we define in each case.




FIGURE 9.3: Probability density plot and cumulative distribution plot of complexity-benefit analysis for a software (CMS and web server) provider-based intervention.

The perspective of search engines. comScore, a company that tracks and periodically reports the search engine market characteristics, reported in March 2014 [40] that 5 search engines handle nearly 100% of searches in the US. Considering the per country population [228] and Internet penetration [105], the US is ranked second in the Internet population. Therefore, we argue that this report offers a solid lower bound for this actor group’s complexity function. We also take into account that in specific geographical regions, certain localized search engines have significant market share. Examples include *Baidu* and *360 Search* in China, and *Yandex* in Russia. Thus, we believe that 20 search engines is a reasonable upper bound for the complexity function. However, given the dominance of the 5 search engines, the complexity function would be represented adequately with a discrete logarithmic distribution ($p = 0.9$). In this assessment we essentially associate the popularity of search engines with their ability to impact the illicit advertising. Due to the disproportionate popularity of this small set of search engines, we argue that the need to require the participation of every search engine is unnecessary; Less popular search engines, even if they are exploited, provide

a rather small potential to redirect web traffic at profitable levels for the criminal activity.

Search engines play—unwillingly—a significant role in the operation of this illicit activity. Through the crime script, we showed that they are a hub for identifying vulnerable websites, and directing web traffic to compromised websites. Consequently, implementing countermeasures at the search engine level would be a relatively centralized and effective way of removing a significant portion of redirected traffic. While there are methodologies for identifying most redirecting websites, the constant refinement of criminal evasion tactics and newly discovered website vulnerabilities can lead to periods of reduced countermeasure effectiveness. In addition, as we discuss in Section 6.3.2 (Table 6.3), on average 3.6% of daily search results are redirecting but we were unable to identify them as redirecting at collection time. Therefore, we estimate that on average, the countermeasures will enable a 96.4% drop of redirected traffic, and U_{Δ} will follow a normal distribution ($\mu = 0.964, \sigma = 0.0964$).

Figure 9.4 presents the characteristics of the impact distribution, when performing a sensitivity analysis with the aforementioned parameters. At the 50th percentile, each search engine is capable of reducing the redirected traffic by 46.4%, with an average of 54.3%. As we discussed earlier, even if displacement to other search engines does occur, we do not expect that the generated traffic will generate enough profit to sustain the illicit activity. Moreover, we highlight the fact that this analysis—in line with the context of this thesis—is done in the context of the US market. Other local markets would need to require the implementation of suggested countermeasures at locally popular search engines, potentially altering the complexity function, and, consequently, the expected impact. Nevertheless, factoring in the market share of each individual search engine in the complexity function, would enable a more fine-grained analysis of their impact.

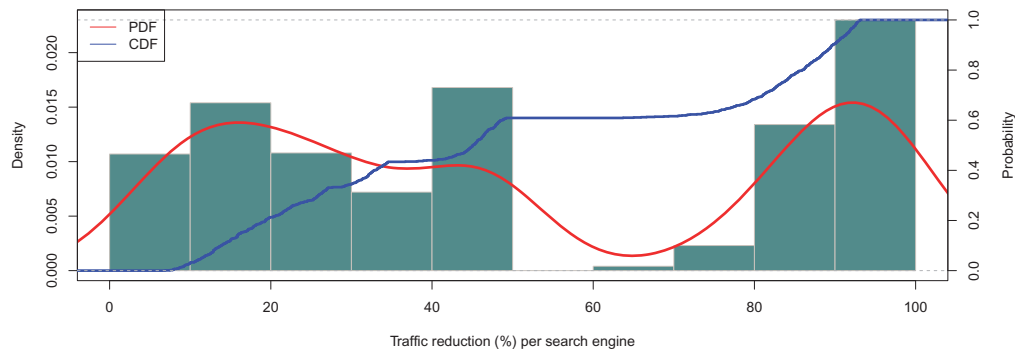


FIGURE 9.4: Probability density plot and cumulative distribution plot of complexity-benefit analysis for a search engine-based intervention.

The perspective of service providers. Based on our earlier discussion, and Table 9.1 it would be counter-intuitive to expect registrars and hosting providers to implement any countermeasure. This is based on the fact that domains and websites, regardless of the nature of their activity (i.e. illicit or not), are the key source of revenue for these service providers. Therefore, any action that would result in a reduction of their customer base, would represent a loss in revenue. However, these service providers are also interested in their reputation, a fact that should provide adequate motivation for them to take action against illicit activities.¹⁰

Registrars, the service providers enabling the registration of domains names, are entities that are usually require accreditation from ICANN to be able to perform their task. ICANN lists about 900 registrars in its accredited registrar list [102]. Web hosting providers on the other hand do not require any such accreditation, making it hard to track their number. A non-authoritative directory of international web hosting providers¹¹ contains 489 such entities. However, the countermeasures appropriate for this actor group do not require for all actors in the group to

¹⁰ <http://blog.legitscript.com/2012/12/internet-bs-domain-name-registrar-does-180-internet-pharmacy-crime/>

¹¹ <http://www.microsoft.com/web/hosting/providers>

simultaneously implement them. This stems from the fact that on any given day, traffic brokers use on average 10 distinct internet service providers, and switching from one provider to another does not happen in a regular fashion (Section 6.4). In the same work, we observe on average a set of 41.3 distinct traffic brokers each day (ranging from 9 to 238), with the set being rather stable over time. At the worst case, where each broker uses a different registrar, the combined average number of actors in the group is 51.3. Based on these observations, we consider the complexity function to be adequately estimated through a normal distribution ($\mu = 51.3, \sigma = 5.13$). We note that, similar to the case of search engines, not all registrars have the same popularity. However, contrary to search engines, the popularity of registrars does not affect their ability to reach the entire Web due to the decentralized nature of the Internet. Therefore, examining the complexity of only popular registrars could possibly lead to target displacement, and we consequently do not take into consideration the popularity of registrars. Still though, the complexity function we define here does not require the involvement of all registrars, which could lead to target displacement. Such effects would need to be assessed empirically post-intervention.

We now turn our attention to identifying the characteristics of the effectiveness function. On an average day, we observe that 74.9% of unlicensed pharmacies receive traffic from traffic brokers (Table 6.4), and we use this as the mean value of a normal distribution to examine the effectiveness of measures. We present the empirical distribution of impact in Figure 9.5. At the 50th percentile of the distribution, each actor is capable of reducing the redirected traffic by 1.46%. However, we argue that this is in the worst-case scenario, as we assume that each traffic broker is using a different registrar.

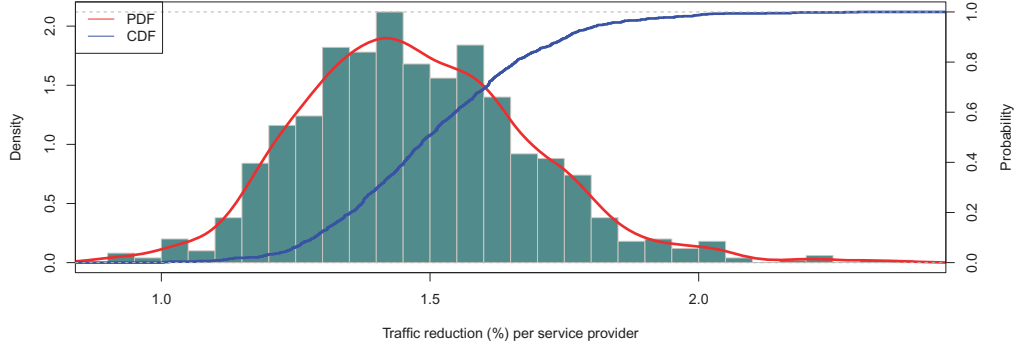


FIGURE 9.5: Probability density plot and cumulative distribution plot of complexity-benefit analysis for a registrar and Internet service provider-based intervention.

Table 9.2: Average reduction of redirected traffic (i.e. effectiveness) per unit of complexity.

Actors	# of actors in group	Expected impact
Software vendors	[200, 500]	0.17%
Search engines	[5, 20]	54.3%
Registrars	$\mathcal{N}(\mu = 51.3, \sigma = 5.13)$	1.48%

Overall assessment

We examined a set of situational prevention measures targeting the criminal operation of illicit online advertising, and in Table 9.2 we summarize our findings on their impact. Overall, we identify four actors capable of implementing the situation measures. However, we reason that only for three of them (i.e. software vendors, search engines, and service providers) we can reasonably expect—from an economic perspective—to be capable of implementing such measures, and we evaluate their impact through a complexity-effectiveness analysis.

We find that the effectiveness per unit of complexity achieved by intervening through software vendors is very limited compared to the other two actors, and especially compare to the search engines. This observation is an artifact of the limited capability of software providers to fix and deploy security patches in a timely

and comprehensive manner. Search engines on the other hand can have high impact against the criminal operation; Due to their limited number, they act as critical components in the criminal infrastructure. However, we argue that situational measures should be implemented in tandem—to the extent possible—for long-lasting disruptive impact.

9.2.2 Unlicensed online pharmacies

In this section, we examine the unlicensed online pharmacies from a procedural perspective, similar to Section 9.2.1. We start with an analysis of the associated criminal processes, and continue by proposing and evaluating appropriate situational prevention measures.

The procedural components of unlicensed pharmacy operation

The operation of unlicensed online pharmacies encapsulates all those processes that enable the illicit online sale of prescription drugs. Figure 9.1 depicts the associated criminal acts, and in the following paragraphs we characterize the operational details of each one separately. This analysis is mainly based on our work in Chapter 5. However, using artifacts from the related work, we also describe the payment processing [147, 148], and shipping infrastructure [77].

Identifying drug suppliers. The drug suppliers are the entities responsible for producing and providing the drug stock of online pharmacies. We argue that each supplier can provide a diverse set of drugs, with distinct differences among them. Therefore, the availability of drugs at the unlicensed online pharmacies can be an estimator of the number of available drug suppliers.

Our empirical examination of the inventories of 256 unlicensed online pharmacies using the search-redirection attack as their advertising technique, has indeed revealed concentrations of drug suppliers (Section 5.3.4). Overall, 50% of the

pharmacies are linked to just 8 drug suppliers. However, this observation is not limited to the specific type of unlicensed pharmacies. A separate set of 256 pharmacies that uses different methods of advertising¹² yields similar concentrations. Similarly, Gelatti et al. [77] using a different methodology, found that orders for prescription drugs placed at different unlicensed online pharmacies, were fulfilled by a small, fixed set of drug manufacturers.

Selecting drugs for sale. While an unlicensed online pharmacy may sell any possible subset of the drugs available through its supplier, being a for-profit business, operating in a shady environment, will make effort to be competitive among both its shady and legitimate counterparts. These unlicensed pharmacies can be competitive through a combination of two strategies: drug selection, and drug pricing. It is noteworthy that licensed pharmacies are rarely able to engage in either strategy; they must fill every prescription for any possible FDA approved drug, and the amount of dispensed drug units is strictly defined in the prescription.

Examining more than 1.02 million drug combinations that appear in 256 unlicensed pharmacies and one licensed online pharmacy (Section 5.3), we identify the drug selection strategies designed to achieve (i) greater variability of available drugs, (ii) greater availability of drug with potency for abuse, and (iii) targeted coverage of medical conditions that generate long-term profit from drug sales.

Defining pricing strategies. The second marketing strategy evolves around drug pricing. Generally, the pharmacy operators engage in a three-tiered approach that makes them competitive compared to licensed pharmacies (Section 5.4). Overall, they offer: (i) generally lower prices, (ii) fake generics, and (iii) volume discounts. In addition, unlicensed pharmacies offer deep discounts for widely used drugs compared to the less popular ones.

¹² The separate set of 256 unlicensed pharmacies appears in the NABP's "not recommended" list.

Deploying pharmacy websites. Online pharmacies are simply e-commerce websites that need to satisfy two prerequisites in order to be operate: (i) host their content on a web server or at a web hosting provider, and (ii) register a domain name. Their choices in both accounts are essential for being and staying operational.

There is a multitude of ways to host a website. For example, one can utilize a web hosting provider as a service, or setup a web server operated from someone's home. For more illicit operations, botnets are commonly utilized to host the questionable content [165]. Using a hosting provider is a common avenue both for legitimate and illicit purposes. In the latter domain, online criminals can benefit from the delayed—or complete lack of—response from service providers to law enforcement requests for taking down illicit content. That is especially important in cases of phishing where the time-to-take-down is critical for the success of the criminal operation [155].

Legitscript and KnujOn reveal that domain name providers (i.e. registrars) can also be considered as enablers of the operation of unlicensed online pharmacies [126]. Registrars have the legal authority to discontinue the operation of domains engaged in illegal activities. However, they do not always have the financial incentive to do so. The authors revealed that four registrars hosting the majority of unlicensed pharmacies at the time, acted as “safe havens” for these illicit operations, by ignoring requests for illicit domain take downs. Levchenko et al. make similar observations, and highlight the capacity of criminals to exploit the systemic weaknesses to their benefit [132].

Receiving web traffic. Once the infrastructure and required collaborations are in place, the online pharmacies are ready to handle incoming web traffic representing potential customers. In this case study, these customers are the outcome of

the illicit advertising discussed in Section 9.2.1. However, we note the availability of additional vectors for traffic acquisition, like email [184] and social networking spam [86]. Our longitudinal analysis of online pharmacies using search-redirection attack to attract potential customers (Table 6.6) has shown that on an average day: (i) 55.9% of unlicensed online pharmacies do not employ an intermediate traffic broker and each is linked to an average of 4.6 compromised websites. (ii) 18.1% receive traffic from dedicated traffic brokers, which in turn are linked to 24.2 compromised websites. (iii) 28.4% of pharmacies receive traffic from shared traffic brokers, which in turn are linked to 5.4 compromised websites.

Therefore, we observe that the unlicensed online pharmacies employ a variety of ways to receive web traffic. Consequently, an efficient preventive measure should be able to tackle all these challenges in parallel.

Processing payments. When customers complete their orders, payments are often processed off-site through affiliate networks [147]. In addition, the payment processors, in 95% of cases, deliver the revenue through popular payment networks like Visa, MasterCard, and American Express [148].

Generally, there are five parties involved in each transaction: (i) the cardholder who issues the payment (i.e. the customer), (ii) the issuing bank (i.e. the customer's bank), (iii) the payment network (e.g. Visa), (iv) the acquiring bank (i.e. the merchant's bank), and (v) the merchant, who receives the payment

McCoy et al. [147] and Levchenko et al. [132] have identified that the acquiring banks are the most crucial component in the payment infrastructure. Only a small number of them are willing to accept the risk of processing high-risk transactions for online pharmaceuticals, especially when there is increased pressure from the payment networks targeting those transactions.

Shipping the merchandise. Legitscript and KnujOn attempted to evaluate the legitimacy of online pharmacies advertising through search engines by placing a number of orders for prescription drugs in 2009 [125]. They found that drugs are shipped directly from the suppliers located mainly in India (via Barbados and Singapore, and packaged in Turkey), in violation of federal laws [227]. More recently, Gelatti et al. performed a similar analysis, ordering prescription drugs online, and having them shipped to Italy [77]. They similarly found that India was the main origin of the received packages.¹³ Other locations of origin included Turkey, the UK, and Vanuatu.

Both analyses point to the fact that online pharmacies ship their merchandise to the US through international locations, in order to exploit the well established jurisdictional (e.g. [91]), and policing (e.g. [218]) limitations. In addition, they present no indication that any of the orders placed originated from within the US.

Situational measures targeting unlicensed online pharmacies

The situational prevention measures targeting the operation of unlicensed online pharmacies are inherently divided in four categories: (i) measures that limit the supply of prescription drugs, (ii) measures that affect the availability of pharmacy websites, (iii) measures that prevent or reduce the network traffic reaching operational pharmacies, and (iv) measures that interfere with the processing and fulfillment of orders placed at unlicensed pharmacies. In the following paragraphs we delve deeper into each of the categories.

Measures limiting prescription drug supply. The set of measures in this category aim at reducing the availability of illicit prescription drugs, and the financial benefits for online criminals. An effective application of such measures should have a severe effect on the operation of unlicensed pharmacies, as detailed in

¹³ Whenever the origin information was available.

the previous section. We identify three specific measures that may achieve these goals:

– **Engage society to increase risk of apprehension.** Given the small number of drug manufacturing labs, provide monetary incentives to report the operation of such locations. These incentives should exceed the expected revenue of the criminal operations, to minimize the potential of bribery.

However, the effectiveness of such measures can be significantly limited if (i) a lab operates in a lawful context, but employees manage to illicitly acquire and sell certain portions of the legally produced drugs, and (ii) the criminal groups controlling the operation of labs are able to provide much stronger—financial or otherwise—incentives to deter potential whistle-blowers.

– **Enable traceability of precursor chemicals.** Enabling proper identification of the well-known set of chemicals used to produce counterfeit drugs, can allow tracing of confiscated drugs back to their producers. This action would potentially increase the risks associated with access to these chemicals, and the costs of illicit drug manufacturing.

– **Enable traceability of specialized equipment.** Being able to identify the owners of specialized equipment used only for production of prescription drugs, would result in (i) an increase in the effort of producing the illicit substances, (ii) a subsequent increase in the operational costs, and (iii) an overall increase in the risk of apprehension.

Considering the small number of “large players” who manufacture the majority of illicitly traded-drugs—eight in total associated with 50% of online pharmacies (Section 5.3.4)—these measures have the potential to be highly effective.

Measures affecting the availability of pharmacy websites. The operation of unlicensed pharmacies has similar characteristics as the traffic brokers discussed in Section 9.2.1. Therefore, the related situational measure (see Section 9.2.1)—namely the use of domain registrars to disrupt the operation of online pharmacies—is also applicable in the present discussion.

Measures reducing number of potential customers. In Section 9.2.1 we extensively address methods of incapacitating the criminal infrastructures sending traffic to unlicensed pharmacies through the search-redirection attack. However, as it is noted elsewhere, unlicensed pharmacies attract potential customers in a variety of additional ways, e.g. through organic search results (Table 6.2), and email spam [116]. While these alternatives can be targeted through rigorous efforts from search engines to exclude such results [136], and though the enforcement of email blacklists [29], they are out of the scope of this analysis.

– **Educate consumers.** While it is well documented that drugs purchased online from unlicensed pharmacies can have severe effects on the health of consumers [18, 19], even people with medical knowledge are evidently unaware of those risks [111], or they choose to ignore such risks for various reasons (e.g. reduced cost, lack of medical insurance) [91, 92]. Therefore, large-scale campaigns providing information about the pitfalls of purchasing drugs online from questionable locations (e.g. [7]) can potentially protect consumers and reduce the profitability of unlicensed pharmacies. However, providing low-cost health care is a much more debatable and tedious task, and recent efforts in this direction [219] are to be evaluated for their long-term effectiveness.

Measures affecting orders placed at unlicensed pharmacies. The purpose of situational measures in this category is to prevent the processing of payments at

unlicensed online pharmacies, and the delivery of their illicit goods. We identify two approaches in this direction:

- **Deny payments.** Payments networks (e.g. Visa) that process credit card payments, have the potential to identify transactions benefiting unlicensed pharmacies, and force merchant banks – through financial disincentives – to sever their business relationships with the illicit pharmacy operators. In this case, there are limited options for the latter party to overcome this hurdle. For example, the offending merchants may have to use an alternative acquiring bank which is not always an option. Also, the merchants may have to fraudulently mislabel the transactions (as non drug-related), in order to avoid detection, by the payment networks. It has been shown that measures in this direction can financially stifle offending enterprises, and provide counter-incentives for banks to cooperate with the online criminals [147].
- **Disrupt the market** by confiscating illicitly imported drugs. Extensive inspection of packages received at international ports of entry under the jurisdiction of US CBP from locations known to ship the illicit merchandise may have a dual effect. While protecting customers from potential health risks [18, 19, 77], this intervention will also cause substantial financial loss to criminals through the unsatisfied requests of refunds from customers [147].

Impact of situational measures

We evaluate the proposed situational measures through a complexity-benefit analysis. Similar to Section 9.2.1 we use the cardinality of the actor group necessary to implement a set of interventions as a non-monetary estimator of the cost of the interventions. Contrary to that analysis though, here we are able to evaluate the benefits of an intervention as an estimated reduction in criminal revenue. In

this effort, we are assisted mostly by the work of Kanich et al. [117] and McCoy et al. [147, 148], who provide an empirical-based insight into the illicit revenues of online pharmacies.

