THE UNIVERSITY OF TULSA
THE GRADUATE SCHOOL

MEASURING BITCOIN-BASED CYBERCRIME

by
Marie Vasek

A dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor of Philosophy
in the Discipline of Computer Science

The Graduate School
The University of Tulsa

2017

THE UNIVERSITY OF TULSA

THE GRADUATE SCHOOL

MEASURING BITCOIN-BASED CYBERCRIME

by
Marie Vasek

A DISSERTATION

APPROVED FOR THE DISCIPLINE OF

COMPUTER SCIENCE

By Dissertation Committee

_____, Chair
Tyler Moore

_____
John Hale

_____
Roger Wainwright

_____
Sal Aurigemma

_____
Nicolas Christin

ABSTRACT

Marie Vasek (Doctor of Philosophy in Computer Science)

Measuring Bitcoin-based Cybercrime

Directed by Tyler Moore

86 pp., Chapter 7: Conclusions

(291 words)

Bitcoin is a decentralized, digital, public currency invented in 2009 by the pseudonymous Satoshi Nakamoto. The decentralized nature of the currency makes it attractive to fraudsters who can transact along with every other user. The digital nature makes it attractive for online businesses. The public nature makes it attractive for businesses who want to imbue trust in their customers as to their cash holdings. Unfortunately, the combination of these features also makes it ripe for cybercriminals. In turn, the public nature of the currency makes it feasible for researchers to be able to measure the prevalence and profits of attacks.

We leverage the public nature of Bitcoin to measure cybercrime. First, we investigate distributed denial of service attacks carried out against various Bitcoin services. We find that Bitcoin currency exchanges, mining pools, gambling operators, online wallets, and financial services are much more likely to be attacked than other services. Next we present the first empirical analysis of Bitcoin-based scams: operations established with fraudulent intent. We find that at least $11 million has been contributed to the scams from 13 000 distinct victims. Furthermore, we present evidence that the most successful scams depend on large contributions from a very small number of victims. We then investigate Ponzi schemes advertised on the Bitcoin forum and the ecosystem that perpetuates them. We find that the more scammers and victims post, the shorter the scam lifetime. Likewise, scams posted by users who register their account on the same day (39% of the total) are found to be much shorter lived. Finally we analyze Bitcoin brain wallets – Bitcoin secured

by the hash of a password or passphrase. We find that most are depleted of money within a day, many within seconds of creation.

ACKNOWLEDGEMENTS

My PhD advisor, Tyler Moore, has been instrumental in the production of this dissertation. I worked with him for six years so far, and hopefully for many more to come. All of the work in this dissertation is coauthored by him. More than that, his advice and leadership has provided me with a solid role model for my research career to come.

I would also like to thank my committee members for their helpful contributions to this dissertation. Thanks John Hale, Roger Wainwright, Sal Aurigemma, and Nicolas Christin. Their feedback has been invaluable.

I would like to thank the other coauthors I have had during this time. Special thanks to Joseph Bonneau, Ryan Castellucci, Jens Grossklags, Ben Johnson, Cameron Keith, Aron Laszka, Micah Thornton, John Wadleigh, and Matthew Weeden. My work is stronger because of the fabulous people I have been privileged to work with.

Many productive (and less than productive) conversations with my lab mates have greatly strengthened my work. Thanks Matt, JT, Andrew, Ali, Gavin, Kyle, and others I have forgotten to mention.

This work could not be done without the support of my friends, albeit from a distance. Thanks Anthea, Linnea, Shannon, Julia, John, and Ryan. Their kind words, helpful distractions, and constructive criticism have been an indispensable source of strength.

My family has supported me through the entirety of my formal education. Thanks Mom and Dad for always being there and reading all of my work. Thanks Reid and Arthur for supporting me and teaching me things I might not have seen the need for, but later have proven to be useful.

This dissertation is dedicated to Eleanor Roosevelt.

PUBLICATIONS

[1] Vasek, M., AND Moore, T. Do malware reports expedite cleanup? An experimental study. In: *Proceedings of the 5th USENIX Workshop on Cyber Security Experimentation and Test*. CSET'12. Bellevue, WA, Aug. 2012.

[2] Vasek, M., AND Moore, T. Empirical Analysis of Factors Affecting Malware URL Detection. In: *8th APWG eCrime Researchers Summit (eCrime)*. Sept. 2013.

[3] Vasek, M., Thornton, M., AND Moore, T. Empirical Analysis of Denial-of-Service Attacks in the Bitcoin Ecosystem. In: *1st Workshop on Bitcoin Research*. Vol. 8438. Lecture Notes in Computer Science. Springer, Mar. 2014, 57–71.

[4] Johnson, B., Laszka, A., Grossklags, J., Vasek, M., AND Moore, T. Game-Theoretic Analysis of DDoS Attacks Against Bitcoin Mining Pools. In: *1st Workshop on Bitcoin Research*. Vol. 8438. Lecture Notes in Computer Science. Springer, Mar. 2014, 72–86.

[5] Vasek, M., AND Moore, T. Identifying Risk Factors for Webserver Compromise. In: *Financial Cryptography and Data Security*. Vol. 8437. Lecture Notes in Computer Science. Springer, Mar. 2014, 326–345.

[6] Vasek, M., AND Moore, T. There's no free lunch, even using Bitcoin: Tracking the Popularity and Profits of Virtual Currency Scams. In: *Financial Cryptography and Data Security*. Ed. by Böhme, R., AND Okamoto, T. Vol. 8975. Lecture Notes in Computer Science. Springer, Jan. 2015, 44–61.

[7] Vasek, M., Bonneau, J., Castellucci, R., Keith, C., AND Moore, T. The Bitcoin brain drain: Examining the Use and Abuse of Bitcoin Brain Wallets. In: *Financial Cryptography and Data Security*. Lecture Notes in Computer Science. Springer, Feb. 2016.

[8] Vasek, M., Wadleigh, J., AND Moore, T. Hacking is not Random: A Case-Control Study of Webserver-Compromise Risk. *IEEE Transactions on Dependable and Secure Computing* 13, 2 (2016), 206–219.

[9] Vasek, M., Weeden, M., AND Moore, T. Measuring the Impact of Sharing Abuse Data with Web Hosting Providers. In: *Workshop on Information Sharing and Collaborative Security*. Oct. 2016.

TABLE OF CONTENTS

ABSTRACT iii

ACKNOWLEDGEMENTS v

PUBLICATIONS vi

TABLE OF CONTENTS ix

LIST OF TABLES x

LIST OF FIGURES xii

CHAPTER 1: **INTRODUCTION** 1
1.1 **Prior Art** 2
1.2 **Structure and Contribution of this Thesis** 6
1.2.1 *Thesis Statement* 6
1.2.2 *Structure* 6
1.2.3 *Contributions* 7

CHAPTER 2: **BITCOIN PRIMER** 9
2.1 **Bitcoin Mining** 10
2.2 **Bitcoin Wallets** 10
2.3 **Bitcoin Services** 11
2.4 **Conclusion** 12

CHAPTER 3: **MEASURING DENIAL-OF-SERVICE ATTACKS IN THE BITCOIN ECOSYSTEM** 13
3.1 **Methodology** 14
3.1.1 *Data Collection* 14
3.1.2 *Classification of Posts Describing Attacks* 15
3.2 **Empirical Analysis** 16
3.2.1 *DDoS Attacks over Time and by Target* 16
3.2.2 *DDoS Attacks on Mining Pools* 20
3.2.3 *DDoS Attacks on Currency Exchanges* 23
3.3 **Related Work** 25
3.4 **Conclusion** 25

CHAPTER 4: **MEASURING THE PROFITS OF BITCOIN SCAMS** 27
4.1 **Methodology for Identifying Scams and Associated Transactions** . 28
4.2 **High Yield Investment Programs** 31
4.2.1 *Traditional HYIPs* 31
4.2.2 *Bridge HYIPs* 32

4.2.3 *Bitcoin-only HYIPs* . . . 35
4.3 **Mining Scams** . . . 36
4.4 **Scam Wallets** . . . 37
4.5 **Bitcoin Exchange Scams** . . . 39
4.6 **Discussion** . . . 40
4.6.1 *Revisiting the Scam Categories* . . . 40
4.6.2 *How are Victim Payments into Scams Distributed?* . . . 41
4.6.3 *Policy Options* . . . 43
4.7 **Related Work** . . . 44
4.8 **Conclusion** . . . 45

CHAPTER 5: **MEASURING THE SUPPLY AND DEMAND FOR BITCOIN SCAMS** 47
5.1 **Methodology** . . . 48
5.2 **Results** . . . 50
5.2.1 *Scammer Interaction and Scam Lifetime* . . . 51
5.2.2 *Victim Behavior* . . . 53
5.2.3 *Proportional Hazards Model* . . . 56
5.3 **Conclusion** . . . 57

CHAPTER 6: **MEASURING THE USE AND ABUSE OF BRAIN WALLETS** 59
6.1 **Data Collection Methodology** . . . 60
6.1.1 *Password Corpora* . . . 61
6.1.2 *Observing Bitcoin Brain Wallet Usage* . . . 63
6.2 **Results** . . . 63
6.2.1 How Prevalent are Brain Wallets? . . . 64
6.2.2 *Draining Brain Wallets* . . . 67
6.2.3 *Network "Stress Test"* . . . 69
6.2.4 *Tracking the Drainers* . . . 70
6.2.5 *Mining Pool Drains* . . . 71
6.2.6 *Impact of Password Strength* . . . 72
6.3 **Conclusion** . . . 74

CHAPTER 7: **CONCLUSION** 75
7.1 **Future Work** . . . 77
7.1.1 *Effects of Technological Counters to DDoS* . . . 78
7.1.2 *Market Responses to DDoS Attacks* . . . 78
7.1.3 *Compounding Effects of Attacks* . . . 78
7.1.4 *Effects of Default Standards on Security* . . . 78
7.1.5 *Effects of Social Behaviors on the Profits of Ponzi Schemes* . . . 79
7.1.6 *Scam Early Detection System* . . . 79
7.1.7 *Cryptocurrency Service Legitimacy Indicators* . . . 79

BIBLIOGRAPHY . . . 80

LIST OF TABLES

3.1 Confusion matrix plus precision, recall and accuracy measures for the word-based classifier. 16

3.2 Prevalence of DoS attacks and anti-DDoS (AD) uptake by service category. 19

3.3 Contingency table comparing the uptake of anti-DDoS protection based on whether or not the service has experienced DDoS attacks. 20

3.4 Contingency table comparing the size of a mining pool to whether or not the pool has experienced DDoS attacks. 21

3.5 Changes in transaction volume on Mt. Gox after a DDoS attack. 24

4.1 For each scam category, we report whether we can directly observe transactions corresponding to what victims pay into scams, what is paid out to victims, and what is paid out to the scammer (indicated by a ✓). 30

4.2 Summary statistics for HYIPs. 32

4.3 Lifetime and payouts for scam wallets and exchanges, plus mining scam payouts. 39

4.4 Recap of Bitcoin scam categories and features. 40

5.1 Bitcointalk forum categories and where scam victims post. Categories are marked as under or overrepresented according to a chi-squared test with 97.5% confidence. Categories with at least 50 000 posts are included. 55

5.2 Cox proportional hazards model: measuring scammer and victim effects on the lifetime of the scam. 56

6.1 Brain wallets and values associated with different password sources. 64

6.2 Top 10 drain addresses from brain wallets, sorted by amount drained in USD. 70

LIST OF FIGURES

3.1 Reported DDoS attacks over time, split up by category of targeted service. 17

3.2 Percentage of DDoS attacks targeting each major category (left); cumulative distribution function of the number of attacks targeting each service (right). 18

3.3 Mining pool hashrate market share (solid line) over time, compared to timing of DDoS attacks (dashed lines). 22

3.4 Daily trade volumes (top) and USD-BTC exchange rate (bottom) at Mt. Gox. Dashed green lines indicate when DDoS attacks on Mt. Gox were reported. 23

3.5 Changes in transaction volume on Mt. Gox after a DDoS attack over time. 24

4.1 Multiple-address transactions in Bitcoin. 30

4.2 Top: Daily volume of all payments into and out of Bridge HYIPs wallet incoming transactions. Bottom: daily volume of incoming payments split by HYIP. 33

4.3 Daily volume of payments into and out of Bridge HYIPs, sorted by total payments received. The green dotted line indicates when the scam is first promoted on `bitcointalk.org`. 35

4.4 Weekly payouts to scam wallets in BTC (top) and USD (bottom). 38

4.5 Lorenz curve for Bridge HYIPs (left) and Bitcoin-only HYIPs (right). . . . 42

4.6 Lorenz curve for total payments into scam categories (left); scatter plot comparing Gini coefficient to the amount of money stolen by scammers (right). 43

5.1 Screenshots of the initial posting for the Ponzi scheme and an example victim response. 49

5.2 Survival analysis of the lifetime of scams. 51

5.3 Lifetime of the scam based on the fraction of the comments about the scam from the scammer. 52

5.4 Lifetime of the scam by interaction by "shill" commenters. 53
5.5 Measuring lifetimes of scams based on attacker accounts. 54
5.6 Number of victim posts after a thread starts. 54
6.1 New brain wallet usage per month (compressed and uncompressed). 67
6.2 CDF and rank-order plot of total value stored in brain wallets. 68
6.3 CDF of the # of hours to drain brain wallets for wallets by value stored . . 68
6.4 How time-to-drain changes over time (median time-to-drain reported per month) 69
6.5 Measuring time-to-drain for passwords versus passphrases 73

CHAPTER 1

INTRODUCTION

Cybercrime is hard to measure. Bitcoin makes it easier. Previous work has depended on criminals accurately reporting on their actual revenues [38], devising clever ways to infer profits [3], or waiting until their databases are publicly leaked [51]. However, with Bitcoin, we can directly measure criminal activities and revenue. Bitcoin is a publicly accessible database of monetary transactions. Given an address, it is trivial for anybody to see all the transactions in and out of the address. This is in stark contrast to the banking system where transaction information is held only by banks and authorized users making abuse hard to externally measure or verify.

In theory, the openness of Bitcoin allows users to garner more trust in their Bitcoin-backed projects. Instead of relying on external auditors to provide assurances of a bank's solvency, Bitcoin exchanges can give a cryptographic proof of solvency for anybody to verify [27]. However, attackers can leverage the same openness to solicit trust in their schemes as well. For example, Ponzi scheme operators can point to specific transactions, demonstrating to victims that they actually pay out.

Currently, Bitcoin resembles the Wild Wild West – there are no rules[1]. We see some Bitcoin currency exchanges that, instead of hardening their services, lay vulnerable to denial of service attacks. In turn, those Bitcoin exchanges stop transacting whenever attackers try to profit off a period of no transactions or decide they do not like the exchange for a political or competitive reason. Large-scale heists of Bitcoin from online services cause anything from temporary loss of service to collapse, not to mention the direct monetary damage. The infrastructure is still vying to be first to the market, not as hardened to these attacks as their counterparts in the traditional banking ecosystem. We see Bitcoin exchanges failing at a high rate (an estimated 45% [55]). Bitcoin still has not established

[1]Some legislation has been approved that attempts to regulate Bitcoin industries (eg see `https://www.loc.gov/law/help/bitcoin-survey/`), but this regulation is relatively new, ad hoc, and not comprehensive. For instance, Bitcoin regulation in New York State is restricted to operations in the state; this mainly functioned to push Bitcoin businesses to other states or to block New York residents from directly accessing their services.

itself; Bitcoin's death has been predicted over 100 times by reputable sources[2], many within the past year. The lack of effective regulation or community norms in this mad rush to the market has contributed to the high rates of collapse and attack.

We also see similar behavior in Bitcoin wallet clients. There are tradeoffs in different methods of storing bitcoin[3] [30]. Since the industry has taken a while to standardize on a set of best practices, we are left with insecure wallets. For example, Android wallets used insecure random numbers, allowing attackers to steal the contents en masse [41]. Scam wallets are advertised on places like the Bitcoin forums and the TOR hidden wiki and usually act like a wallet when a small amount of Bitcoin is deposited, stealing all the funds when a user subsequently put in a larger amount. We also find a number of users put their bitcoin in brain wallets which secures their bitcoin with only a password. This trend enabled a sequence of attackers who scanned the Bitcoin network for such addresses and drained them.

Bitcoin is a tool that allows us to measure criminal behavior. We will outline subsequent behaviors in this dissertation, but note that the interesting takeaways are not the particulars about this vehicle for crime, but rather the methods in which attackers wield the network and its users.

1.1 Prior Art

There is a large body of research on improving understanding of how cybercriminals operate so that we might more effectively dismantle their networks. This thesis contributes to the knowledge on these threats by studying their prevalence in the Bitcoin ecosystem. The transparency of Bitcoin enables us to glean more information on scams than was previously possible. In other cases, the greater prevalence of attacks lets us study their effects more closely.

More broadly, the work contained in this thesis measures cybercrime incidents with an eye towards disrupting attacks using economic approaches, rather than purely technical ones. Bitcoin here is a tool that we can use to directly measure cybercrime incidents

[2] `http://bitcoinobituaries.com`

[3] Following convention, lowercase "bitcoin" refers to a unit of currency within the uppercase Bitcoin system.

without having to trust the attackers or rely on outside reports. A large body of work has been conducted in this space, known as security economics, that complements the underlying premise of our work. Anderson and Moore review economic mechanisms that cause suboptimal security outcomes [4]. They note that computer security often imposes negative externalities. A user might not notice or care that their computer is infected with malware, but when that malware puts their computer in a botnet and is used to take down large parts of the Internet in a denial of service attack (e.g. the recent Mirai botnet [39]), the Internet at large is harmed by that user's insecurity. Because the user is not harmed directly, they are less likely to invest in secure outcomes. Other work in security economics has addressed the strategic interaction between attackers and defenders. For instance, Böhme and Moore looked at the effect of domain takedowns on phishing [15]. They noticed that criminals were registering domains in a certain country from a particular domain registrar. Once the domain registrar caught onto the scam, the criminals merely moved to a different country's infrastructure and started over again. Even though each country's domain registrar eventually wisened to the scam, the abuse still lingered in another place. Here the technical solution – identifying the phishing domains and removing them – was only a small part of the overall solution. Rather, the economic solution – sharing information about the criminal group and raising the price of a country's domains – has proven to be more effective.

Other work in this field measures similar systems as presented here, albeit without considering Bitcoin. Moore et al. documented the online high yield investment program (HYIP) ecosystem [56]. They monitored over 1 000 such scams through an aggregator service. They estimated the profits of these scams by estimating visitors to the scams, the percentage that they invest, and the amount of money each investor invests in the scam. Neisius and Clayton built upon that work, noting the monetization of HYIP kits that made HYIPs easy for anybody to set up [64]. They enumerate the profit model of these schemes using numbers the criminals publish along with estimated numbers. Drew and Moore found clusters of replicated HYIP websites, pointing to the high use of HYIP kits in creating Ponzi scheme websites [29].

Moore et al. measured the worldwide prevalence of denial of service attacks [54]. Mirkovic and Reiher devise a taxonomy to separate out different types of DDoS attacks and different ways of defending against them [53]. Nazario analyzes politically motivated DDoS attacks [63]. He walked through the history of these attacks and how these politically motivated criminals leveraged the same infrastructure as profit motivated ones. Parts of this work also fits into the greater literature of reputation mechanisms. Resnick et al. provide a general overview for reputation systems as well as drawbacks in them [66]. Shen et al. provide analysis of reviewers posting about products on online retailers [75]. They found that popular reviewers post about popular products that have few reviews and also tend to provide similar reviews to the existing ones about the product.

Research in this dissertation also relates to other work measuring parts of the Bitcoin network. As interest in Bitcoin has exploded, researchers have undertaken a number of measurement studies to improve our understanding of how Bitcoin is used and abused in practice. Our blockchain analysis techniques are similar to that of others. Ron and Shamir reconstruct a transaction graph from the Bitcoin blockchain in order to find out how money changes hands and identify suspicious transactions (e.g., attempts to launder identity) [68][69]. Other researchers spend bitcoin and trace their money through the Bitcoin blockchain. For example, Meiklejohn et al. measure the traceability of transactions initiated at many Bitcoin service providers [52] by spending money at each of them. Möser et al. systematically analyze the traceability of three popular Bitcoin mixing services by sending some bitcoin through each one [58]. They found that most were low volume and effectively broken.

Other researchers have performed measurement studies on other nefarious uses for Bitcoin. We draw a distinction between Bitcoin-*based* cybercrime and Bitcoin-*facilitated* cybercrime. Our work concentrates on Bitcoin-based cybercrime: crimes that can happen only because Bitcoin exists. Bitcoin-based cybercrime targets Bitcoin users, frequently uses Bitcoin-related ruses, as well as predominantly using Bitcoin as a payment mechanism. Examples of Bitcoin-based cybercrime include:

- Denial of Service attacks on Bitcoin services (Chapter 3)

- Bitcoin scam services (Chapter 4)
- Bitcoin brain wallet drainers (Chapter 6)
- Bitcoin currency exchange thefts [55]
- Spam transaction attacks against the Bitcoin network [8]

The methodology and analysis framework that we lay out in this dissertation is directly applicable to Bitcoin-based crimes such as these.

Bitcoin-facilitated crime merely uses Bitcoin as a payment mechanism or monetization strategy. If Bitcoin did not exist, then these criminals would move to use different payment mechanisms, like Western Union, cash, or Perfect Money (a centralized digital currency). Examples of Bitcoin-facilitated cybercrime include:

- Ransomware [47, 43]
- Online illicit marketplaces [24]
- Bitcoin mining malware [40]
- High-yield investment programs (a form of online Ponzi scheme) that accept Bitcoin and other currencies (Chapter 4)

Bitcoin has a small community of actors [16]. Maurer et al. associated the distributed network of Bitcoin nodes with the distributed network of conversations, like those found on the Bitcoin forums [48]. We do not disagree that the "sociality of trust" that Bitcoin offers seems to be both ingrained in the code and the community. We use this small network of trust ingrained in code and in people to more easily measure communications and outcomes.

Work has been done looking at inherent security vulnerabilities in the Bitcoin network. Barber et al. describe a Doomsday, "51%", attack where miners edit the past transaction history of the blockchain [9]. Eyal and Sirer further refine the attack assuming colluding miners, lowering the threshold from 50% to 33% of total mining hashrate needed to control the blockchain [31]. Heilman et al. devise an eclipse attack where attackers with

enough IP addresses can control all the incoming connections to a Bitcoin node and control their information about the Bitcoin network [36]. Nayak et al. then combine selfish mining and the eclipse attack for their stubborn mining attack [62]. Similar attacks inspired by selfish mining include Bonneau's bribery attack, Sapirshtein et al.'s optimal selfish mining refinements, and Teutsch et al.'s alternative puzzle attack [19][71][77].

Kroll et al. model whether a miner should join a mining pool using game theory. They expand their model to describe a "Goldfinger" attack on the Bitcoin network [44]. Rosenfeld describes a double-spending attack [70]. Andrychowicz et al. study malleability of bitcoin transactions [6] and Decker and Wattenhofer measured the attackers targeting MtGox with these transactions [28]. Transactions are malleable when an attacker can transform an original transaction to a different one where all the inputs and outputs are the same, but it hashes to a different value. These transactions cause issues for primitive software that rely heavily on transaction hashes. McCorry et al. consider refund attacks which BIP70 (a Bitcoin payment protocol standard) enables [50]. These attacks against merchants using this protocol enables customers to both receive items and their money back with plausible deniability that such an attack occurred.

1.2 Structure and Contribution of this Thesis

1.2.1 Thesis Statement

We measure cybercrime activity in the Bitcoin ecosystem to better understand attacker motivation and the efficacy of various crimes and countermeasures.

1.2.2 Structure

Chapter 2 gives an overview about the Bitcoin network. We detail different aspects of the Bitcoin ecosystem that are needed to understand the rest of this dissertation.