The situational measures designate the following actor groups having the capacity to implement them: (i) federal law enforcement agencies (e.g. FBI, and DEA), that have the tools and resources to limit the manufacturing and supply of illicit prescription drugs, either directly or through international collaboration, (ii) domain registrars, capable of disrupting the access to and operation of unlicensed online pharmacies [108], (iii) payment networks like Visa and MasterCard, capable of interfering with and interrupting the realization of payments, and (iv) federal agencies (e.g. CBP, USPS) and private companies (e.g. UPS, FedEx), capable of intercepting counterfeit pharmaceutical goods while in transit for delivery to customers.

Estimating the monetary effects of intervention The set of interventions targeting the supply of drugs can result in either (i) a reduction of sales, leading up to a complete halt, when unlicensed pharmacies forfeit their access to drug suppliers, or (ii) to a reduction in demand due to price increases. To estimate the reduction in demand, it is essential to have a good approximation of the *price elasticity of demand*, and the percentage of resulting price increases.

Increase in drug prices. While we cannot accurately predict the effects of intervention in prices, we will evaluate the effects of resulting price increases with an upper bound equal to the difference of prices between unlicensed pharmacies and their licensed counterparts. The price difference has been measured to be statistically-significant, with unlicensed online pharmacies offering lower prices by a median of 56% (Section 5.4). We use this upper bound as potential customers looking for better prices at unlicensed pharmacies, would have no economically-

rational incentive to purchase from these online stores if they could get the same deal, while avoiding shady transactions.

We do not discount the competitive advantage of being able to purchase drugs without valid prescriptions from unlicensed pharmacies. However, due to the lack of empirical data to characterize the customer population in this regard, we state this as a limitation of the present analysis.

Price elasticity of demand. Rhodes et al. [192] empirically measured the price elasticity of demand E_d for marijuana as being within the relatively elastic range, i.e. between $-2.79 \leq E_d \leq -2.65$.¹⁴ This means that for every percentage point of change in prices, the change in demand would be between 2.65% and 2.79% in the opposite direction. Due to the unavailability of information on the price elasticity of prescription drugs, we work with the assumption that products sold through unlicensed online pharmacies have a comparable price elasticity of demand as marijuana. This drug is the only one from the specific study closely related to the market we examine, as it is the least addictive drug the authors examine.

Given the previous discussion on the estimated upper bound in price increases, we will use a slightly lower range of elasticity to match the expectation that a 56% increase in prices would completely diminish the customer base. Therefore, we examine a demand elasticity in the range $-1.79 \leq E_d \leq -1.65$.

We now proceed with the complexity-benefit evaluation of the countermeasures, grouped according to the implementing actors, and using Monte Carlo simulations. Consequently, at each of the following paragraphs we will attempt to identify the distributions of complexity and benefit, whenever possible.

¹⁴ Elastic demand (i.e. $E_d < -1$) means that for a change in price by $S\%$, the demand will change by $D\%$, where $|S| < |D|$, and the changes are negatively correlated.

The perspective of law enforcement agencies. DEA’s Office of Diversion Control maintains a list of specialized equipment (e.g. tablet presses) and of 28 widely accessible (e.g. as over-the-counter medication) chemicals (e.g. ammonia gas) dubbed “Special Surveillance List”, that can be used for the production of counterfeit drugs [232]. This list is complemented by two additional lists of 40 chemicals (List I and II), designated through the Control Substances Act [227]. Entities trading chemicals and equipment in those lists are required to use caution when the quantities sold indicate a potential for illicit drug manufacturing.

While the existence of these lists has the potential to allow for rigorous monitoring of attempts for illicit drug manufacturing, there are two aspects that limit this potential. First, there is no expectation of or requirement for formal surveillance, leaving it up to the trading entities to report suspect transactions, with violators facing only civil penalties in the form of fines (up to \$250,000). Second, any enforcement is limited by the jurisdictional reach of the enforcing agencies (i.e. only within the US). Given the international locations of clandestine laboratories, and even if the traceability of offending transactions was fully automated, reasonable effects could be achieved only through international collaboration. For example, the latest Operation Pangea VII, with 111 participating countries, required the collaboration of nearly 200 enforcement agencies—two per country on average [108].

In this analysis, we make the assumption that existent standards—termed e-pedigree—allowing for automated traceability of drugs and their chemicals can be internationally implemented.¹⁵ However, enforcement needs to be constant and persistent at the locations where clandestine drug labs operate [28].

We have previously found that about $L_{critical} = 8$ labs provide the supply for about 50% of 256 unlicensed pharmacies (Section 5.3.4), while we also observe 82 additional, small-scale labs (therefore $L_{total} = 90$). In the worst case, each of

¹⁵ In the US the related laws are part of the Prescription Drug Marketing Act of 1987 [226].

those labs may operate in a different jurisdiction, requiring the cooperation of law enforcement agencies from 90 countries. Combining the previous observations, we estimate the number of continuously engaged enforcement agencies through a normal distribution with a mean value of $2 \times L_{total} = 180$ ($\mu = 180, \sigma = 18$). However, complete obliteration of clandestine labs may not be necessary to disrupt the market, as shown by Baveja et al. [15]. Therefore, and given the dominance of 8 labs, we will also consider the option of taking action only at the locations of those “big players”, approximating the complexity function through a normal distribution ($\mu = 2 \times L_{critical} = 16, \sigma = 1.6$). We argue that proper implementation of the countermeasures will inadvertently lead to a complete halt of sales for the pharmacies that depend on the affected drug labs, ranging from 50% of the 256 pharmacies (for $L_{critical}$) to 100% (for L_{total}).

As unlicensed online pharmacies usually operate under the umbrella of affiliate networks, we argue that we can estimate the financial loss caused through any relevant situational measure, by examining affiliate network revenues. McCoy et al. [148] using ground-truth data on GlavMed, one of the largest affiliate networks in the illicit online prescription drug market, found that the average weekly revenue per affiliate as being around \$2,000 (i.e. \$9,000 on an average 4.5 week-long month).¹⁶ In addition, examining 699,428 billed orders over a period of 40 months, the authors found the average purchase valued at \$115. Therefore, each affiliate—which in this case we equate to a single unlicensed pharmacy for simplicity—processes 78.3 orders per month on average (normal distribution, $\mu = 78.3, \sigma = 7.83$).

Unfortunately, there is no empirical data to quantify the effect on drug prices, for each drug manufacturer that is shut down due to the discussed interventions.

¹⁶ However, the authors note that the top 10% of affiliates—in terms of generated revenue—in the affiliate networks they examined, account for 75-90% of total revenue. For example, the affiliate reported as the largest overall earner, generated \$4.6 million in commissions [148].




FIGURE 9.6: Probability density plot and cumulative distribution plot of complexity-benefit analysis for law enforcement based intervention.

Therefore, we take a rather simplistic approach to estimate this effect. We consider the effect of taking down the top 8 labs that supply 50% of pharmacies, as resulting in a 50% price increases (i.e. equal to the number of affected pharmacies). We estimate the effect of each of the 82 remaining labs, as a fixed percentage point increase in prices equal to $\frac{56\% - 50\%}{82} = 0.07\%$, where 56% is the median price difference of drugs between unlicensed and licensed pharmacies, as discussed earlier.

In Figure 9.6 we present the estimated reduction of average revenue per month at each pharmacy, following a Monte Carlo sensitivity analysis. Based on this analysis, we observe that the intervention would initially result in an increased revenue per affiliate by \$285.50 on average, due to increase in prices, regardless of the reduction in the customer base, as a consequence of the elastic market. However, due to the debilitating impact on the ability of pharmacies to find drug suppliers, we estimate a significant reduction in monthly revenue for the set of 256 pharmacies by \$65,386 per unit of complexity.

The perspective of registrars. Domain registrars are the entities providing domain name services, and are naturally an inherent part of the operation of unli-

censed online pharmacies. While registrars have the legal ability and responsibility to discontinue their services to websites that are evidently unlawful, LegitScript and KnujOn have been regularly publishing reports on the uncooperative nature of some registrars. In a 2010 report examining more than 10,000 pharmacies linked to the EvaPharmacy affiliate network—the largest affiliate network at the time—the authors found a total of 16 registrars associated with those domains, an average of 531 per registrar [126]. In the authors' attempts to inform the registrars regarding the illicit domains under their realm, 11 took action disabling 9,803 such domains, while the remaining five rejected similar requests. The takeaway point is that while shutting down unlicensed pharmacies through cooperative registrars have a positive immediate effect, as long as some registrars remain willing to support their illicit business, this effect is also very short-lived. This is a consequence of the relative inexpensive and speedy nature of acquiring a new domain name. Therefore, interventions implemented through this set of actors, should require the persistent and full cooperation of all ICANN-accredited registrars, and of their affiliates.

As we discussed in Section 9.2.1, there are currently 900 ICANN-accredited domain registrars [102]. Contrary to the previous analysis though, countermeasures in this case need to be actively implemented by all such registrars to prevent pharmacy operators from quickly switching registrars once their domain name stops functioning. Since the number of registrars is relatively fixed, we examine the sensitivity of the complexity function through a normal distribution ($\mu = 900, \sigma = 90$).

Assuming that each unlicensed pharmacy is operated by a different affiliate, each pharmacy domain take-down would result in a \$9,000 monthly revenue loss (per [148], and the earlier discussion). In Figure 9.7 we present the distribution




FIGURE 9.7: Probability density plot and cumulative distribution plot of complexity-benefit analysis for registrar based intervention.

characteristics of the benefit over complexity ratio. On average, the discussed set of interventions results in a reduced revenue of \$2,595 per unit of complexity.

The perspective of payment networks. McCoy et al. [147] have identified three payment networks as being primarily used to process payments on behalf of unlicensed online pharmacies. The set of interventions discussed here are capable of denying the complete number of related payments. Consequently, with an average generated monthly revenue of \$9,000 per each of the 256 affiliates, the average loss of monthly revenue would be \$768,000 per unit of complexity.

The perspective of shipping and inspection actors. According to McCoy et al., 75% of the 78.3 orders per month (58.7), placed at each GlavMed affiliate, are shipped to customers in the US [148] from abroad [77, 125]. Based on the average value of \$115 per order, and the case 256 unlicensed pharmacies, each month 15,027 shipments of illicit drugs enter the country on average.

An operation coordinated by CBP in 2000, evaluated the agency's ability to inspect shipments for illicitly imported drugs [218]. During this operation, the agency identified that 11.7% of the 16,500 shipments that should have been inspected over a period of one week, were actually inspected. In addition, 37.8% of the in-




FIGURE 9.8: Probability density plot and cumulative distribution plot of complexity-benefit analysis for a US Customs and Border Protection-based intervention.

spected shipments (4.4% of the total) contained illicitly imported drugs valued at \$82,915,¹⁷ or at \$373,117 on a monthly basis. Lacking more accurate data, we use these proportions as the base case characterizing the success rate. In other words, for full identification of illicitly imported drugs, CBP needs not only to increase its inspection capacity by 81.3%, but also to extend its capacity to identify suspicious shipments by 62.2%.

While so far we have used the number of actors as an estimator of the intervention complexity, in this case one unit of complexity represents a success rate of 4.4%. The reason for taking this slightly different approach is because the one actor we examine here has many distributed components (i.e. inspection sites) which need to be improved in order to achieve higher success rates. Therefore, a 100% success rate is equal to 22.7 units of complexity, and we analyze the complexity function through a normal distribution of the current success rate ($\mu = 0.044, \sigma = 0.0044$).

In Figure 9.8 we present the distribution of reduced revenue per unit of complexity following a Monte Carlo sensitivity analysis after 1,000 iterations. On aver-

¹⁷ Based on the average purchase total of \$115.

Table 9.3: Average reduction of revenue from illicit online sales of prescription drugs (i.e. benefit) per unit of complexity.

Actors	# of actors in group	Expected impact
Law enforcement	[16, 180]	\$65,386
Registrars	$\mathcal{N}(\mu = 900, \sigma = 90)$	\$2,595
Payment networks	3	\$768,000
Shipping and inspection	[1,22.7]	\$75,990

age, intervention at this level would result in a \$75,990 direct revenue loss per unit of complexity. As a reminder, this loss would take effect when customers request refunds after their order gets confiscated at the ports of entry.

Overall assessment

We examine the various complexity-benefit analyses through a comparative lens, and Table 9.3 summarizes our findings presented in the previous paragraphs. Focusing specifically on the average criminal revenue reduction per unit of complexity, we find that interventions implemented through payment networks are by far the most effective, reducing revenues by \$768,000 per month. In terms of the second position, interventions targeting shipments at the ports of entry (\$75,990) fare better than interventions targeting clandestine drug labs (\$65,386). However, examining the related cumulative distribution functions in Figure 9.9, we observe that in fact the latter intervention has Second Order Stochastic Dominance [89] over the former. In other words, the distributions reveal that targeting drug labs is a more effective solution. In addition, considering that this solution is designed to work on a global scale—compared to the solution targeting shipments only at the US ports of entry—has the potential to have a more severe impact.

We also show the ineffectiveness of intervention at the registrar level. With an average reduction of revenues by \$2,595 per unit of complexity, the measure’s impact is one to two orders of magnitude lower than the alternatives. This obser-

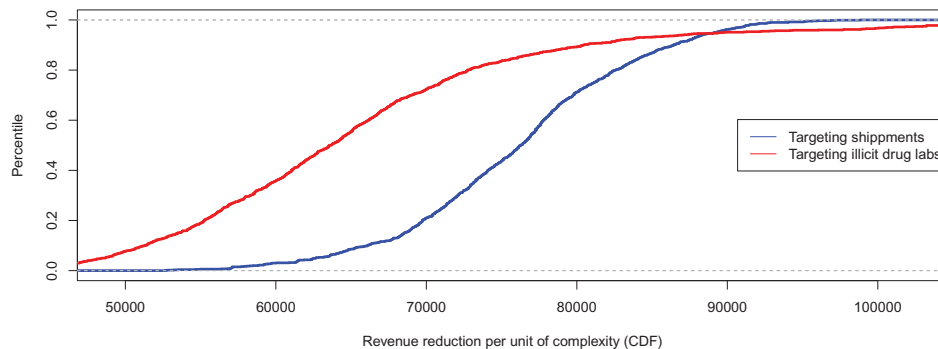


FIGURE 9.9: CDFs of the benefits of two interventions per unit of complexity to identify the stochastic dominant.

vation highlights our argument in Section 3.2.3 that the focus of law enforcement on taking down online pharmacy storefronts [230, 231] is rather futile and short-sighted.

We conclude this analysis restating the observation that interventions that can be relatively centralized as a consequence of the small number of implementing actors are overall more effective in reducing the criminal revenues. However, this comparative analysis does not suggest that interventions should be implemented in isolation, even if a specific intervention is much more effective than others. For example, only disrupting payment processing, may eventually lead to the use of alternative payment networks like PayPal or even decentralized ones like Bitcoin [147]. On the contrary, the situational measures should be considered in combination, for a long-lasting disruptive impact on the illicit online activities.

9.3 The case of trending term exploitation

In this Section we examine the procedural components enabling online criminals to exploit trending terms, and profit either through ad-filled websites, or by infecting the computers of visitors with malicious software. This examination is informed by

our empirical analysis on trending term exploitation (Chapter 7), and implemented through CSA. Further on, we consider the identified criminal processes to formulate appropriate countermeasures through situational prevention. Finally, we assess the effectiveness of the proposed countermeasures based on their implementation complexity, and the expected societal benefits.

9.3.1 A procedural analysis of trending term exploitation

We identify the components of this online crime case study based on data we collected over a period of 9 months, between July 2010 and March 2011, and the associated empirical analysis (Chapter 7). In Figure 9.10 we provide a graphical representation of the crime script, and in the following paragraphs we discuss each component separately. This criminal operation uses two separate methods of traffic monetization: (i) one based on malware (e.g. fake antivirus), and (ii) one based on interaction with advertisements. Therefore, we provide a separate discussion for each monetization path that appears after the search engine manipulation component.

We will be using two key terms to characterize the targeted trending terms: *popularity*, and *monetary value*. Popularity is defined as the number of times a term is queried for over a time period, and this estimation is often provided by the search engines. Similarly, we define the monetary value of a term as the price one would have to pay to a search engine to promote their websites at the top of search engine results—a type of results often termed as sponsored or as advertisement.

Identify trending terms. Online criminals can use a variety of sources to identify and target popular search terms that can drive traffic into their illicit operation. Examples of such sources include Google's hot trends [82], Yahoo!'s buzz log [251], Twitter's Trends [216], and Microsoft Bing Trending News [149]. In addition, adver-

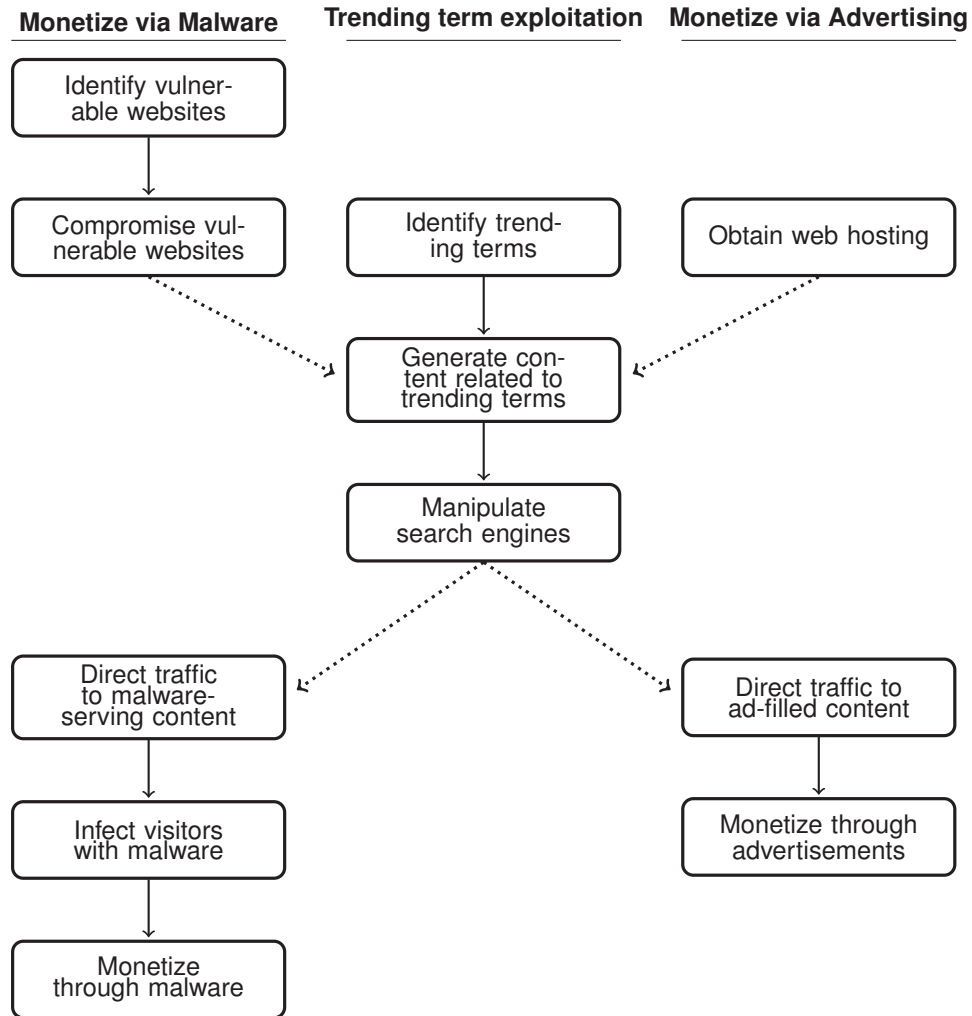


FIGURE 9.10: Components of the crime commission process in the case of trending term exploitation.

saries may use traditional media outlets—like TV-based news—to identify emerging search topics. However, John et al. [115] have provided empirical evidence that over 95% of terms used for search engine manipulation are retrieved automatically from Google’s service [82].

Generate content. John et al. reverse-engineered a commonly used script which generates, in an automated way, content highly relevant to the trending terms identified in the previous stage. Specifically, the process works as follows [115]:

[T]he script generates content that is relevant to [a] keyphrase [...] with the help of search engines. It queries `google.com` for the keyphrase, and fetches the top 100 results, including the URLs and snippets. It also fetches the top 30 images from `bing.com` for the same keyphrase. The script then picks a random set of 10 URLs (along with associated snippets) and 10 images and merges them to generate the content page.

Furthermore, the auto-generated content can be hosted either at websites compromised by the miscreants (i.e. in the case of malware-based monetization), or at regular hosting providers (i.e. in the base of ad-based monetization).

Manipulate search engines. The purpose of this action is to make the auto-generated content appear at the top of the targeted search engine results in order to attract as many “clicks” as possible, replacing legitimate results with original content. Both John et al. [115] and our own analysis (Section 7.2.2) have shown that online criminals achieve this goal through link farming [88]. Examining search results associated with trending terms over a period of nine months, we found that on average 4.7 out of the top 10 results are malicious or abusive (Table 7.2).

Manipulating trending terms to serve malware

Next, we examine the crime script actions that use search engine manipulation to infect computers with malicious software. In essence, this operation first requires the miscreants to identify and compromise vulnerable websites, which in turn host the malicious software infecting traffic coming from search engines.

Based on the work of John et al. [115], the first two script actions (i.e. identifying and compromising vulnerable websites) are identical to the ones we analyze

in Section 9.2.1. Therefore, we will only be discussing the script actions that are unique to this operation.

Directing traffic to malware-serving content. We have previously shown that miscreants do not target a specific type or category of trending terms, but they rather attempt to profit through all possible terms (Section 7.2.4). The motivation for this approach is that miscreants can maximize the popularity of the fraudulent content only within the scope of their reach, but in the end, other content can be more relevant to the searched terms, and they will unequivocally have the preference of search engines.

However, we have identified the characteristics of trending terms, in terms of popularity and monetary value, that make them prone to abuse. Overall (Section 7.2.4), 38% of relatively unpopular terms (i.e. being queried 1,000 times per day at most) contain results linked to malware, compared to 6.2% of the most popular terms (i.e. greater than 100,000 queries per day). In regards to the value of terms, we observe a similar pattern, with more “expensive” terms containing fewer malicious results.

Infecting visitors with malware. While there are various types of malware in the wild (e.g. ransomware,¹⁸ and spyware¹⁹), the related work has shown that trending term exploitation is most often monetized through fake antivirus software²⁰ [188, 209]. In this regard, Stone-Gross et al., using ground truth data, have shown that 2.16% of users exposed to this illicit activity end up getting infected [209].

¹⁸ Ransomware is a type of malicious software that usually takes control of the data on infected computers, and requests a payment (i.e. ransom) from the owners of those computers (i.e. victims) in return to restoring their access to the data.

¹⁹ Spyware is a type of malicious software that logs sensitive information entered by victims through infected computers (e.g. e-banking credentials, social security numbers) and transmit them back to the miscreants.

²⁰ Fake antivirus software (FakeAV) is a type of malicious software that claims to be free legitimate antivirus software to persuade users to install them. Once installed, they usually deny access to the computers of the victims, demanding “licensing” payments.

Monetizing infections. Once the online criminals victimize their targets through malware installation, it is straightforward for them to demand and receive payments from their victims (\$58 on average per victim [209]). Payments are submitted through credit cards, processed via major payment processors (i.e. Visa, MasterCard, and American Express), and availed to the miscreants through the complicit or cooperative banks [147].

Manipulating trending terms to serve advertisements

We now examine the alternative monetization path that involves the use of ad-filled websites. Contrary to malware-serving websites, miscreants do not need to use compromised websites for their illicit operation. The fact that this operation is clearly fraudulent, but not always clearly illegal, removes the need to obfuscate the websites existence or functionality, and they can therefore be deployed at any (e.g. free) hosting provider. The script action describing the methods used to obtain web hosting is similar to the one describing the operation of traffic brokers in Section 9.2.1, and we will not be discussing it further here. Therefore, we continue with the analysis of the actions that are characteristic to the specific operation.

Directing traffic to ad-filled content. While, as previously mentioned, miscreants do not target specific types of trending terms in order to maximize the incoming traffic, we have empirically measured that they are more successful with less popular terms, and terms that have low monetary value. The differentiating aspect of this monetization path though is that there are also specific categories of terms that are more or less prone to exploitation. For example, trending terms related to shopping or science are positively correlated to ad-based misuse, while terms in the automotive and health categories have a negative correlation.

Monetizing through ads. With online advertising, anyone can place advertisements on websites under their control, and profit based on the interaction of visitors with the ads. For the purpose of this discussion we refer to the entities that directly profit from advertisements as ad-hosts. Ad-hosts do not have to make the decision of what advertisements they host, but they rather outsource this task to advertising networks which provide the back-end business intelligence while profiting at the same time on a commission basis. Similarly, an advertisee does not directly choose where they advertise, and depend on the ad networks to choose the right medium and audience. Our analysis has revealed that ad-filled websites primarily make use of PPC ads (83%), followed by banners (66%), and affiliate marketing(16%), grossing \$100,000 on average per month (Section 7.3.2).

9.3.2 Situational measures targeting trending term exploitation

In this Section we are building on top of the trending term exploitation crime script, in an effort to identify relevant and applicable situational prevention measures. These measures are tailored to affect the criminal operation at each crime script action either by increasing the risk of apprehension or by reducing the criminal profits.

Measures affecting criminal infrastructures

The crime script actions that describe the criminal infrastructures supporting the two monetization paths, namely (i) identifying and compromising vulnerable websites for malware-based revenue, and (ii) obtaining web hosting for ad-based revenue, have similar characteristics as with the case of the illicit prescription drug trade. We argue that the respective situational measures discussed in Section 9.2.1 are also applicable in the present case. Therefore, we only name the relevant situation measures here for the sake of completeness: (i) Utilize webmas-

ters for website hardening, (ii) CMS and web server hardening, (iii) utilize search engines to increase the effort and risks of compromise, (iv) conceal vulnerable websites, (v) extend guardianship for high-value targets, (vi) reduce anonymity for suspect queries, (vii) utilize search engines to conceal victimized targets, and (viii) utilize webmasters to identify compromise.

Measures affecting trending term exploitation

The following situational measures target the crime script actions involving the identification of trending terms, and their use for automated content generation and search engine manipulation.

Reduce anonymity or access to trending terms. The list of trending terms is currently a publicly available resource not requiring any special permission to access. This unvetted availability offers the miscreants an opportunity to identify the appropriate “baits” for driving traffic to their illicit operations. A mere requirement for authenticated access (i.e. requiring users to log-in, or use a digital certificate) would significantly increase the efforts of adversaries in obtaining this information. In addition, this would increase their risks of being detected and identified as abusers of such services.

Deny benefits of auto-generated content. John et al. [115] have suggested an effective method for identifying content automatically generated based on trending terms. Search engines could use a similar approach to identify offending websites and take action upon them, effectively countering the associated manipulation.

Given the ever-changing attack methodologies, search engines can take a two-staged approach to minimize the negative effect on false-positives. *Depreferencing*, as suggested by Edwards et al. [62], can be used as a temporary first-stage countermeasure to demote potentially harmful or otherwise offending websites in

the search engine result, leading to significantly lower traffic headed on their way. At the next stage, content analysis of the depreferenced websites would result in either blacklisting offending websites, or whitelisting false-positives. The task of identifying malware at websites has been an on-going effort for the past years, achieving a negligible false-negative rate [84]. Similarly, in terms of identifying ad-filled websites, we have previously defined a machine learning method for automated classification yielding a 87.3% success rate (Section 7.1.3).

Measures reducing web traffic abuse

Here we examine the script actions directing web traffic from the search engines either to malware or to ad content.

Extend guardianship to exploitable trending terms. For both methods of trending term monetization, we have shown that less popular and “cheaper” trending terms are the ones driving the profit for the online miscreants. Therefore, we suggest that countermeasures at the search engine level, providing rigorous protection of results for the specific types of trending terms can reduce the overall profitability of the illicit operation. Such measures would be similar to the ones suggested above (i.e. denying the benefits of auto-generated code), but targeted to the specific trending terms.

Measures affecting malware-based monetization

The following set of countermeasures aim at reducing the exposure of web traffic to the malicious websites, and the access of miscreants to payment processors.

Alert potential victims and provide instructions. The purpose of the this measure is to alert Internet users that are about to visit or are visiting a malicious web location through search results about its maliciousness. These alerts can be

placed either at the search result pages [71], or at the browser level, once the user navigates to the malicious website [67].

It is important to note the significance of being able to detect a malicious website as soon as it comes to existence. Our analysis in Section 7.2 has shown that this is not always the case. About 0.015% of top 10 results at any given point in time (Table 7.2), are malicious but unfortunately undetected at the time of their appearance. However, proposed methods for malicious website identification that depend of the URL structure instead of the content of potentially malicious websites [144], may allow for faster detection.

Deny payments. This intervention would require payment networks (e.g Visa) to discontinue their business relations with banks that deliver payments to the online criminals. However, we do not suggest that Visa or other similar networks block payments from victims to the miscreants. This approach could be potentially damaging to the victims, as they would have no way of escaping their victimized state. We rather suggest the suspension of the merchant accounts, and the blacklisting of complicit banks, similar what McCoy et al. describe in [147].

Measures affecting ad-based monetization

Miscreants monetize web traffic by diverting users to their abusive, ad-filled websites. Therefore, this set of interventions aims at identifying such abusive behavior, and preventing miscreants from receiving payments.

Disrupt the market of ad fraud. Related work has revealed the possibilities for automated identification of click-based (pay-per-click) [56], impression-based (pay-per-view) [208], and commission-based [61] fraudulent ad monetization. Such interventions may be implemented within the ad networks (e.g. Google AdSense), by creating heuristics of illegitimate activities, based on known legitimate

use patterns. Large scale implementation of the suggested prevention measures have the potential of significantly reducing the profitability of the specific online criminal activity. Specifically, for pay-per-view networks, Springborn and Barford estimate a reduction in abusive traffic between 46% and 99.51% [208], and Dave et al. estimate a reduction by 23.6% for PPC networks [56]. Moreover, Edelman suggested that using a model of delayed payment in commission-based ad networks, fraudulent activity can be reduced by 71% [61].

9.3.3 Impact of situational measures targeting trending term exploitation

We examine the impact of the suggested situational measures through a complexity-effectiveness analysis, following the same methodology as in Section 9.2.1. As a reminder, we argue that the cardinality of the actor set implementing a set of countermeasures, can be effectively used as an estimator of the implementation complexity C . Intuitively, when fewer actors are needed to implement a countermeasure, it has the potential of having a higher impact (i.e. being more efficient) and being easier or faster to implement and maintain.

Despite the availability of revenue estimations for the trending term exploitation [162], we assess the effectiveness E of the countermeasures through the estimated reduction of web traffic (in %) reaching malware-serving and ad-filled websites. We take this approach for two reasons: (i) in order to limit the number of assumptions we have to make when estimating revenues from either one of the illicit or abusive criminal operations, and (ii) to generate findings directly comparable with the impact analysis of situational measures targeting abusive advertising in the illicit prescription drug trade (Section 9.2.1). We have measured, as base case of the number of visitors (i.e status quo), a total of $A = 4,284,458$ monthly visits for ad-filled websites, and $M_{total} = 203,815$ for malware-serving

websites. In addition the number of visitors exposed to *undetected* malware is $M_{exposed} = 48,975$ on average (Table 7.5).

Considering the crime script in Section 9.3.1, we examine the following actor groups, capable of implementing the suggested countermeasures: (i) software vendors, (ii) registrars and hosting providers, (iii) search engines, (iv) payment networks, and (v) advertisement networks.

The perspective of software vendors. Briefly, the software vendors are responsible for implementing countermeasures that prevent attackers from compromising websites and using them, in this case, to serve malware. The relevant impact analysis in Section 9.2.1 is also applicable here, revealing an expected reduction in infections by 0.17% on average (Figure 9.3).

The perspective of registrars and of hosting providers. Countermeasures associated with this actor set, are designed to affect the operation of domains with ad-filled content. However, and considering the fact that this activity, while abusive, is rarely deemed illegal, we do not have any reasonable expectation for this actor group to implement the related situational prevention measures.

The perspective of search engines. This is a much more powerful set of actors, as they currently provide opportunities for identifying trending terms, and directing search traffic to the abusive or malicious destinations. At the same time, their number, and therefore the implementation complexity is small as discussed in Section 9.2.1) (discrete logarithmic distribution, $p = 0.9$).


Completely removing access to the list of trending terms would prevent miscreants from generating content that is tailored directly for the specific terms, leading to a significant reduction in acquired traffic. Nevertheless, we expect that miscreants would be able to find alternative sources of this information, but without doubt of lesser quality and with additional effort (e.g. based on TV news). However,

search engines are also capable of reducing access to auto-generated content and its effects. In 2011 we were able to measure the effects of an intervention from Google [203] which demoted “low quality” search results like the ones we discuss here. We found that this intervention effectively reduced access to ad-filled websites, with the reduction ranging between 41% and 93% (Section 7.4, Table 7.6). Therefore, we assess the effectiveness of situational measures in the present actor group with a uniform distribution in this range.


In Figure 9.11(a) we present the distribution of the impact function of measures. On average, each search engine can reduce the overall traffic headed to abusive destinations by 37.5% (median: 30.8%). However, we note that not all search engines have the same impact, as we assumed in this analysis. To achieve higher accuracy, we would need to examine the effectiveness as a probability conditional to each search engine’s market share.

The perspective of payment networks. While denying payments through payment network intervention would not directly reduce the amount of traffic landing on malware-serving websites, this approach would make the illicit operation unsustainable as it has been documented by McCoy et al. [147]. The authors have noted that many organized criminal networks engaged in this illicit activity have completely ceded their operations after a relevant intervention, and others have been mostly unsuccessful in using alternative payment networks (e.g. PayPal). Therefore, countermeasures at this level can result in complete obliteration or traffic ending up at malware-serving websites, with an impact of $\frac{100\%}{3} = 33.3\%$ per payment network.

The perspective of ad networks. Similar to the case of payment networks, interventions at the ad network level aim at making this abusive operation less profitable, and consequently reduce the incentives to engage in such activity.



(a) Search engines.



(b) Ad networks.

FIGURE 9.11: Plots of impact (Equation 9.3) probability density and cumulative distribution functions of measures targeting trending term exploitation, when considering different actor sets. Analyzed through Monte Carlo simulations with 1,000 iterations.

In assessing the effectiveness of the related measures, we consider the various estimated reductions in fraudulent activities for the different types of ad networks [56, 61, 208], and we examine the effectiveness as a continuous distribution with parameters $min = 23.6\%$ and $max = 99.5\%$. Furthermore, we analyze the complexity of measures, by considering ad networks large enough to reach the majority of Internet users. To this end, we consult the list of top 20 ad networks in April 2014 [41], which identifies 14 of them as reaching at least 50%—[50.3%, 93.8%]—of the 228 million Internet users in that month.

Table 9.4: Average reduction of traffic being subject to trending term exploitation (i.e. effectiveness) per unit of complexity.

Actors	# of actors in group	Expected impact
Software vendors	[200, 500]	0.17%
Search engines	[5, 20]	37.5%
Payment networks	3	33.3%
Advertising networks	[1,14]	12.5%

In Figure 9.11(b) we present the characteristics of the impact distribution in the case of ad network-based situational measures. On average, and assuming that each network has equal impact (which is not necessarily the case based on [41]), the measures can result in a 12.5% reduction in affected traffic per unit of complexity (median: 8.2%).

9.3.4 Overall Assessment

In Table 9.4 we summarize the findings of the impact analysis. Overall, two actor groups—namely the search engines and the payment networks—can have a notable impact against trending term exploitation by introducing hardships in the functionality and profitability of the illicit operation. We also find that advertising networks can also have a significant impact through the detection of fraudulent activities.

Once again, we see two key factors contributing to higher levels of impact. First, a low degree of complexity for implementing a given set of countermeasures allows for faster and more flexible solutions. When this aspect is paired with the ability of each involved actor to have a high degree of control over the critical components of the criminal infrastructures, we are able to maximize the resulting impact.

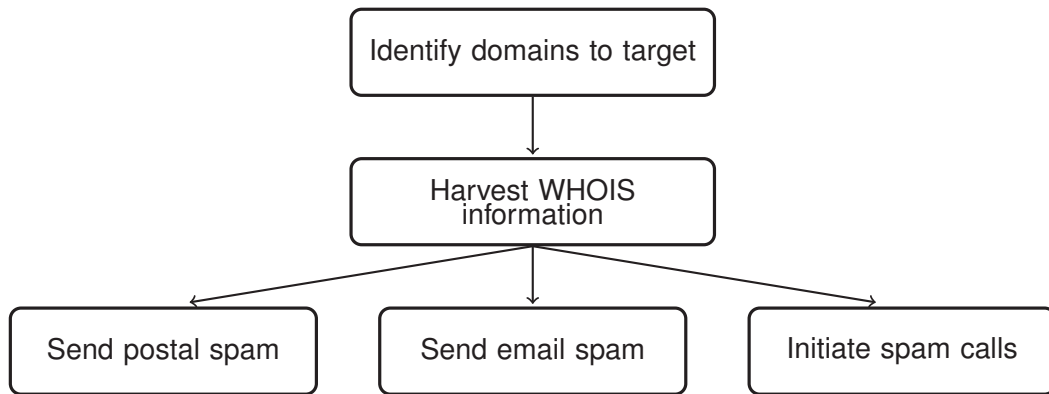


FIGURE 9.12: Components of the crime commission process in the case of WHOIS misuse.

9.4 The case of WHOIS misuse

In this section we examine the case-study of WHOIS misuse from a situational crime prevention perspective, based on our analysis in Chapter 8. We will show that WHOIS misuse is a rather simple case of online crime in terms of its structure, operation, and prevention, especially when compared to two other cases in Sections 9.2 and 9.3. However, the primary purpose of the analysis in this section is to highlight the fact that our methodological approach is not applicable only in complex cases of online crime, but also in more straightforward circumstances.

We start by briefly revisiting the high-level characteristics of WHOIS misuse, in the form of a crime script. We then use the defined script to suggest appropriate situational prevention measures, capable of reducing the opportunities to engage in this fraudulent activity. Finally, we conclude with a qualitative assessment of the problem of WHOIS misuse and of the suggested measures, in an effort to highlight the significance of opportunities in fighting online crime when such activity is relatively simple from a technical standpoint.

9.4.1 *A procedural analysis of WHOIS misuse*

In our empirical analysis of WHOIS misuse in Chapter 8 we have identified its primary forms affecting the majority of registrants: (i) email, (ii) phone, and (iii) postal spam. Consequently, in this analysis we focus exclusively on these types of WHOIS misuse, and in Figure 9.12 we present the related criminal processes through a crime script. In the remainder of the section we will visit each component separately, in an effort to identify their specific procedural characteristics.

Identify domains to target. In order to misuse the information on WHOIS, online criminals first need to identify the set of domain names to use for querying WHOIS. In this regard, there are various sources online for acquiring domain lists either for free or as a paid service. For example, Verisign provides the complete list of `.com` and `.net` domains (among others), in the form of DNS zone files. Other services, like `http://www.dailychanges.com` provide lists of newly registered domains on a daily basis. We also note a technique called “zone file enumeration”, which allows an adversary to effectively reconstruct a zone file containing all registered domains under a given top level domain in an iterative way.

Harvesting WHOIS information. Generally, anyone wanting to get the WHOIS information associated with a domain, does not need to do any research in terms of identifying the authoritative source of this information. Instead, the WHOIS protocol [53] enables services like `www.internic.net/whois.html`²¹ to identify the authoritative source (i.e. the registrar or registries that maintain the specific domain’s registration information), and provide the WHOIS information to the requestor. However, a WHOIS request can also target a specific registrar or registry,

²¹ “InterNIC is a registered service mark of the US Department of Commerce. It is licensed to the ICANN, which operates this web site.”

which can in turn respond with the requested information only when the queried domain is under their realm (i.e. has been registered with the specific registrar or registry).

Different registrars and registries take various approaches in safeguarding WHOIS information from harvesting. Notably, we have found that 37.5% of the 16 largest domain registrars, and one out of five most populous registries do not provide any protective mechanisms (Section 8.4, Table 8.3). Consequently, online criminals can have a higher success rate in their harvesting efforts, whenever a WHOIS request is handled by servers lacking any anti-harvesting measures.

Misusing WHOIS information

Once online criminals have harvested the registrant information of targeted domains, they can use it in a variety of ways. In the next paragraphs we outline the three key types of WHOIS misuse affecting registrants.

Sending postal spam. We have found that online criminals misuse registrants' postal addresses either to advertise commercial products, or to fraudulently request payments (by issuing fake invoices) for services registrants have not requested (Section 8.3.1).

Initiating spam voice calls. Miscreants initiate unsolicited phone calls using the harvested registrants' phone numbers in an effort to sell (usually) Internet-related services (Section 8.3.2).

Sending email spam. This type of WHOIS misuse is the most prevalent one, possibly because of its rather low cost when compared to the other forms of misuse. Our empirical measurement of WHOIS-attributed email misuse showed that it affects 70.5% of domains registered in the five most populous gTLDs (Section 8.3.3).