Chapter 3 details the distributed denial of service (DDoS) attacks that plague various parts of the Bitcoin ecosystem. These profit and political-minded attacks change targets based on profitability of different Bitcoin sectors as well as political changes (such as the acceptance of alternate cryptocurrencies). We note that DDoS attacks on Bitcoin services have similar motivations as DDoS attempts on the rest of the web (money, making a polit-

ical statement) but with Bitcoin, the small number of attacks and actors allows us to more easily quantify their proliferation and impact.

Chapter 4 analyzes Bitcoin scams across different ruses and the profitability of these scams. Here we use the Bitcoin blockchain to directly measure the current profitability of any given scam. We find the scams by inspecting the venues that they are advertised on – the Bitcoin forum and subreddit. Chapter 5 then looks at the behavior of scammers and victims and how that contributes to the lifetime of individual scams. We find the scams on the Bitcoin forums and then look at how the users interact with the scams and consequently how some scams to profit and others to die. We also look at scammer behavior and the differences between scams run by those with reputation and those without.

Chapter 6 looks at bad passwords that users select to control their bitcoin via brain wallets. Here, Bitcoin facilitates a platform to analyze bad password design as well as attacker behavior. We find that users picking bad passwords to secure their Bitcoin wallet fuels an ecosystem of attackers trying to drain these wallets as fast as they are filled.

1.2.3 Contributions

The research contributions for this work are found both in the data collection methodology and the analysis of gathered data. As explained above, Bitcoin's transparency and high rate of attack creates an opportunity to gather novel datasets in new ways. The first step for all the chapters involves identifying candidates to potentially measure. In Chapter 3, we gather reports of DDoS attacks from user reports on forums. In Chapter 4, we find candidate scams using aggregated defender data. In Chapter 5, we gather our candidate scams directly from scammer advertising venues. We generate our candidate brain wallets for Chapter 6 by gathering a large corpus of passwords and passphrases. Then we confirm the data we collect. In Chapter 3, we construct a rule-based classifier to identify that posts in fact discuss DDoS attacks. We use manual inspection to confirm data for Chapter 4 and automatically identify payout mechanisms to confirm data for Chapter 5. Finally, we use this confirmed data to measure attacker behavior. Chapter 3 provides a reliable estimation of attack targets and date. Similarly, for Chapter 5 we collect usage, performance and demographic indicators from forum posts. Chapter 4 uses the public Bit-

coin blockchain to measure money in and out of scams. We use the blockchain also in Chapter 6 to measure attacker draining behavior on our brain wallets.

The methodological contributions just outlined in turn enable novel analysis of these new cybercrime datasets. We provide summary statistics of attack prevalence over time for Chapter 3, dividing the attacks based on target. We document prevalence and revenues of scams over time for Chapter 3 quantifies the impact of DDoS attacks on differing targets. Chapter 4, dividing the scams based on a new scam taxonomy. Chapter 5 quantifies the effects of scammer behavior and victimology on the lifetime of the scam. Chapter 4 also quantifies for the first time HYIP cash flows and victims losses over time. Chapter 6 quantifies users' password selection and attacker draining behavior, which provides a direct insight into how users cannot pick strong passwords, even to directly secure their money.

CHAPTER 2

BITCOIN – A PRIMER

Bitcoin [60], launched in 2009, is the first decentralized cryptographic currency and has recently attracted considerable research [14][20]. The Bitcoin system is a public distributed database (**blockchain**) where the records are transactions (bitcoin). Bitcoin uses **proof of work** to ensure that these records are both correct and universally accepted.

Bitcoin is attractive to techno-libertarians as well as those who, for a variety of reasons, do not have access to or do not trust their financial system. The appeal to libertarians partially lays in the fact that there is a fixed amount of bitcoin that will ever be created. This will eventually make Bitcoin a deflationary currency,[1] and this is seen as a good thing – a foil to most inflationary currencies where entities are incentivized to be constantly in debt. The distributed nature of the network ensures that one entity cannot gain power over Bitcoin. This decentralization appeals to a wide range of people, from US citizens who perceive the US banking system to be corrupt, to Afghani women who are not legally allowed to have a bank account in their country.

Most of what we currently consider money gets value from a governmental backing. In contrast, Bitcoin is backed by the Bitcoin network. Very few people are paid in bitcoin. There is no central Bitcoin organization that people trust to back their bitcoin. No nation state uses bitcoin. Instead, bitcoins are worth money because people are willing to trade cash for bitcoin.

The rest of this chapter will discuss some basic mechanics of how Bitcoin works, eschewing the more technical details. For a more detailed explanation of the technical underpinnings, we refer the reader to [61].

[1] The finite money supply will make the currency deflationary if the currency will continue to be used. Fernández-Villaverde and Sanches discuss deflationary virtual currencies more in depth, particularly how they interact with each other and potential government issuance [33].

2.1 Bitcoin Mining

Bitcoin **miners** perform proof of work by attempting to produce a partially matching preimage of a hash function. Hashcash, the proof of work mechanism used by Bitcoin, was originally proposed by Back as a denial-of-service countermeasure [7]. They run software that repeatedly hashes different nonces combined with metadata about the state of the Bitcoin system. If a miner happens to guess a nonce that hashes to a number lower than the target value, they satisfy the proof of work and have the right to add a block of transactions to the network. As an incentive to mine, the first transaction is a **block reward** from the network to the miner. The current block reward is 12.5 bitcoin, or about 12 715 USD as of March 2017. The block reward programmatically decreases with time until it reaches zero. At that time, no new bitcoin will ever be created and the network will rely on fees attached to transactions to incentivize miners.

This block will be accepted into the network if and only if it contains only valid transactions, it has the appropriate proof of work attached, and it is accepted before another block with the same parent. Note that the target value needed to satisfy the proof of work is recalculated approximately every two weeks to try to ensure that a new block is released on average every ten minutes.

Not all miners mine alone. While this is what the original creator of Bitcoin intended, eventually miners realized that they could have more reliable returns if they pooled together and, in return, shared the reward. Bitcoin **mining pools** split work (aka the nonce search space) between pool participants and then, if a block is won, split the reward. Mining pools reduce slightly the expected block reward for individual miners (since some money is usually taken off the top for the pool operator), but increase the expected probability of earning any money during a short time period [73]. In the beginning, this trade-off was not worth it. Currently, however, almost all of the Bitcoin mining power comes from mining pools.

2.2 Bitcoin Wallets

Holding bitcoin is reliant on public key cryptography. Transactions which transfer control of bitcoins are authorized by ECDSA digital signatures. Most bitcoin are tied to a

single private key, public key pair. A bitcoin **address** is a readable encoding of a hash of a bitcoin public key or script.

Bitcoin **wallets** store private keys corresponding to a users' public keys. Some software stores that key information on a users' local machine. This can be printed out to be put in a safe as a backup. Other software stores the key information on the hosted platforms' machines. These software services act as de-facto banks and are a target of attacker interests.

2.3 Bitcoin Services

If a user wants some bitcoin but does not want to mine their own, they can trade their other currency for bitcoin at a Bitcoin **currency exchange**. These are very similar to other currency exchanges in some ways. One fundamental difference is that many bitcoin currency exchanges allow users to store their money on their platform, acting more like a bank.

Other ways to acquire bitcoin are peer-to-peer marketplaces and bitcoin ATMs. Peer-to-peer marketplace such as btc-otc[2] or Local Bitcoins[3] let users trade bitcoin directly with other users. Established in 2010, btc-otc uses a web of trust, a rating system which allows users to trust or distrust other users before trading with them. LocalBitcoins was started in 2012 and allows users to meet with other users in person or online to trade bitcoin. Bitcoin ATMs, such as Robocoin and Skyhook, are machines that allow users to insert their cash and deposits bitcoin into their wallet (and sometimes visa-versa as well).

The dominance of a few large bitcoin exchanges has centralized activities, which is attractive to attackers. Indeed many bitcoin exchanges have been hacked [55]. On the other hand, many attackers *stealing* bitcoin or earning bitcoin through nefarious purposes cash out their bitcoin on Bitcoin exchanges. Sometimes criminals directly cash out their bitcoin on a public exchange. Other times they run their money through **mixing** services that try to provide anonymity by creating a transaction with many people paying in and

[2] `https://bitcoin-otc.com/`
[3] `https://localbitcoins.com/`

hiding which of the pay outs go to which particular person. These services attempt allow users to hide their transaction trail in the public blockchain.

Bitcoins can be spent online on various marketplaces. OpenBazaar is a peer-to-peer market place started in 2014 which enables users to buy a large variety of goods from sellers. BTC-OTC also facilitates trades of goods as do various Bitcoin forums. Traditional retailers such as `Overstock.com` and `Newegg.com` have also adopted Bitcoin. The Silk Road was a marketplace designed specifically for illicit goods started in 2011. The Silk Road had many privacy protections to try to evade governmental control, including the exclusive use of Bitcoin for payment. The Silk Road was most well known for selling illegal drugs, but they also sold other illicit goods such as ebooks [24]. This particular marketplace was shut down in October 2013 by the FBI. Other similar marketplaces have since tried to take its place [76].

Gambling sites are also popular, accounting for a disproportionate amount of traffic on the network. The leader, SatoshiDice, accounted for about 60% of the Bitcoin network traffic in 2013 [72][52]. SatoshiDice uses the transparency of the Bitcoin network to gain a user's trust. Users bet on a number and the system rolls a die. If the user guesses a number greater than the system, then they win; on expectation, the system takes a 1.9% overhead on all the bets. SatoshiDice is currently blocked in the US.

2.4 Conclusion

Bitcoin is a platform we use to measure cybercrime. Bitcoin is decentralized which enables cybercriminals to transact along with everybody else. The openness of the platform allows us to directly measure transaction history. These two key features of Bitcoin allow us to study basic human behavior through the lens of the seemingly chaotic Bitcoin network.

CHAPTER 3

MEASURING DENIAL-OF-SERVICE ATTACKS IN THE BITCOIN ECOSYSTEM

Distributed denial-of-service (DDoS) attacks are inexpensive to carry out and quite disruptive. In Bitcoin, there are many motivators for launching DDoS attacks: competing services could launch them in order to improve market share, traders could target exchanges to buy or sell at favorable prices [46], and miners outgunned in the rush to increase computational power could try to cripple larger pools in order to increase their odds of solving the hash puzzle first [42]. Despite their apparent frequency, very little is known about the true prevalence of service-denial attacks on Bitcoin. To that end, we carry out an empirical analysis of reports of such attacks made on the popular `bitcointalk.org` discussion forum. We begin in Section 3.1 by outlining how we gather reports of DDoS attacks from public sources. We employ a simple rule-based classifier that distinguishes between the discussion of those experiencing attacks from other messages mentioning DDoS attacks.

We present our analysis in Section 3.2. We identify 142 distinct DDoS attacks taking place between May 2011 and October 2013. We first explain how these attacks vary over time and by category of service affected (e.g., currency exchanges, mining pools, gambling websites). We present evidence that those services that have suffered DDoS attacks are much more likely to now take steps to prevent future DDoS-es. We examine the relationship between a mining pool's size and its susceptibility to attacks, and we look at how attacks relate to the trading volumes and exchange rate at Mt. Gox, the largest currency exchange during this time period. We review related work in Section 3.3, and conclude in Section 3.4.

The research contributions for this chapter are both in the data collection methodology and in the analysis of the gathered data. Our data collection contributions are inferring DDoS attacks from user reports and providing a reliable estimation of attack targets and date. Our analysis contributions are summary statistics of prevalence over time and quantifying impact of attacks (on mining pools and exchanges).

3.1 Methodology

We first set out our approach to data collection in Section 3.1.1. Then we describe and evaluate our method for identifying posts that report DDoS attacks in Section 3.1.2. The collected data and analysis scripts are publicly available for replication purposes at `doi:10.7910/DVN/25541`.

3.1.1 Data Collection

Identifying when a denial-of-service attack has taken place can be difficult. If we knew in advance the websites to monitor, we could run a regular script that attempts to visit the websites. However, simply because we can connect to a website does not mean that others are being blocked. Furthermore, some services (e.g., mining pools) are not run as websites, so non-standardized means of connecting would be required. Finally, it would be desirable to peer back further into the past to check for historical reports of DDoS attacks.

To that end, we decided to inspect reports of DDoS attacks posted to the popular `bitcointalk.org` forum. Using the Google Custom Search API, we identified all posts including the term "ddos" on the website appearing between February 2011 and October 2013. Because the Google API limits the results to the top 100 results, we issued queries restricted to week-long intervals. In only 3 weeks (during April and May 2013) did the API return the maximum 100 results. In those cases we shortened the time interval further to ensure that we obtained all results including "ddos".

In total, we identified 2 940 distinct pages on `bitcointalk.org` that mentioned "ddos". However, many duplicates existed in these pages, such as when a single thread spans multiple pages. Consequently, we identified 1 355 distinct pages comprised of the first page of the thread. For each page, we then fetched a local copy of the page and automatically extracted the thread title, plus the first post's text, URLs, poster handle and date. We also extracted the forum title. Not all posts actually described DDoS attacks, however. In Section 3.1.2 we explain how to distinguish between discussion of perceived DDoS attacks and other DDoS-related threads.

We collected additional information to complement the information gathered on DDoS reports. For instance, we fetched a directory of 1 240 online services supporting

Bitcoin [13] and 32 mining pools [12]. We extracted category and subcategory information for these services from parsing the directory. We threw out any services that did not resolve after an automatic and manual check.

Subsequently, we identified the use of anti-DDoS providers by resolving the websites of all known Bitcoin services and comparing against known IP ranges for CloudFlare [25], Incapsula [35], and Amazon Web Services [2]. CloudFlare and Incapsula are content distribution networks (CDNs), whereas Amazon hosts material. All three are identifiable by IP range. For services not resolving to these networks, we looked up their AS number using the IP address. We did not find any other content distribution networks serving more than two Bitcoin services. Therefore, we are confident we found all significant network-based anti-DDoS protections. Other forms of protection, such as DDoS detection built in to security appliances, could not be identified and are beyond this chapter's scope.

Finally, we identified historical market share of mining pools from 22 Internet Archive snapshots of `http://blockchain.info/pools` dating to October 2011.

3.1.2 Classification of Posts Describing Attacks

As noted above, many of the posts mentioning "ddos" do not actually describe experiences with denial-of-service attacks. Instead, users discussed ways to defeat DDoS attacks, posted advertisements for services with built-in protections against attacks, and speculated on the motivations behind prior attacks.

We built a simple word-based classifier to identify just those threads describing DDoS attacks currently in progress. Of course, we cannot confirm that what the posters describe is actually a DDoS attack rather than a server overloaded with demand. Nonetheless, user reports do provide a useful indication of when such attacks most likely occur. We flagged all posts with the following words and phrases in the title as DDoS attacks: "unreachable", "offline", "online", "down","flooding","attack","ddos","unavailable","blocking", and "connect". Any posts including the words "anti-ddos" or "vote" in the title were marked as not describing attacks.

To evaluate the classifier's accuracy, we compared it against a manually labeled set of 207 posts. The results are given in Table 3.1. Overall accuracy is 75%. The false negative

	Actual	
	DDoS	Not DDoS
Predicted DDoS	42	36
Predicted Not DDoS	15	114
Precision 54%, Recall 74%, Accuracy 75%		

Table 3.1: Confusion matrix plus precision, recall and accuracy measures for the word-based classifier.

rate is modest (26%), but false positives are problematic. Thus the classifier does a pretty good job at finding DDoS reports, whereas many posts flagged as DDoS in fact are not.

Consequently, we manually inspected the 362 posts identified by the classifier as describing attacks from the full dataset. We found that 200 posts actually described attacks. We use these posts in the analysis that follows below. Based on the observed recall rates, we expect that there are around 70 more posts describing attacks not included in our analysis. However, we defer improving the classifier further and identifying those posts to future work.

There is one final subtlety in the data collection that bears mentioning. Sometimes multiple posts discuss the same DDoS event. To account for that, we define distinct DDoS attacks as any post mentioning a service on a given day. For instance, if three posts describe an attack on Mt. Gox on April 26, 2013, we count that as a single attack. If however, a single post mentions a DDoS on three different services, we count that as three attacks. Using this approach, the 200 posts correspond to 142 distinct DDoS attacks.

3.2 Empirical Analysis

We first discuss how DDoS attack targets have changed over time in Section 3.2.1, along with an examination of which service categories are targeted more and less often. We then study attacks on mining pools in Section 3.2.3, followed by attacks on currency exchanges in Section 3.2.3.

3.2.1 DDoS Attacks over Time and by Target

We begin by examining how reports of DDoS attacks on Bitcoin services have evolved over time. Figure 3.1 plots the number of reported DDoS attacks per month since May 2011.

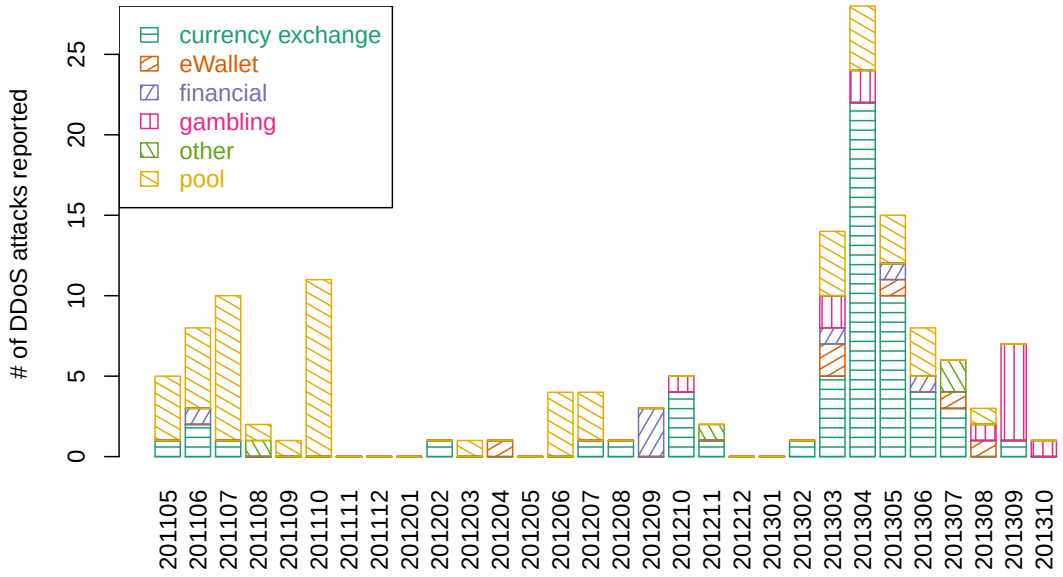

Figure 3.1: Reported DDoS attacks over time, split up by category of targeted service.

We can see that the number and target of reported attacks varies greatly over time. Initially, in the second half of 2011, most DDoS reports concerned mining pools. Then there were very few reported attacks of any kind during the first half of 2012. During the second half of 2012, DDoS attacks picked up again, initially targeting pools, but more frequently targeting currency exchanges and other websites. During 2013, attacks on pools continued, but they were joined by DDoS on gambling websites, online wallets, and currency exchanges. Attacks on currency exchanges dominated the totals from March–June 2013, coinciding with rising exchange rates and unprecedented interest in Bitcoin. While we expect that some of these reported DDoSes were in fact triggered by customer demand, it is nonetheless interesting to see the rise in reported abuses. Finally, DDoS on exchanges fell sharply in August. However, Bitcoin-based gambling websites experienced a surge of DDoS activity in its place.

Figure 3.2 (left) shows how DDoS attacks stack up by category over all time. The most targeted service category is currency exchanges (41%), followed closely by mining pools (38%). These were trailed by gambling (9%), finance (5%), and online wallets (4%). DDoS attacks on other services accounted for 3% of the total.

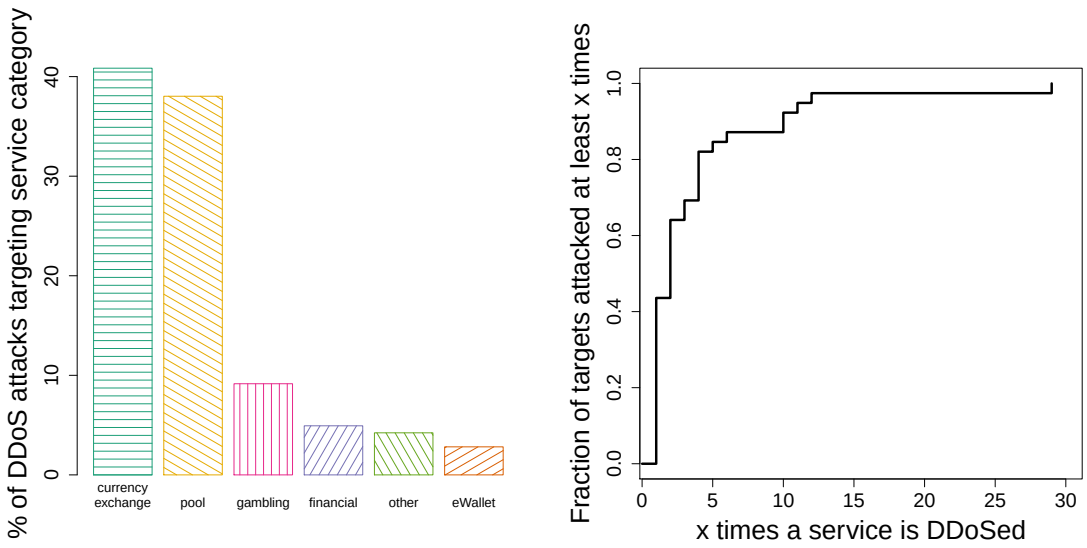

Figure 3.2: Percentage of DDoS attacks targeting each major category (left); cumulative distribution function of the number of attacks targeting each service (right).

While some services are targeted only once by DDoS attacks, others are repeatedly hit by them. Figure 3.2 (right) plots a CDF of the number of times a service is DDoSed. Out of the services targeted by a DDoS attack, 44% are only attacked once, while 15% are attacked on at least five occasions. One service, the Mt. Gox currency exchange, suffered 29 DDoS attacks on different days. We study the timing of attacks on Mt. Gox in greater detail in Section 3.2.3 below.

Table 3.2 shows another way to look at the breakdown of DDoS attacks by category. The first column lists the number of services for each category that are still operational (i.e., their listed websites resolve), followed by the percentage of services in each category that have suffered DDoS attacks. Overall, 7.3% of services actually experienced a DDoS attack. The variation across categories is substantial: 27% of pools have experienced DDoS attacks compared to just 0.7% of shops selling physical products. Currency exchanges, mining pools, financial services and online wallets are targeted more frequently than other categories. These differences compared to the average are statistically significant with 95% confidence according to a χ^2 test. One surprise is that Bitcoin payment systems are not targeted by DDoS attacks any more than average.

Given the very real threat of DDoS attacks on Bitcoin services, it is not surprising that many services take steps to defend against these attacks. Moving over to the next

Category	#	Suffer DDoS %	Suffer DDoS Sig.?	Use AD %	Use AD Sig.?	AD + DDoS	AD Only	DDoS Only
Material/physical products	295	**0.7**	−	**10.5**	−	2	29	0
Internet & mobile services	225	1.8		16.9		0	38	4
Online products	185	3.8		14.6		3	24	4
Professional services	137	0		10.2		0	14	0
Currency exchanges	119	**10.9**	+	**36.1**	+	10	33	3
Travel/tourism/leisure	78	0		10.3		0	8	0
Commerce & community	71	1.4		12.7		1	8	0
Getting started	31	0		12.9		0	4	0
Financial	26	**15.4**	+	26.9		1	6	3
Pool	41	**26.8**	+	**34.1**	+	5	9	6
Bitcoin eWallets	17	**17.6**	+	35.3		2	4	1
Bitcoin payment systems	11	9.1		18.2		1	1	0
Average		*7.3*		*19.9*				

Table 3.2: Prevalence of DoS attacks and anti-DDoS (AD) uptake by service category.

column grouping, we report for each category the percentage of services that use anti-DDoS services (either Amazon, Incapsula, or CloudFlare). Overall, around 20% of online Bitcoin services have anti-DDoS protection.

Anti-DDoS protection is more popular in some categories than others. Around one third of exchanges and pools have anti-DDoS protection. This difference in proportion (compared to the 20% average) is statistically significant according to a χ^2 test. Shops selling material and physical products and accepting Bitcoin were substantially less likely to be protected from DDoS attacks – only 10.5% rely on these services. Financial firms and online wallets also frequently employ anti-DDoS protection, but the differences are not statistically significant.

Finally, the last grouping in Table 3.2 shows for each category how many services have anti-DDoS protection and have been attacked, how many have anti-DDoS and have not been attacked, and how many have been DDoSed but do not have anti-DDoS protection from Amazon, Incapsula, or CloudFlare. It is noteworthy that across categories it is far more common to have anti-DDoS protection than it is to have actually experienced a DDoS attack. Even in categories where no service has experienced a DDoS attack (e.g., travel and professional services), there is substantial uptake of anti-DDoS protection.

We can also answer a related question: Are services that have experienced DDoS in the past more likely to get anti-DDoS protection afterwards? Table 3.3 helps to answer the question for all services.

	Use Anti-DDoS		No Anti-DDoS	
	#	%	#	%
Suffered DDoS	25	54%	21	46%
No DDoS	178	15%	1 012	85%

Table 3.3: Contingency table comparing the uptake of anti-DDoS protection based on whether or not the service has experienced DDoS attacks.

Of the 46 distinct services that have experienced DDoS attacks, more than half now have anti-DDoS protection. It is impossible to tell whether or not they had such service at the time of attack. Among services that have not yet experienced a DDoS attack, only 15% have anti-DDoS protection. The difference in proportion (15% vs. 54%) is statistically significant, according to a χ^2 test ($p \ll 0.0001$ with χ^2 value of 47.232). We conclude that providers are much more likely to obtain anti-DDoS protection if they are targeted by DDoS attacks.

3.2.2 DDoS Attacks on Mining Pools

Given that mining pools are frequently targeted by DDoS attacks, we now study them in greater detail. We first investigate whether the size of a mining pool affects its chances for being DDoSed. Mining pool size constantly changes, sometimes in response to DDoS attacks. Hence, we needed a historical record of mining pool market shares. Using the Internet Archive, we accessed 22 historical copies of `blockchain.info/pools` that breaks down hashrate by pool. We deem a pool to be "big" if it is observed to have at least a 5% share of the hashrate during two or more observations. All other pools are deemed "small".

Table 3.4 shows how the incidence of DDoS attacks vary by pool size. 5 out of 8 big pools (63%) have suffered DDoS attacks, compared to just 7 out of 41 small pools (17%). These percentage differences are statistically significant, according to a χ^2-test with a p-value of 0.022. Why would large pools be targeted for DDoS attacks more than small

pools? Attackers gain more by targeting large pools, since taking one out can substantially increase the odds of winning the round.

	Small Pools		Big Pools	
	#	%	#	%
Suffered DDoS	7	17.1%	5	62.5%
No DDoS	34	82.9%	3	37.5%

Table 3.4: Contingency table comparing the size of a mining pool to whether or not the pool has experienced DDoS attacks.

Figure 3.3 examines the historical hashrate-based market share for six of the larger pools. DDoS reports are indicated by the vertical dashed lines. Some pools seem unfazed by DDoS attacks (e.g., Slush's Pool, Eclipse MC, and Eligius). BTC Guild actually increased its market share following a DDoS attack in mid-2012. However, substantial declines followed a later attack in mid-2013. Furthermore, one can see that sometimes DDoS attacks target multiple pools simultaneously. For example, DeepBit was targeted by attacks at the same time as BTC Guild and Eclipse MC. DeepBit's share of the hashrate tumbled, while it appears that Eclipse MC and BTC Guild benefited as a result. Later attacks in 2013 on BTC Guild and Eclipse MC reduced their own shares, with Eligius benefiting this time even though it too had been hit by DDoS attacks.

Based on this analysis, we reject the notion that DDoS attacks always trigger a decline in market share for affected mining pools. Instead, we see that DDoS attacks often precede shakeups in pool market share. However, at this point we cannot reliably predict who the winners and losers will be as a result.

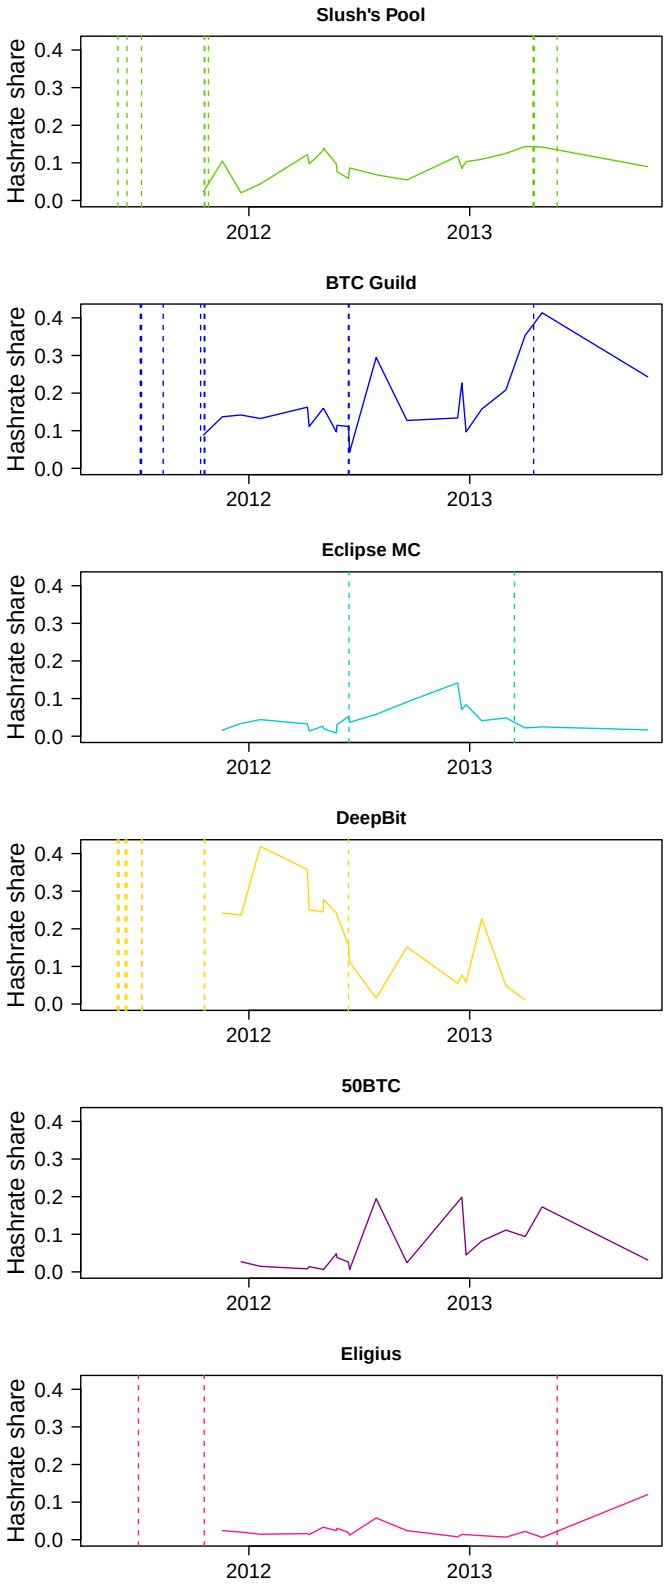

Figure 3.3: Mining pool hashrate market share (solid line) over time, compared to timing of DDoS attacks (dashed lines).

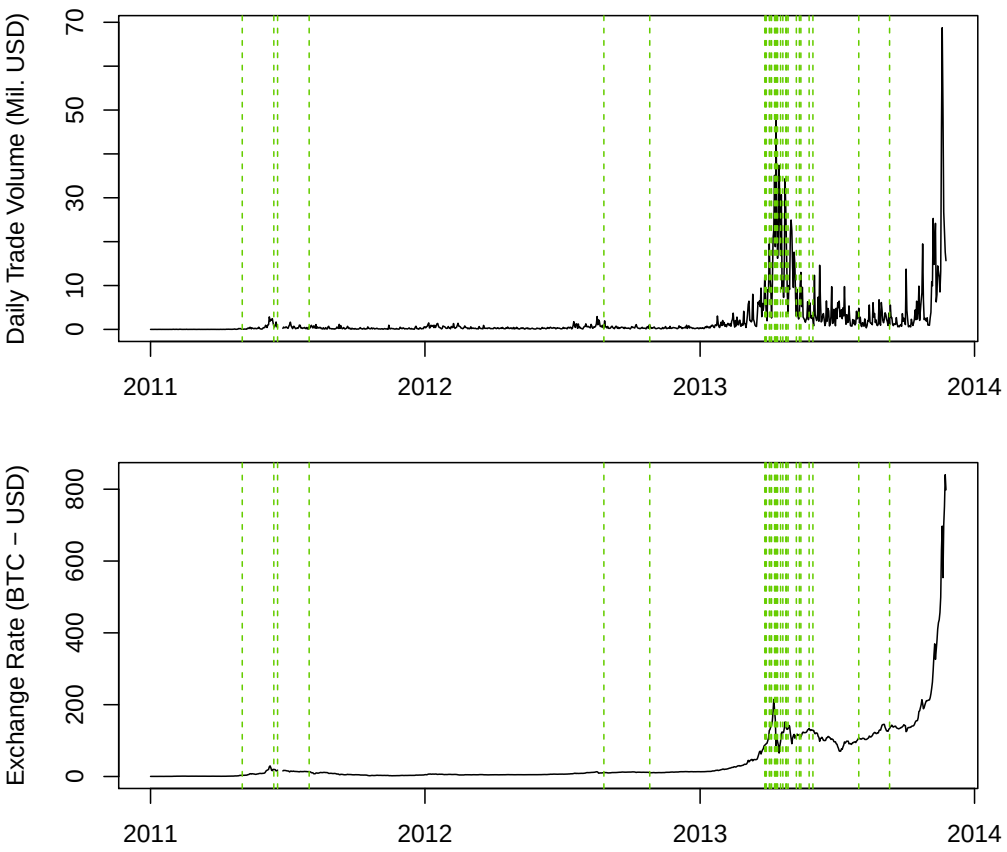

Figure 3.4: Daily trade volumes (top) and USD-BTC exchange rate (bottom) at Mt. Gox. Dashed green lines indicate when DDoS attacks on Mt. Gox were reported.

3.2.3 DDoS Attacks on Currency Exchanges

Currency exchanges are the most frequent target of DDoS attacks. We defer to future work a more detailed analysis of how DDoS attacks affect exchange operations in general. Instead, we take a closer look at attacks targeting Mt. Gox, the largest currency exchange during the time of our study and most frequent attack target.

Figure 3.4 plots trade volumes and USD-BTC exchange rates at Mt. Gox, along with DDoS attacks as dashed green lines. We can see that Gox suffered some DDoS attacks in 2011 shortly after experiencing unprecedented peaks in trading volume. (It can be difficult to see on the current graph since trading has exploded so much since early 2013.) Note that these early attacks, plus one in late 2012, came shortly after a fall from a new peak in the

Δ Transaction Vol.	**# of Attacks**	**% Attacks**	**% Change** (median)
Increase	12	41.4%	53.3%
Decrease	17	58.6%	34.2%

Table 3.5: Changes in transaction volume on Mt. Gox after a DDoS attack.

exchange rate. This behavior is consistent with the modus operandi of blocking exchanges in order to slow down a panicked sell-off.

When Bitcoin's exchange rate shot up in spring 2013, trading volume also soared to unprecedented heights. Dozens of DDoS claims were made in April and May 2013, eventually subsiding. Two more reports were made later in 2013, but these were one-off reports rather than a chorus as in the spring. Doubtless, some reports were caused by surging demand rather than by a botnet. The blogger *organofcorti* observed a drop in trading volume at Mt. Gox after Mt. Gox's Dwolla account was seized in spring 2013 [65], which could explain some of the reported attacks in times of lower trading volume.

In the (slight) majority of cases, we observe a decrease in transaction volume in the week following a DDoS attack compared to the week prior, as seen in Table 3.5. We also notice that the median size of the transaction volume change is greater when the transaction volume increases. Figure 3.5 show this trend over time. We observe that the increases and decreases tend to be clustered together in time. This suggests that certain DDoS attack campaigns can be recovered from quickly while others cannot.

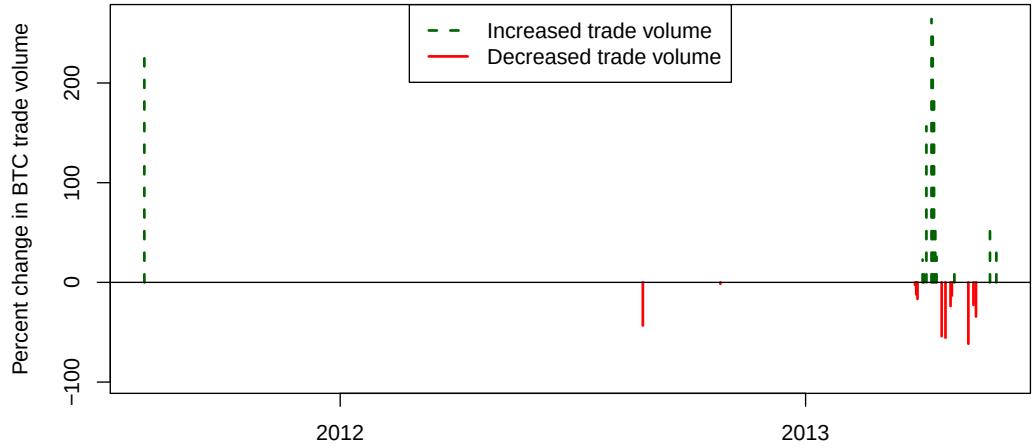

Figure 3.5: Changes in transaction volume on Mt. Gox after a DDoS attack over time.

3.3 Related Work

This work continues in the vein of previous measurement studies, in that it collects publicly-available data to better explain the a phenomena in the bitcoin ecosystem. Our work gathers this data from `bitcointalk.org` to infer DDoS attacks, in line with other work that generates situational awareness of natural disasters from tweets [78] and work that determines whether a service is up or down from tweets [59]. We are not aware of any prior work measuring the occurrence of DDoS attacks on Bitcoin. There has been a large-scale study that measures how prevalent DDoS attacks are in the context of websites and blogs [82]. But there are several reasons why we believe Bitcoin DDoS attacks are worth studying on their own. First, there are unique incentives at play that reward DDoS attacks, such as traders who benefit by blocking others' transactions. Second, Bitcoin's unregulated environment has facilitated criminality in pursuit of profits, with DDoS an attractive tool for unscrupulous operators. Indeed, the most closely related work to our own is that of Johnson et al., who present a game-theoretic model of the trade-offs mining pools face between investing in upgrades to computing infrastructure and engaging in DDoS attacks [42]. Their model nicely complements the empirical work undertaken in this chapter.

Work that has built upon ours has occurred after this original paper was published. Laszka et al. expand their game-theoretic model of Bitcoin mining pools the look at both the short term and long term effects of attacks [45]. They find two long term equilibria – one with no attacks and another where one mining pool is attacked by the others. Ritter et al. use Twitter to infer security events such as data breaches and DDoS attacks [67]. However, their approach leaves many false positives[1]. Feder et al. expand upon our data, particularly on the currency exchange, Mt. Gox [32]. They found that DDoS attacks caused fewer big trades in the day following the attack.

3.4 Conclusion

We have presented an empirical study of DDoS attacks targeting a wide range of operators in the Bitcoin ecosystem. Using posts to the popular `bitcointalk.org` forum,

[1] `https://web.archive.org/web/20160315132927/http://kb1.cse.ohio-state.edu:8123/events/ddos`

we identify and analyze 142 distinct DDoS attacks. We find that 7.4% of Bitcoin-related services have experienced DDoS attacks. Currency exchanges are targeted most often, followed by mining pools, gambling operators, financial service providers, and online wallet operators. Attack frequency is highly variable: pools were targeted most often back in 2011, followed by a wave of attacks targeting currency-exchanges in Spring 2013. DDoS on gambling operators, nonexistent until December 2012, have picked up considerably in the latter part of 2013.

We also carried out preliminary analysis into the effects of DDoS attacks on mining pools and currency exchanges. One striking finding is that over 60% of large mining pools have been DDoSed, compared to just 17% of small ones. This suggests that the large pools are big targets for unscrupulous miners hoping to increase their odds of winning freshly minted Bitcoins.

Our results indicate that Bitcoin DDoS attacks merit further investigation. Nonetheless, the findings often raise more questions than they answer. To get those answers, a richer and more robust dataset is needed. Our dataset is based on circumstantial evidence of DDoS attacks reported on a single, albeit popular, web forum. Such reports do not constitute definitive evidence that a DDoS has taken place. Future investigations could corroborate reports with supplementary evidence, such as directly measuring inaccessibility from probes and incorporating reports from additional sources besides `bitcointalk.org`.

CHAPTER 4

MEASURING THE PROFITS OF BITCOIN SCAMS

As more people have been drawn to Bitcoin, frequently out of a desire to get rich quickly, more hucksters have appeared to take advantage of these eager new targets. Because Bitcoin is so new, the newly emerging scams are frequently poorly understood. The goal of this chapter is to systematically investigate different types of Bitcoin scams, explain how they work, and measure their prevalence. It is hoped that by understanding how these scams work we will identify ways to arrest their rise.

To that end, we identify four types of scams currently plaguing Bitcoin: high-yield investment programs, mining investment scams, scam wallet services and scam exchanges. Using reports obtained from discussion forums and tracking websites, we study 41 distinct scams operational between 2011 and 2014 where we could find the associated Bitcoin address(es). So while the study is by no means comprehensive, we are able to analyze the blockchain and provide a lower bound estimate of the prevalence and criminal profits associated with these scams.

We find that $11 million worth of bitcoin has been contributed to the scams, and that at most $4 million has been returned to the victims. For the HYIPs and mining scams, we estimate that about 13 000 victims contributed funds. We also show that the most successful scams draw the vast majority of their revenue from a few victims, presenting an opportunity for law enforcement to track down and prosecute the scammers.

Section 4.1 describes the methodology for identifying scams, as well as how we examine the blockchain to identify payments into and out of scams. Section 4.2 reports on high-yield investment programs (HYIPs), online Ponzi schemes where existing investors are paid lucrative returns from the contributions of new investors. Section 4.3 examines mining-investment scams, which is a form of advanced-fee fraud that exploits people's interest in Bitcoin mining by promising a way to profitably mine without making large up-front investments in expensive hardware. Sections 4.4 and 4.5 cover scam wallets and exchanges, respectively. Here, the scammers provide sought-after services such as mixing at a seemingly

affordable price, only to steal incoming transfers from customers. Section 4.6 compares the different scam categories and considers what the appropriate response, if any, should be from the Bitcoin community and policymakers. Finally, we review related work in Section 4.7 and conclude in Section 4.8.

The research contributions for this chapter are both in the data collection methodology and in the analysis of the gathered data. Our data collection contributions are finding candidate scams using aggregated defender data, confirming said scams through manual inspection, and measuring money in/out of scams using data from the public Bitcoin blockchain. Our data analysis contributions are constructing a taxonomy of scam categories, documenting each categories' prevalence and revenues, and providing a first longitudinal analysis of HYIP cash flows and victim losses.

4.1 Methodology for Identifying Scams and Associated Transactions

We compile a list of 349 distinct candidate scams from an aggregated thread on `bitcointalk.org`[1], a blacklist of suspected fraudulent services maintained at `http://www.badbitcoin.org/thebadlist/index.htm`, and a website tracking Bitcoin-based HYIPs called `cryptohyips.com`[2]. We manually inspected all services on the list to identify only those operations established with fraudulent intent. For instance, we exclude Hashfast, a mining company that recently filed for Chapter 11 bankruptcy protection, as well as losses from Mt. Gox, a bitcoin exchange that failed. We also removed a number of false positives with no clear connection to cryptocurrencies, such as `unclechiens.com` (a Chinese restaurant in Texas). In total, this sheds 26% of our candidate list.

We also exclude from consideration all efforts beyond the purview of this chapter, such as phishing websites, malware websites, and pay-for-click websites. We are left with 192 scams to investigate further, 55% of the candidates. We categorize each scam's type by inspecting the website through the Internet Archive (since many scams have since disappeared) and targeted Google searches on the domain.

[1] `https://bitcointalk.org/index.php?topic=576337`

[2] Data and analysis scripts are publicly available at `doi:10.7910/DVN/28561`.

We next seek out associated Bitcoin addresses for each scam using threads on `bitcointalk.org`, `reddit.com/r/bitcoin`, and named addresses and transactions on `blockchain.info`. We exclude any "dual-use" addresses that are also used for other purposes. In all, we find usable Bitcoin addresses for 20% of the scams.

The next goal is to identify payments made into and out of the scam. To that end, we download the Bitcoin blockchain using the Bitcoin Core client on August 25, 2014. Using znort987's Bitcoin blockparser [81] we query for all transactions involving our set of scam addresses. This gives us traffic levels for incoming transactions to each scam. We then take a complete SQL dump of the Bitcoin blockchain and query for all the transactions where the input or output address match one of our scam addresses. This gives us the Bitcoin addresses of the victims as well as the outgoing transactions from the scam. To separate out transactions made by scammers, we omit all outgoing transactions going to other addresses associated with the same scam. We also omit transactions occurring before and after the first incoming transaction to the scam.

One challenge for researchers inspecting a blockchain is dealing with multiple sources and destinations in transactions. Figure 4.1 demonstrates the three cases where these transactions arise. We deal with multiple source–single destination transactions (Figure 4.1 (left)) as follows. If the destination is a scam address and the source addresses are not also identified as being part of the scam, we group the source addresses together as a single victim.[3] In general, two addresses are assigned to the same *address group* if they ever paid into the same scam during the same transaction. For multiple-source transactions involving a scam address, we only count the scam address's contribution towards the total payout from the scam.

For transactions with a single source and multiple destinations (Figure 4.1 (center)), we attribute only the source amount to the scam. For instance, suppose Fig is a victim address and Honeydew is the scam. Even though Fig pays 0.4 BTC, we tally only the 0.32 BTC transferred to Honeydew as part of the scam's total incoming payments.

[3]Note that we deliberately make no attempt to deanonymize the actual victims beyond identifying that the addresses participated in the scam.

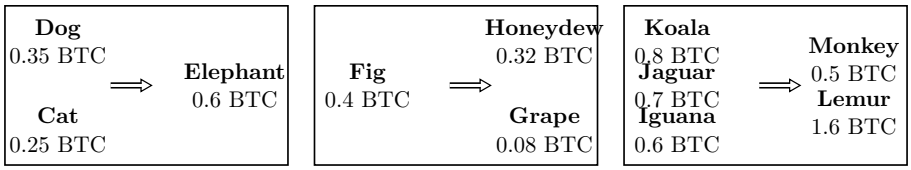

Figure 4.1: Multiple-address transactions in Bitcoin.

	Victim pay in	Payout to victim	Payout to scammer
HYIPs	✓	✓	derived
Mining scams	✓	derived	derived
Scam wallets			✓
Exchange scams			✓

Table 4.1: For each scam category, we report whether we can directly observe transactions corresponding to what victims pay into scams, what is paid out to victims, and what is paid out to the scammer (indicated by a ✓).