9.4.2 *Situational measures targeting WHOIS misuse*

In this section we discuss countermeasures against WHOIS misuse from the perspective of SCP, taking into consideration the criminal processes outlined in Section 9.4.1. On a high level, we identify two types of countermeasures: (i) measures against misusing WHOIS to acquire registrant information for illicit purposes, and (ii) measures aiming at protecting registrants after their information has been acquired by miscreants through WHOIS misuse.

Measures affecting WHOIS misuse

The following set of situational prevention measures suggests that reducing the availability of opportunities to misuse WHOIS can result in a reduction of victimized registrants though an increase in the associated criminal efforts.

Reduce anonymity or access to domain lists. The availability of domain name lists is the primary source of information for online criminals before engaging in WHOIS misuse. Consequently, limiting access to this information can be a key approach for reducing misuse. However, ICANN's RAA [101] requires registrars to provide domain zone files to any third party requesting access and agreeing to a specific set of guidelines for proper use of the acquired information. Therefore, measures limiting further publication of WHOIS data can be implemented through the RAA and enforced by the registrars. In terms of limiting zone file enumeration, we note that a technical solution (called *NSEC3* [21]) is already part of the DNS protocol, and is capable of preventing such attempts.

Harden WHOIS. In the current state of affairs, WHOIS, by definition, lacks any security mechanisms enabling authenticated access, leaving it up to the various registrars and registries to implement anti-harvesting measures. Given that a big portion of the registrars do not provide any protective measure, various authors

have suggested updating the WHOIS protocol to incorporate authentication provisions [98, 173, 211]. On a short-term basis, we have empirically shown that various query rate limiting techniques at the registrar and registry level have a statistically significant impact in reducing the occurrence of WHOIS misuse by 2.3 times, compared to when no such measure is in place (Section 8.5).

Measures affecting misuse of information of registrants

Once miscreants gain access to the contact details of registrants, there is little to be done in preventing the misuse of the information. Therefore the following measure suggests limiting the amount of registrant information that is publicly available.

Remove the targets. ICANN's expert working group on directory services has suggested the complete abandonment of WHOIS in favor of service that allows access to registrant information only for permissible purposes [65]. This would effectively remove public access to this information altogether, making WHOIS misuse, as we know it, inapplicable.

9.4.3 Overall assessment of situational measures targeting WHOIS misuse

The problem of WHOIS misuse, and the possible solutions are apparently simple. We have experimentally shown that implementing a query rate limiting mechanism at the registrar and registry level is a straightforward way of significantly reducing the occurrence of WHOIS misuse. Other, more radical approaches like eliminating WHOIS altogether have the potential of similar or greater impact against this illicit operation. Engaging in a more detailed complexity-effectiveness analysis of the suggested countermeasures would essentially require a similar analysis of intervention at the registrar level as in Sections 9.2.2 and 9.3.3, adjusting only for the

different effectiveness factor. Therefore, we argue that such analysis will provide only marginal benefits in our understanding.

We highlight instead the fact that allowing unrestricted access to WHOIS makes it easy for online criminals to harvest registrant contact information for illicit or fraudulent purposes. Therefore, online criminals do not necessarily need elaborate technical skills to engage in their illicit activities, as with the previously discussed case studies; Misusing WHOIS requires a few lines of code. Therefore we argue that the criminal efforts are inversely proportional to the extent of available opportunities.

9.5 Concluding remarks: Towards a generalizable methodology for online crime analysis and prevention

In the previous sections we examined in detail the criminal processes enabling three cases of online crime: (i) the online prescription drug trade, (ii) the trending term exploitation, and (iii) WHOIS misuse. This analysis is heavily based on empirical examination and measurements of the illicit operations, and it enhances our understanding of the online criminal infrastructures and the interactions among their components. This understanding is of paramount importance as it enables the subsequent identification and critical evaluation of situational prevention measures capable of increasing the criminal efforts and risks in engaging in such activity, while concurrently reducing the associated criminal profits.

Examining online crime on a case-by-case basis is important in itself to prevent further victimization of Internet users. However, we now take a holistic approach in studying online crime, in an effort to define the components of a generalizable methodology for online crime analysis and prevention. We start by creating a canvas of the common aspects and characteristics of the three case studies,

to identify the components of the criminal infrastructures being critical from an operational, an economic, and a preventive perspective.

9.5.1 Commonalities in criminal infrastructures

Online criminals often exploit insecure software and poorly maintained websites to achieve their goals. Nevertheless, the economic incentives that could potentially drive efforts to reduce the criminal opportunities are either minimal or nonexistent. Therefore, the following discussion is focused on the more actionable elements of the criminal infrastructures.

Search engines and payment networks. In the majority of the criminal operations we examine, search engines and payment networks are big part of the problem, but of the solution as well. Search engines enable online criminals in terms of identifying their victims (e.g. vulnerable websites), and of funneling web traffic into illicit businesses. Further on, payment networks allow online criminals to monetize this stream of potential customers, giving them further incentive to continue operating their illicit business.

Both types of actors share important characteristics. They are limited in number (at least the most popular ones), and they are overly important for the function of the illicit operations. These characteristics make them very effective whenever they take an action that limits the opportunities for offending. In addition, they suggest that search engines and payment networks are, in a sense, part of the critical infrastructure of online crime.

Registrars and internet service providers. Equally essential are the various service providers that enable online criminal activity by providing necessary resources to the miscreants. Examples of such resources are the Internet locations

(i.e. websites) actively engaging in illicit activity, and personal information of potential victims.

However, the greater size of this group of actors, has the potential to impact the implementation complexity of effective countermeasures. While specific actors in these groups may be more powerful in terms of their market share, we cannot argue that employing only the specific subset of actors for implementing a set of countermeasure will have similar effectiveness as with, for example, search engines. In such case, online criminals can move to a different service provider, and continue with their illicit operation. On the other hand, the miscreants do not have the option of choosing which search engine they will manipulate as their only option is to make every effort to target the most popular ones, in order to maximize their expected profit.

Law enforcement. The global scale of online criminal operations is evidently a significant hardship from the perspective of law enforcement. International cooperation is necessary for targeting criminal operations taking place beyond the jurisdiction of the victimized population. Consequently, online criminals have the incentive to diversify the physical locations of their infrastructures, whenever this is applicable. However, empirical analysis of online crime can inform the decision-making process. This way, targeted enforcement can have a detrimental effect, especially whenever a physical relocation of criminal resources imposes a significant financial burden to online criminals (e.g. clandestine drug manufacturing)

9.5.2 Designing effective solutions

This discussion naturally leads to the following question: How can we deal with online crime in a unified, methodical, and efficient manner? Historically, the different components of online crime have been targeted in isolation, either by law

enforcement, or through technical solutions. We have showed that this approach has only short-term or superficial effects, as it usually does not affect the critical components of the criminal infrastructures. The overall problem is not that there is no incentive to target those components, but the fact that they require often complex methods to bring them out of obscurity. The methodology we suggest here takes instead an empirical approach in studying online crime, looking for the processes most vulnerable to intervention. We have identified the following common criminal processes that have such characteristics:

- Identifying potential victims
- Scaling the attacks, and
- Monetizing the attacks

The commonalities across cases of online crime we presented in Section 9.5.1 serve as indicators and guidelines for potential intervention points. However, we do not expect that they are applicable in all cases of online crime. Nevertheless, there are cases that may involve identical or similar criminal structures, and, in such cases, identifying potential solutions can be rather trivial. For example, the cases of the illicit online trade of counterfeit software, watches, and books employs the same methodological aspects of abusive advertising examined in Section 9.2.1. In addition, while monetizing trending terms through ad-filled websites (Section 9.3.1) is distinctly different than profiting from illicit sales of prescription drugs (Section 9.2.2), they both share characteristics that allow online criminals to profit from their illicit operations, by taking advantage of the poor (regulatory or operational) protection of the specific monetization paths.

9.5.3 Limitations and future work

Our definition of the impact metric necessarily makes a set of assumptions that have been discussed throughout this chapter. We briefly reexamine them here with the intent to provide insights into possible ways of improving our assessments.

Specifically, the complexity metric incorporates the assumption that all actors within a group of actors (e.g. search engines) have the same potential of effectiveness. This assumption is largely dependent both on the specific type of actors, and on the temporal granularity of our estimation. For example, considering the impact of interventions against illicit advertising that service providers are capable of undertaking (Section 9.2.1), our complexity estimation assumes that all service providers have the same potential of effectiveness. However, examining the immediate impact of such intervention, service providers with high market share, or with high concentration of traffic brokers have a potential for higher immediate effectiveness. In this case, though, we argue that online criminals can start using other services providers who are not participating in this effort, considerably limiting the mid- and long-term effectiveness of this small-scale intervention.

At the same time, the same observations are not applicable to the group of search engines. For example, requiring just Google—with a market share of about 67.6% [40]—to take action against traffic redirection would result in the same short-, mid- and long-term reduction in the specific illicit activity. A key fact in understanding the difference between those two groups of actors, is that online criminals are not capable of choosing which search engine to exploit, contrary to the case of service providers. The criminals are rather dependent on the search engine that handles the most search traffic, towards keeping their illicit operation profitable.

With those considerations in mind, our interpretation of impact could become notably richer if it incorporates such temporal characteristics, as well as an expected “switching cost” among actors. This cost should be characteristic of the ability of the criminals to use an alternative actor, following an intervention that disrupts a previously established exploitation of a different actor within the same group.

As highlighted in the previous discussion, some of the limitations of the impact estimations are an artifact of the limitations in the complexity estimations. To this end, we argue that including the aspect of market share of each actor within a group of actors may potentially enhance our understanding of immediate impact. In addition, considering the monetary estimates of the actual implementation costs of countermeasures can provide deeper insights both in terms of the associated private costs—e.g. how much does Google have to invest towards protecting vulnerable websites?—and also in terms of the societally acceptable levels of enforcement of such measures—i.e. the cost of the deterrents should not exceed the cost of the targeted criminal activities.

Summary and conclusions

The fear of punishment, as a deterrent, serves as the cornerstone of modern justice systems. While laws deem specific actions and behaviors as illegal, without the deterrent effect of punishment, a risk-seeking individual will inevitably choose to profit by victimizing a vulnerable target. Online crime is an obvious artifact of such opportunistic behavior, as online criminals are rarely brought to justice.

Contrary to traditional crime, the characteristic features of online crime often make it immune to traditional intervention approaches that would normally act as deterrents. Specifically, it is performed within a globalized virtual environment, the Internet, which allows for a certain degree of anonymity—or at least perceived anonymity. In this case, anonymity enables miscreants to profit illicitly or fraudulently without the fear of attribution, prosecution, and punishment [13]. In addition, even when a criminal action can be attributed to specific actor(s), jurisdictional complications often allow such actions to remain unpunished.

In this thesis we examine online crime and evaluate present deterrent efforts. We build upon three contemporary cases of online criminal activity, finding that

such efforts are often misapplied or poorly coordinated, leading into nullified long-term effects. We combine extensive empirical measurements and economic analyses to offer an understanding as to why current interventions, while increasingly using more resources, remain largely ineffective. Along this process, we uncover the existence of often complex infrastructures supporting the illicit online activities. Our empirical analysis shows that within these infrastructures, there exist procedural components with similar characteristics across different types of online crime. Such components depend on resources that are limited in number (e.g. search engines, payment processors), but essential for the criminal operations. Most importantly, due to the obscure nature of criminal online operations, the actors controlling the critical resources, while being capable of implementing effective interventions, are often unaware of their role in enabling the illicit activities. Consequently, this lack of awareness is translated into an opportunity for their victimization. In fact, whenever serendipity allowed us to measure the effects of an intervention affecting—only by chance—the availability of such opportunities, we were able to observe their detrimental effect in the criminal profitability.

We consider this renewed understanding of opportunities—allowing online criminals to exploit limited resources to the benefit of their illicit operations—towards suggesting more effective countermeasures. To this end, the theory of Situational Crime Prevention (SCP), rooted in the domain of criminology, appears as a plausible approach towards reducing the availability of those opportunities. SCP suggests that reducing such opportunities forces miscreants to reevaluate their methods; For example, they will either try to overcome the introduced hurdles by accepting more risk—leading to an increased risk of apprehension—or they will simply accept the new status quo—leading to a reduction in their profits.

In either case, countermeasures compatible to SCP can act as efficient, long-lasting deterrents.

Informed by our empirical analysis, we break down three cases of online crime into their procedural components using Crime Script Analysis (CSA). CSA, originating from the domain of cognitive psychology, allows us to identify and suggest situational prevention measures applicable at every step of the criminal processes. In addition, we assess the measures' impact using a novel index we define as complexity-effectiveness. This impact measure effectively incorporates the notion of a critical resource by considering countermeasures through the degree of complexity of their implementation. In essence, a measure is characterized by the extent of reduction of an illicit activity per unit of complexity. At a given level of reduction in crime, countermeasures requiring the involvement of a large number of actors fair worse than countermeasures requiring the participation of fewer actors. For example, an intervention employing search engines to reduce the amount of hijacked traffic landing at unlicensed online pharmacies has more impact than a measure achieving the same degree of reduction while requiring the participation of software vendors (Section 9.2.1).

The previous discussion naturally leads to the following question: How applicable or extensible is this analysis to online crime in general, beyond the case studies examined in this thesis? While we put significant effort in measuring and analyzing (i) the case of the illicit online prescription drug trade, (ii) the case of trending term exploitation, and (iii) that of WHOIS misuse, we argue that neither their traits are necessarily characteristic of online crime in general, nor are the suggested situational prevention measures generally applicable. Alternatively, we could have used various other cases of online crime—like the high-yield Ponzi investment programs [161], and typosquatting [160]—to derive similar insights.

The cases we selected are, however, adequate to provide the insights we set out gather, and important from a societal perspective.

Our contribution is rather in the methodological approach we follow to (i) understand the structure of online criminal networks, (ii) identify the associated critical resources providing opportunities to profit illicitly, and (iii) suggest and evaluate appropriate countermeasures. We show that measures lacking empirical support, or not targeting critical resources are often futile. We further argue that policy makers and technology providers need to work in tandem, to finally get the upper hand in incapacitating online criminals. In addition, through this work, we suggest that the research community engaged in measurements of online crime, can receive significant gains by combining their work with well-established concepts from different scientific domains. Indeed, in this thesis have been able to use traditional crime prevention concepts from criminology, and effectively adapt them to the unique characteristics of digital crime.

This work should be received as the beginning of a fascinating journey towards making Internet more secure. Considering our methodological contributions in the fight against online crime, future research work should attempt to apply this methodology to the various other cases of illicit online activity. There is already significant effort from the research community in providing an empirical-based understanding for a wide range of online criminal processes. However, such work often falls short in assessing and evaluating countermeasures that target the availability of criminal opportunities. Adopting a comprehensive approach—like the one we outline in this dissertation—to characterize the illicit online activity, and use this understanding to identify and reduce such opportunities can lead to a more effective, scientific, and result-oriented approach against online crime.

Furthermore, we highlight the need to expand our understanding that characterizes the impact and the complexity of suggested countermeasures. The model

we propose in this dissertation takes a rather simplistic approach of estimating complexity as the number of actors capable of undertaking a set of countermeasures. However, per our discussion in Chapter 9, our understanding of complexity could be enhanced through the inclusion in this model of the specific characteristics of those actors. Such characteristics are potentially the actual monetary costs of implementing a countermeasure, and the effectiveness—from the criminal’s perspective—of switching to a different actor who is not participating in an intervention. For example, a criminal cannot “simply” switch from manipulating search engine A to search engine B once A has implemented a situational prevention measure, because the criminal cannot control which search engine his potential victims use. However, the same argument does not apply for Internet service providers enabling the operation of traffic brokers (Section 9.2.1). In addition, a version of the impact metric that considers the immediate and the mid-term effect of an intervention can allow us to gain deeper insights into what an effective strategy for deploying a countermeasure would look like—e.g. prioritize intervention at search engine A and then expand to other search engines.

Such enhancements in the evaluation of impact would also better inform the public policy-making process. For example, in answering a question like “What is a tolerable level of online crime?”, we need to understand first the actual direct costs in implementing countermeasures. Accepting a certain degree of online crime means that we have reached an equilibrium between the benefits to the society (i.e. reduction in crime) and the costs of interventions. In such occasion, increasing the level of intervention to reach, for example, zero tolerance, may be a financially irrational effort. However, such evaluations need a deeper understanding of the intervention costs. This dissertation introduces tools in this direction, by providing evidence-based guidance as to which interventions have the potential to be more cost-effective.

Appendix A

Surveying registrants on their WHOIS misuse experiences

In this appendix we provide the fine details—in terms of methodology and findings—of the pilot registrant survey [127] that guided our measurements on WHOIS misuse in Chapter 8 (Section 8.1.2). In essence, we surveyed a representative sample of registrants with domains in the 5 most populous gTLDs to gain a better understanding of their direct experiences with WHOIS misuse. The details on the registrant sample selection are available in Section 8.1.1.

We start with a discussion on the methodology and design details of the survey in Section A.1. In Section A.2, we describe issues presented during the survey, which affected the degree of representation of our findings. Then, in Section A.3, we then present the demographics of the registrant sample, and our discoveries related to the ways registrants experience misuse of their personal information as a consequence of its public availability in WHOIS. We conclude in Section A.4, where we also provide a discussion of the limitations of the present analysis.

A.1 Methodology

We used email messages to invite registrants to participate in the survey. We acquired their contact information through the WHOIS entries associated with the domains in our sample (Section 8.1.1). The invitation contained a short description of the study, information about the principal investigator, and links to either participate in the survey or opt out from any future messages and reminders from us. Because this survey was designed to be taken by non-Internet-savvy registrants, the invitation (i) briefly described domain registration and the role of WHOIS data in simplified language, (ii) included the name of the sampled domain name included in our survey, and (iii) suggested that invitees check the information available through WHOIS on the domains they own. We also offered the option to download the questionnaire and email the responses to us. The content of the invitation is available in Section B.1 of Appendix B.

When participants clicked on the link to participate they were presented with a consent form that describes briefly the procedures, requirements, risks, benefits, associated compensation (entry into a random prize drawing), and privacy assurances we offered. The text is available in Section B.2 of Appendix B.

The survey lasted three and a half months—from September 2012 until December 2012. The invitations were sent out in stages, and each group of invitees was offered a period of five weeks to complete the survey. We also scheduled the distribution of weekly reminders to non-respondents, increasing the response rate. The survey was implemented with SurveyMonkey,¹ and all connections to the service were encrypted. Invitees were assured that all responses would be treated as confidential, with survey data published only in aggregate, anonymized form.

¹ <http://www.surveymonkey.com>

A.1.1 Survey translations

Our sample of 1619 registrants covers 81 countries, requiring a significant effort to translate the survey in all associated languages. Given that many of the 81 countries were mapped to a handful of participants, and the expected low response rate (15%), we decided to not produce all required translations. We observed that 90% of registrants was located in just 18 countries, with the remaining 10% spread across 63 countries. Hence, we provided the survey in the native language of the top 90% of the participants.

In all, the survey questions were available in the following languages: English, Chinese, French, Japanese, Spanish, Italian, and Portuguese. The 18 countries included Germany and Turkey, but were not able to secure proper translations. Therefore, we offered the English version of the survey to participants from those two countries. This effectively reduced the portion of participants surveyed in their native language to 84.9%.

We relied on native speakers of various languages from Carnegie Mellon University for the translations with a background in computer networks or computer security. These characteristics allowed them to have the technical background to produce meaningful translations, and to integrate nuances of the different cultures. In addition we also offered definitions for key terms used in the survey questions to accommodate participants unfamiliar with the technical jargon. These definitions are available in Section B.4.

A.1.2 Types of questions

The survey was divided in three parts. The first set of questions was designed to collect data on the demographics of the participants. The second part asked questions about seven different types of misuse of WHOIS: postal spam, email

spam, voice spam, identity theft, unauthorized intrusion to servers, denial of service, Internet blackmailing. In addition, we included an open-ended section for any other type of misuse a registrant may have experienced. We requested that the participants optionally provide a detailed description of their experiences in any of the previous categories.

Due to the length of the survey—which could take up to 30 minutes to complete—we assessed that a large portion of participants would abandon the survey before completion. In an effort to avoid biases related to the design of the survey (i.e. the order of the questions) we randomized the sequence of questions for the different types of misuse.

The third and final part of the survey collected information related to the actions taken by the participants in response to their WHOIS information being misuse.