With multiple sources and destinations (Figure 4.1 (right)), we assign the amount paid in or out of the scam to the corresponding address group. For example, suppose Lemur is the scam address. Here, the victim group Koala–Jaguar–Iguana contributes 1.6 BTC to Lemur's scam. While in theory services such as CoinJoin [49] could account for many such transactions, in practice we do not observe very many transactions of this type.

Finally, we note that when identifying victim groups we could mistakenly identify online web wallets that pay out multiple users from the same address as a single address group. To check for this, we inspected all multiple-destination transactions whose source address appeared more than three times. In all cases, we did not find that the source addresses corresponded to web wallets. One potential explanation for this is that many scams prohibit using web wallets as a method of payment.

In addition to gathering data directly from the blockchain, we also analyze scams that raise funds through selling shares. We gather the share holdings from BitFunder and cross list that with cost of the shares from announcements on `bitcointalk.org`. For each scam, we omit the top holding who we verify is the scammer in all instances.

Ideally, we would analyze payments from victims into scams, payments back to victims, and scammer profits. For some scams, we can observe all such payments, whereas for

others we can only observe certain categories. Table 4.1 summarizes the types of observable transactions for each scam type. Full details are given in subsequent sections.

Finally, due to high volatility of the bitcoin exchange rate, it makes sense to also report scam revenues in terms of its dollar equivalent. In order to convert BTC to USD, we gathered the daily closing USD-BTC exchange rate from the four highest-volume USD exchanges during the period of our study (Mt. Gox, Bitstamp, Bitfinex and BTC-E), as reported to `http://www.bitcoincharts.com`. We then converted any transactions into USD using the average exchange rate on the day of the transaction.

4.2 High Yield Investment Programs

Moore et al. first described high-yield investment programs (HYIPs) in [56]. HYIPs are online Ponzi schemes where people are promised outlandish interest rates on deposits (e.g., 1–2% interest per day). Unsurprisingly, the schemes eventually collapse, and they are replaced by new programs often run by the same criminals. Moore et al. observed that these HYIP schemes relied on virtual currencies such as Liberty Reserve, Perfect Money, and EuroGoldCash for deposits and withdrawals. The centralized nature of these particular currencies has left them vulnerable to countermeasures by law enforcement. For example, Liberty Reserve was taken down by the US government in 2013 for money-laundering activities. In response, some programs have begun accepting decentralized digital currencies such as Bitcoin and Litecoin. Furthermore, most HYIPs directly advertise Bitcoin addresses in order to accept incoming payments, as opposed to using a payment processor such as BitPay or Coinbase.

We observe a number of different types of HYIPs that accept Bitcoin: HYIPs that stay in the traditional HYIP ecosystem; HYIPs that bridge the traditional HYIP ecosystem and the Bitcoin community; and HYIPs that originate in the Bitcoin ecosphere.

4.2.1 Traditional HYIPs

We first investigated the extent to which traditional HYIPs have begun to embrace Bitcoin. To our surprise, we found that most HYIPs do not accept bitcoin as payment. We believe the reason why is that the leading kit for developing HYIP websites, Gold Coders,

does not support payments in Bitcoin or other cryptocurrencies. Neisius and Clayton analyzed the HYIP ecosystem, and they estimated that between 50–80% of HYIP websites they observed used the Gold Coders kit [64].

When we observed several "aggregator" websites that track HYIPs, we found some traditional HYIPs that accept BTC or LTC. We then inspected HYIPs with a publicly-accessible incoming address but had never been mentioned on `bitcointalk.org`. All of these programs had insignificant transaction volume. Based on these findings, we do not consider traditional HYIPs further in our analysis.

4.2.2 Bridge HYIPs

	Bridge HYIPs	Bitcoin-only HYIPs
# Scams	9	23
Median lifetime (days)	125	37
# still operational	1	0
Victim pay in		
# address groups (total)	9 410	3 442
# address groups (median)	298	157
Amount paid (total)	$6 456 593	$842 909
Payout to victim		
Amount paid (total)	$3 464 476	$802 655
Payout to scammer		
Amount paid (total)	$2 992 117	$40 254

Table 4.2: Summary statistics for HYIPs.

Some scams first appear in the traditional HYIP ecosystem before being brought over to the Bitcoin world through posts on `bitcointalk.org`. In these cases we frequently find a high volume of BTC transactions. For example, Leancy claimed to have received over $5M in investments[4] from a variety of currencies. From observing payments into its Bitcoin address, we estimate $1 674 270 came from bitcoin deposits.

Overall, we observe a total of nine such scams that brought in 12 622 BTC ($6.5M) from September 2, 2013 through September 9, 2014. Table 4.2 reports key summary statis-

[4] `https://web.archive.org/web/20140322111925/https://leancy.com/`

tics for the nine bridge HYIPs observed. Median lifetime of the bridge HYIPs is 125 days, with one HYIP still in operation at the time of writing.

The $6.5M in contributions came from 9 410 distinct address groups, which provides an upper bound for the number of victims contributing to these scams. The scams in turn paid at most $3.5M back to the victims, leaving $3M in profit to the operators. It is likely that at least some of the $3.5M in payouts went to addresses controlled by scammers, so we expect the actual profit rate to be much higher.

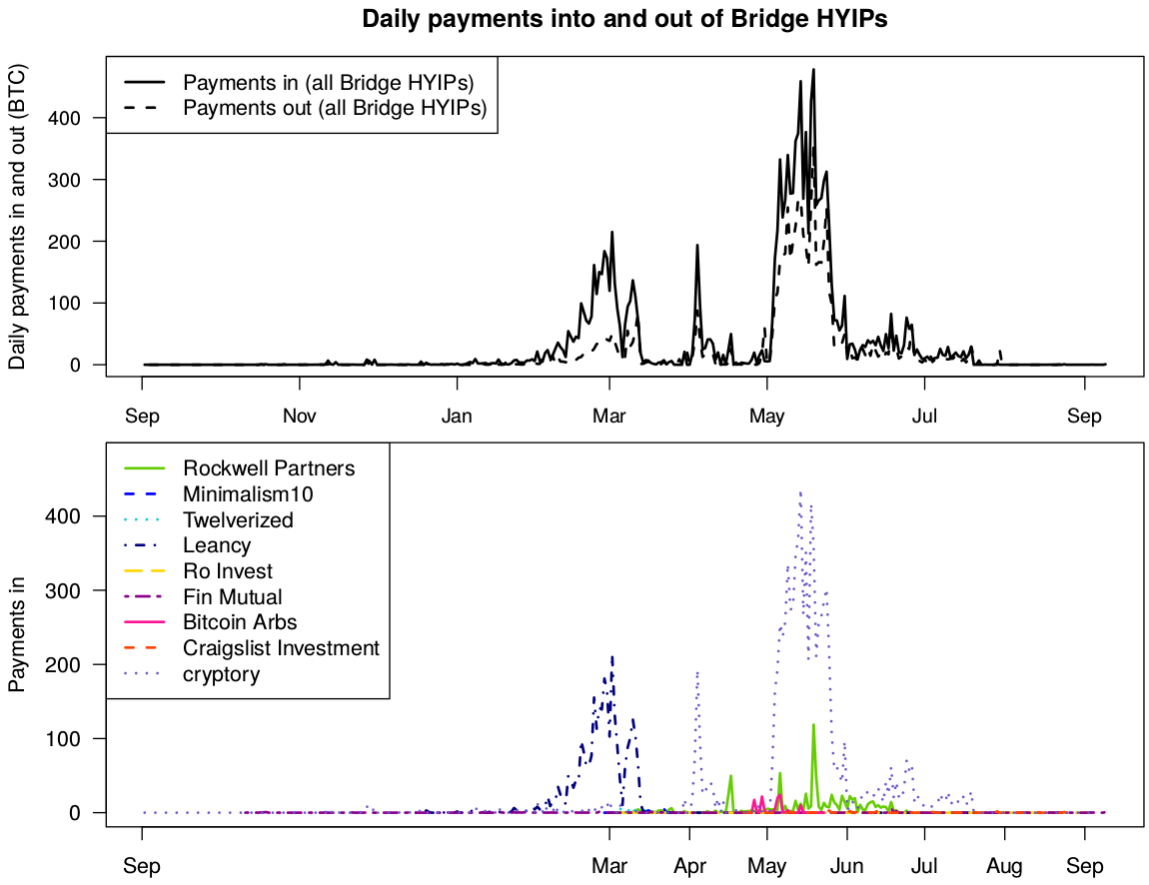

Figure 4.2: Top: Daily volume of all payments into and out of Bridge HYIPs wallet incoming transactions. Bottom: daily volume of incoming payments split by HYIP.

These summary statistics obscure the details of how individual scams performed over time. Figure 4.2 (top) plots the aggregate payments into and out of the nine bridge HYIPs. We can see that, in aggregate, the payments flowing into the scams always keep pace with the payments flowing out. We also see huge spikes in the money flowing in at different points throughout the period, with nearly all of the activity taking place in 2014. Figure 4.2

(bottom) breaks out the incoming payments to the associated scams. We can see that the first big spike is due to the rise of Leancy, the second Cryptory, the third Rockwell Partners and the fourth Cryptory (with a small contribution form Rockwell Partners). Hence the overall burstiness observed in the scam contributions can be attributed to different scams receiving a surge of investment before falling rapidly.

Figure 4.3 compares the transactions in and out for the top 8 performing bridge HYIPs. The graphs are presented in decreasing order of scam size, and the graph also includes a green dotted line indicating the day the scam first appeared on the `bitcointalk.org` forum.

For example, for Leancy (top right) we see the first BTC transaction on December 16, 2013, but the volume picks way up on February 4, 2014 when a user, `LeancyBTC`, posted an advertisement for the scheme in the Bitcoin forums[5]. Most reports precede spikes in investment, though the jump is not always as immediate as in the case of `LeancyBTC`'s post.

The other key conclusion that can be drawn from these graphs is that the most successful scams manage to pay out far less than they take in, and they do so consistently over time. In theory, Ponzi schemes need not collapse until withdrawal requests overwhelm the cash reserves of the scammer. In practice, for Leancy and Cryptory, the scheme stopped paying out as soon as the funds stopped flowing in. These operators could have kept up the appearance of legitimacy by honoring withdrawal requests after new deposits stopped, but they chose not to. Instead, they found it more profitable to simply disappear once the deposits did.

For the less successful scams (bottom of graph), the outgoing payments often exceed the incoming payments. Hence, in these cases it does appear that the scammer gave up once the scam failed to take off, even after honoring withdrawal requests that exceeding available deposits.

[5] `https://bitcointalk.org/index.php?topic=448250`

4.2.3 Bitcoin-only HYIPs

In addition to HYIPs that happen to accept bitcoin, many shady operators have set up Ponzi schemes using bitcoin as a method of payment. We term these frauds *Bitcoin-only HYIPs* because they operate like HYIPs even if they do not share the same heritage as traditional HYIPs.

The premise behind Bitcoin-only HYIPs varies considerably. Some purport to be legitimate investment vehicles. The biggest of these is Bitcoin Savings and Trust (first

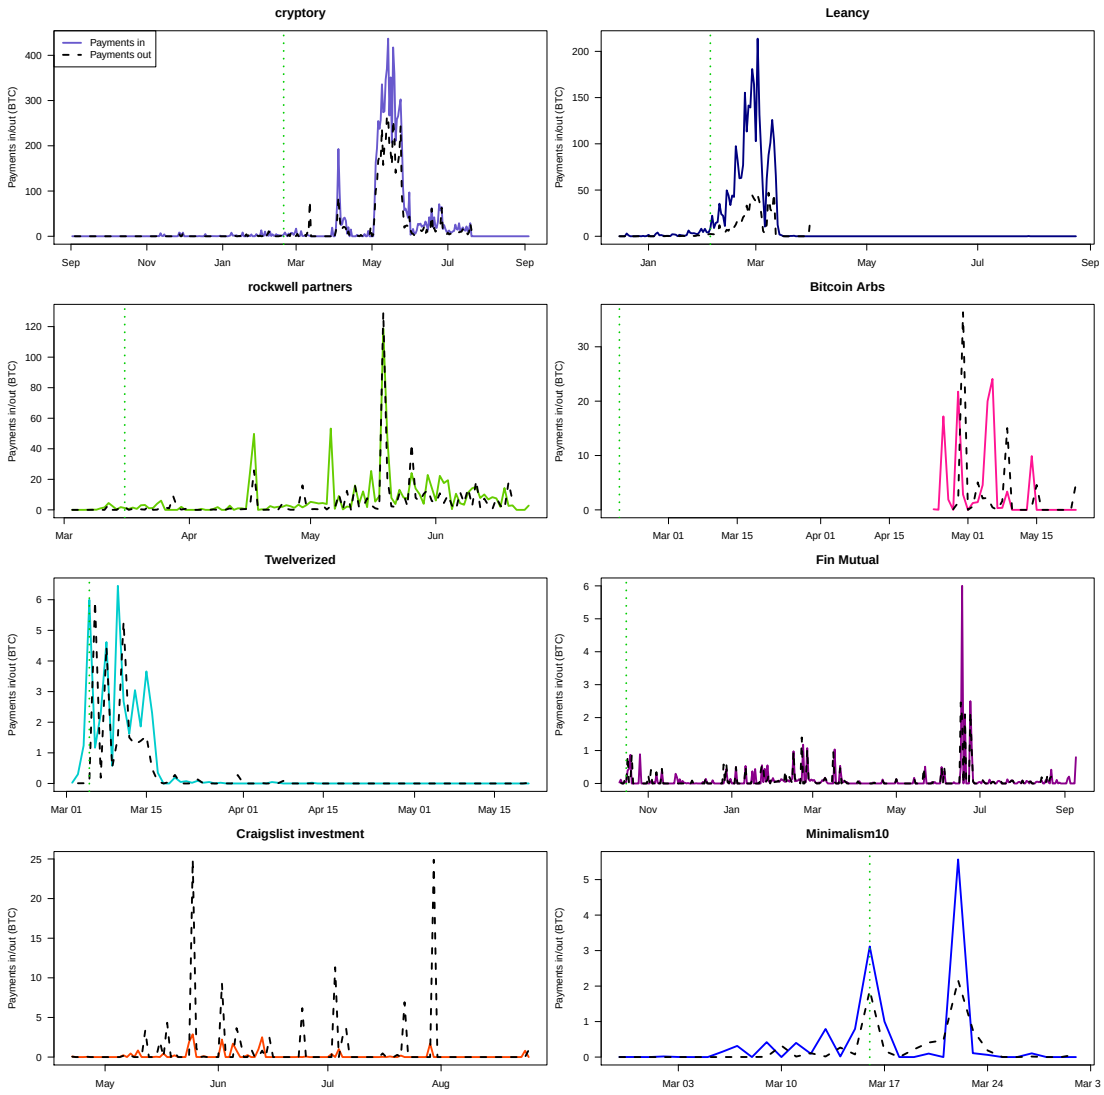

Figure 4.3: Daily volume of payments into and out of Bridge HYIPs, sorted by total payments received. The green dotted line indicates when the scam is first promoted on `bitcointalk.org`.

launched under the name "First Pirate Savings and Trust") which allegedly raised 4.5 million USD [74]. (Unfortunately, since the address used for this Ponzi was also used for a legitimate Bitcoin marketplace, we do not include it in our analysis. Reported estimates in volume vary greatly[6].) Others purport to be online Bitcoin wallets offering an outlandishly high daily rate of return on the money kept in the wallet. While these schemes are fraudulent by design, they lure in unsuspecting, naïve victims as well as those fully aware that they are investing in a Ponzi scheme. The rest were transparently Ponzis. Some of these offer an "hourly" rate of return and purport to deposit that return back hourly. Others offer an increased payout upon a subsequent pay in. Some schemes just offer a lump payout after a period of time.

In total, we observed 23 Bitcoin-only Ponzi schemes, which earned 1 562 BTC (843K USD) from January 2, 2013 through September 9, 2014. Table 4.2 reports the key summary statistics. Compared to Bridge HYIPs, Bitcoin-only HYIPs are shorter-lived and less profitable. The schemes collapse within 37 days (median) and the scammers have collectively netted only $40K during that time. Again, we expect that some of the payouts to victims are actually addresses controlled by scammers, so the scammer's profit is likely higher.

4.3 Mining Scams

Since virtually every operation that sells mining equipment has been accused of being a scam, we adopt the narrower definition of scams as those mining operations that take payments from "investors" but never deliver product. Note that "cloud mining" operations that are transparently Ponzi schemes are considered in our HYIP discussion in Section 4.2. Furthermore, we also exclude the many "cloud mining" operations that have not been shown to be Ponzi schemes but are dubious in nature.

We analyze five mining scams (Labcoin, Active Mining Corporation, Ice Drill, `AsicMiningEquipment.com`, `Dragon-Miner.com`). We consider Labcoin here instead of Section 4.2 since it did not promise outrageous returns and it did purport to deliver hashing output to some degree[7]. Similarly, Active Mining and Ice Drill are operations

[6] `https://bitcointalk.org/index.php?topic=576337#post_toc_38`
[7] `https://bitcointalk.org/index.php?topic=263445.msg3417016`

that raised money to purportedly make ASICs and share the profits but never delivered. `AsicMiningEquipment.com` and `Dragon-Miner.com` are fraudulent mining e-commerce websites.

Relevant summary statistics are presented in Table 4.3. Notably, due to the nature of the scam, none of this contributed money is returned to the victims.

4.4 Scam Wallets

We now consider fraudulent services that masquerade as Bitcoin wallets. Note that we categorize wallets that purport to offer a daily return on savings as Ponzi schemes and discuss them in Section 4.2. Scam wallets, by contrast, offer many of the features of online wallets, but with a key difference: the operators siphon some or all of the currency transferred to the wallet.

The basic ruse goes as follows:

1. Victim deposits bitcoin into scam wallet.
2. If the amount of money falls below the threshold, the money stays.
3. If the amount of money is above the threshold, the scammer moves the money into her own wallet.

We identified this process by examining 15 threads on the `bitcointalk.org` forums and 7 threads on the Bitcoin subreddit (`reddit.com/r/bitcoin`) where users complained of losing money once they began depositing larger amounts. Bitcointalk users `drgonzo`[8] and `Artificial`[9]put over 10 bitcoin into their respective Easy Coin accounts in early 2013 but were each left with 0.099 bitcoin (0.1 bitcoin minus their mixing fee) immediately following. Whereas Bitcointalk user `BitcoinOnFire` [10] reports that the first Easy Coin transaction he made worked, but when he moved over a few bitcoin in early 2014, that was quickly drained. Bitcointalk user `Kazimir`[11] reports that putting in less than 0.1 bitcoin into Bitcoinwallet.in in late 2013 which was fine. Reddit user `LutherForThePeople`[12] reports putting in a small

[8] `https://bitcointalk.org/index.php?action=profile;u=106769`
[9] `https://bitcointalk.org/index.php?action=profile;u=109912`
[10] `https://bitcointalk.org/index.php?action=profile;u=323407`
[11] `https://bitcointalk.org/index.php?action=profile;u=58460`
[12] `https://www.reddit.com/user/LutherForThePeople`

amount of bitcoin into Easy Coin in 2013 which was fine and then upon putting in more bitcoin, the scammers drained his account.

We were able to analyze three of these services (Onion Wallet[13], Easy Coin[14], and Bitcoinwallet.in[15]), in which all transfers from the victims were ultimately delivered to the same address held by the scammer. These particular scams advertise themselves as offering a mixing service that enhances transaction anonymity for customers. In fact, all three services appear to be operated by the same scammer, because the siphoning transfers all go directly to the same Bitcoin address. The wallets do in fact operate a mixing service, which makes it impractical to trace back incoming transfers from victims into the service. However, since the scammer sends all stolen bitcoins to the same address, we are able to track the ill-gotten gains for these three scams collectively.

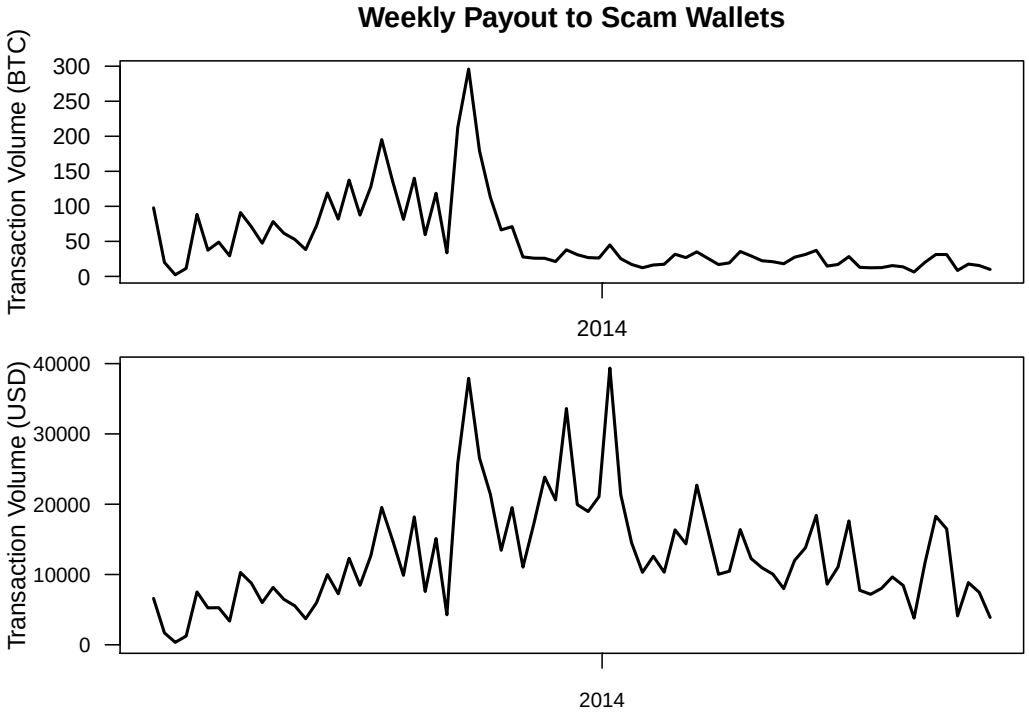

Figure 4.4: Weekly payouts to scam wallets in BTC (top) and USD (bottom).

Figure 4.4 (top) plots the amount of Bitcoin drained out of victim accounts each week. The highly volatile trend suggests that the scam had more success in 2013 compared

[13] `http://ow24et3tetp6tvmk.onion/`

[14] `http://easycoinsayj7p5l.onion/` and `https://web.archive.org/web/20130905204338/https://easycoin.net/`

[15] `https://web.archive.org/web/20140213235218/https://bitcoinwallet.in/`

to 2014. However, normalizing the scammer intake against the BTC–USD exchange rate, as in Figure 4.4 (bottom), tells a different story. It suggests that the scammer drains off an amount of BTC corresponding to a steady USD-denominated wage. Compared to the Bitcoin HYIPs and mining scams, these wallet scams offer a much steadier stream of between $10–40K in ill-gotten gains each week. In total, this scammer's revenue (through 11 September 2014) was about 4 100 BTC, which corresponds to nearly $1 million. Finally, we note that the scam appeared to fold when the original paper for this chapter was originally published at the end of November 2014.

Scam	Lifetime		Payout to scammer	
	Days	Alive?	BTC	USD
Scam wallets	535	yes	4 105	$359 902
Scam exchanges				
BTC Promo	98	yes	44	$22 112
btcQuick		no	929	$73 218
CoinOpend	29	no	575	$264 466
Ubitex	91	no	30	$96[16]
Mining scams	Data Source			
Labcoin	Blockchain		241	$48 562
AMC	BitFunder		18 041	$1 327 590
Ice Drill	BitFunder		14 426	$1 558 008
Asic Mining	Blockchain		12.6	$5 532
Dragon Miner	Blockchain		1.63	$1 019

Table 4.3: Lifetime and payouts for scam wallets and exchanges, plus mining scam payouts.

4.5 Bitcoin Exchange Scams

We look at four scams purporting to be Bitcoin exchanges: BTC Promo, btcQuick, CoinOpend, and Ubitex. Most of these scams entice victims by offering features that many other exchanges do not offer such as PayPal/Credit Card processing, or a better exchange rate than established players. Unfortunately for the customer, they never actually receive the bitcoin or cash after making payment. Ubitex purported to be an in-person exchange, but never got off the ground. Speculation exists as to whether Ubitex is a scam or just a flopped business, but we treat it as a scam here.

[16] 20.189BTC corresponding to $15 515 reported invested on GLBSE, but not trackable on blockchain. Address is from `bitcointalk` forum post asking for Ubitex donations.

Table 4.3 reports the key figures for the scam exchanges. The longer-lived scam exchanges survived for approximately three months, but they also drew in the least amount of money from victims. CoinOpend and btcQuick each operated for less than one month, but during that time drew in hundreds of thousands of dollars from victims.

4.6 Discussion

4.6.1 Revisiting the Scam Categories

Scam category	Scam revenue		Hook	Victim awareness	Trackability
Bridge HYIPs	$6.5M	(in)	Greed	low–high	med.
BTC-only HYIPs	$840K	(in)	Risk appetite, greed	high	high
Mining scams	$2.9M	(in/out)	Advanced-fee fraud	low	low
Wallet scams	$360K	(out)	Information asymmetry	low	low
Exchange scams	$455K	(out)	Information asymmetry	low	low

Table 4.4: Recap of Bitcoin scam categories and features.

The scams presented differ in several key ways, as summarized in Table 4.4. First, we can see that Bridge HYIPs have taken in the most revenue from victims. This may reflect the more mature nature of these scams, as traditional HYIPs have been operating for years. Thus, they already have an established base of users and extensive advertising. The Bitcoin-based schemes, by contrast, are much newer and so we would expect that the scams are not as refined. A less optimistic interpretation, therefore, is that there is considerable room for growth in the magnitude of these frauds as Bitcoin increases in popularity. Furthermore, we note that true total of scammer profits could be much higher, given that we could only track revenues for 21% of the reported scams.

The scams also differ in the way they "hook" victims. HYIPs exploit people's greed, or more precisely, their susceptibility to the narrative that it is easy to get rich quick just by using Bitcoin. Mining scams exploit this same desire, but wrap it in more measured promises of future riches. Mining scams are classic advanced-fee fraud: victims pay money in hopes of getting larger sums down the line, but that day never comes.

Wallet and exchange scams, by contrast, exploit the difficulty people have in judging the legitimacy of web services. Thus, the scammers take advantage of an information

asymmetry that naturally exists. So long as it is difficult to distinguish between good services and bad ones, there will remain an opening for scammers to profit.

User awareness to the scams also varies considerably. Some participants in HYIPs know that they are likely investing in a Ponzi scheme, but they hope to cash out before the scheme collapses. Most Bitcoin-based HYIPs, however, are transparent about the dodgy nature of the service. For example, Bit Twin offers to double your bitcoins within 48 hours. Hence, some scams might even be considered a form of gambling. However, investors in mining, exchange and wallet scams are usually completely unaware that anything untoward is going on with the service until they have lost their money.

Finally, we can distinguish between how inherently trackable these scams are. Some bridge HYIPs can be readily tracked, since they publish a single incoming payment address online. Others use a service such as `blockchain.info` which generates a new incoming address for each visitor. Many require investors to sign up first in order to receive the incoming payment address, which could be changed for different investors. Most Bitcoin-only HYIPs can be readily tracked, since the service usually posts the address in order to signal trustworthiness in the service. Any service that attempts to hide the payment addresses would be viewed with suspicion.

Mining, exchange and wallet scams need not be trackable. The ones we observed happened to make their addresses publicly available, but there is no reason that this should always be. Hence, we anticipate these frauds to remain difficult to track via the blockchain moving forward.

4.6.2 How are Victim Payments into Scams Distributed?

We now examine how the size of payments into scams are distributed. This is an important question, because it influences how successful scammers select targets. A relatively even distribution of payments into scams would indicate that scammers must recruit lots of victims who each contribute a small but substantial amount. By contrast, an uneven distribution suggests that scammers should focus on the small number of marks who will give away the vast majority of the money contributed to the scheme.

To answer this question, we compute measurements typically used in assessing income inequality. Figure 4.5 plots Lorenz curves for each of the HYIP scams we identified. Perfect equality would be indicated by a diagonal line with slope equal to 1, while curves appearing further down and to the right indicate greater inequality in payments from address groups. The left graph plots Bridge HYIPs while the right plots Bitcoin-only Ponzis. We see considerable variation, but with a small number of victims contributing much of the payments in most cases. For instance, in Leancy approximately 20% of the victim population contributed 90% of the payments to scammers. We see even greater variation in the Bitcoin-only HYIPs.

Next, we consider variations across scam categories. Figure 4.6 (left) plots the Lorenz curves for all payments into the 3 scam categories. Payments into mining scams are the most skewed: nearly all of the total contributions come from less than 10% of the victims. While still very skewed, Bridge HYIPs rely on contributions from more victims than do the Bitcoin-only HYIPs: the smallest 80% of address groups account for around 5% of the scammer's haul for Bitcoin-only HYIPs, compared to 15% for Bridge HYIPs.

Figure 4.6 (right) examines the relationship between inequality of payments into scams and the total money drawn into the scams. The graph plots the Gini coefficient for each scam (where 0 indicates all incoming payments are equal and 1 indicates complete

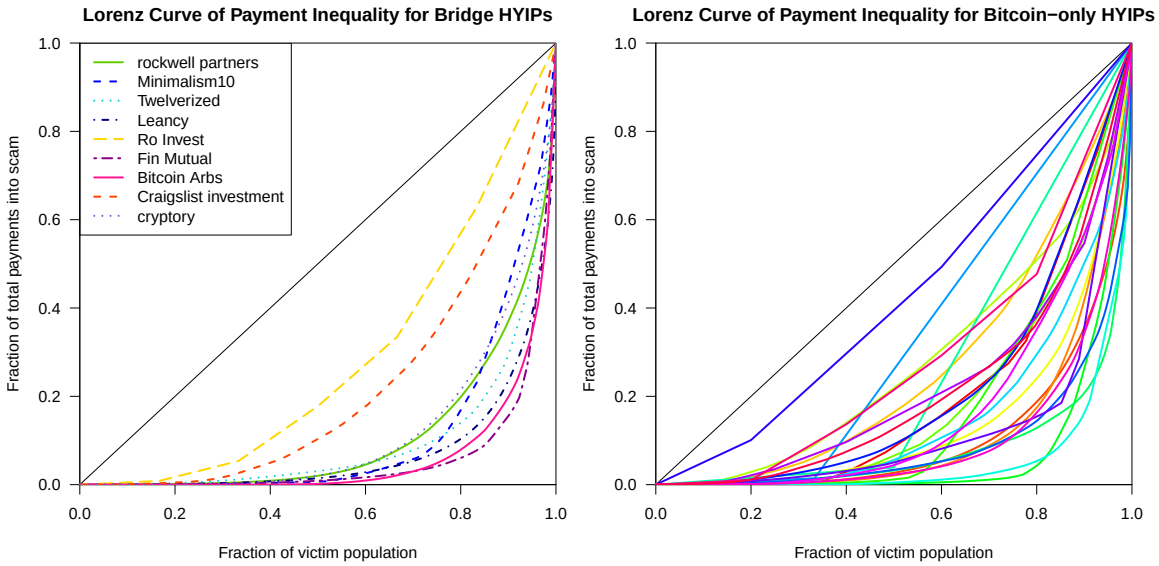

Figure 4.5: Lorenz curve for Bridge HYIPs (left) and Bitcoin-only HYIPs (right).

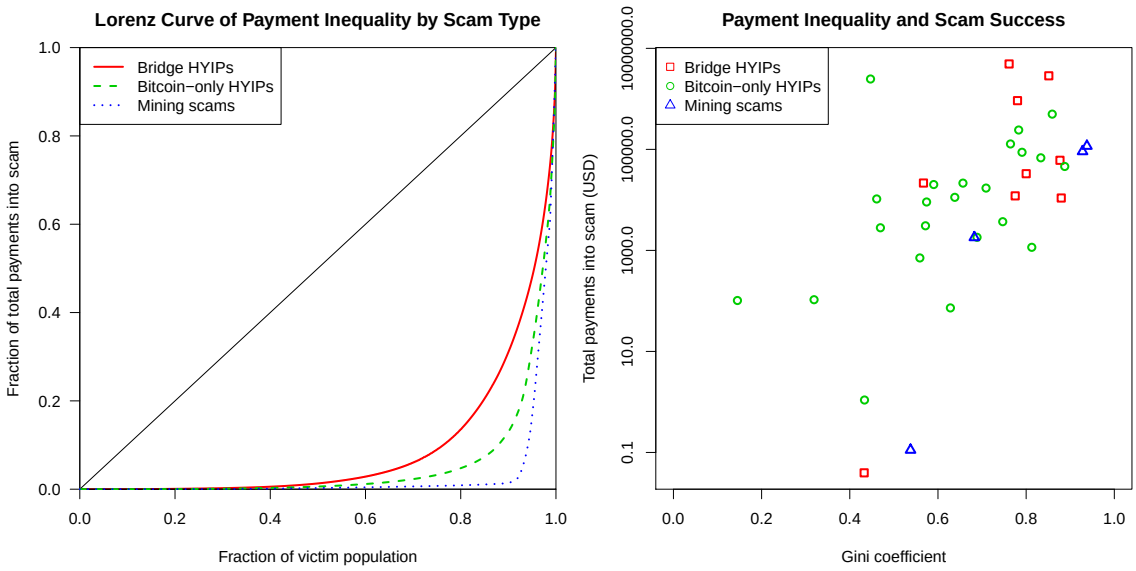

Figure 4.6: Lorenz curve for total payments into scam categories (left); scatter plot comparing Gini coefficient to the amount of money stolen by scammers (right).

inequality) against the total payments paid into each scam. We can see that the least successful scams tend to be the most equal, whereas the most successful scams are more unequal. Hence, for a scam to be successful, it appears that it must catch the few "big fish" who will pay the bulk of the money into the scam.

The high concentration in payment size into scams has implications for law enforcement actions against the scammers. Most successful scams have a few big contributors, who might be more willing to assist with in an investigation. Furthermore, the individual losses suffered by these victims are more likely to meet the threshold required to get the attention of high-tech crime units.

4.6.3 Policy Options

We have already established that different types of Bitcoin scams exist, and that many are growing in popularity. But there are many issues with Bitcoin, as well as cybercrime in general. Given that context, why might Bitcoin scams matter? Here are three plausible reasons: (i) if there are many victims, (ii) if substantial amounts of money is being lost, or (iii) if the scams undermine trust in the ecosystem.

This chapter has established a lower bound on answers to the first two reasons. The number of victims and magnitude of their losses, while considerable, is substantially smaller than those afflicted by failures elsewhere in Bitcoin, such as the Mt. Gox collapse. So on the current figures alone, we cannot conclude that eradicating these scams should take priority.

However, there are two counterarguments that suggest a more robust response is warranted. First, the scams are growing substantially in popularity and profitability. Rooting out the scams at this early stage may be more feasible, and doing so we could avoid the substantial indirect costs imposed by exposing many new Bitcoin users to such a negative experience. The second counterargument is that, for the wallet and exchanges scams at least, their continued prevalence threatens to undermine trust in the overall ecosystem. If people cannot determine whether the service they are interacting with is legitimate due to an information asymmetry, then everyone in the ecosystem, even legitimate exchanges and wallets, suffers.

4.7 Related Work

High-yield investment programs were first documented in the research literature by Moore et al. [56]. They documented over 1 000 such scams, provided a primer on the ecosystem's operation, and established that tracking websites accurately monitor the scam's operation. Neisius and Clayton also investigated HYIPs, focusing on the profits accrued by support organizations in setting up and monitoring HYIP scams [64]. Both papers focused on traditional HYIPs that have operated with impunity for several years using centralized virtual currencies such as Liberty Reserve and Perfect Money. In this work, we have instead focused on HYIPs that use cryptocurrencies as payment. The blockchain has enabled us to accurately measure, for the first time, the amount of money transferred into HYIPs by victims and out by the scam operators.

The Bitcoin Foundation surveyed prominent Bitcoin participants about different hypotheticals that could affect the Bitcoin ecosystem [11]. While they did not explicitly ask about Bitcoin scams, they found that mismanaged Bitcoin businesses was a top threat to Bitcoin's success. They also found people feared Bitcoin getting a "bad reputation" for being a haven for wicked behavior. This includes a concern over Bitcoin being used for

gambling (e.g., many Bitcoin-only HYIPs). The scams presented in this chapter doubtless could harm Bitcoin's reputation if they are not eradicated.

Since our original paper was published, other researchers have been doing work on other cryptocurrency scams. Bartoletti et al. analyzed Ponzi schemes using the cryptocurrency Ethereum and found similar results as to our Bitcoin-based Ponzi scams [10]. Soska and Christin looked at online black marketplaces and found that some would "exit scam" or run the marketplace legitimately for a time and then take all the money deposited in it and leave [76]. They found that this behavior lowered users' confidences in these marketplaces for a couple months, but long term, the online drug market was resilient to these scams.

4.8 Conclusion

Scams – operations established with fraudulent intent – pose serious dangers to the Bitcoin ecosystem. First, there is the direct harm imposed on the victims who pass money to the scammers, never to see it again. Second, and perhaps more substantially, there is indirect harm imposed on all users, even those who don't fall victim to scams. This harm manifests in damage to the reputation of legitimate operations and the undermined trust of users who become more reticent to try out new services.

Fortunately, the blockchain creates an opportunity in that transactions may often be tracked, which could make it easier to assess the true risk posed by scams and make it harder for scammers to hide. To that end, in this chapter we have presented the first systematic, empirical analysis of Bitcoin scams. We identified four categories of scams: Ponzi schemes, mining scams, scam wallets and fraudulent exchanges. By analyzing transactions into and out of 42 such scams, we estimate that approximately $11 million has been contributed to scams by at least 13 000 victims, much of it within the past year.

We found that Bridge HYIPs, an established scam that predates Bitcoin, take in 60% of the total revenue. The blockchain has enabled us to more accurately estimate the financial success of these scams than in previous work, by directly measuring money flowing into HYIPs for the first time. We also worry that the other scam categories may soon rise to the level of HYIPs as scammers wise up to what is possible.

To combat any future rise, continued measurement of the threat as outlined in this chapter is essential. Furthermore, by investigating losses from victims contributing the largest amounts, there may be an opportunity for law enforcement to crack down on scams more effectively.

CHAPTER 5

MEASURING THE SUPPLY AND DEMAND FOR BITCOIN SCAMS

Bitcoin draws out risk-seeking individuals. The exchange rate is volatile; many businesses built on top of it are speculative in nature; the currency is anonymous and distributed. Consequently, it is perhaps unsurprising that many Bitcoin users have taken to Ponzi schemes (and Ponzi scheme runners to Bitcoin).

In this chapter, we look at the ecosystem around Ponzi schemes advertised to Bitcoin users. The previous chapter established a lower bound for the amount of money earned by criminals through Bitcoin scams. Here we more comprehensively study the scams by gathering data where they are promoted. As well as shedding light on the "supply" side of Ponzi schemes, we also look at the "demand" side by gathering data on victim interactions with the scams. People keep falling for Bitcoin scams, but why? Bitcoin users like to purport themselves as particularly technologically savvy, but does that help or hinder their susceptibility to scams? How do the steps taken by scammers, such as engaging shills to promote their products, affect their success? Ultimately, our goal is to shed light on why criminals are able to prosper in this ecosystem.

Even with the improved coverage relative to Chapter 4, our results are necessarily incomplete. There are inevitably scams which use Bitcoin and we do not measure. There are also scammers which create multiple accounts to talk about their scam and we only are able to extricate the obvious cases of this behavior. Despite these limitations, we provide a large-scale analysis of this online Ponzi scheme ecosystem.

The research contributions for this chapter are both in the data collection methodology and in the analysis of the gathered data. Our data collection contributions are gathering candidate scam data directly from scammer advertising venues, automatically confirming scams by inspecting payout mechanisms, and, for confirmed scams, collecting usage, performance and demographic indicators from forum posts. This yields a richer dataset on Ponzi schemes than has been collected before in prior work. Our data analysis contributions lever-

age this novel dataset to describe supply-side characteristics of scams and scammers as well as describe demand-side characteristics of victims.

5.1 Methodology

We aimed to measure scams by collecting data from the places they were advertised. This helps us generate a comprehensive list of advertised scams. For the purposes of this study, we elected to focus on Ponzi schemes exclusively. Of course, there are many different types of scams affecting Bitcoin, as shown in Ch. 4. We focus on Ponzi schemes due to their reliance on public advertising and the consistency of locations for such advertising. Since Ponzi schemes must advertise to stay in business, we are relatively secure in the comprehensiveness of our approach.

In order to collect information about the scams, we crawled the entire history of three subforums of `bitcointalk.org`: Scam accusations, Gambling: Games and Rounds, and Gambling: Investment Games. Investment games is a subforum where users submit Ponzi schemes or moderators move threads on Ponzi schemes. However, the previous chapter found a number of Ponzi schemes advertised in other subforums of bitcointalk. We chose the two most popular subforums for Ponzi schemes from that work, scam accusations and games and rounds. In total, we crawled 11 424 threads on these three subforums from June 2011 through November 2016. We considered all the subforums of bitcointalk where data was found for Chapter 4. We then looked at the forums and looked for Ponzi schemes. We omitted subforums like the gambling subforum which predominantly contained posts about online card games and other non-Ponzi scheme activity.

Since threads on these forums covered other topics than just promoting Ponzi schemes, we refined this further to threads that referenced "ponzi" or "hyip" in the first 10 comments. We then processed these further to only consider threads which contained a URL or bitcoin address for the scam. This left us with 1 810 scams advertised through 1 804 Ponzi-registered domains as well as 1 448 Bitcoin addresses collated from 2 617 threads. We merged threads containing the same domain or Bitcoin address, since many scams were advertised multiple times or in multiple places. Note that we threw out threads containing

Figure 5.1: Screenshots of the initial posting for the Ponzi scheme and an example victim response.

a whitelist of legitimate gambling domains[1]. We also did not consider popular domains, removing from consideration any URLs in the Alexa top 10 000 domains such as `google.com` and `wikipedia.org`.

Our objective is to extract as much information about reflecting supply and demand for scams by examining threads discussing the schemes. In particular, we are interested in measuring the lifetime of the scam, the profiles of the scammers and their victims, and how interactive the threads on scams are. We considered the opening time a scheme was operational to be the first time it was advertised on bitcointalk and the closing time to be the last comment time on threads relating to the scheme. The difference between these times is the lifetime of the scam. We looked at 10 different scams for which we had ground truth on the lifetime of the scam, and found that this method was reasonably accurate within a couple days of the length of the scam.

We identify three distinct categories of posters: scammers, victims, and shills. We consider the scammer to be the original poster about the scheme and the victims to be the commenters who were not the scammer or a shill. For each scammer and victim, we analyzed their most recent posting history (maximum 20 posts). We looked at the other subforums they posted in as well as the number of times they posted on any given Ponzi-

[1]This list was curated by bitcointalk user **mem** here: `https://bitcointalk.org/index.php?topic=75883.0`.

related thread. For scammers, we looked at public interaction with victims; similarly, for victims, we looked at public interactions with scammers. We also tried to find evidence of public history on the forums.

We classified shills as victims who post only about a single scam and nowhere else on the forum. We devised this rule upon looking through scam threads and finding users who were extremely positive. Some of these users posted about multiple threads, seemingly different content, and largely had corroborating evidence, such as transaction information. Others only posted about one or a few scams with similar content. We tried to identify these posters automatically, and the most straightforward was by number of threads posted on. While not all shills only post about one particular scam and not all posters with history on only one scammer thread are shills, we have conlcluded that this simple approach provides an effective rough cut to study this effect.

Finally, we sought a way to measure the effort the scammer made to imbue trust in his scheme from the Bitcoin forum. The markers of trust and reputation used include the time between registration and posting about a scam (with shorter gaps seemingly less trustworthy) and the overall posting history of the scammer including frequency and topics.

5.2 Results

We found 1 780 scams from 1 956 scammers on 2 625 forum posts. Scams with multiple scammers had multiple threads about the scam originating from different usernames. By randomly inspecting 20 such instances of this, we found that in most cases, both usernames appear to be the same scammer or at least operating the same scam. We found 11 990 users who posted in response to these posts.

Figure 5.2 shows the lifetimes of the scams. About a quarter of the attempted scams did not last a day and half only lasted a week. However, some scams lasted a long time, with the longest lasting scam lasting over three years. From manual inspection, many of the scams lasting a day were shut down by the moderators or other entities. The rest of this section will break down this vast difference in lifetimes between these scams and quantify the differences both in attacker strategies and in victimology.

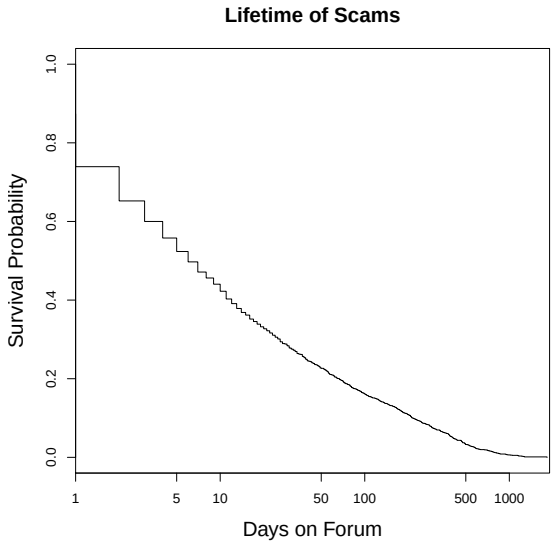

Figure 5.2: Survival analysis of the lifetime of scams.

5.2.1 Scammer Interaction and Scam Lifetime

Figure 5.3 shows the difference in lifetime based on the amount of scammer interaction. Out of the 344 threads that only had one post by the scammer on them, less than 50% lasted longer than a day – 19 of them only consisted of one post total. We found that more scammer posting helped enliven the scam – whereas an average scam lasted about a week, the average scam where the scammer posted at least half of the posts lasted about three weeks. Scammers interacting with their victims seem to prop up their scam, at least in the short term. The difference in these curves, measured by running the survival curve difference test, is statistically significant at the p=0.01 level.

We can see if we can see the same effect for shills as well as scammers, since most of the postings by scammers seems rather overt. Figure 5.4 shows the average lifetime of a scam based on the percentage of posts by shills. Scams where more than 10% of the posts are from shills last longer than those where more than 10% of the posts are from scammers. Furthermore, more shill posts seems to be more effective than the combined strategy, considering both shill posts and scam posts to contribute to the lifetime. Running a survival curve differences test, the effect of the differing shill interaction percentages on the lifetime of a scam are statistically significantly different at the p=0.1 level.

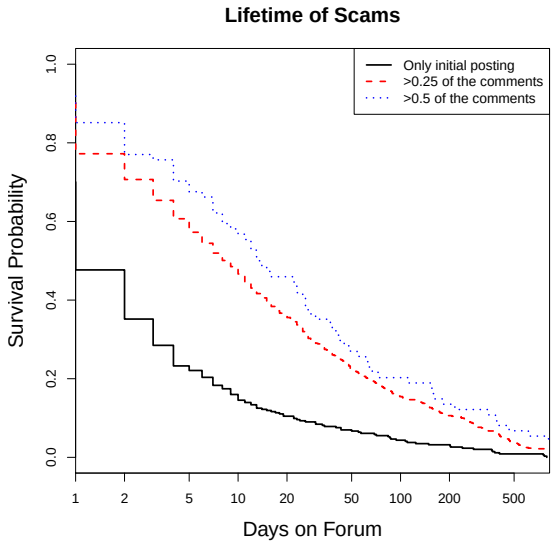

Figure 5.3: Lifetime of the scam based on the fraction of the comments about the scam from the scammer.