A.2 Response and error rates

Between May and August of 2012, we ran two pilots of the registrant survey to assess possible issues with the design and/or implementation of the survey. One pilot involved tech-savvy colleagues at CMU with great experience in user surveys. This pilot helped us identify and fix a number of design issues. The second pilot targeted a broader audience of randomly-selected English speaking registrants, and was intended to assess the expected response rate. In this second pilot, we did not receive any responses out of the 48 invitations sent. We identified as a possible problem the excessive length of the survey, which apparently discouraged participation. Therefore, we attempted to remedy this by offering entry into a random prize drawing to participants that would complete the survey in its entirety. Note there was no incentive to report having encountered misuse; re-

spondents were only required to complete survey sections that pertained to their experiences.

Overall, we sent out 1619 invitations and had 57 participants: 52 in English, 3 in Japanese, and 2 in Spanish, achieving a response rate of 3.6%. Out of these 57 participants, 41 were complete responses. Such a low number in collected responses impacts our targeted levels of significance, namely the error rate. The resulting error rate for the statistic we are measuring (is there observed WHOIS misuse?) is 12.7%. This means that for 95% of the population, the measured misuse deviates from the actual misuse in 12.7% of registrants. For the remaining 5% of the population, the measured misuse can deviate by more than 12.7% from the actual value (i.e. far more or far less misuse).

A.3 Analysis of responses

We start the analysis of the collected responses by first giving an overview of the characteristics of the sample in terms of the demographics, and the reported knowledge on WHOIS. We then delve into the specific types of reported WHOIS misuse.

A.3.1 Characteristics of the participants

From a demographic standpoint, the participants were mainly from English speaking countries (92%) even though we made efforts—as previously discussed—to include a wide geographical range of participants. We collected responses from the following countries (in descending order of the number of participants): USA, Japan, United Arab Emirates, Australia, Canada, Switzerland, Germany, Spain, UK, India, and Mexico. There were also respondents that did not disclose their location.

Although each registrant was surveyed just once, the majority of the participants (60%) have more than 10 domains registered, with 9% of the participants operating a single domain. Additionally, the domains in our sample are mainly registered by self-described for-profit businesses or organizations (49%), followed by the domains registered by individuals (33%), and domains registered by non-for-profit organizations (14%). Moreover, respondents reported that most of the domains (46.5%) in our sample are used for commercial activities. Finally, the great majority of the participants (93%) indicated they were aware of the existence and purpose of WHOIS.

Comparing the self-reported demographics of our sample with the WHOIS-based findings of the WHOIS registrant identification study [175], we see that the top two categories are occupied by similar entities in both studies; Individual—natural—person registrants appeared roughly with the same frequency (30% vs. 33%). In our sample though, the combined share of categories representing legal entities was 62% compared to 39% in [175].

A.3.2 Reported WHOIS misuse

We now present our findings for each specific type of WHOIS misuse we evaluate. In each set of questions, we first asked the participants to report if they have experienced misuse affecting a specific type of information supplied when registering their domains. If the answer was yes, we then asked more specific questions about those misuse incidents. 25 of the respondents (43.9%) reported experiencing some kind of misuse of their WHOIS information, mainly affecting postal addresses, email addresses, and phone numbers.

Postal address misuse

38.6% of surveyed registrants (22) have received postal spam mailed to an address published in WHOIS, and 29.8% (17) believed the unsolicited mail resulted from misuse of their WHOIS postal address. As a proof of their suspicion, participants provided details of the unsolicited mail; it was either directly related to one of their domains, or it advertised web services. Moreover, 21.1% (12) of the participants reported that their WHOIS postal address was not published in any other public directory (e.g. phone book, website, etc.).

The majority of the respondents having received postal spam (14% of total, 8) experience this misuse a few times a year, with 11% (6) receiving postal spam a few times a month, and 5% (3) less than once a year. The reported subjects of the unsolicited correspondence were mainly related to fake domain name renewals and transfers, followed by messages related to website hosting, and search-engine optimization (SEO) services.

Email address misuse

25 registrants (43.9%) reported receiving spam email at an account associated with a WHOIS email address. 29.8% (17) of those associate the misuse of their email address to WHOIS because the topics of the spam emails specifically targeted domain name registrants (e.g. domain name transfer offers, SEO offers). 14% (8) of the registrants stated they have not listed the misused email address in any other public directory.

The majority of the respondents (10%, 6 registrants) identifying WHOIS misuse as a cause for email spam reported receiving such emails a few times a day, followed by 9% of responses (5 registrants) receiving unsolicited email a few times

a week. The topics of the unsolicited messages are similar to the ones reported for postal spam.

Phone number misuse

22.8% (13) of registrants reported receiving voicemail spam, with 12.3% (7) attributing the spam to WHOIS misuse. They were able to associate the voicemails with WHOIS because the caller either explicitly mentioned a domain name under the registrants' control or they were offering Internet-related services. 9% (5) of the registrants who claimed to have experienced this type of misuse, had reportedly not listed their number in any other public directory.

Identity theft

Two of the participants reported having experienced identity theft, but none could tie these events to WHOIS misuse.

Unauthorized intrusion to servers

In order to measure the extent of misuse of WHOIS information to gain unauthorized access to servers, we first asked the participants if they were the system administrators of Internet servers associated with one of their registered domains. The number of participants that have this role is very small (7%, 4), with just one person experiencing unauthorized intrusion. That respondent, however, could not link this intrusion to WHOIS misuse.

Blackmail

One participant reported being a victim of blackmail as a result of their information being published in the WHOIS directory. The registrant was allegedly accused by a third-party company of violating the terms of domain registration because of the

name the registrant chose for the domain. The registrant was offered the option to settle in exchange for a fee, and—after consulting with lawyers—decided to not take any action. After a few months, and a series of emails from the third party, the latter stopped communicating with the registrant.

Other

Although we gave registrants an opportunity to describe types of WHOIS misuses not otherwise covered, no participant claimed to have experienced any other type of WHOIS misuse.

A.4 Discussion

Getting registrants to communicate their experiences in terms of the possible misuse of their personally identifiable information listed in WHOIS proved to be a challenging task. Even with an incentive to participate—a raffle at the end of the survey—we were able to collect responses only from a small portion of invitees (57 out of 340, or 17%). However we managed to get a clear insight into the prevalence of WHOIS misuse, and the specific types of information usually targeted.

43.9% of registrants claim to have experienced some type of WHOIS misuse. Given the margin of error (12.7%) this observation neither confirms or disproves that WHOIS-misuse is affecting the majority of registrants. It does confirm though the hypothesis that public access to WHOIS data leads to a measurable and statistically-significant degree of misuse.

Email addresses are mostly targeted, followed closely by the postal addresses. Phone numbers are also misused, but with a much smaller occurrence and higher adverse impact per incident.

A.4.1 Potential survey biases

We contemplate the biases the survey design introduced to evaluate the possibility of over or under-reporting of WHOIS misuse. First, by not providing translated versions of the survey to 15% of the sample, we may have missed some incidents of misuse experienced by registrants that do not speak English. However, given the observed response rate (3.6%), the expected response rate of that portion of the sample (15%) is less than 1%—3.6% of 15%. In retrospect, even if provided all the possible translations, we would not receive a statistically-significant number of responses from this group.

Another possible bias is that registrants may be more willing to report a harmful act (e.g. experience with misuse) rather than a lack of harmful incidents, over-representing WHOIS misuse. In addition, we did not attempt to verify or corroborate any WHOIS misuse incident, which could lead to false representation of the extent of WHOIS misuse. However, the strong economic incentive we provided—entry into a random prize drawing—should mitigate this potential source of bias.

Finally, the great majority of the survey participants originated from North America. This fact affects our findings in the following ways; first, we are unable to analyze the geographical distribution of misuse, as the survey suffers from coverage bias. Consequently, findings are also descriptive of a narrower portion of the world population than we intended.

Appendix B

Registrant survey supplemental material

B.1 Invitation to participate in registrant survey

Dr. Nicolas Christin

Carnegie Mellon University – CyLab

4720 Forbes Avenue, CIC Rm 2108

Pittsburgh, PA 15213 USA

<http://www.andrew.cmu.edu/user/nicolasc/>

Please click here to verify authenticity of this email:

<http://dogo.ece.cmu.edu/whois-study/>

Dear [FirstName], Sampled Domain Name: [CustomData]

Interested in winning the new Apple iPad 4G or an Apple iPod Shuffle? Read on.

We are computer security researchers in Carnegie Mellon University’s Cyber Security Lab (CyLab) (<http://www.cylab.cmu.edu>). We are conducting a study that may help reduce Internet-based crimes, and we need your help!

At some point — perhaps when you created a website or an email account — you registered a domain name. During registration, you were asked to provide contact details (name, email, phone number, address). These details are published in a public Internet directory called "WHOIS." ANYONE, including us, can look up this directory to find out registration information. By sharing your experience as a domain name Registrant, you can help us better understand potential misuses of WHOIS registration data.

The results of this study will help the Internet community to fight various forms of online crime. We will NOT collect your personal information, unless you specifically give us permission to contact you to discuss this survey. Information about this option is available at the end of the survey. The survey should take about 30 minutes to complete, and will ask questions about the domain name you have registered and your experience using it.

You can complete the survey in two ways:

- Complete and submit an on-line survey form by clicking [SurveyLink] (PREFERRED)
- Download survey questions from http://dogo.ece.cmu.edu/whois-study/WHOIS_Misuse_Survey_Registrant_Printable.pdf and email answers to whois-study@andrew.cmu.edu.

We aim to complete this survey by [closing date here]. Please click on the link below if you do not wish to participate or receive further communication from us. You will not be contacted further. [RemoveLink]

If you fully complete the survey, you will be entered in a drawing for a chance to win one new iPad ("iPad 3") 16GB with 4G, or one of four 2GB iPod Shuffle.

Thank you very much for your time and consideration. We look forward to hearing from you.

Sincerely,

—

Nicolas Christin, Ph.D

Carnegie Mellon University CyLab

B.2 Consent form

This survey is part of a research study conducted by Prof. Nicolas Christin at Carnegie Mellon University.

The purpose of the research is to investigate the extent to which public availability of certain information online leads to the information being misused by unauthorized parties.

Procedures.

Participants are expected to answer a survey. The expected duration of participation is 30 minutes.

Participant Requirements.

Participation in this study is limited to individuals age 18 and older.

Risks.

The risks and discomfort associated with participation in this study are no greater than those ordinarily encountered in daily life or during other online activities.

Benefits.

There may be no personal benefit from your participation in the study, but the knowledge received may be of value to humanity.

Compensation & Costs.

By fully completing the survey, you will be entered in a drawing for a chance to win an Apple iPad 4G, or one of four Apple iPod Shuffle. There will be no cost to you if you participate in this study.

Confidentiality.

By participating in this research, you understand and agree that Carnegie Mellon may be required to disclose your consent form, data and other personally identifiable information as required by law, regulation, subpoena or court order. Otherwise, your confidentiality will be maintained in the following manner:

Your data and consent form will be kept separate. Your consent form will be stored in a locked location on Carnegie Mellon property and will not be disclosed to third parties. By participating, you understand and agree that the data and information gathered during this study may be used by Carnegie Mellon and published and/or disclosed by Carnegie Mellon to others outside of Carnegie Mellon. However, your name, address, contact information and other direct personal identifiers in your consent form will not be mentioned in any such publication or dissemination of the research data and/or results by Carnegie Mellon.

Right to Ask Questions & Contact Information.

If you have any questions about this study, you should feel free to ask them by contacting the Principal Investigator now at

Dr. Nicolas Christin
Carnegie Mellon INI & CyLab
4720 Forbes Avenue, CIC Room 2108
Pittsburgh, PA 15217 USA
Phone: 412-268-4432
Email: nicolasc@cmu.edu

If you have questions later, desire additional information, or wish to withdraw your participation please contact the Principal Investigator by mail, phone or e-mail in accordance with the contact information listed above.

If you have questions pertaining to your rights as a research participant; or to report objections to this study, you should contact the Research Regulatory Compliance Office at Carnegie Mellon University. Email: irb-review@andrew.cmu.edu. Phone: 412-268-1901 or 412-268-5460.

The Carnegie Mellon University Institutional Review Board (IRB) has approved the use of human participants for this study.

Voluntary Participation.

Your participation in this research is voluntary. You may discontinue participation at any time during the research activity.

I am age 18 or older.	Yes	No
I have read and understand the information above.	Yes	No
I want to participate in this research and continue with the survey.	Yes	No

B.3 Survey questions

1. How many domain names have you currently registered?

- 1
- 2-10
- More than 10

2. Please list all of the domain names that you have registered. If you registered more than one name, please separate them with commas (,) – for example, “mycorp1.com, mycorp2.com.”

[Open ended]

2.1 Please tell us the “sampled domain name” that appears in your survey invitation letter.

[Open ended]

When answering questions that follow, please think about your experiences as the Registrant of this sampled domain and communication sent to addresses that you supplied when registering that domain. Before continuing, you may find it helpful look up your own domain in WHOIS using <http://whois.domaintools.com>.

3. Thinking about why you registered this domain name and how you use it, please indicate which of the following categories best describes you as this domain name’s Registrant:

- I registered the domain for my own use as an Individual
- I registered the domain for use by a For-profit business or organization
- I registered the domain for use by a Non-profit organization
- I registered the domain for use by an informal interest group (e.g., tennis club)
- Other (please specify)

3.1 Is this domain name used for any commercial activities – for example, to sell or advertise goods or services or to collect donations?

- Yes

- No

- Not sure or prefer not to answer

4. Please indicate the country that you identified when you registered this domain name. Note: WHOIS identifies several contacts for each domain name, including an administrative contact (usually you) and a technical contact (may be your Internet service provider). Here, we are interested in the country identified in YOUR contact details.

(Drop down list)

5. Please identify the Registrar (that is, the registration service provider) from whom you obtained this domain name. If you do not know or recall, you may leave this blank.

[Open ended field]

6. Before taking this survey, did you know that the contact details which you provided during domain registration would be publicly available on the Internet through "WHOIS"?

[Yes/No]

7. Since registering this domain name, have you ever received unsolicited *postal mail* at any of the *postal addresses* that you specified in contact details during domain registration?

[YES/NO]

7.1 [If yes to Q7] Do you have reason to suspect that you received this unsolicited postal mail because your postal address was published in WHOIS?

[YES/NO]

7.1.1 [If yes to Q7.1] Why do you think so?

[Open ended field]

7.1.2 [If yes to Q7.1] Is the postal address published in another public directory or Internet source (for example, a phone book, a website, your email signature)?

[Yes/No]

7.1.3 [If yes to Q7.1] How often do you receive unsolicited postal mail at the postal addresses published in WHOIS?

- A few times in a week
- A few times in a month
- A few times in a year
- Less than once in a year

7.1.4 [If yes to Q7.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

7.1.5 [If yes to Q7.1] Please describe reasons for which you were contacted in these cases (e.g., a domain name hosting services offer)

[Open ended]

7.1.6 [If yes to Q7.1] If you know or can recall who contacted you in a recent case, please tell us more about that entity (e.g., sender's name, type of company)

[Open ended]

7.1.7 [If yes to Q7.1] Did this unsolicited postal mail have any adverse impact on you?

- Yes (describe)
- No

7.2 [If no to Q7.1] Could the postal address have been obtained from another public directory or Internet source (for example, a phone book, a website, your email signature)?

[Yes/No]

7.2.1 [If no to Q7.2] How do you think your postal address was obtained?

[Open ended]

8. Since registering this domain name, have you ever received *unsolicited electronic mail* at any of the *email addresses* that you specified in contact details during domain registration?

[YES/NO]

8.1 [If yes to Q8] Do you have reason to suspect that you received those emails because your email address was published in WHOIS?

[YES/NO]

8.1.1 [If yes to Q8.1] Please specify why you think so.

[Open ended field]

8.1.2 [If yes to Q8.1] Is the misused email address published in another public directory or Internet source (for example, a website, your email signature, Facebook, Twitter)?

[Yes/No]

8.1.3 [If yes to Q8.1] How often do you experience misuse of your email address published in WHOIS?

- A few times a day
- A few times in a week

- A few times in a month
- A few times in a year
- Less than once in a year

8.1.4 [If yes to Q8.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

8.1.5 [If yes to Q8.1] Please describe the reasons for which you were contacted in these cases (e.g., a domain name hosting services offer, targeted phishing email)

[Open ended]

8.1.6 [If yes to Q8.1] If you know or can recall who contacted you in a recent case, please tell us more about that entity (e.g., sender's name, type of company)

[Open ended]

8.1.7 [If yes to Q8.1] Did this unsolicited email have any adverse impact on you?

- Yes (describe)
- No

8.2 [If no to Q8.1] Could the email address have been obtained from another public directory or Internet source (for example, a website, your email signature, facebook, twitter)?

[Yes/No]

8.2.1 [If no to Q8.2] How do you think your email address was obtained?

[Open ended]

9. Since registering this domain name, have you ever received *unsolicited voice calls* at the *phone number(s)* that you specified in contact details during domain registration?

[YES/NO]

9.1 [If yes to Q9] Do you have reason to suspect that those unsolicited voice calls happened because your phone number(s) are published in WHOIS?

[YES/NO]

9.1.1 [If yes to Q9.1] Please specify why you think so.

[open ended]

9.1.2 [If yes to Q9.1] Is the misused phone number(s) published in another public directory or Internet source (for example, a phone book, a website, your email signature)?

[Yes/No]

9.1.3 [If yes to Q9.1] How often do you experience misuse of your phone number(s) published in WHOIS?

- A few times a day
- A few times in a week
- A few times in a month
- A few times in a year
- Less than once in a year

9.1.4 [If yes to Q9.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

9.1.5 [If yes to Q9.1] Please describe the reasons for which you were contacted in these cases (e.g., a domain name hosting services offer)

[Open ended]

9.1.6 [If yes to Q9.1] If you know or can recall who contacted you in a recent case, please tell us more about that entity (e.g., sender's name, type of company).

[Open ended]

9.1.7 [If yes to Q9.1] Did these unsolicited calls have any adverse impact on you?

- Yes (describe)

- No

9.2 [If no to Q9.1] Could the phone number have been obtained from another public directory or Internet source (for example, a phone book, a website, your email signature)?

[Yes/No]

9.2.1 [If no to Q9.2] How do you think your phone number(s) was obtained?

[Open ended]

10. Since registering this domain name, have you ever had your *identity* (e.g. name, address, phone number) abused or stolen? An example would be fraudulent use of your identity (without your knowledge) to apply for a credit card or receive financial services.

[YES/NO]

10.1 [If yes to Q10] Was this identity specified in contact details during domain registration?

[Yes/No]

10.1.1 [If yes to Q10.1] Do you have reason to suspect that the identity abuse happened because your identity details are published in WHOIS?

[YES/NO]

10.1.1.1 [If yes to Q10.1.1] Please specify why you think so.

[Open ended]

10.1.1.2 [If yes to Q10.1.1] Are the misused identity details published in another public directory or Internet source (for example, your email signature, a workplace directory, Facebook)?

[Yes/No]

10.1.1.3 [If yes to Q10.1.1] How many times have been your identity published in WHOIS abused or stolen?

- Once
- Twice
- Three times
- More than three times (please indicate)

10.1.1.4 [If yes to Q10.1.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

10.1.1.5 [If yes to Q10.1.1] Please describe how your identity details were misused (e.g. issuing of a loan, credit card)

[Open ended]

10.1.1.6 [If yes to Q10.1.1] If you know or suspect who is responsible for this identity abuse/theft please tell us more about that entity (e.g., name, relationship to you if any).

[Open ended]

10.1.1.7 [If yes to Q10.1.1] Please describe the adverse impact of this identity abuse/theft on you. For example, would you rate the impact as minor, major, or severe?

[Open ended]

10.1.2 [If no to Q10.1.1] Could the identity details have been obtained from another public directory or Internet source (for example, your email signature, a workplace directory, Facebook)?

[Yes/No]

10.1.2.1 [If no to Q10.1.2] How do you think identity details were obtained?

[Open ended]

11. Are there any Internet servers (web, email, etc.) now reachable using the domain name that you registered?

[YES/NO]

11.1 [If yes to Q11] Are you the system administrator of these servers? That is, do you own and operate the computer on which the server runs? (If your servers are hosted by a web or email services provider, the answer to this question should be NO. If you're not sure about the answer, chances are good it should be NO.)

[YES/NO]

11.1.1 [If yes to Q11.1] Since registering this domain name, have you ever experienced unauthorized intrusion into servers within this domain for which you have administrative rights?

[YES/NO]

11.1.1.1 [If yes to Q11.1.1] Do you have reason to suspect that the unauthorized intrusion(s) happened because your identity details are published in WHOIS?