We indirectly measure scammer reputation in two ways: by examining where scammers post and by measuring the time between registration and scam posting. Figure 5.5 shows the breakdown in the efficacy of the scam by the reputation of the scammer. On the left we look at the other posts/comments made by the user who first posted the scam. We distinguish between only posting on one scam, only posting on (multiple) scam posts, and those scammers who post in other parts of bitcointalk. We see that scammers that only post on one scam have a lower lifetime compared to scammers that post outside of just one scam. The difference in these survival lifetimes are significant at the p=0.01 level. Figure 5.5b shows the lifetime based on if the scammer account was created on the same day as the scam or not. 39% of scammer accounts were created within a day as the corresponding bitcointalk post. We find that scams advertised by scammer with newly created accounts die quicker than those with older accounts. Half of scams that have been created at least a day prior to posting end within 26 days compared to only 4 days for those created the same day. The difference in these survival plots is statistically significant at the p=0.01 level.

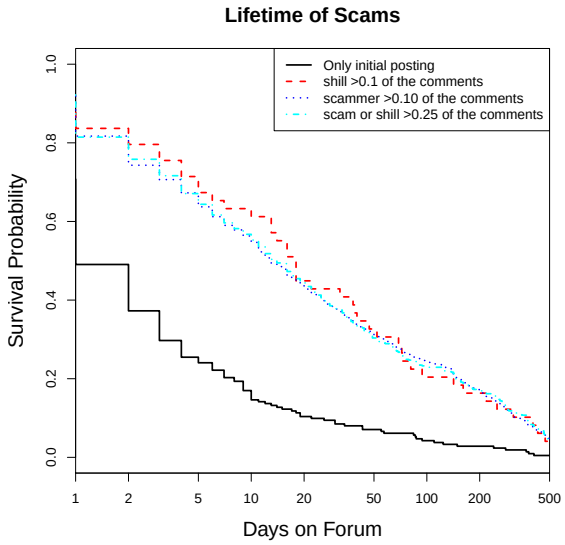

Figure 5.4: Lifetime of the scam by interaction by "shill" commenters.

5.2.2 Victim Behavior

We measured the responses from 11 902 victims from 89 439 comments on 2 629 threads on 1 779 scams. In this section, we examine characteristics of the user accounts that post in threads about Ponzi schemes.

In Figure 5.4 we separated out shills from the victims and the scammers. We can see that shill and scammer activity are associated with longer lifetimes. Active shills do appear to survive slightly longer than active scammers for the first couple of months, but the overall effect is indistinguishable between shills and scammers.

Table 5.1 shows how Ponzi scheme victims' post history compares to that of other users active on bitcointalk. For this, we scraped bitcointalk's aggregated posts statistics for ground truth and categorized each post using bitcointalk's categories. The Ponzi victims' post history was statistically significantly different (at the p=0.01 level) than the general post history, both aggregating by thread and by overall topic. Ponzi victims are overrepresented in the "economy" section, which is unsurprising since this is the section where Ponzi scheme advertisements are located. Ponzi victims are also overrepresented in the "other" section. When we look further into this forum category, we find that Ponzi victims are overrepresented in the "Off Topic" and "meta" board commenters and under represented in "Politics & Society" and "Beginners". We also see Ponzi victims underrepresented in

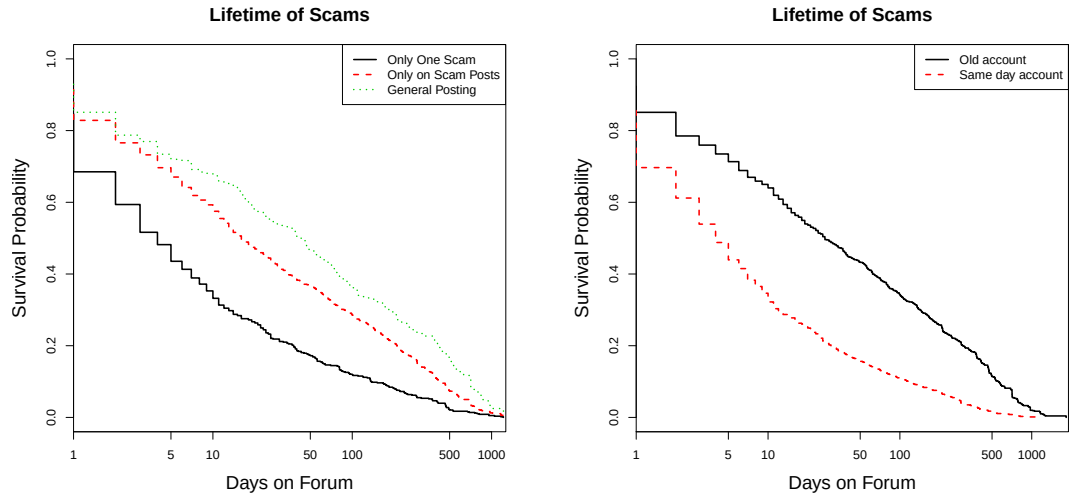

(a) Lifetime of scams, distinguishing between post history.

(b) Lifetime of scams, distinguishing between newly created accounts and older accounts.

Figure 5.5: Measuring lifetimes of scams based on attacker accounts.

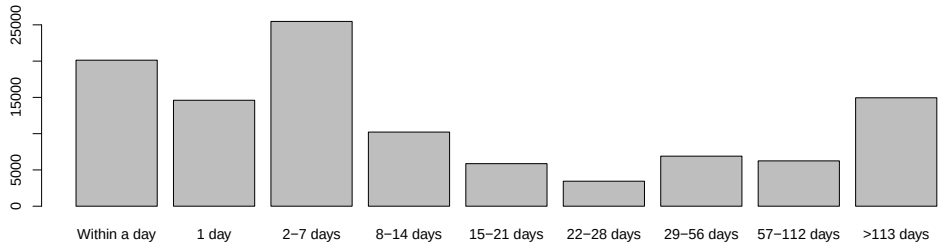

Figure 5.6: Number of victim posts after a thread starts.

many technical boards, like "Development & Technical Discussion" and "Mining" but are overrepresented in "Mining Speculation".

We can also look at what time these victims posted on threads about the scheme. The median time for victims to comment on a thread is about 5 days after the initial post. Figure 5.6 analyzes this effect further. While most victims post within a week, there is quite a long tail of victim posts. We see victims posting over half a year after the start of the initial scam posting.

Category	# Victim Posts	# Other Posts	
Altcoins (all)	32 536	5 429 022	(−)
Alternative Clients	106	54 159	(−)
Bitcoin Discussion	8 872	998 246	(+)
Development & Technical Discussion	683	162 405	(−)
Group Buys	498	84 734	
Hardware	2 730	518 728	(−)
Mining	427	1 044 148	(−)
Mining software (miners)	274	67 561	(−)
Mining speculation	616	63 071	(+)
Pools	885	177 985	(−)
Press	696	74 437	(+)
Project Development	1 526	137 245	(+)
Technical Support	586	58 952	(+)
Auctions	1 865	108 048	(+)
Collectibles	1 063	60 745	(+)
Computer hardware	1 462	118 584	(+)
Currency exchange	3 124	138 264	(+)
Digital goods	7 303	277 903	(+)
Economics	3 692	1 204 450	(−)
Gambling	12 070	1 297 038	(+)
Gambling discussion	5 677	340 593	(+)
Games and rounds	23 331	388 689	(+)
Goods	1 251	587 681	(−)
Investor-based games	15 402	115 454	(+)
Lending	3 230	138 108	(+)
Marketplace	517	5 372 844	(−)
Micro Earnings	3 694	144 797	(+)
Scam Accusations	4 643	116 151	(+)
Securities	1 338	202 813	
Service Announcements	2 338	288 993	(+)
Service Discussion	3 692	330 535	(+)
Services	8 528	407 342	(+)
Speculation	5 058	883 584	(−)
Trading Discussion	1 678	257 930	
Local (all)	14 932	4 454 405	(−)
Archival	1 026	147 836	
Beginners & Help	3 923	564 720	
Meta	1 960	134 319	(+)
Off-topic	8 309	563 710	(+)
Politics & Society	2 181	290 782	

Table 5.1: Bitcointalk forum categories and where scam victims post. Categories are marked as under or overrepresented according to a chi-squared test with 97.5% confidence. Categories with at least 50 000 posts are included.

5.2.3 Proportional Hazards Model

To distill the varying effects on the lifetime of a Ponzi scheme, we run a Cox proportional hazards model. Our dependent variable is the lifetime of the scam, measured in days. For independent variables, we have used:

daily # victim comments This measures the number of *victim* comments over the lifetime of the scam. We use a daily count, since the overall count is, unsurprisingly, highly correlated with the lifetime of the scam.

daily # scammer comments This measures the number of *scammer* comments over the lifetime of the scam. Again, we use a daily count to control for the correlation between this variable and the lifetime of the scam.

shill has posted? This is true if a "shill" (described more thoroughly in Section 5.2.1) has posted anywhere in the thread. This accounts for their presence, since the number of comments by these users is so low.

same day account This is true if the scammers' bitcointalk account was registered the same day as the original post for the scam.

	coef	exp(coef)	95% CI	p value
Daily # victim comments	0.028	1.029	(1.022 , 1.036)	$\ll 0.0001$
Daily # scammer comments	0.022	1.022	(1.002 , 1.043)	0.034
Shill has posted?	-0.846	0.429	(0.385 , 0.479)	$\ll 0.0001$
Same day account	0.374	1.453	(1.320 , 1.599)	$\ll 0.0001$
Log-rank test: $Q = 489.2$, $p \ll 0.0001$, $R^2 = 0.218$.				

Table 5.2: Cox proportional hazards model: measuring scammer and victim effects on the lifetime of the scam.

Table 5.2 shows the results of running this regression. We note that all the variables are statistically significant to at least the $p = 0.05$ level, with three of the variables highly significant. The best way to interpret the table is to focus on the exp(coef) column. Values greater than one correspond to an increase in the hazard rate, while those less than one correspond to a decrease. The hazard rate captures the instantaneous probability that a scam will shut down, so an increased hazard rate means a greater risk of shutdown.

Each additional daily comment by a victim correlates to a 2.9% increase in the hazard rate. The effect is similar, though slightly weaker, for additional posts by the scammer. The result is somewhat counterintuitive; one might have expected scams with more active participation to be longer-lived, yet the opposite is true. One possible explanation is that victims are more likely to post when there are problems, and so are scammers.

By contrast, a shill posting on a thread is correlated with a massive 57% decrease in the hazard rate. This indicates that shills may play a significant role in prolonging the lives of scams, helping to draw in more victims and settle the nerves of existing investors.

Unsurprisingly, a scammer creating an account on the same day as the initial post correlates with a longer scam lifetime. This confirms the intuition from Figure 5.5, which suggests that no post history shortens the lifetime of the scam. The Cox model shows that scams created by newly registered posters face a 45% increase in the hazard rate.

Reflecting on the overall model, we conclude that posts by shills may prolong a scam's lifetime dramatically, whereas posts made by victims and scammers have the opposite effect. Finally, the reputation of posters as indicated by posting history also appears to significantly affect the scam's expected lifetime.

5.3 Conclusion

Bitcoin Ponzi schemes are alluring. The victims of these scams enjoy the thrill of the risk and the opportunity to earn a windfall. The scammers are seduced by the opportunity to earn hard-to-trace money with seemingly little effort.

To measure this, we crawl 11 424 threads on three subforums of the Bitcoin forums from June 2011 through November 2016 to find 1 780 scams from 1 956 scammers on 2 625 forum posts targeting 11 990 users. We find that more daily scammer and victim interaction shortens the life of the scam. Furthermore, we find that shill interaction, or users that only post in one thread, lengthens the life of the scam. We show that having a reputation on the Bitcoin forum matters: posting a scam the same day as an account was created is associated with a quicker demise.

In addition to investigating perpetrators of these frauds, we also analyze the users who fall victim to them. We compare the post history of scam victims to overall Bitcoin

forum statistics and find that scam victims disproportionately post in other forums like "Off-Topic" and "Mining Speculation". We find that most victims post within the first five days of a scam post, with a long tail that post even over a year after the initial posting.

CHAPTER 6

MEASURING THE USE AND ABUSE OF BRAIN WALLETS

In this chapter we study the use of *brain wallets*, or private keys which are deterministically derived from passwords or passphrases. Compared to other paradigms for managing Bitcoin keys, such as storing them on a personal computer or a dedicated hardware device, this approach is convenient as the user can spend their bitcoins simply by typing their password. Because their private keys are not permanently stored on devices, brain wallets cannot be exfiltrated by malware [9].

However, there is a big downside: anyone who guesses a user's password can immediately steal their funds. Worse, attackers can perform unthrottled (offline) guessing to test candidate passwords. Attackers guessing a password can quickly test whether it matches *any* user's brain wallet by scanning for use of the derived public key's hash on the Bitcoin blockchain, a public ledger of all transactions. We replicate this password-guessing attack in a research setting by non-invasively testing candidate passwords for historical use as a Bitcoin brain wallet address.

Others have investigated brain wallets. Eskandari et al. studied bitcoin wallet software and found that while brain wallets are supported across platforms and require little trust in devices, the threat of weak passwords eclipses those benefits [30]. BIP 38 specifies a format for password-protected private key encryption as a second factor [21]. Our work also builds upon work on passwords for financial systems. While there is little evidence that users choose significantly stronger passwords to protect financial online accounts [18], Herley argues that users rationally choose weak passwords for online accounts [37] as they are protected by anti-fraud systems.

In this work we report on the first large-scale attempt to measure brain wallet use and abuse in the wild. Surprisingly, we identified a relatively small number of brain wallets in use: fewer than 2 000 total. This is despite a significant amount of interest in the concept and the existence of several software tools for creating and using brain wallets.

Our results are necessarily incomplete in that password-derived public keys are indistinguishable from pseudorandomly-generated public keys without knowledge of the password. Put another way, we do not know how many brain wallets are in use for which we were not able to guess the password. Nonetheless, given that we tried over 3.9 trillion passwords and passphrases from over twenty customized word lists, we are confident that the use of brain wallets remains quite rare.

Our results reveal the existence of an active attacker community that rapidly steals funds from vulnerable brain wallets in nearly all cases we identify. In total, approximately $261K worth of bitcoin has been loaded into brain wallets, with the ten most valuable wallets accounting for over three quarters of the total value. Many brain wallets are drained within minutes and nearly all wallets are drained within 24 hours.

We present evidence that the time required to drain has rapidly shortened, with median times that were measured in hours through mid-2013 now measured in minutes or seconds. We document how a dozen or so drainers have emptied multiple brain wallets. Finally, upon examining the cracked passwords, we find no evidence that users pick stronger passwords when they load more money into the wallets. However, we do observe that the addresses created with passphrases rather than passwords are drained slower.

The research contributions for this chapter are both in the data collection methodology and in the analysis of the gathered data. Our data collection contributions are generating a large corpus of candidate passwords and efficiently using the blockchain to study attacker draining behaviors. Our data analysis contributions are using Bitcoin to quantify password and passphrase selection as well as measuring attacker draining performance using timing data.

6.1 Data Collection Methodology

We first review how the candidate passwords[1] were constructed and then explain how we checked for their usage in brain wallets.

[1]Technically these are passwords and passphrases (sequence of words). We use password for simplicity of presentation.

6.1.1 Password Corpora

We have constructed an extensive set of passwords derived from publicly available sources. This includes prior password leaks (e.g., Rockyou, Yahoo!, LinkedIn) word and derived phrase lists (e.g., English Wikipedia, Wikiquote), and information gleaned from Bitcoin discussion forums. Coming up with comprehensive, representative passphrase lists was hard – most of the other datasets do not include them. Rather, we had to generate them by scraping websites and looking for media dumps. In total, we tested approximately 3.9 trillion passwords and passphrases for usage in brain wallets. Testing was carried out using the open-source project Brainflayer[2].

Word lists were tried directly unless otherwise specified. The following word lists were used:

1. **xkcd:** Lists obtained on July 10^{th}, 2014 from three sources[3]. Combinations up to four words with and without spaces, both lowercase and initial caps. All words used for three word combinations; words common to all three lists used for four word combinations.

2. **Urban Dictionary:** Terms and phrases from the crowd-sourced slang dictionary[4].

3. **Password dumps:** LinkedIn, MySpace, RockYou, Rootkit.com, Stratfor, eHarmony, Nvidia, Gamigo, Adobe, Project Whitefox, LinkedIn, and found dumps from Pastebin in 2012[5]

4. **Security industry lists:** CrackStation, Naxxatoe, Uniqpass (combination of 2012-01-01 and 2012-04-01 lists), Oclhashcat (medium), Skull Security[6] (RockYou list excluded) all with permutations (unchanged, initial caps, force lowercase, initial caps and strip spaces, lowercase and strip spaces).

[2] `https://github.com/ryancdotorg/brainflayer`

[3] `https://xkpasswd.net/s/`, `http://correcthorsebatterystaple.net/`, and `http://preshing.com/20110811/xkcd-password-generator/`.

[4] List was sourced from `https://github.com/inieves/urban-dictionary-scraper/blob/4a86fd9ef4c2f8812dc78f5862c327912213436a/dict/UrbanDictionary.txt`.

[5] List sourced from `https://blog.thireus.com/look-back-on-2012s-famous-password-hash-leaks-wordlist-analysis-and-new-cracking-techniques/`.

[6] `https://wiki.skullsecurity.org/Passwords`

5. **Facebook names:** Names from Facebook obtained from Skull Security[7], lower case and initial caps, with and without spaces

6. **BitSig:** Data scraped from `bitsig.io`, a website that allows users to post brain wallet plaintexts as a timestamping scheme.

7. **Bitcoin IRC:** Channel chat messages from Bitcoin-related IRC channels from 2011 through 2017.

8. **Reddit:** Reddit commenters and comments, sourced from `https://archive.org/details/2015_reddit_comments_corpus`.

9. **WikiQuote:** English, Spanish, Russian and German quotes from 3/2013 with permutations (with/without spaces, initial caps/all lower case, with/without trailing punctuation, with/without commas).

10. **BrainyQuote:** Quotes scraped from BrainyQuote.

11. **Wiki/Brainy:** Combined from WikiQuote and BrainyQuote.

12. **Lyrics:** Lyrics and song titles purchased from `https://andymoore.info/mysql-lyrics-database/` with some permutations applied (unchanged, all caps, all lower, with/without punctuation)

13. **English Wikipedia:** Words and phrases scraped from `en.wikipedia.org`.

14. **Openwall:** Password cracking wordlists from 20+ languages purchased from Openwall[8].

15. **Purdue:** Wordlists from the information security department at Purdue[9]

16. **Keyboard Patterns:** Keyboard patterns from oclhashcat

In addition to the aforementioned word lists, we tested the following:

1. **Brute Force:** All numbers up to 11 digits, printable ASCII up to 6 characters, and 7 lowercase with spaces.

[7] `https://wiki.skullsecurity.org/Passwords#Facebook_lists`
[8] `http://www.openwall.com/wordlists/`
[9] `http://ftp.cerias.purdue.edu/pub/dict/wordlists/`

2. **Misc:** Known public brain wallets and iterative attacks on found brain wallets by taking plaintexts cracked via other methods and prepending/appending them with up to three additional printable ascii characters or additional words from the Ubuntu American English wordlist[10].

Table 6.1 in Section 6.2 details the number of brain wallet passwords obtained from each source, along with the total amount drained.

6.1.2 Observing Bitcoin Brain Wallet Usage

We use the SHA256 hash of the password as the private key. We then generate the corresponding public key using a few speedups to the secp256k1 curve library[11] [26]. We download the Bitcoin blockchain using Bitcoin core software[12] and extract all the unique Bitcoin addresses using znort987's block parser[13]. We then add all the addresses to a bloom filter for quick lookup and a sorted list for false positive detection. We compare all the addresses generated from candidate passwords against the bloom filter and confirm positive results against the sorted list. After we find all of the used brain wallet addresses, we supplement this information by querying all our brain wallet addresses against the `blockchain.info` API to obtain precise timestamps for all transactions. Transactions with brain wallets as recipients are incoming payments and transactions with brain wallets as sources are outgoing payments.

6.2 Results

We investigate brain wallet usage by examining all blockchain transactions through the end of February 2017. We excluded 17 784 brain wallets that were suddenly assigned a tiny amount of bitcoin from 36 linked input addresses within a few hours on August 31, 2013. We strongly suspect these brain wallets were set up by a researcher. We also excluded transactions included in a Bitcoin network "stress test" as detailed in Section 6.2.3 and transactions drained by Bitcoin mining pools detailed in Section 6.2.5. The reason we

[10] Particularly, `/usr/share/dict/words`, which contains 99 171 entries and has the sha256sum 126a4ef38493313edc50b86f90dfdaf7c59ec6c948451eac228f2f3a8ab1a6ed

[11] `https://github.com/bitcoin-core/secp256k1`

[12] `https://github.com/bitcoin/bitcoin`

[13] `https://github.com/znort987/blockparser`

Source	# Wallets	(non-empty)	Unique	90% # drains	Total BTC	Total USD
Word lists						
xkcd	155	3	8	8	126.94	8 857.49
Urban Dictionary	244	0	1	3	51.01	54 41.56
Password dumps	815	0	44	3	199.20	39 155.22
Industry lists	876	0	32	3	364.91	37 096.71
Facebook names	364	0	23	4	107.78	14 425.13
BitSig	235	0	71	8	1 586.78	63 818.81
Bitcoin IRC	454	1	17	6	777.52	25 355.79
Reddit	843	8	120	3	2 175.42	99 089.43
WikiQuote	281	0	3	7	113.60	17 700.88
BrainyQuote	61	0	0	6	85.12	14 037.94
Wiki/Brainy	83	0	0	6	101.05	14 481.48
Lyrics	438	0	17	4	270.28	19 257.41
Wikipedia	176	0	5	6	565.77	15 645.48
Openwall	456	0	0	3	60.69	14 097.56
Purdue	424	0	0	3	118.95	14 983.66
Keyboard Patterns	19	0	0	5	0.96	246.96
Non-word lists						
Brute Force	586	2	84	2	96.44	23 796.09
Misc	268	7	268	1	73.67	26 941.39
Overall	1 730	21	488	3	2 846.23	260 792.30

Table 6.1: Brain wallets and values associated with different password sources.

exclude such activity is that our interest is in studying legitimate use of brain wallets. We report on their prevalence, draining, and password strength.

6.2.1 How Prevalent are Brain Wallets?

We have found 1 730 distinct brain wallets using 1 686 different passwords. The slight difference is from to the small number of instances where compressed and uncompressed wallets were used for the same password. In total, these brain wallets received 2 846 BTC (approx. \$261K[14]).

Table 6.1 reports the brain wallets identified, broken down according to the password sources. The most popular source, yet not a very unique one, was from our industry password lists. This was also the second largest source of passwords (the largest being the brute force list). While these lists contained 876 different brain wallets, only 32 of those wallets were not found using all of our other methods. The single most uniquely popular

[14]All USD calculations presented here are normalized by the corresponding day's exchange rate on Bitstamp, as reported by `bitcoincharts.com`.

source is the miscellaneous category. Note that here we only counted number of unique wallets, since the input to this list is all the brain wallets cracked from other methods. This list contained 268 wallets not found in other lists, 7 of which had not been emptied as of early March 2017. Many of these wallets shared prefixes and had digits at the end. The second most uniquely popular list was sourced from Reddit comments and commenters and had 120 wallets not found in other lists, 8 of which were non-empty. We only considered brain wallets using the sha256 hash of the password/passphrase. We did not look for brainwallets using different key derivation formulas, like `brainwallet.io` and WarpWallet which use the scrypt function on the password.