[YES/NO]

11.1.1.1.1 [If yes to Q11.1.1.1] Please specify why you think so.

[Open ended]

11.1.1.1.2 [If yes to Q10.1.1.1] Are the misused identity details published in another public directory or Internet source (for example, your email signature, a workplace directory, Facebook)?

[Yes/No]

11.1.1.1.3 [If yes to Q11.1.1.1] How many times have you observed intrusions into your server(s) that you can relate to your identity details published in WHOIS?

- Once
- Twice
- Three times
- More than three times (please indicate)

11.1.1.1.4 [If yes to Q11.1.1.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

11.1.1.1.5 [If yes to Q11.1.1.1] Please describe the adverse effect and severity of the unauthorized intrusion (e.g. web site defacement)

[Open ended]

11.1.1.1.6 [If yes to Q11.1.1.1] If you know or suspect who was behind a recent intrusion, please tell us more about that entity (e.g., source IP address or domain name).

[Open ended]

11.1.2 [If yes to Q11.1] Have any of the servers in your domain(s) been a victim of denial of service (DoS) attack? (If unsure, the answer should be NO.)

[YES/NO]

11.1.2.1 [If yes to Q11.1.2] Do you think the DoS attack happened because your identity details are published in WHOIS?

[YES/NO]

11.1.2.1.1 [If yes to Q11.1.2.1] Why do you think so?

[Open ended]

11.1.2.1.2 [If yes to Q11.1.2.1] Are the misused identity details published in another public directory or Internet source (for example, your email signature, a workplace directory, Facebook)?

[Yes/No]

11.1.2.1.3 [If yes to Q11.1.2.1] How many times have you have you experienced a DoS attack against one or more of the servers within this domain that you attribute to WHOIS misuse?

- Once
- Twice
- Three times
- More than three times (please indicate)

11.1.2.1.4 [If yes to Q11.1.2.1] When is the last time that you experienced this?

- Within this week

- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

11.1.2.1.5 [If yes to Q11.1.2.1] Please describe the adverse impact of the attack (e.g. unable to provide services to customers, etc)

[Open ended]

11.1.2.1.6 [If yes to Q11.1.2.1] If you are know or suspect who was behind a recent attack, please tell us more about that entity (e.g., caller's name, type of company)

[Open ended]

12. Since registering this domain name, have you ever been a victim of blackmail or intimidation?

[YES/NO]

12.1 [If yes to Q12] Was the identity (e.g., name, address, phone number, etc) that was the target of blackmail or intimidation specified in contact details during domain registration?

[Yes/No]

12.1.1 [If yes to Q12.1] Do you have reason to suspect that the blackmail or intimidation was related to the fact that your identity details are published in WHOIS?

[YES/NO]

12.1.1.1 [If yes to Q12.1.1] Please specify why you think so.

[Open ended]

12.1.1.2 [If yes to Q12.1.1] Are the misused identity details published in another public directory or Internet source (for example, email signature, workplace directory, Facebook)?

[Yes/No]

12.1.1.3 [If yes to Q12.1.1] How many times have you have you been blackmailed or intimidated using your identity details published in WHOIS?

- Once
- Twice
- Three times
- More than three times (please indicate)

12.1.1.4 [If yes to Q12.1.1] When was the last time that you experienced this?

- Within this week
- Within this month
- Within the past three months
- Within this year
- More than a year ago (please specify)

12.1.1.5 [If yes to Q12.1.1] Please describe a recent incident (e.g., how you got blackmailed or intimidated).

[Open ended]

12.1.1.6 [If yes to Q12.1.1] If you know or suspect who was behind a recent incident, please tell us more about that entity (e.g., name, relationship to you if any)

[Open ended]

12.1.1.7 [If yes to Q12.1.1] Please describe the adverse impact this incident had on you. For example, would you rate the incident's impact as minor, major, or severe?

[Open ended]

13. Have you received any other type of harmful Internet communication or experienced any other harmful acts that you have reason to believe may represent WHOIS data misuse?

[Yes/No]

13.1 [If yes to 13] Please tell us what you experienced, why you believe WHOIS contact details for this domain name might have played a role, and whether the contact details misused in this incident were available from any other source.

[Open ended]

14. If you believe that the information you used for domain name registration has been misused in any way, and you have indicated this in any one of the previous questions, did you subsequently take any measures to avoid WHOIS misuse in the future?

[I have experienced misuse and taken measures/I have experienced misuse and not taken measures/I have not experienced misuse]

14.1 [If yes to Q14] Please tell us about the measures that you took. Check all steps that you tried and explain any additional strategies you tried that are not listed below:

- Cancelling your domain name's registration or moving it to a different Registrar.
- Changing your email address or domain name or any other misused WHOIS data.
- Replacing your own WHOIS contact addresses with forwarding addresses supplied by a service provider (such as your domain's Registrar).
- Replacing your WHOIS contact names and addresses with the names and addresses of a service provider (for example, someone registering the domain name on your behalf).
- Supplying partially incorrect or incomplete information when re-registering the domain name or updating its WHOIS contact details (e.g., using a fake street number with everything else valid)
- Supplying completely fake information when re-registering the domain name or updating its WHOIS contact details.
- Applying a spam filter or registering with an identity theft protection service or some other step to deal with the consequences of WHOIS misuse (as opposed to reducing misuse itself).

- Other (please describe)

[Important note: As previously stated, your individual answer to this question is completely confidential and will NOT be shared with your Registrar or ICANN.]

15. Are you aware of any strategies that your domain name's current Registrar may be taking to reduce or protect against WHOIS data misuse?

[YES/NO]

15.1 [If yes to Q15] Please describe: [open ended field]

16. Do you grant us permission to contact you further in case we need clarifications about your answers to this survey?

[YES/NO]

16.1 [If yes to Q16] If yes, please enter your email here.

[Open ended]

B.4 Definitions of terms

The following are the descriptions for the technical terms provided as part of the registrant survey.

Identity theft.

Identity theft occurs when someone uses your personally identifiable information, like your name, address, phone number, Social Security number (or national identification number), or credit card number, without your permission, to commit fraud or other crimes. Some examples of identity theft include renting an apartment, obtaining a credit card, or establishing a telephone account in your name, without your permission.

Identity thieves steal information by going through trash looking for bills or other paper with your personally identifiable information, soliciting your information by sending emails pretending to be your bank (see also Phishing), calling your financial institution while pretending to be you, etc. Thieves may also be able to get some personally identifiable information by searching WHOIS for domain name contact names and addresses.

Blackmail.

In common usage, blackmail is a crime involving threats to reveal substantially true and/or false information about a person to the public, a family member, or associates unless a demand is met. Blackmail can include coercion involving threats of physical harm, criminal prosecution, or taking the person's money or property. In the context of WHOIS misuse, blackmailers may use some personally identifiable information by searching WHOIS for domain name contact names and addresses.

Email spam.

Spam email is an unsolicited mail message, sent to your email address without your permission. The sender of spam is commonly called a "spammer" Spammers send the same email to a large number of email addresses. They may obtain email addresses from many different sources such as websites and chat forums. It is also possible for spammers to search WHOIS for domain name contact email addresses.

Spam email is often used to advertise (or sell) legal and illegal products and to attempt to steal sensitive information like credit card numbers (see also Phishing). Products commonly advertised by spam include prescription drugs, herbal medications, replica watches, online gambling and pornography.

Postal spam.

Postal spam is unsolicited postal mail sent to a residential or commercial postal mailbox or another postal address, and is similar in concept to email spam (see Email Spam). Postal spammers may obtain postal addresses from many different sources, both offline and on-line, including searching WHOIS for domain name contact postal addresses.

Phishing.

Phishing attacks attempt to steal your personally identifiable information (see also Identity Theft) and financial account information. A common tactic used during phishing attacks is sending spam emails that contain links to counterfeit websites (see also Email Spam). Phishing emails may contain details about recipients, obtained from many different sources, including searching WHOIS for domain name contact names, addresses and phone numbers.

The attacker can use techniques to hide the identity of the phishing message's true sender and make the email look like someone else sent it. For example, a phishing email may appear to come from a legitimate bank, but when you click on

the link, you may be taken to a website designed to look like the bank's website. This may trick you into divulging sensitive data such as banking or other website account usernames and passwords.

Alternatively, when you click on a phishing email link, you may be taken to a website that attempts to automatically install malicious software on your computer without your permission or knowledge. For example, a key-logger program may be installed to send everything that you type (e.g., passwords) to a remote attacker.

Vishing.

Vishing attacks attempt to steal your personally identifiable information (see also Identity Theft) and financial account information. Vishing attacks are similar to phishing attacks (see Phishing), but are conducted using voice or telephone calls instead of email messages. The attacker can use techniques to hide the vishing caller's true caller identification number and make the caller's number appear to be another party's number. Vishing attack victims may be tricked into revealing sensitive information.

For example, the attacker may call you, claiming to be a representative of a bank, and request your banking information for administrative purposes. Alternatively, upon receiving a vishing call, you may hear an automated voice message requesting you to immediately call a specified number to verify account details. But that number reaches the attacker, not your bank.

Email virus.

The most generic definition of an email virus is malicious software (also called "malware") delivered as an email file attachment. When the recipient opens the attached file, the malicious software is installed or otherwise activated. The malicious software may damage data or services on the recipient's computer. It may also carry out harmful actions on behalf of the attacker. Common examples in-

clude deleting files, sending spam emails (see Email Spam) on the attacker's behalf, tracking the user's actions, and downloading and installing additional malicious software. Mail messages that carry viruses may be sent to email addresses obtained from many different sources, including searching WHOIS for domain name contact addresses.

Denial of Service (DoS).

In a denial-of-service attack, an attacker attempts to prevent legitimate users from accessing or making use of information or services. By targeting your computer and its network connection, or the computers and network of Internet sites that you are trying to use, an attacker may be able to prevent you from accessing email, websites, online service provider accounts (banking sites, etc.), or other services that rely on the computers or networks that are under DoS attack.

Not all disruptions to service are the result of a DoS attack. There may be technical problems with a particular network, or system administrators may be performing maintenance. However, the following symptoms could indicate a DoS attack: (i) unusually slow performance when opening files or accessing websites, (ii) unavailability of a particular website, (iii) inability to access any website, or (iv) a dramatic increase in the amount of spam that you receive.

DoS attacks may be launched against targets identified in many different sources, including searching WHOIS for domain name contact names and addresses.

Unauthorized intrusion.

Unauthorized intrusion occurs when an attacker gains access to services or information on a computer system without the owner's permission. It is also possible that the attacker is a legitimate user of the computer system, but has managed to gain access to an access level higher than she is authorized to access.

Unauthorized intrusion can happen in many ways. Some common techniques used by intruders are sending malicious messages to the targets computer through the network, tricking the administrator of the computer system in to installing malicious software (see also Phishing), and guessing the administrator's account user-name and password. Unauthorized intrusions may be launched against targets identified in many different sources, including searching WHOIS for domain name contact names and addresses.

B.4.1 Document information

This document was prepared to help users completing surveys being conducted by computer security researchers at Carnegie Mellon University - Cylab. This document is for research and education purposes only, and is not for commercial or business purposes. Anyone can use this document in part or whole by citing all the sources cited in this document, and adhering to the terms of use specified by the sources cited in this document. All queries regarding this document should be directed to whois-study@cmu.edu.

B.4.2 Acknowledgment of sources

All sources used to create this document are specified below. Some sentences have been quoted verbatim or with slight modifications to assist readers with limited knowledge of computer terminology. Further, certain references to United States specific terminology (e.g., Social Security Number) have been reduced as this document is intended for use by an international audience.

Identity Theft <http://www.ftc.gov/bcp/edu/microsites/idtheft/consumers/about-identity-theft.html#Whatisidentitytheft>

Denial of Service <http://www.us-cert.gov/cas/tips/ST04-015.html>

Phishing <http://www.icann.org/en/general/glossary.htm#P>

Blackmail <http://en.wikipedia.org/wiki/Blackmail>

Email spam <http://www.spamhaus.org/definition.html>

Email Viruses <http://www.mysecurecyberspace.com/encyclopedia/index/intrusion.html>

Phishing <http://www.ftc.gov/bcp/edu/pubs/consumer/alerts/alt127.shtm>

Vishing http://www.fbi.gov/news/stories/2010/november/cyber_112410

Unauthorized intrusion <http://www.mysecurecyberspace.com/encyclopedia/index/intrusion.html>

Bibliography

- [1] “Adblock easy list,” <https://easylist-downloads.adblockplus.org/easylist.txt>, Adblock, last accessed August 18, 2014.
- [2] “Adify vertical gauge shows steady growth in seven of eleven critical verticals,” Adify, <http://www.smartbrief.com/news/aaaa/industryMW-detail.jsp?id=732F69A7-9192-4E05-A261-52C068021634>. Last accessed May 5, 2011.
- [3] S. Afroz, V. Garg, D. McCoy, and R. Greenstadt, “Honor Among Thieves: A Common’s Analysis of Cybercrime Economies,” in *eCrime Research summit*. San Francisco, CA: IEEE, 2013.
- [4] A. Aizcorbe and N. Nestoriak, “Price indexes for drugs: A review of the issues,” Bureau of Economic Analysis, BEA Working Papers, 2010. [Online]. Available: <http://econpapers.repec.org/paper/beawpaper/0050.htm>
- [5] “India’s top court dismisses drug patent case,” <http://www.aljazeera.com/news/asia/2013/04/2013412275825670.html>, Al Jazeera, last accessed August 18, 2014.
- [6] “Alexa Web Information Service,” <http://aws.amazon.com/awis/>, Amazon Web Services.
- [7] American Medical Association, “Illicit online pharmacies resort to hacking to gain customers,” <http://www.amednews.com/article/20110905/business/309059964/7/pdf>, Sep. 2011, last accessed August 18, 2014.
- [8] D. Anderson, C. Fleizach, S. Savage, and G. Voelker, “Spamscatter: Characterizing internet scam hosting infrastructure,” in *Proceedings of 16th USENIX Security Symposium*. USENIX Association, 2007, pp. 1–14.
- [9] R. Anderson, C. Barton, R. Böhme, R. Clayton, M. J. van Eeten, M. Levi, T. Moore, and S. Savage, “Measuring the cost of cybercrime,” in *The Economics of Information Security and Privacy*. Springer, 2012, pp. 265–300.

- [10] Anti-Phishing Working Group, "Phishing attack trends report - Q2 2010," January 2010.
- [11] "Open Directory project," <http://www.dmoz.org/>, AOL Inc.
- [12] "SpamAssassin," <http://spamassassin.apache.org/>, The Apache Software Foundation.
- [13] H. L. Armstrong and P. J. Forde, "Internet anonymity practices in computer crime," *Information management & computer security*, vol. 11, no. 5, pp. 209–215, 2003.
- [14] A. Austin and L. Williams, "One technique is not enough: A comparison of vulnerability discovery techniques," in *International Symposium on Empirical Software Engineering and Measurement (ESEM)*. IEEE, 2011, pp. 97–106.
- [15] A. Baveja, R. Batta, J. P. Caulkins, and M. H. Karwan, "Modeling the response of illicit drug markets to local enforcement," *Socio-Economic Planning Sciences*, vol. 27, no. 2, pp. 73–89, 1993.
- [16] C. Beccaria, "Dei delitti e delle pene (On crimes and punishments)," *Il Caffè*, 1764.
- [17] G. S. Becker, "Crime and punishment: An economic approach," in *Essays in the Economics of Crime and Punishment*. UMI, 1974, pp. 1–54.
- [18] T. L. Bessell, J. N. Anderson, C. a. Silagy, L. N. Sansom, and J. E. Hiller, "Surfing, self-medicating and safety: buying non-prescription and complementary medicines via the internet." *Quality & safety in health care*, vol. 12, no. 2, pp. 88–92, Apr. 2003. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1743681/>
- [19] T. L. Bessell, C. A. Silagy, J. N. Anderson, J. E. Hiller, and L. N. Sansom, "Quality of global e-pharmacies: can we safeguard consumers?" *European journal of clinical pharmacology*, vol. 58, no. 9, pp. 567–72, Dec. 2002. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/12483449>
- [20] "Black market reloaded," <http://5onwnspjvuk7cwvk.onion/>. Last accessed November 25, 2012.
- [21] D. Blacka, B. Laurie, G. Sisson, and R. Arends, "RFC 5155: DNS security (DNSSEC) hashed authenticated denial of existence," *Request for Comments*, vol. 5155, 2008.

- [22] K. Borgolte, C. Kruegel, and G. Vigna, "Delta: automatic identification of unknown web-based infection campaigns," in *Proceedings of ACM CCS 2013*, Berlin, Germany, Nov. 2013, pp. 109–120.
- [23] S. Brown, P. Elkin, S. Rosenbloom, C. Husser, B. Bauer, M. Lincoln, J. Carter, M. Erlbaum, and M. Tuttle, "VA national drug file reference terminology: a cross-institutional content coverage study," *Medinfo*, vol. 11, no. Pt 1, pp. 477–81, 2004.
- [24] J. M. Buchanan and W. C. Stubblebine, "Externality," *Economica*, pp. 371–384, 1962.
- [25] L. Camp and S. Lewis, *Economics of Information Security*, ser. Advances in Information Security. Springer, 2004.
- [26] D. Canali, D. Balzarotti, and A. Francillon, "The role of web hosting providers in detecting compromised websites," in *Proceedings of the 22nd international conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2013, pp. 177–188.
- [27] J. Castronova, "Operation Cyber Chase and other agency efforts to control internet drug trafficking the "virtual" enforcement initiative is virtually useless," *The Journal of Legal Medicine*, pp. 1–17, 2006.
- [28] J. P. Caulkins, "The distribution and consumption of illicit drugs: some mathematical models and their policy implications," Ph.D. dissertation, Massachusetts Institute of Technology, 1990.
- [29] N. Chachra, D. Savage, and G. Voelker, "Empirically characterizing domain abuse and the revenue impact of blacklisting," in *Workshop on the Economics of Information Security*, 2014.
- [30] Y.-N. Chiu, B. Leclerc, and M. Townsley, "Crime script analysis of drug manufacturing in clandestine laboratories implications for prevention," *British journal of criminology*, vol. 51, no. 2, pp. 355–374, 2011.
- [31] N. Christin, S. Egelman, T. Vidas, and J. Grossklags, "It's all about the Benjamins: Incentivizing users to ignore security advice," in *Proceedings of IFCA Financial Cryptography'11*, Saint Lucia, Mar. 2011.
- [32] N. 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. International World Wide Web Conferences Steering Committee, 2013, pp. 213–224.

- [33] N. Christin, S. S. Yanagihara, and K. Kamataki, “Dissecting one click frauds,” in *Proceedings of the 17th ACM conference on Computer and communications security*, ser. CCS ’10. Chicago, Illinois, USA: ACM, 2010, pp. 15–26.
- [34] R. V. Clarke, “Situational crime prevention: Its theoretical basis and practical scope,” *Crime and justice*, pp. 225–256, 1983.
- [35] —, *Situational crime prevention: Successful case studies*. Harrow and Heston (Guilderland, NY), 1997.
- [36] R. V. Clarke and D. B. Cornish, “Modeling offenders’ decisions: A framework for research and policy,” *Crime and justice*, pp. 147–185, 1985.
- [37] R. V. Clarke and J. E. Eck, “Crime analysis for problem solvers,” *Washington, DC: Center for Problem Oriented Policing*, 2005.
- [38] R. Clayton, “How much did shutting down McColo help?” in *Proceedings of the Sixth Conference on Email and Antispam (CEAS)*, Jul. 2009.
- [39] R. Clayton and T. Mansfield, “A study of Whois privacy and proxy service abuse,” in *Proceedings of the 13th Workshop on Economics of Information Security*, State College, PA, Jun. 2014.
- [40] “February 2014 US search engine rankings,” https://www.comscore.com/Insights/Press_Releases/2014/3/comScore_Releases_February_2014_U.S._Search_Engine_Rankings, comScore Inc., 2014, last accessed August 18, 2014.
- [41] “Top 20 ad networks,” http://www.comscore.com/Insights/Market_Rankings/comScore_Media_Metrix_Ranks_Top_50_US_Desktop_Web_Properties_for_April_2014, comScore Inc., Apr. 2014, last accessed August 18, 2014.
- [42] G. Cooper and E. Herskovits, “A Bayesian method for the induction of probabilistic networks from data,” *Machine learning*, vol. 9, no. 4, pp. 309–347, 1992.
- [43] D. B. Cornish, “The procedural analysis of offending and its relevance for situational prevention,” *Crime prevention studies*, vol. 3, pp. 151–196, 1994.

- [44] D. B. Cornish and R. V. Clarke, "Understanding crime displacement: An application of rational choice theory," *Criminology*, vol. 25, no. 4, pp. 933–948, 1987.
- [45] —, "Opportunities, precipitators and criminal decisions: A reply to Wortley's critique of situational crime prevention," *Crime prevention studies*, vol. 16, pp. 41–96, 2003.
- [46] A. Costin, J. Isacenkova, M. Balduzzi, A. Francillon, and D. Balzarotti, "The role of phone numbers in understanding cyber-crime schemes," Technical Report, EURECOM, RR-13-277, Tech. Rep., 2013.
- [47] Court of Appeals, 8th Circuit, "US v. Birbragher," in *F. 3d*, vol. 603, 2010, p. 478, no. 08-4004.
- [48] D. T. Courtwright, "The controlled substances act: how a "big tent" reform became a punitive drug law," *Drug and Alcohol Dependence*, vol. 76, no. 1, pp. 9 – 15, 2004.
- [49] M. Cova, C. Leita, O. Thonnard, A. Keromytis, and M. Dacier, "An analysis of rogue AV campaigns," in *Proc. RAID 2010*, Ottawa, ON, Canada, Sep. 2010.
- [50] D. R. Cox, "Regression models and life-tables," *Journal of the Royal Statistics Society, Series B*, vol. 34, pp. 187–220, 1972.
- [51] "The center for safe internet pharmacies," <http://www.safemedsonline.org/>, CSIP, last accessed August 18, 2014.
- [52] C. J. Dahlman, "The problem of externality," *Journal of law and economics*, pp. 141–162, 1979.
- [53] L. Daigle, "Rfc 3912: Whois protocol specification," *Request for Comments*, vol. 3912, 2004.
- [54] J. Damron, "Identifiable fingerprints in network applications," *USENIX; login*, vol. 28, no. 6, pp. 16–20, 2003.
- [55] V. Dave, S. Guha, and Y. 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*. ACM, 2012, pp. 175–186.

- [56] —, “ViceROI: catching click-spam in search ad networks,” in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 765–776.
- [57] G. Del Pino, “The unifying role of iterative generalized least squares in statistical algorithms,” *Statistical Science*, vol. 4, no. 4, pp. pp. 394–403, 1989.
- [58] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” in *Proceedings of the 13th USENIX Security Symposium*, Aug. 2004.
- [59] J. E. Dunn, “Srizbi grows into world’s largest botnet,” *CSO*, May 2008, <http://www.csoonline.com/article/356219/srizbi-grows-into-world-s-largest-botnet>.
- [60] E. Eckholm, “Abuses are found in online sales of medication,” *New York Times*. *July*, vol. 9, 2008.
- [61] B. Edelman, “Deterring online advertising fraud through optimal payment in arrears,” *Financial Cryptography and Data Security*, vol. 5628, p. 17, 2009.
- [62] B. Edwards, T. Moore, G. Stelle, S. Hofmeyr, and S. Forrest, “Beyond the blacklist: modeling malware spread and the effect of interventions,” in *Proceedings of the 2012 workshop on New security paradigms*, 2012, pp. 53–65.
- [63] K. Elliott, “The who, what, where, when, and why of WHOIS: Privacy and accuracy concerns of the WHOIS database,” *SMU Science & Technology Law Review*, vol. 12, p. 141, 2008.
- [64] “Experian hitwise reports bing-powered share of searches reaches 30 percent in march 2011,” Experian Hitwise, Apr. 2011, <http://www.hitwise.com/us/press-center/press-releases/experian-hitwise-reports-bing-powered-share-of-s/>.
- [65] Expert Working Group on gTLD Directory Services, “A next generation registration directory service,” <https://www.icann.org/en/groups/other/gtld-directory-services/initial-report-24jun13-en.pdf>, ICANN, 2013, last accessed August 18, 2014.
- [66] M. Felson and R. V. Clarke, “Opportunity makes the thief,” *Police research series, paper*, vol. 98, 1998.

- [67] I. Fette, “Understanding phishing and malware protection in google chrome,” <http://blog.chromium.org/2008/11/understanding-phishing-and-malware.html>, Google Inc., Nov. 2008, last accessed August 18, 2014.
- [68] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee, “RFC 2616: Hypertext Transfer Protocol–HTTP/1.1,” *Request for Comments*, vol. 2616, 1999.
- [69] M. Finifter, D. Akhawe, D. Wagner, J. Mickens, and M. Finifter, “An empirical study of vulnerability rewards programs,” in *Proceedings of the 22nd USENIX Security Symposium*. USENIX Association, 2013, pp. 13–25.
- [70] M. Finifter and D. Wagner, “Exploring the relationship between web application development tools and security,” in *2nd USENIX Conference on Web Application Development*, 2011.
- [71] O. Fisher, “Malware? We don’t need no stinking malware!” <http://googleonlinesecurity.blogspot.com/2008/10/malware-we-dont-need-no-stinking.html>, Google Inc., Oct. 2008, last accessed August 18, 2014.
- [72] D. Florêncio and C. Herley, “Where do all the attacks go?” in *Economics of Information Security and Privacy III*. Springer, 2013, pp. 13–33.
- [73] J. Franklin, A. Perrig, V. Paxson, and S. Savage, “An inquiry into the nature and causes of the wealth of internet miscreants.” in *ACM conference on Computer and communications security*, 2007, pp. 375–388.
- [74] K. Fung, C. McDonald, and B. Bray, “Rxterms—a drug interface terminology derived from rxnorm,” in *AMIA Annual Symposium Proceedings*, vol. 2008. American Medical Informatics Association, 2008, p. 227.
- [75] D. Garland, “Ideas, institutions and situational crime prevention,” *Ethical and social perspectives on situational crime prevention*, 2000.
- [76] U. Gasser, “Regulating search engines: Taking stock and looking ahead,” *Yale Journal of Law and Technology*, vol. 8, no. 1, p. 7, 2006.
- [77] U. Gelatti, R. Pedrazzani, C. Marcantoni, S. Mascaretti, C. Repice, L. Filipucci, I. Zerbini, M. Dal Grande, G. Orizio, and D. Feretti, “‘You’ve got m@il: Fluoxetine coming soon!’: Accessibility and quality of a prescription drug sold on the web,” *International Journal of Drug Policy*, vol. 24, no. 5, pp. 392–401, 2013.

- [78] Generic Names Supporting Organization, "Motion to pursue WHOIS studies," <http://gnso.icann.org/en/council/resolutions#20100908-3>, ICANN, 2010, last accessed August 18, 2014.
- [79] I. P. Gomez, "Beyond the neighborhood drugstore: US regulation of online prescription drug sales by foreign businesses," *Rutgers Computer & Technology Law Journal*, vol. 28, p. 431, 2002.
- [80] "Google insights for search," <http://www.google.com/insights/search/>, Google Inc.
- [81] "Google traffic estimator," <https://adwords.google.com/select/TrafficEstimatorSandbox>, Google Inc.
- [82] "Google trends," <http://www.google.com/trends/>, Google Inc., last accessed August 18, 2014.
- [83] "Google Web Search API," <https://code.google.com/apis/websearch/>, Google Inc., last accessed August 18, 2014.
- [84] "Making the web safer," <http://www.google.com/transparencyreport/safebrowsing/>, Google Inc., last accessed August 18, 2014.
- [85] "Quality guidelines: Cloaking," <https://support.google.com/webmasters/answer/66355>., Google Inc., last accessed August 18, 2014.
- [86] C. Grier, K. Thomas, V. Paxson, and M. Zhang, "@spam: the underground on 140 characters or less," in *Proceedings of the 17th ACM conference on Computer and communications security*. ACM, 2010, pp. 27–37.
- [87] O. H. Griffin, "Is the government keeping the peace or acting like our parents? rationales for the legal prohibitions of ghb and mdma," *Journal of Drug Issues*, vol. 42, no. 3, pp. 247–262, 2012.
- [88] Z. Gyöngyi and H. Garcia-Molina, "Link spam alliances," in *Proceedings of the 31st international conference on Very large data bases*. VLDB Endowment, 2005, pp. 517–528.
- [89] J. Hadar and W. R. Russell, "Rules for ordering uncertain prospects," *American economic review*, vol. 59, no. 1, pp. 25–34, 1969.
- [90] G. Hardin, "The tragedy of the commons," *Science*, vol. 162, no. 3859, pp. 1243–1248, 1968.

- [91] J. E. Henney, "Cyberpharmacies and the role of the US Food And Drug Administration." *Journal of medical Internet research*, vol. 3, no. 1, p. E3, 2001. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1761882/>
- [92] J. E. Henney, J. E. Shuren, S. L. Nightingale, and T. J. McGinnis, "Internet purchase of prescription drugs: buyer beware," *Annals of internal medicine*, vol. 131, no. 11, pp. 861–862, 1999.
- [93] C. Herley, "So long, and no thanks for the externalities: the rational rejection of security advice by users," in *Proceedings of the 2009 workshop on New security paradigms workshop*. ACM, 2009, pp. 133–144.
- [94] C. Herley and D. Florêncio, "Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy," in *Economics of Information Security and Privacy*. Springer, 2010, pp. 33–53.
- [95] R. Hesseling, "Displacement: A review of the empirical literature," *Crime prevention studies*, vol. 3, pp. 197–230, 1994.
- [96] K. J. Higgins, "Google, godaddy help form group to fight fake online pharmacies," *Dark Reading*, Dec. 2010, <http://www.darkreading.com/security/privacy/228800671/google-godaddy-help-form-group-to-fight-fake-online-pharmacies.html>.
- [97] D. L. Hoffman and T. P. Novak, "How to acquire customers on the web," *Harvard business review*, vol. 78, no. 3, pp. 179–188, 2000.
- [98] S. Hollenbeck, K. Ranjbar, A. Servin, A. Newton, N. Kong, S. Sheng, B. El-lacott, F. Obispo, and F. Arias, "Using HTTP for RESTful WHOIS services by internet registries," *Internet-Draft*, 2012.
- [99] D. W. Hosmer Jr and S. Lemeshow, *Applied logistic regression*. John Wiley & Sons, 2004.
- [100] "gTLD-specific monthly registry reports," [http://www.icann.org/sites/default/files/mrr/\[gTLD\]/\[gTLD\]-transactions-201102-en.csv](http://www.icann.org/sites/default/files/mrr/[gTLD]/[gTLD]-transactions-201102-en.csv), ICANN, Feb. 2011, last accessed August 18, 2014.
- [101] "2013 Registrar Accreditation Agreement," <http://www.icann.org/en/resources/registrars/raa/approved-with-specs-27jun13-en.htm>., ICANN, 2013, last accessed August 18, 2014.

- [102] ICANN, "ICANN-accredited registrars," <http://www.icann.org/registrar-reports/accredited-list.html>, 2014, last accessed August 18, 2014.
- [103] ICANN. Security and Stability Advisory Committee, "Is the WHOIS service a source for email addresses for spammers?" <http://www.icann.org/en/committees/security/sac023.pdf>, 2007, last accessed August 18, 2014.
- [104] —, "Advisory on registrar impersonation phishing attacks," <http://www.icann.org/en/committees/security/sac028.pdf>, 2008, last accessed August 18, 2014.
- [105] International Telecommunications Union, "Percentage of individuals using the internet 2000-2012," http://www.itu.int/en/ITU-D/Statistics/Documents/statistics/2013/Individuals_Internet_2000-2012.xls, 2013, last accessed August 18, 2014.
- [106] Interpol, "Operation Cobra," <http://www.interpol.int/Crime-areas/Pharmaceutical-crime/Operations/Operation-Pangea>, last accessed August 18, 2014.
- [107] —, "Operation Mamba," <http://www.interpol.int/Crime-areas/Pharmaceutical-crime/Operations/Operation-Mamba>, last accessed August 18, 2014.
- [108] —, "Operation Pangea VII," <http://www.interpol.int/News-and-media/News/2014/N2014-089>, last accessed August 18, 2014.
- [109] —, "Operation Storm," <http://www.interpol.int/Crime-areas/Pharmaceutical-crime/Operations/Operation-Storm>, last accessed August 18, 2014.
- [110] —, "Operations Pangea (I-VII)," <http://www.interpol.int/Crime-areas/Pharmaceutical-crime/Operations/Operation-Pangea>, last accessed August 18, 2014.
- [111] L. Ivanitskaya, J. Brookins-Fisher, I. O Boyle, D. Vibbert, D. Erofeev, and L. Fulton, "Dirt cheap and without prescription: how susceptible are young US consumers to purchasing drugs from rogue internet pharmacies?" *Journal of medical Internet research*, vol. 12, no. 2, p. e11, Jan. 2010. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2885783/>
- [112] J. Jang, D. Brumley, and S. Venkataraman, "Bitshred: feature hashing malware for scalable triage and semantic analysis," in *Proceedings of the 18th*

ACM conference on Computer and communications security. ACM, 2011, pp. 309–320.

- [113] K. E. Jerian, “What’s a legal system to do-the problem of regulating internet pharmacies,” *Albany Law Journal of Science & Technology*, vol. 16, p. 571, 2006.
- [114] T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay, “Accurately interpreting clickthrough data as implicit feedback,” in *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval.* ACM, 2005, pp. 154–161.
- [115] J. P. John, F. Yu, Y. Xie, A. Krishnamurthy, and M. Abadi, “deSEO: combating search-result poisoning,” in *Proceedings of the 20th USENIX conference on Security.* USENIX Association, 2011, pp. 20–20.
- [116] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage, “Spamalytics: An empirical analysis of spam marketing conversion,” in *Proceedings of the 15th ACM conference on Computer and communications security.* ACM, 2008, pp. 3–14.
- [117] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. M. Voelker, and S. Savage, “Show me the money: Characterizing spam-advertised revenue,” in *Proceedings of the 20st USENIX conference on Security symposium.* USENIX Association, 2011.
- [118] E. Kao, “Making search more secure,” <http://googleblog.blogspot.com/2011/10/making-search-more-secure.html>., Google Inc., Oct. 2011.
- [119] E. Kaplan and P. Meier, “Nonparametric estimation from incomplete observations,” *Journal of the American Statistical Association*, vol. 53, pp. 457–481, 1958.
- [120] R. Kohavi, “A study of cross-validation and bootstrap for accuracy estimation and model selection,” in *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, vol. 14, no. 2, 1995, pp. 1137–1145.
- [121] J. Lacoste and P. Tremblay, “Crime and innovation: A script analysis of patterns in check forgery,” *Crime Prevention Studies*, vol. 16, pp. 169–196, 2003.
- [122] M. Lee, “Who’s next? identifying risks factors for subjects of targeted attacks,” in *Proceedings of the Virus Bulletin Conference*, 2012, pp. 301–306.

- [123] “Setting the record straight,” <http://www.legitscript.com/about/setting-the-record-straight>, LegitScript, last accessed August 18, 2014.
- [124] “The Leading Source of Internet Pharmacy Verification,” <http://www.legitscript.com/>, LegitScript.
- [125] “Yahoo! Internet pharmacy advertisements,” <http://www.legitscript.com/download/YahooRxAnalysis.pdf>, LegitScript and Knujon, 2009, last accessed August 18, 2014.
- [126] “Rogues and registrars,” <http://www.legitscript.com/download/Rogues-and-Registrars-Report.pdf>, LegitScript and Knujon, 2010, last accessed August 18, 2014.
- [127] N. Leontiadis and N. Christin, “Empirically measuring WHOIS misuse,” in *Proceedings of the 19th European Symposium on Research in Computer Security*. Wroclaw, Poland: Springer, Sep. 2014.
- [128] N. Leontiadis, T. Moore, and N. Christin, “Measuring and analyzing search-redirect attacks in the illicit online prescription drug trade,” in *Proceedings of the 20th USENIX Security Symposium*, San Francisco, CA, Aug. 2011, pp. 281–298.
- [129] —, “Pick your poison: pricing and inventories at unlicensed online pharmacies,” in *Proceedings of the 14th ACM conference on Electronic Commerce*, ser. EC ’13. Philadelphia, Pennsylvania, USA: ACM, Jun. 2013, pp. 621–638.
- [130] —, “A nearly four-year longitudinal study of search-engine poisoning,” in *Proceedings of the 21st ACM conference on Computer and Communications Security*, ser. CCS ’14. Scottsdale, Arizona, USA: ACM, Nov. 2014.
- [131] J. Leskovec, L. Backstrom, and J. Kleinberg, “Meme-tracking and the dynamics of the news cycle,” in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2009, pp. 497–506.
- [132] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu *et al.*, “Click trajectories: End-to-end analysis of the spam value chain,” in *IEEE Symposium on Security and Privacy*. IEEE, 2011, pp. 431–446.

- [133] M. Levi and M. Maguire, "Reducing and preventing organised crime: An evidence-based critique," *Crime, Law and Social Change*, vol. 41, no. 5, pp. 397–469, 2004.
- [134] Z. Li, S. Alrwais, X. Wang, and E. Alowaisheq, "Hunting the red fox online: Understanding and detection of mass redirect-script injections," in *Proceedings of the 35th IEEE Symposium on Security and Privacy*. San Jose, CA, USA: IEEE, May 2014.
- [135] Z. Li, S. Alrwais, Y. Xie, F. Yu, M. S. Valley, and X. Wang, "Finding the linchpins of the dark web: a study on topologically dedicated hosts on malicious web infrastructures," in *IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 112–126.
- [136] B. Liang and T. Mackey, "Searching for safety: addressing search engine, website, and provider accountability for illicit online drug sales." *American journal of law & medicine*, vol. 35, no. 1, pp. 125–84, Jan. 2009. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/19534258>
- [137] Z. Lin, D. Li, B. Janamanchi, and W. Huang, "Reputation distribution and consumer-to-consumer online auction market structure: an exploratory study," *Decision Support Systems*, vol. 41, no. 2, pp. 435–448, 2006.
- [138] M. Lincoln, S. Brown, V. Nguyena, T. Cromwella, J. Carter, M. Erlbaum, and M. Tuttle, "US department of veterans affairs enterprise reference terminology strategic overview," in *Proceedings of the 11th World Congress On Medical Informatics*, ser. Medinfo, vol. 107, 2004, pp. 391–395.
- [139] C. Littlejohn, A. Baldacchino, F. Schifano, and P. Deluca, "Internet pharmacies and online prescription drug sales: a cross-sectional study," *Drugs: Education, Prevention, and Policy*, vol. 12, no. 1, pp. 75–80, 2005.
- [140] S. Liu, W. Ma, R. Moore, V. Ganesan, and S. Nelson, "RxNorm: prescription for electronic drug information exchange," *IT professional*, vol. 7, no. 5, pp. 17–23, 2005.
- [141] J. Long, *Google hacking for penetration testers*. Syngress, 2011, vol. 2.
- [142] L. Lu, R. Perdisci, and W. Lee, "Surf: detecting and measuring search poisoning," in *Proceedings of the 18th ACM conference on Computer and communications security*. ACM, 2011, pp. 467–476.

- [143] C. Lumezanu and N. Feamster, "Observing common spam in twitter and email," in *Proceedings of the 2012 ACM conference on Internet measurement conference*. ACM, 2012, pp. 461–466.
- [144] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond blacklists: learning to detect malicious web sites from suspicious URLs," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2009, pp. 1245–1254.
- [145] T. K. Mackey and B. A. Liang, "Global reach of direct-to-consumer advertising using social media for illicit online drug sales," *Journal of medical Internet research*, vol. 15, no. 5, 2013.
- [146] "Mapping the Mal Web," McAfee, 2010, http://us.mcafee.com/en-us/local/docs/Mapping_Mal_Web.pdf.
- [147] D. McCoy, H. Dharmdasani, C. Kreibich, G. Voelker, and S. Savage, "Priceless: The role of payments in abuse-advertised goods," in *Proceedings of the 19th ACM conference on Computer and communications security*. Raleigh, NC: ACM, Oct. 2012, pp. 845–856.
- [148] D. McCoy, A. Pitsillidis, G. Jordan, N. Weaver, C. Kreibich, B. Krebs, G. M. Voelker, S. Savage, and K. Levchenko, "Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs," in *Proceedings of the 21st USENIX conference on Security symposium*. USENIX Association, 2012.
- [149] "Bing trending news," <http://www.bing.com/news>, Microsoft Inc., last accessed August 18, 2014.
- [150] "Bing webmaster guidelines: Cloaking," <http://www.bing.com/webmaster/help/webmaster-guidelines-30fba23a>., Microsoft Inc., last accessed August 18, 2014.
- [151] "Microsoft, yahoo! change search landscape," Microsoft Inc., <http://www.microsoft.com/presspass/press/2009/jul09/07-29release.mspx>.
- [152] C. C. Miller, "Google is said to have broken internal rules on drug ads," *New York Times*, May 2011, article appeared in print on May 14, 2011, on page B2 of the New York edition. Available online at <http://www.nytimes.com/2011/05/14/technology/14google.html>.