The password sources used for our study can of course also be used by attackers. These attackers scan the Bitcoin network looking for brain wallets, taking all the funds from anything with a balance. We call this behavior "draining". One way to estimate the popularity of password sources among attackers is to compare how often repeated drains occur. The fifth column shows the 90^{th} percentile for number of drains observed on passwords identified by each source. Larger numbers indicate that more attackers are using the source. Perhaps unsurprisingly, passwords derived from xkcd are drained repeatedly the most, tied with passwords found from BitSig messages.

The last two columns provide an alternative way to value the passwords obtained from different sources. Each represents the total value put into brain wallets whose passwords are identified by these sources (in BTC and USD, respectively). By this measure, the Reddit word list is still the most valuable at $99K, followed by BitSig, password dumps, and industry lists.

Figure 6.1 plots when wallets were first used over time, beginning with the first brain wallet established in July 2011. Monthly totals of new wallets are reported, and the bar chart breaks down the use of compressed and uncompressed brain wallets. We can see that the number of new brain wallets has increased since Bitcoin's early days, though the total remains small.

We consider two different types of wallets (encoded hashed public keys): compressed and uncompressed. Each Bitcoin public key corresponds to a point on an elliptic curve. An uncompressed public key stores the `x` and `y` coordinate and almost twice the size of the

compressed public key which only contains the `x` coordinate and a flag for the half of the curve the point is on [23][15]. Relatively speaking, uncompressed wallets are more prevalent. We found 1 566 uncompressed wallets compared to 164 compressed. Note that the brain wallet service `bitaddress.org` offers only uncompressed brain wallets whereas the (defunct) `brainwallet.org` defaulted to uncompressed brain wallets (though it supported both). We note that the number of newly created brain wallets dramatically decreased after August 2015. On August 7, 2015 Ryan Castellucci gave a talk on brain wallets at the DEFCON conference. Right after his talk, `brainwallet.org` went offline. We note that his talk had seemingly no effect on attacker behavior; rather, the effects were concentrated on users creating fewer brain wallets.

Compressed keys are only supported in versions of Bitcoin clients released after March 30, 2012; we observed 30 brain wallets before then, the first being "This string contains 0.25 BTC hiding in plain sight." seen with in July 2011. This wallet was created by Mike Caldwell and was offered to the first person to figure out his puzzle[16]. It was subsequently drained within four hours by an anonymous bitcointalk user. Unsurprisingly, the second brain wallet found in our collection was "Here is another 0.08 BTC waiting to be claimed." and also from the same user in July 2011.

Figure 6.1 also plots the USD value of the brain wallets each month. We can see that this is quite volatile. Most months, the total value hovers around a few thousand dollars, but frequently the amount stored spikes greatly, including to a peak of over $45K in July 2015. Notably, there is no discernible relationship between the number of new wallets created and the value stored.

The left plot in Figure 6.2 gives the CDF of brain wallet value in USD. While most brain wallets store little money (just 5% of the brain wallets received the equivalent of $100 or more), the bulk of the total value in brain wallets is associated with a small number of addresses. The right plot of Figure 6.2 presents a rank-order plot, which reveals that just 10 wallets account for approximately 85% of the total dollar value placed into all brain wallets.

[15]This is described further here: `https://bitcoin.org/en/developer-guide#public-key-formats`.
[16]https://bitcointalk.org/index.php?topic=28877

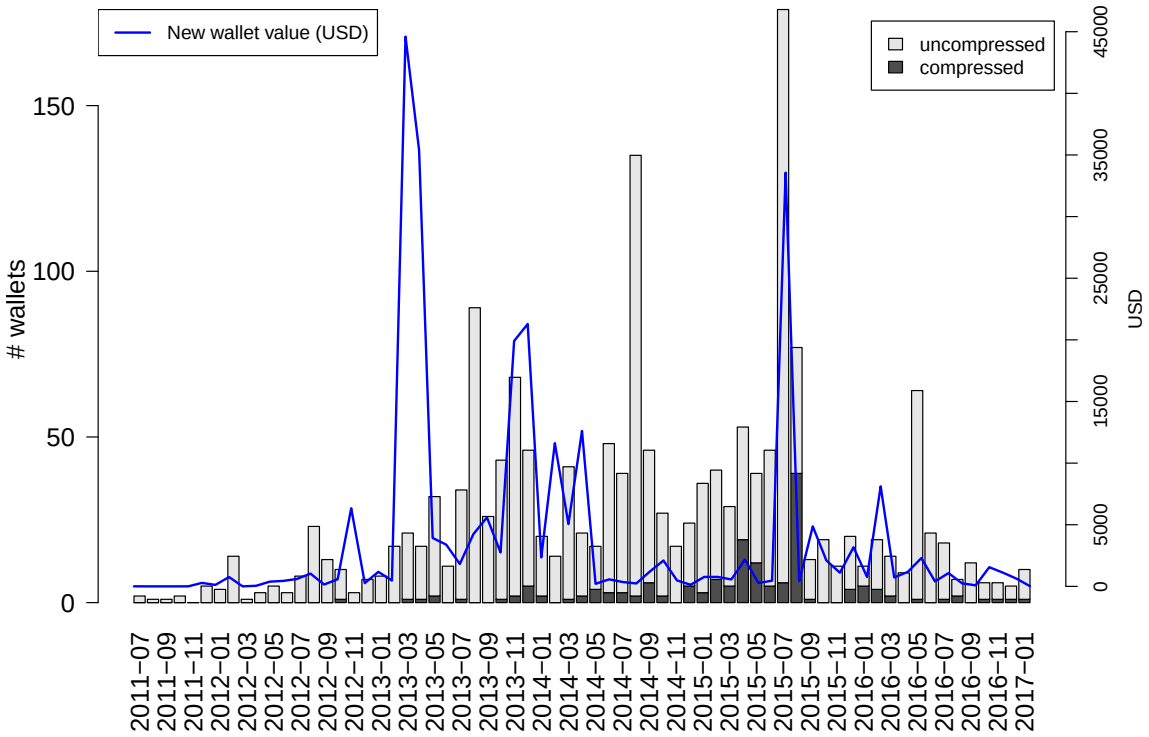

Figure 6.1: New brain wallet usage per month (compressed and uncompressed).

6.2.2 Draining Brain Wallets

As explained at the beginning of this chapter, because the addresses used by brain wallets are deterministically computed from passwords, there is a risk that attackers might guess the password and drain the wallet's value. Many users select brain wallets with the intention of keeping their bitcoin there for a long time, analogous to hiding cash under a mattress. Therefore, when bitcoins are drained from these addresses (i.e., the account balance falls to zero), it strongly suggests that an attack may have taken place.

Perhaps the best way to quantify brain wallet insecurity is to examine the time required to drain wallets. Figure 6.3 plots a CDF of the observed time-to-drain. The solid black line shows the distribution for all wallets. Half of the wallets are drained in 32 minutes or less. Subsequently, the rate of draining slows, but nearly all brain wallets are drained within 24 hours. While some of these drains are initiated by the brain wallet owners, it is likely that most are not.

We can also see the difference in draining speed when wallets are loaded with large or small amounts of money. The red dashed line plots the cumulative distribution for

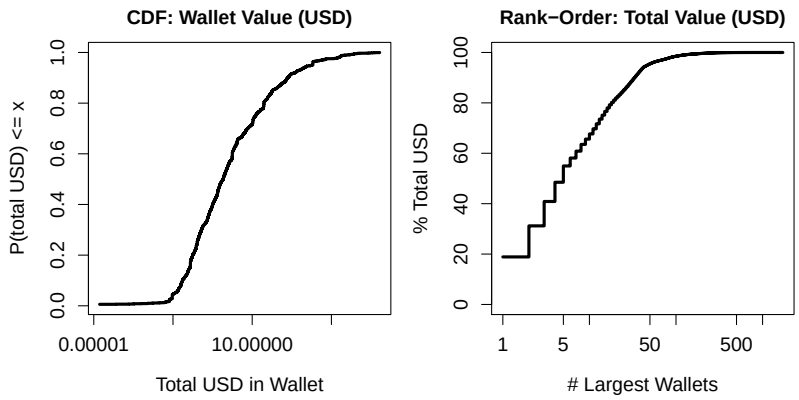

Figure 6.2: CDF and rank-order plot of total value stored in brain wallets.

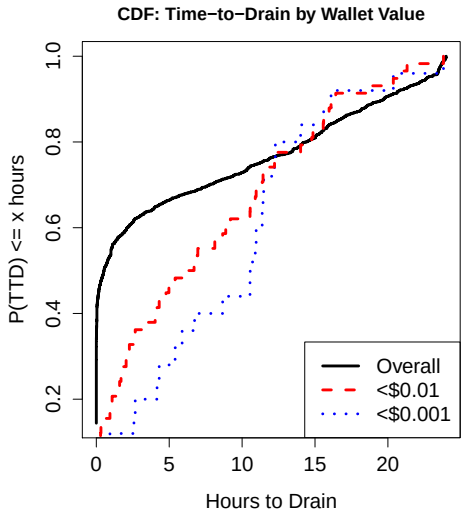

Figure 6.3: CDF of the # of hours to drain brain wallets for wallets by value stored

wallets loaded with a penny or less. The blue dotted line plots the cumulative distribution for wallets loaded with a tenth of a penny or less. We see that brain wallets with such a low amount of money are drained slower than other brain wallets, presumably because the amount in the wallet is lower than the fee accompanied with the transaction. However, these low value brain wallets are eventually drained within a day of creation.

How often are brain wallets drained? 98% of the brain wallets have been drained at least once. We observed 3 153 distinct draining events on 1 714 brain wallets. 76% of wallets are drained exactly once, while 14% are drained twice, and 1.9% are drained ten or more times. Figure 6.4 plots the median time-to-drain by month. While this is always

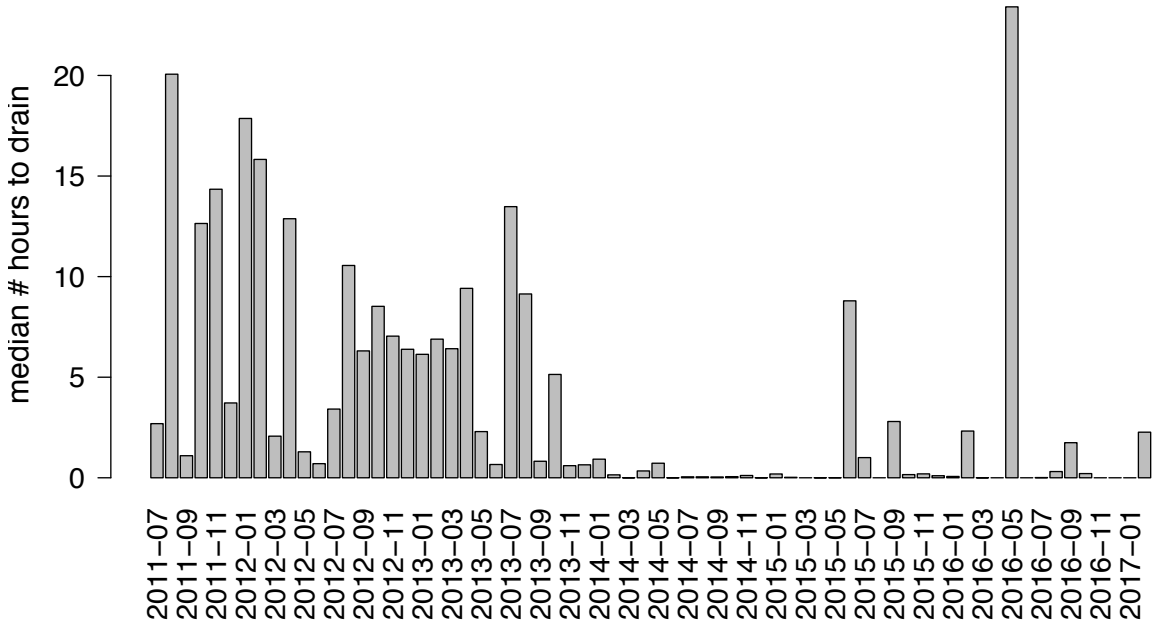

Figure 6.4: How time-to-drain changes over time (median time-to-drain reported per month)

brief (less than one day), by September 2013 it becomes measured in minutes and seconds rather than hours, with spikes as transaction fees rise.

6.2.3 Network "Stress Test"

On May 4, 2015, Gavin Andresen published his first in a series of blog posts aimed at convincing the Bitcoin community to adopt a larger block size [5]. This sparked a heated debate about the Bitcoin blocksize that broke past the technical developers of the Bitcoin protocol into the broader sphere of people that care about Bitcoin.

In order to test the appropriateness of the current Bitcoin network capacity, a group of users, most notably from CoinWallet, decided to broadcast a large quantity of small value transactions with the stated purpose of creating a 30-day backlog of transactions [1].

We inadvertently measured some aftermath of this contentious test, since the attackers chose to send large quantities of small amounts of Bitcoin to select brain wallets. From June 13 to August 28, 2015, 15 brain wallet addresses were used in 20 172 transactions using only 6.6 bitcoin. Many of these transactions originated from or drained to an address

Rank (USD)	Drained # pwd	Drained (USD)	Drains	Description
1	1	22 466	1	woodchuck drain (unintentionally done by researcher Castellucci, https://rya.nc/dc23)
2	1	15 267	1	woodchuck drain (done by owner)
3	13	14 561	31	drainer https://bitcointalk.org/index.php?topic=1138273.40
4	208	13 968	345	drainer https://bitcoin.stackexchange.com/questions/35971/double-spend-bitcoin-bots
5	1	13 766	1	Antidisestablishmentarianismistic drain
6	2	11 528	2	drainer https://archives.somethingawful.com/showthread.php?threadid=3606857&userid=0&perpage=40&pagenumber=971
7	1	10 526	23	1 9 9 2 1 1 2 4 1 2 6 drainer
8	1	10 009	1	drainer https://redd.it/1j9p2d
9	1	9 963	3	1BQmbdHdtdJnGbhNLgnr5w5pKJ4aFghdLp drainer
10	1	6 597	1	1PsenWrxazHNrEC9pR7JESb37aogZZFWUW drainer

Table 6.2: Top 10 drain addresses from brain wallets, sorted by amount drained in USD.

that starts with "1aa", which is tied to CoinWallet [17]. We removed these transactions from our analysis.

6.2.4 Tracking the Drainers

How can these drains occur so fast? Many bots monitor for new transactions depositing into known brain wallets. These *drainers* quickly send the money to their own addresses, often with a sizable fee to encourage miners to pick up the transaction quickly. In contrast to many criminals who take steps to cover their tracks (e.g., by funneling transactions through many addresses), drainers are proud of their achievements. Consequently, they make it easy for all to see that they have done the draining, such as by using the same address for all drains. This makes it easier for researchers to document their activities.

How many drainers did we find? The graph in Figure 6.4 also plots in red the number of drainers actively receiving money from brain wallets. Overall their number seem to be fairly steady, except for a large spike in July 2015. We believe that this is related to the stress test on the network during this time. However, we are not able to always reliably distinguish stress test orchestrators from normal users and so we keep this data in our analysis.

Digging deeper, we manually inspected the top 10 receivers of money, shown in Table 6.2 which is sorted by the total amount drained in USD. The table indicates how

[17] https://bitcointalk.org/index.php?topic=1175321.500

many distinct brain wallets were drained, the associated value in USD, and the number of drain events that occurred. Most of the popular addresses were associated with a single password drain. In a few cases, the owner explicitly confirmed that they drained their own by online postings. However, a number of these one password drains were reported publicly as stolen. It would not be surprising if the spikes seen in Figure 6.4 were due to one drainer that used multiple Bitcoin wallets to drain their money.

A few drainers are very successful while the rest do not make very much. The top 4 drainers have netted the equivalent of $51 000 between them. The drainer who has emptied the most brain wallets – 208 in all – has earned $13 967 for the effort. But other drainers have stolen very little money. For example, one drainer stole from 60 different brain wallets but netted only $5.25 worth of bitcoin. Why is this? Looking back at Figure 6.4 at the money flowing into brain wallets indicates this amount has diminished as Bitcoin's overall popularity has risen.

We also investigated the behavior of successful drainers. Some have claimed that drainers purposely avoid emptying brain wallets with small stores of value [34]:

> Another example is brainwallets, we have clear evidence that people who crack brainwallets intentionally avoid sweeping small amounts (And even coordinate among each other) in order to avoid alerting users prematurely.

We did not find any evidence for this practice among the most successful drainers. The median value of a drained brain wallet among each of the most successful drainers was under $1 (typically a few cents). We do find that wallets holding less than a cent are drained at a slightly slower rate than other wallets (as seen in Figure 6.3) but are still drained within a day. We attribute this to relatively higher transaction fees rather than attacker coordination.

6.2.5 Mining Pool Drains

We found evidence that 8 mining pools had 157 710 drains accounting for 88 708 transactions since September 2013 from 15 popular brain wallets amounting to 1.58 BTC or 437 USD. These are not included in our other measures of draining activity. This is

because the rest of our measures assume that drainers are not miners and only include the money that the drainers take themselves and not the transaction fee. However, when mining pools drain brain wallets, they instead leave the full value as a transaction fee. Since transaction fees are taken by the mining pool that put the corresponding transaction in a block, this effectively does the same thing.

Rather than making a profit, this seems to be motivated by cleaning up the unclaimed transaction set. The network stress test (explained further in Section 6.2.3 was trying to create a backlog of transactions and one of the ways was through creating a large number of brain wallets. It is telling that none of the mining pool drains was over 0.14 USD. The Chinese mining pool F2Pool mined 97% of these drains which contained 99.6% of the total money we measured which was drained this way (in USD).

6.2.6 Impact of Password Strength

Measuring the "strength" (or resistance to guessing) of an individual password is a hard problem. Many standard metrics, such as the NIST "entropy" formula, have been shown to be poor predictors of actual cracking time [79]. In practice, many websites use inconsistent and poorly specified methods for giving users feedback on password strength [22]. The gold-standard of non-parametric statistics requires very large sample sizes and is hence impractical in our setting [17]. Instead, we use the `wheelerzxcvbn` formula as a rough measure of password strength [80]. While it produces an integer value for the estimated cracking time of any string, we conservatively use the value only to induce an ordinal ranking on the strength of our cracked passwords.

Using this metric, we are able to test several hypotheses about the impact of factors such as the time a brain wallet was created or the total amount stored on the strength of the passwords chosen. For each hypothesis we computed the (non-parametric) Spearman's rank-correlation coefficient (ρ) against a null hypothesis of no correlation. We did not observe statistically significant correlations ($p > 0.1$ in all cases) between the estimated password strength and the date the brain wallet address was initially used or the total amount ever sent to the address or the time it took for the wallet to initially be drained of funds. This suggests that, consistent with previous password research, we find no evidence that users

are able to pick stronger passwords when protecting a larger quantity of money. A major limitation of using `wheelerzxcvbn` as a password metric is that it's not as accurate with passphrases. Since many of the brain wallets are secured by passphrases, rather than single passwords, `wheelerzxcvbn` overestimates the strength of some of the weaker passphrases, muddling our results.

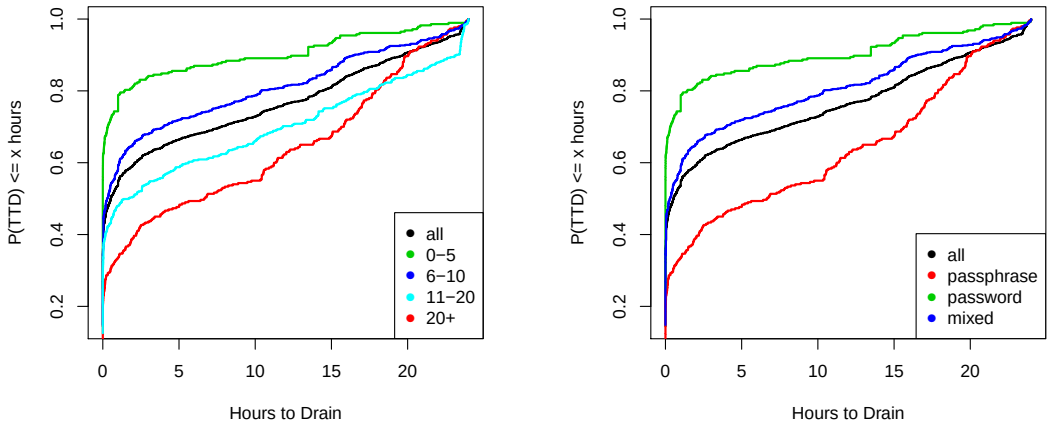

(a) CDF of the # of hours to drain brain wallets for wallets by length of password/passphrase

(b) CDF of the # of hours to drain brain wallets for wallets by type of list

Figure 6.5: Measuring time-to-drain for passwords versus passphrases

Since standard password strength metrics are not appropriate for passphrases, we used other metrics to discern the impact of using passphrases over passwords on the time attackers take to drain the wallet. Figure 6.5a measures this by looking at the length of the password. This chart shows that shorter length passwords (presumably one word) take a shorter amount of time to drain. The green line and dark blue lines (passwords of 0-5 and 6-10 characters) are drained substantially faster than the black line (the average).

We can also look at this by looking at the time to drain based on which password list we found the corresponding brain wallet. Figure 6.5b shows the time to drain based on this distinction. Again, we see that brain wallets found in password lists (like our Industry lists and password dumps) were drained faster than those found in our passphrase lists (like our lyrics scrape and quotes). Drains of mixed lists (like messages scraped from BitSig.io

and our generated misc list based on permutations on found brain wallets) were drained at a similar rate to an average brain wallet.

6.3 Conclusion

The idea behind brain wallets is elegant and alluring: remembering a password is surely easier than a private key. Unfortunately, as this chapter makes clear, it is also an extremely insecure way to store bitcoin. Drainers lurk over the blockchain, ready to pounce as soon as new brain wallets are established.

By examining 3.9 trillion candidate passwords, we found 1 730 brain wallets that were active at some point in time. Unfortunately, we also found that nearly all were drained – usually quickly. Brain wallets created with passphrases were drained slower than those created with a password, but both were usually drained within a day. While our findings are necessarily incomplete, they certainly suggest that brain wallets are not a secure method for using bitcoin. Perhaps the most surprising result of our analysis is the relative scarcity of brain wallets in use today. This is actually quite encouraging, because it means that fewer users are at risk to these attacks than has previously been supposed.

CHAPTER 7

CONCLUSION

Bitcoin is a public, decentralized currency that allows anybody to trace all transactions made using the currency. The decentralized nature of Bitcoin makes it particularly attractive to cybercriminals. These features allow us to use Bitcoin as a tool to measure cybercriminal activity.

We concentrate on Bitcoin-*based* cybercrime, as opposed to Bitcoin-*facilitated* cybercrime. Bitcoin-based cybercrime is crime that can happen only because Bitcoin exists. It is spread and enabled within the Bitcoin ecosystem; the criminals and victims are both actors in this ecosystem which helps perpetuate the crime. Bitcoin-facilitated cybercrime uses Bitcoin solely for its payment mechanism or as a monetization strategy. For instance, Ponzi schemes that accept payments in Bitcoin and use the Bitcoin forums to rally behind their scheme are Bitcoin-based. However, Bitcoin malware that uses victim computers to mine Bitcoin is a Bitcoin-facilitated crime; here Bitcoin is a tool to surreptitiously earn money rather than a place to garner victims from as well as monetize them. Other examples of Bitcoin-facilitated cybercrime include ransomware and anonymous online marketplaces.