- [153] P. Mockapetris, “Domain names – Implementation and specification (RFC 1035),” *Information Sciences Institute*, 1987.
- [154] N. Mohan, “The AdSense revenue share,” May 2010, <http://adsense.blogspot.com/2010/05/adsense-revenue-share.html>.
- [155] T. Moore and R. Clayton, “Examining the impact of website take-down on phishing,” in *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*. ACM, 2007, pp. 1–13.
- [156] —, “The consequence of non-cooperation in the fight against phishing,” in *eCrime Researchers Summit*. IEEE, 2008, pp. 1–14.
- [157] —, “Evil searching: Compromise and recompromise of internet hosts for phishing,” in *Financial Cryptography and Data Security*. Springer, 2009, pp. 256–272.
- [158] T. Moore, R. Clayton, and R. Anderson, “The economics of online crime,” *The Journal of Economic Perspectives*, vol. 23, no. 3, pp. 3–20, 2009.
- [159] T. Moore, R. Clayton, and H. Stern, “Temporal correlations between spam and phishing websites,” in *Proceedings of the 2nd USENIX conference on Large-scale exploits and emergent threats*. USENIX Association, 2009.
- [160] T. Moore and B. Edelman, “Measuring the perpetrators and funders of typosquatting,” in *Financial Cryptography and Data Security*. Springer, 2010, pp. 175–191.
- [161] T. Moore, J. Han, and R. Clayton, “The postmodern ponzi scheme: Empirical analysis of high-yield investment programs,” in *Financial Cryptography and Data Security*. Springer, 2012, pp. 41–56.
- [162] T. Moore, N. Leontiadis, and N. Christin, “Fashion crimes: trending-term exploitation on the web,” in *Proceedings of the 18th ACM conference on Computer and communications security*, ser. CCS ’11. Chicago, Illinois, USA: ACM, 2011, pp. 455–466.
- [163] C. Morselli and J. Roy, “Brokerage qualifications in ringing operations,” *Criminology*, vol. 46, no. 1, pp. 71–98, 2008.
- [164] J. Mueller, “Upcoming changes in Google’s HTTP referrer,” <http://googlewebmastercentral.blogspot.com/2012/03/upcoming-changes-in-googles-http.html>, Google Inc., Mar. 2012.

- [165] Y. Nadji, M. Antonakakis, R. Perdisci, D. Dagon, and W. Lee, "Beheading hydras: performing effective botnet takedowns," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 121–132.
- [166] Y. Nadji, M. Antonakakis, R. Perdisci, and W. Lee, "Connected colors: Unveiling the structure of criminal networks," in *Research in Attacks, Intrusions and Defenses*, 2013.
- [167] D. S. Nagin, "Deterrence in the twenty-first century," *Crime and Justice*, vol. 42, no. 1, pp. 199–263, 2013.
- [168] S. Nakamoto, "Bitcoin: a peer-to-peer electronic cash system," <http://bitcoin.org/bitcoin.pdf>, Oct. 2008, last accessed August 18, 2014.
- [169] "NABP list of "not recommended" sites," <http://www.nabp.net/programs/consumer-protection/buying-medicine-online/not-recommended-sites/>, National Association of Boards of Pharmacies, last accessed August 18, 2014.
- [170] "Verified Internet Pharmacy Practice Sites," <http://vipps.nabp.net/>, National Association of Boards of Pharmacies, last accessed August 18, 2014.
- [171] J. Nazario and T. Holz, "As the net churns: Fast-flux botnet observations," in *Proceedings of the 3rd International Conference on Malicious and Unwanted Software*, Oct. 2008, pp. 24–31.
- [172] J. A. Nelder and R. W. M. Wedderburn, "Generalized linear models," *Journal of the Royal Statistical Society. Series A*, vol. 135, no. 3, pp. pp. 370–384, 1972.
- [173] A. Newton, D. Piscitello, B. Fiorelli, and S. Sheng, "A restful web service for internet names and address directory services," *USENIX; login*, pp. 23–32, 2011.
- [174] Y. Niu, H. Chen, F. Hsu, Y.-M. Wang, and M. Ma, "A quantitative study of forum spamming using context-based analysis," in *Proceedings of the 14th Annual Network and Distributed System Security Symposium (NDSS)*, 2007.
- [175] NORC, "Proposed design for a study of the accuracy of WHOIS registrant contact information," University of Chicago, Tech. Rep., 2009.

- [176] A. Ntoulas, M. Najork, M. Manasse, and D. Fetterly, “Detecting spam web pages through content analysis,” in *Proceedings of the 15th international conference on World Wide Web*, ser. WWW ’06. Edinburgh, Scotland: ACM, 2006, pp. 83–92.
- [177] Office of the Deputy US Attorney General, “Google forfeits \$500 million generated by online ads & prescription drug sales by canadian online pharmacies,” *Justice News*, Aug. 2011.
- [178] G. Palla, I. Derényi, I. Farkas, and T. Vicsek, “Uncovering the overlapping community structure of complex networks in nature and society,” *Nature*, vol. 435, pp. 814–818, Jun. 2005.
- [179] D. Pauli, “Srizbi botnet sets new records for spam,” *PCWorld*, May 2008, http://www.pcworld.com/businesscenter/article/145631/srizbi_botnet_sets_new_records_for_spam.html.
- [180] J. Pearl, *Bayesian Networks: A Model of Self-Activated: Memory for Evidential Reasoning*. Computer Science Department, University of California, 1985.
- [181] “Free and open source forum software,” <http://www.phpbb.com>, PhpBB.
- [182] A. C. Pigou, *The economics of welfare*. Transaction Publishers, 1924.
- [183] A. Pitsillidis, K. Levchenko, C. Kreibich, C. Kanich, G. M. Voelker, V. Paxson, N. Weaver, and S. Savage, “Botnet Judo: Fighting spam with itself,” in *Proceedings of the 17th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, USA, Mar. 2010.
- [184] A. Pitsillidis, C. Kanich, G. M. Voelker, K. Levchenko, and S. Savage, “Taster’s choice: a comparative analysis of spam feeds,” in *Proceedings of the 2012 ACM conference on Internet measurement conference*. ACM, 2012, pp. 427–440.
- [185] N. Provos, P. Mavrommatis, M. Rajab, and F. Monroe, “All your iFrames point to us,” in *Proceedings of the 17th USENIX Security Symposium*, Aug. 2008.
- [186] N. Provos, D. McNamee, P. Mavrommatis, K. Wang, and N. Modadugu, “The ghost in the browser: Analysis of web-based malware,” in *Proceedings of the First USENIX Workshop on Hot Topics in Understanding Botnets*, ser. HotBots’07, Cambridge, MA, Apr. 2007.

- [187] U. N. Raghavan, R. Albert, and S. Kumara, "Near linear time algorithm to detect community structures in large-scale networks," *Physical Review E*, vol. 76, no. 3, p. 036106, 2007.
- [188] M. Rajab, L. Ballard, P. Mavrommatis, N. Provos, and X. Zhao, "The nocebo effect on the web: an analysis of fake anti-virus distribution," in *Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats*. USENIX Association, 2010, p. 3.
- [189] A. Ramachandran and N. Feamster, "Understanding the network-level behavior of spammers," in *Proceedings of ACM SIGCOMM'06*, Pisa, Italy, Sep. 2006.
- [190] "The Metasploit penetration testing framework," <http://www.metasploit.com/pdf>, Rapid7, 2014, last accessed August 18, 2014.
- [191] J. Reichardt and S. Bornholdt, "Statistical mechanics of community detection," *Physical Review E*, vol. 74, no. 1, p. 016110, 2006.
- [192] W. Rhodes, P. Johnston, S. Han, Q. McMullen, and L. Hozik, *Illicit drugs: Price elasticity of demand and supply*. Abt Associates, 2000.
- [193] G. Salton and M. J. McGill, "Introduction to modern information retrieval," *McGraw-Hill computer science series*, 1983.
- [194] D. Samosseiko, "The partnerka – what is it, and why should you care?" in *Virus Bulletin Conference*, 2009.
- [195] E. U. Savona, "Infiltration of the public construction industry by italian organised crime," *Situational Prevention of Organized Crimes*, pp. 130–150, 2010.
- [196] J. J. Schlesselman and M. A. Schneiderman, "Case control studies: design, conduct, analysis," *Journal of Occupational and Environmental Medicine*, vol. 24, no. 11, p. 879, 1982.
- [197] B. Schwarz, "Google adwords click through rates: 2% is average but double digits is great," Jan. 2010, <http://www.seroundtable.com/archives/021514.html>. Last accessed May 3, 2011.
- [198] F. Scott and A. Yelowitz, "Pricing anomalies in the market for diamonds: evidence of conformist behavior," *Economic Inquiry*, vol. 48, no. 2, pp. 353–368, 2009.

- [199] D. Segal, “A bully finds a pulpit on the web,” *New York Times*, Nov. 2010, article appeared in print on November 28, 2010, on page BU1 of the New York edition. Available online at <http://www.nytimes.com/2010/11/28/business/28borker.html>.
- [200] —, “The dirty little secrets of search,” *New York Times*, Feb. 2011, article appeared in print on February 13, 2011, on page BU1 of the New York edition. Available online at <http://www.nytimes.com/2011/02/13/business/13search.html>.
- [201] L. W. Sherman, “The rise of evidence-based policing: Targeting, testing, and tracking,” *Crime and Justice*, vol. 42, no. 1, pp. 377–451, 2013.
- [202] Silk Road, “Silk Road anonymous marketplace,” <http://silkroad6ownowfk.onion>. Last accessed August 18, 2014.
- [203] A. Singha, “Finding more high-quality sites in search,” <http://googleblog.blogspot.com/2011/02/finding-more-high-quality-sites-in.html>, Google Inc., Feb. 2011, last accessed August 18, 2014.
- [204] R. R. Sokal, “A statistical method for evaluating systematic relationships,” *Univ Kans Sci Bull*, vol. 38, pp. 1409–1438, 1958.
- [205] A. Sorensen, “Equilibrium price dispersion in retail markets for prescription drugs,” *Journal of Political Economy*, vol. 108, no. 4, pp. 833–850, 2000.
- [206] K. Soska and N. Christin, “Automatically detecting vulnerable websites before they turn malicious,” in *23rd USENIX Security Symposium*. San Diego, CA: USENIX Association, Aug. 2014.
- [207] J. F. Spillane, “Debating the Controlled Substances Act,” *Drug and Alcohol Dependence*, vol. 76, no. 1, pp. 17 – 29, 2004.
- [208] K. Springborn and P. Barford, “Impression fraud in on-line advertising via pay-per-view networks,” in *Proceedings of the 22nd USENIX Security Symposium*. Washington, D.C.: USENIX Association, 2013, pp. 211–226.
- [209] B. Stone-Gross, R. Abman, R. A. Kemmerer, C. Kruegel, D. G. Steigerwald, and G. Vigna, “The underground economy of fake antivirus software,” in *Proceedings of the 10th Workshop on the Economics of Information Security*, Fairfax, VA, Jun. 2011.

- [210] B. Stone-Gross, C. Kruegel, K. Almeroth, A. Moser, and E. Kirda, "Fire: Finding rogue networks," in *Computer Security Applications Conference, 2009. ACSAC'09. Annual*. IEEE, 2009, pp. 231–240.
- [211] A. Sullivan and M. S. Kucherawy, "Revisiting WHOIS: Coming to REST," *IEEE Internet Computing*, vol. 16, no. 3, 2012.
- [212] "The General Store," <http://xqz3u5drneuzhaeo.onion/users/generalstore/>. Last accessed November 25, 2012.
- [213] "Wayback machine," <https://archive.org/web/>, The Internet Archive.
- [214] "The definition of spam," <http://www.spamhaus.org/consumer/definition/>, The Spamhaus Project, last accessed August 18, 2014.
- [215] R. Thomas and J. Martin, "The underground economy: priceless," *USENIX; login*, vol. 31, no. 6, pp. 7–16, 2006.
- [216] "Twitter developers trends resources," <http://dev.twitter.com/doc/get/trends/>, Twitter, last accessed August 18, 2014.
- [217] United States. 106th Congress. Senate. Committee on Health, Education, Labor, and Pensions, *E-drugs: who regulates Internet pharmacies?*, ser. S. hrg. US Government Printing Office, Mar. 2000, vol. 4.
- [218] United States. 107th Congress. House. Committee on Energy and Commerce. Subcommittee on Oversight and Investigations, *Continuing concerns over imported pharmaceuticals*, ser. H. hrg. US Government Printing Office, June 2001, vol. 4. [Online]. Available: <http://www.fda.gov/NewsEvents/Testimony/ucm115214.htm>
- [219] United States. 111th Congress, *The Patient Protection and Affordable Care Act*, ser. Public Law 111-148. US Government Printing Office, 2010. [Online]. Available: <http://www.gpo.gov/fdsys/pkg/PLAW-111publ148/html/PLAW-111publ148.htm>
- [220] United States. 113th Congress, "Safeguarding America's Pharmaceuticals Act of 2013," <http://www.govtrack.us/congress/bills/113/hr1919>, last accessed August 18, 2014.
- [221] United States. 118th Congress. Senate. Committee on Governmental Affairs. Permanent Subcommittee on Investigations, *Buyer beware: the danger of purchasing pharmaceuticals over the Internet*, ser. S. hrg. US

Government Printing Office, June 17 and July 22 2004, vol. 108. [Online]. Available: <http://www.fda.gov/NewsEvents/Testimony/ucm113635.htm>

- [222] United States. Congress. House, *Proceedings of Congress and General Congressional Publications*, ser. Congressional Record. US Government Printing Office, November 2002, vol. 148.
- [223] United States. Congress. House. Committee on Energy and Commerce, *Ryan Haight Online Pharmacy Consumer Protection Act*, ser. Report. US Government Printing Office, 2008.
- [224] United States. Congress. House. Committee on Interstate and Foreign Commerce, *Federal food, drug, and cosmetic act*. US Government Printing Office, 1938.
- [225] United States. Congress. House. Committee on the Judiciary. Subcommittee on Courts, the Internet, and Intellectual Property, *Internet Domain Name Fraud: The U.S. Government's Role in Ensuring Public Access to Accurate WHOIS Data*, ser. H. hrg. US Government Printing Office, September 2003.
- [226] United States. Congress. Senate. Committee on Finance. Subcommittee on International Trade, *Prescription Drug Marketing Act of 1987: hearing before the Subcommittee on International Trade of the Committee on Finance*, ser. S. hrg. US Government Printing Office, 1988, no. v. 4.
- [227] United States. Congress. Senate. Committee on Government Operations. Subcommittee on Executive Reorganization and Government Research and United States. Congress. Senate. Committee on Government Operations. Subcommittee on Intergovernmental Relations, *Drug Abuse Prevention and Control*, ser. Drug Abuse Prevention and Control. US Government Printing Office, 1971, vol. 74-76.
- [228] US Census Bureau, "International data base country rankings," <http://www.census.gov/population/international/data/idb/rank.php>, last accessed August 18, 2014.
- [229] US Department of Justice. Drug Enforcement Administration, "Controlled substances - alphabetical order," http://www.deadiversion.usdoj.gov/schedules/orangebook/c_cs_alpha.pdf, last accessed August 18, 2014.
- [230] —, "Operation Cyber Chase," <http://www.justice.gov/dea/pubs/pressrel/pr042005.html>, last accessed August 18, 2014".

- [231] —, “Operation Cyber X,” <http://www.justice.gov/dea/pubs/pressrel/pr092105b.html>, 2005, last accessed August 18, 2014.
- [232] US Department of Justice. Drug Enforcement Administration. Office of Diversion Control, “Special surveillance list of chemicals, products, materials and equipment used in the clandestine production of controlled substances or listed chemicals,” http://www.dea.gov/diversion/chem_prog_advisories/surveillance.htm, May 1999, last accessed August 18, 2014.
- [233] US Food and Drug Administration, “Current drug shortages index,” <http://www.fda.gov/Drugs/DrugSafety/DrugShortages/ucm050792.htm>. Last accessed August 18, 2014.
- [234] —, “Imported drugs raise safety concerns,” <http://www.fda.gov/Drugs/ResourcesForYou/Consumers/ucm143561>, last accessed August 18, 2014.
- [235] —, “National drug code directory,” Nov. 2010, <http://www.fda.gov/Drugs/InformationOnDrugs/ucm142438.htm>.
- [236] M. Vasek and T. Moore, “Identifying risk factors for webserver compromise,” in *Financial Cryptography and Data Security*, 2014.
- [237] “The domain industry brief,” Verisign, 2010, http://www.verisigninc.com/assets/Verisign_DNIB_Nov2010_WEB.pdf.
- [238] “Free online virus, malware and URL scanner,” <https://www.virustotal.com/>, VirusTotal, last accessed August 18, 2014.
- [239] D. Wang, M. Der, M. Karami, L. Saul, D. McCoy, S. Savage, and G. Voelker, “Search + seizure: The effectiveness of interventions on seo campaigns,” in *Proceedings of ACM IMC’14*, Vancouver, BC, Canada, Nov. 2014.
- [240] D. Wang, G. Voelker, and S. Savage, “Juice: A longitudinal study of an SEO botnet,” in *Proceedings of NDSS’13*, San Diego, CA, Feb. 2013.
- [241] D. Y. Wang, S. Savage, and G. M. Voelker, “Cloak and dagger: dynamics of web search cloaking,” in *Proceedings of the 18th ACM conference on Computer and communications security*, ser. CCS ’11. Chicago, Illinois, USA: ACM, 2011, pp. 477–490.
- [242] Y.-M. Wang, M. Ma, Y. Niu, and H. Chen, “Spam double-funnel: connecting web spammers with advertisers,” in *Proceedings of the 16th international*

conference on World Wide Web, ser. WWW '07. Banff, Alberta, Canada: ACM, 2007, pp. 291–300.

- [243] P. A. Watters, A. Herps, R. Layton, and S. McCombie, “Icann or icant: Is whois an enabler of cybercrime?” in *Fourth Cybercrime and Trustworthy Computing Workshop*. IEEE, 2013, pp. 44–49.
- [244] “Drugs and medications a–z,” <http://www.webmd.com/drugs/index-drugs.aspx>, WebMD, last accessed August 18, 2014.
- [245] WHOIS Task Force 3, “Improving accuracy of collected data,” <http://gnso.icann.org/en/issues/whois-privacy/tor3.shtml>, ICANN, 2003, last accessed August 18, 2014.
- [246] R. Willison, “Understanding the perpetration of employee computer crime in the organisational context,” *Information and organization*, vol. 16, no. 4, pp. 304–324, 2006.
- [247] T. Wilson, “Researchers link storm botnet to illegal pharmaceutical sales,” *Dark Reading*, Jun. 2008, <http://www.darkreading.com/security/security-management/211201114/index.html>.
- [248] “Blog tool, publishing platform, and CMS,” <http://www.wordpress.org>, Wordpress, last accessed August 18, 2014.
- [249] Y. Xie, F. Yu, K. Achan, R. Panigrahy, G. Hulten, and I. Osipkov, “Spamming botnets: signatures and characteristics,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 171–182, Aug. 2008.
- [250] “Content quality guidelines: Cloaking,” <https://help.yahoo.com/kb/search/content-quality-guidelines-sln2245.html>, Yahoo! Inc., last accessed August 18, 2014.
- [251] “Yahoo buzzlog,” <http://buzzlog.yahoo.com/overall/>, Yahoo! Inc., last accessed August 18, 2014.
- [252] “Yahoo site explorer,” <http://siteexplorer.search.yahoo.com/>, Yahoo! Inc.
- [253] F. Yarochkin, V. Kropotov, Y. Huang, G.-K. Ni, S.-Y. Kuo, and I.-Y. Chen, “Investigating DNS traffic anomalies for malicious activities,” in *43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop*. IEEE, 2013, pp. 1–7.

- [254] F. Ye and D. Lord, "Comparing three commonly used crash severity models on sample size requirements: multinomial logit, ordered probit and mixed logit models," *Analytic methods in accident research*, vol. 1, pp. 72–85, 2014.
- [255] J. Zhuge, T. Holz, C. Song, J. Guo, X. Han, and W. Zou, "Studying malicious websites and the underground economy on the chinese web," *Managing Information Risk and the Economics of Security*, pp. 225–244, 2009.