This work falls in the general category of security economics: research that tackles computer security issues by examining the incentives of attackers and defenders rather than by purely technical means. Using the Bitcoin ecosystem, we measure crimes in greater detail than previous researchers have been able to achieve. Consequently, we also increase understanding of particular cybercriminal activities beyond when they involve Bitcoin.

The methodology and analysis in this dissertation enables researchers to measure Bitcoin-based crimes by leveraging various aspects of the Bitcoin ecosystem. For example, we first gather information about a crime of interest using Bitcoin forums and other conversation websites like the Bitcoin subreddit. We then use that information to start a deeper investigation throughout the Bitcoin ecosystem using sources like the forum posts and also external information like the Bitcoin blockchain, third-party defender websites, and third-party password and passphrase dumps. We then analyze this information, keeping in

mind information learned at every step. We remove false positive information in line with false positives known in the data sources and the Bitcoin ecosystem at large. A similar methodological approach can be used to measure other forms of cybercrime that rely on small ecosystems such as Bitcoin.

We investigate three types of Bitcoin-based crimes: distributed denial of service (DDoS) attacks, scams, and brain wallets (Bitcoin stored using the hash of a weak password or passphrase). Chapter 3 quantifies the effect of DDoS attacks on Bitcoin mining pools and currency exchanges. Chapter 4 analyzes the revenues from various scams in the Bitcoin ecosystem, and Chapter 6 measures the Bitcoin lost when stored with weak passwords. But, while this dissertation focuses primarily on direct costs of crime, most of the harm from these crimes is indirect. For instance, while a mining pool might lose money when they are subject to a DDoS attack, the greater harm is a Bitcoin miner deciding to stop mining due to the incessant attacks. Users might not buy Bitcoin if they are unable to tell the difference between a legitimate currency exchange and a fraudulent one. Bitcoin users might stop transacting in Bitcoin after their money, tied to a weak password, is stolen.

This dissertation also advances our understanding of how these particular Bitcoin-based cybercrimes operate and the impact they have. The analysis also often sheds new light on cybercrimes writ large, due to the additional information newly available by studying how the crimes happen in Bitcoin.

Chapter 3 analyzes DDoS attacks on various Bitcoin services. We quantify these attacks over time as well as various countermeasures, like the use of anti-DDoS proxy services, that defenders use to counter the attacks. We find that up until late 2013, at least 7.4% of Bitcoin-related services have experienced DDoS attacks. Bitcoin currency exchanges were the most frequent DDoS target during this time period. We investigate Bitcoin mining pools further and find that over 60% of large mining pools have been DDoSed whereas only 17% of small ones have been DDoSed. This implies that big mining pools are targeted for DDoS attacks, potentially by other pools. Johnson et al. analyze this from a game theoretic lens, showing that it is cheaper for many pools to DDoS a larger pool rather than invest more in pool mining resources [42].

We systematically measure the revenues of scams using Bitcoin in Chapter 4. We measure transactions in and out of 42 scams and estimate that at least $11 million has been lost to these scams from 13 000 victims. Most of this revenue was earned by Ponzi schemes that also accept currencies other than Bitcoin and are advertised in other non-Bitcoin-related places. We find more Ponzi schemes that exclusively accept Bitcoin and are only advertised on Bitcoin forums, but each scheme on average lasts a shorter amount of time and earns less money. Chapter 5 measures Ponzi schemes advertised on the Bitcoin forums in greater depth. This chapter focuses on the ecosystem around Ponzi schemes by studying the forums on which they are advertised. We find over 1 700 of these scams, half of which end within a week of being started. Using a Cox proportional hazards model, we find that the frequency of victim and scammer posts is negatively correlated with scam survival. Furthermore, scams also collapse more quickly if the associated scammer registers their Bitcoin forum account on the same day that they first post about the scam. By contrast, scams that are promoted by shills tend to survive for longer.

Chapter 6 looks at brain wallets, or Bitcoin wallets secured by the hash of a password or passphrase. We use 3.9 trillion candidate passwords and passphrases to find 1 730 brain wallets using 1 686 different passwords which have received 2 846 BTC (approx. $261K). Almost $100K of these passwords were found using Reddit usernames and comments, eight of which still have a balance as of March 2017. Nearly all of the identified 1 730 brain wallets were drained, many within seconds. Brain wallets created with passphrases were drained slower than those created with a password, but both were nearly always drained within a day. We do not find evidence that attackers wait to drain wallets, though sometimes the amount in the brain wallet is too low to justify the transaction fee. We find that mining pools, particularly F2Pool, empty brain wallets by claiming the balance as a transaction fee, as a way to reduce the amount of unspent transaction outputs on the Bitcoin network.

7.1 Future Work

The work done in this dissertation lends itself to future work both in measuring Bitcoin-based cybercrime but also in other areas.

7.1.1 Effects of Technological Counters to DDoS

Chapter 3 showed how platforms subject to DDoS attacks were more likely to use anti-DDoS proxy services. However, we did not investigate the effects of other anti-DDoS countermeasures. For instance, different mining pool schemes influence the prevalence and success of DDoS attacks. It is unclear if schemes like peer-to-peer pools more prone to DDoS attacks, or if there are other explanatory factors for their lack of prominence.

7.1.2 Market Responses to DDoS Attacks

Moore and Christin found that transaction volume mattered more than attack susceptibility when predicting the future viability of a Bitcoin exchange [55]. It is unclear if this model carry over to Bitcoin mining pools. The case study of DeepBit which lost its market dominance due to repeated DDoS attacks would suggest not.

7.1.3 Compounding Effects of Attacks

Decker and Wattenhofer measured the effects of transaction malleability attacks on the now-defunct currency exchange Mt. Gox [28]. These attacks were happening around the same time that DDoS attacks were supposedly being perpetuated on Mt. Gox. More work needs to be done to investigate the ties between these contemporaneous attacks, to see if the same people are profiting off of both types of attack.

7.1.4 Effects of Default Standards on Security

The popular Bitcoin brain wallet generating website `brainwallet.org` (now defunct) had two default passwords: the empty string and "correct horse battery staple". Those two passwords were also two of the most popular brain wallet passwords. This website also defaulted to uncompressed keys. While most wallets spent recently are compressed, most brain wallets are still uncompressed. Möser and Böhme saw that default transaction fees were the most commonly used fees [57]. We can measure how default options affect security. For instance ether.camp implements an Ethereum (a different cryptocurrency) wallet which defaults to using a brain wallet with the Keccak hash of a users' login password as the private key. We can also measure how address reuse (a known privacy and security risk) and how default software options contribute to this behavior.

7.1.5 Effects of Social Behaviors on the Profits of Ponzi Schemes

Chapter 5 looks at the social behaviors that influence the lifetime of various Ponzi schemes. Some of these schemes are included in the analysis on the profits of Ponzi schemes in Chapter 4. We can then measure the effects of various measures (like scammer interaction, scammer reputation, and victim posting) studied in depth in Chapter 5 on the profits brought in by the scam. We want to ascertain whether there are different lessons learned for the lifetime of the scam as measured by comments rather than the profits of the scam.

7.1.6 Scam Early Detection System

It is hard for users to know which cryptocurrency services are legitimate or not, since none have been around for very long. Chapter 4 found that scammers swindled over $11 million worth of bitcoin through fraudulent services.A scam early detection system could help correct this market failure of asymmetric information. This would feed in data from the public forums as well as transaction data from the public blockchain. Such a system would also be able to detect scams that initially behave as legitimate services, but then take all the accumulated money and run after they encounter a large enough sum of money. These so-called "exit scams" currently plague the Bitcoin economy, and there are currently no mechanisms in place to discourage or otherwise distinguish them.

7.1.7 Cryptocurrency Service Legitimacy Indicators

Many cryptocurrency-related services are advertised on various forums, but no clear, consistent ways to measure the relative safety of each service. In Chapter 6, we found that users' lack of understanding on how to secure their bitcoin let criminals steal $260 000 worth of bitcoin. A sequence of indicators of legitimacy for cryptocurrency services could correct this information asymmetry. These indicators range from technical factors, such as the techniques the service uses to secure their customers' currency, to other factors, such as the service's interactions on social media. Such a system would help correct the *moral hazard* inherent in these services that have little regulatory oversight. It would also encourage legitimate services to enact features that help protect their and their customers' bitcoin.

BIBLIOGRAPHY

[1] Allison, I. *CoinWallet says Bitcoin stress test in September will create 30-day backlog.* `http://www.ibtimes.co.uk/coinwallet-plans-bitcoin-dust-attack-september-create-30-day-transaction-backlog-1515981`. Aug. 2015.

[2] Amazon Web Services *Announcement: Amazon EC2 Public IP Ranges.* `https://forums.aws.amazon.com/ann.jspa?annID=1701`. Last accessed November 21, 2013.

[3] Anderson, R., Barton, C., Böhme, R., Clayton, R., Eeten, M. van, Levi, M., Moore, T., AND Savage, S. Measuring the Cost of Cybercrime. In: *11th Workshop on the Economics of Information Security (WEIS)*. 2012.

[4] Anderson, R., AND Moore, T. The Economics of Information Security. *Science* 314, 5799 (2006), 610–613.

[5] Andresen, G. *Time to roll out bigger blocks.* `http://gavinandresen.ninja/time-to-roll-out-bigger-blocks`. May 2015.

[6] Andrychowicz, M., Dziembowski, S., Malinowski, D., AND Mazurek, Ł. On the malleability of Bitcoin transactions. In: *Financial Cryptography and Data Security*. Springer. 2015, 1–18.

[7] Back, A. *Hashcash - a denial of service counter-measure.* `http://hashcash.org/papers/hashcash.pdf`. 2002.

[8] Baqer, K., Huang, D. Y., McCoy, D., AND Weaver, N. Stressing Out: Bitcoin Stress Testing. In: *3rd Workshop on Bitcoin and Blockchain Research*. Springer, 2016, 3–18.

[9] Barber, S., Boyen, X., Shi, E., AND Uzun, E. Bitter to better: How to make Bitcoin a better currency. In: *Financial Cryptography and Data Security*. Springer, 2012, 399–414.

[10] Bartoletti, M., Carta, S., Cimoli, T., AND Saia, R. Dissecting Ponzi schemes on Ethereum: identification, analysis, and impact. *arXiv preprint arXiv:1703.03779* (2017).

[11] Bitcoin Foundation *Removing Impediments to Bitcoin's Success: A Risk Management Study*. `https://bitcoinfoundation.org/static/2014/04/Bitcoin-Risk-Management-Study-Spring-2014.pdf`. 2014.

[12] Bitcoin Wiki *Category: Pool Operators*. `https://en.bitcoin.it/wiki/Category:Pool_Operators`. Last accessed November 21, 2013.

[13] Bitcoin Wiki *Trade*. `https://en.bitcoin.it/wiki/Trade`. Last accessed November 21, 2013.

[14] Böhme, R., Christin, N., Edelman, B., AND Moore, T. Bitcoin: Economics, Technology, and Governance. *Journal of Economic Perspectives* 29, 2 (2015), 213–38.

[15] Böhme, R., AND Moore, T. The iterated weakest link - a model of adaptive security investment. In: *8th Workshop on the Economics of Information Security*. 2009.

[16] Bohr, J., AND Bashir, M. Who uses Bitcoin? An exploration of the Bitcoin community. In: *Twelfth Annual International Conference on Privacy, Security and Trust*. IEEE. 2014, 94–101.

[17] Bonneau, J. Statistical metrics for individual password strength. In: *20th International Workshop on Security Protocols*. Cambridge, UK, Apr. 2012.

[18] Bonneau, J. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In: *2012 IEEE Symposium on Security and Privacy*. San Francisco, CA, USA, May 2012.

[19] Bonneau, J. Why buy when you can rent? Bribery attacks on Bitcoin-style consensus. In: *3rd Workshop on Bitcoin and Blockchain Research*. Springer, 2016.

[20] Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., AND Felten, E. W. Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In: *IEEE Symposium on Security and Privacy*. San Francisco, CA, USA, May 2015.

[21] Caldwell, M., AND Voisine, A. *BIP 38: Passphrase-protected private key*. Nov. 2012.

[22] Carné de Carnavalet, X. de, AND Mannan, M. From very weak to very strong: Analyzing password-strength meters. In: *Network and Distributed System Security Symposium (NDSS 2014)*. Internet Society. 2014.

[23] Certicom Research *SEC 1: Elliptic Curve Cryptography*. `https://www.secg.org/sec1-v2.pdf`. Jan. 2010.

[24] Christin, N. Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In: *22nd International World Wide Web Conference*. 2013, 213–224.

[25] CloudFlare *CloudFlare IP Ranges*. `http://www.cloudflare.com/ips`. Last accessed November 21, 2013.

[26] Courtois, N., Song, G., AND Castellucci, R. *Speed Optimizations in Bitcoin Key Recovery Attacks*. `http://eprint.iacr.org/2016/103.pdf`.

[27] Dagher, G. G., Bünz, B., Bonneau, J., Clark, J., AND Boneh, D. Provisions: Privacy-preserving proofs of solvency for Bitcoin exchanges. In: *22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM. 2015, 720–731.

[28] Decker, C., AND Wattenhofer, R. Bitcoin transaction malleability and MtGox. In: *European Symposium on Research in Computer Security*. Springer. 2014, 313–326.

[29] Drew, J., AND Moore, T. Automatic identification of replicated criminal websites using combined clustering. In: *International Workshop on Cyber Crime*. IEEE. 2014, 116–123.

[30] Eskandari, S., Barrera, D., Stobert, E., AND Clark, J. A First Look at the Usability of Bitcoin Key Management. In: *Workshop on Usable Security*. 2015.

[31] Eyal, I., AND Sirer, E. G. Majority is not Enough: Bitcoin Mining is Vulnerable. In: *Financial Cryptography and Data Security*. Lecture Notes in Computer Science. Springer, Mar. 2014.

[32] Feder, A., Gandal, N., Hamrick, J., AND Moore, T. The Impact of DDoS and Other Security Shocks on Bitcoin Currency Exchanges: Evidence from Mt. Gox. In: *15th Workshop on the Economics of Information Security (WEIS)*. 2016.

[33] Fernández-Villaverde, J., AND Sanches, D. *Can currency competition work?* NBER Working Paper No. 22157. 2016.

[34] gmaxwell *#bitcoin-wizards*. https://botbot.me/freenode/bitcoin-wizards/2015-09-22/. 2015.

[35] Harel, U. *Restricting direct access to your website (Incapsula's IP addresses)*. http://support.incapsula.com/hc/en-us/articles/200627570-Restricting-direct-access-to-your-website-Incapsula-s-IP-addresses-. Last accessed January 15, 2014.

[36] Heilman, E., Kendler, A., Zohar, A., AND Goldberg, S. Eclipse attacks on Bitcoin's peer-to-peer network. In: *24th USENIX Security Symposium*. 2015, 129–144.

[37] Herley, C. So long, and no thanks for the externalities: The rational rejection of security advice by users. In: *Workshop on New Security Paradigms*. ACM. 2009, 133–144.

[38] Herley, C., AND Florêncio, D. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In: *Workshop on the Economics of Information Security*. Springer, 2009, 33–53.

[39] Hilton, S. *Dyn Analysis of Friday October 21 Attack*. http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/.

[40] Huang, D. Y., Dharmdasani, H., Meiklejohn, S., Dave, V., Grier, C., McCoy, D., Savage, S., Weaver, N., Snoeren, A. C., AND Levchenko, K. Botcoin: Monetizing stolen cycles. In: *Network and Distributed System Security Symposium*. 2014.

[41] Hutchinson, L. *All Android-created Bitcoin wallets vulnerable to theft*. http://arstechnica.com/security/2013/08/all-android-created-bitcoin-wallets-vulnerable-to-theft/.

[42] Johnson, B., Laszka, A., Grossklags, J., Vasek, M., AND Moore, T. Game-Theoretic Analysis of DDoS Attacks Against Bitcoin Mining Pools. In: *1st Workshop on Bitcoin Research*. Vol. 8438. Lecture Notes in Computer Science. Springer, Mar. 2014, 72–86.

[43] Kharraz, A., Robertson, W., Balzarotti, D., Bilge, L., AND Kirda, E. Cutting the Gordian knot: A look under the hood of ransomware attacks. In: *International Confer-*

ence on Detection of Intrusions and Malware, and Vulnerability Assessment. Springer. 2015, 3–24.

[44] Kroll, J., Davey, I., AND Felten, E. The Economics of Bitcoin Mining, or Bitcoin in the Presence of Adversaries. In: *Twelfth Annual Workshop on the Economics of Information Security (WEIS'13)*. Washington, DC, June 2013.

[45] Laszka, A., Johnson, B., AND Grossklags, J. When Bitcoin mining pools run dry. In: *Bitcoin Workshop*. Springer. 2015, 63–77.

[46] Leyden, J. How mystery DDoSers tried to take down Bitcoin exchange with 100Gbps crapflood. *The Register* (Oct. 2013). `http://www.theregister.co.uk/2013/10/17/bitcoin_exchange_ddos_flood/`.

[47] Liao, K., Zhao, Z., Doupé, A., AND Ahn, G.-J. Behind Closed Doors: Measurement and Analysis of a CryptoLocker Ransoms in Bitcoin. In: *Eleventh APWG eCrime Researcher's Summit*. June 2016.

[48] Maurer, B., Nelms, T. C., AND Swartz, L. "When perhaps the real problem is money itself!": the practical materiality of Bitcoin. *Social Semiotics* 23, 2 (2013), 261–277.

[49] Maxwell, G. *CoinJoin: Bitcoin privacy for the real world*. `https://bitcointalk.org/index.php?topic=279249`. Last accessed 29 August 2014.

[50] McCorry, P., Shahandashti, S. F., AND Hao, F. Refund attacks on Bitcoin's Payment Protocol. In: *Financial Cryptography and Data Security*. 2016.

[51] McCoy, D., Pitsillidis, A., Grant, J., Weaver, N., Kreibich, C., Krebs, B., Voelker, G., Savage, S., AND Levchenko, K. Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs. In: *USENIX Security Symposium*. 2012, 1–16.

[52] Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., AND Savage, S. A Fistful of Bitcoins: Characterizing Payments Among Men with No Names. In: *Conference on Internet Measurement Conference*. IMC '13. ACM, Barcelona, Spain, 2013, 127–140.

[53] Mirkovic, J., AND Reiher, P. A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM SIGCOMM Computer Communication Review* 34, 2 (2004), 39–53.

[54] Moore, D., Shannon, C., Brown, D. J., Voelker, G. M., AND Savage, S. Inferring Internet Denial-of-Service activity. *ACM Transactions on Computer Systems* 24, 2 (2006), 115–139.

[55] Moore, T., AND Christin, N. Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk. In: *Financial Cryptography and Data Security*. Vol. 7859. Lecture Notes in Computer Science. Springer, Apr. 2013, 25–33.

[56] Moore, T., Han, J., AND Clayton, R. The Postmodern Ponzi Scheme: Empirical Analysis of High-Yield Investment Programs. In: *Financial Cryptography and Data Security*. Vol. 7397. Lecture Notes in Computer Science. Springer, 2012, 41–56.

[57] Möser, M., AND Böhme, R. Trends, tips, tolls: A longitudinal study of Bitcoin transaction fees. In: *2nd Workshop on Bitcoin and Blockchain Research*. Springer. 2015, 19–33.

[58] Möser, M., Böhme, R., AND Breuker, D. An Inquiry into Money Laundering Tools in the Bitcoin Ecosystem. In: *8th APWG eCrime Researchers Summit*. IEEE, 2013.

[59] Motoyama, M., Meeder, B., Levchenko, K., Voelker, G. M., AND Savage, S. Measuring Online Service Availability Using Twitter. In: *Workshop on Online Social Networks*. Usenix. 2010.

[60] Nakamoto, S. *Bitcoin: A Peer-to-Peer Electronic Cash System*. `http://www.bitcoin.org/bitcoin.pdf`. 2009.

[61] Narayanan, A., Bonneau, J., Felten, E., Miller, A., AND Goldfeder, S. *Bitcoin and Cryptocurrency Technologies*. 2016.

[62] Nayak, K., Kumar, S., Miller, A., AND Shi, E. Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In: *IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE. 2016, 305–320.

[63] Nazario, J. Politically motivated denial of service attacks. *The Virtual Battlefield: Perspectives on Cyber Warfare* (2009), 163–181.

[64] Neisius, J., AND Clayton, R. Orchestrated Crime: The High Yield Investment Fraud Ecosystem. In: *Eighth APWG eCrime Researcher's Summit*. Birmingham, AL, Sept. 2014.

[65] organofcorti MTGOX volume post Dwolla: A single statistical test. *Neighbourhood Pool Watch* (July 2013). `http://organofcorti.blogspot.com/2013/07/114-mtgox-volume-post-dwolla-single.html`.

[66] Resnick, P., Kuwabara, K., Zeckhauser, R., AND Friedman, E. Reputation systems. *Communications of the ACM* 43, 12 (2000), 45–48.

[67] Ritter, A., Wright, E., Casey, W., AND Mitchell, T. Weakly supervised extraction of computer security events from Twitter. In: *24th International Conference on World Wide Web*. ACM. 2015, 896–905.

[68] Ron, D., AND Shamir, A. Quantitative analysis of the full Bitcoin transaction graph. In: *Financial Cryptography and Data Security*. Vol. 7859. Lecture Notes in Computer Science. Springer, 2013, 6–24.

[69] Ron, D., AND Shamir, A. How Did Dread Pirate Roberts Acquire and Protect His Bitcoin Wealth? In: *1st Workshop on Bitcoin Research*. Vol. 8438. Lecture Notes in Computer Science. Springer, Mar. 2014.

[70] Rosenfeld, M. *Analysis of hashrate-based double-spending*. 2012. URL: `https://bitcoil.co.il/Doublespend.pdf`.

[71] Sapirshtein, A., Sompolinsky, Y., AND Zohar, A. Optimal selfish mining strategies in Bitcoin. In: *Financial Cryptography and Data Security*. 2016.

[72] *Satoshi Dice: Bitcoin Gambling and Casino Games*. `https://satoshidice.com/`.

[73] Schrijvers, O., Bonneau, J., Boneh, D., AND Roughgarden, T. Incentive Compatibility of Bitcoin Mining Pool Reward Functions. In: *Financial Cryptography and Data Security*. 2016.

[74] Securities and Exchange Commission *SEC v. Trendon T. Shavers, et al.* `http://www.sec.gov/litigation/complaints/2013/comp-pr2013-132.pdf`.

[75] Shen, W., Hu, Y. J., AND Ulmer, J. R. Competing for Attention: An Empirical Study of Online Reviewers' Strategic Behavior. *Mis Quarterly* 39, 3 (2015), 683–696.

[76] Soska, K., AND Christin, N. Measuring the Longitudinal Evolution of the Online Anonymous Marketplace Ecosystem. In: *USENIX Security Symposium*. 2015, 33–48.

[77] Teutsch, J., Jain, S., AND Saxena, P. When cryptocurrencies mine their own business. In: *Financial Cryptography and Data Security*. 2016.

[78] Vieweg, S., Hughes, A. L., Starbird, K., AND Palen, L. Microblogging during two natural hazards events: What Twitter may contribute to situational awareness. In: *Conference on Human Factors in Computing Systems*. ACM. 2010, 1079–1088.

[79] Weir, M., Aggarwal, S., Collins, M., AND Stern, H. Testing metrics for password creation policies by attacking large sets of revealed passwords. In: *17th ACM conference on Computer and communications security*. ACM. 2010, 162–175.

[80] Wheeler, D. L. zxcvbn: Low-budget password strength estimation. In: *USENIX Security*. 2016.

[81] znort987 *blockparser*. `https://github.com/znort987/blockparser`.

[82] Zuckerman, E., Roberts, H., McGrady, R., York, J., AND Palfrey, J. G. *2010 Report on Distributed Denial of Service (DDoS) Attacks*. Tech. rep. 2010-16. `http://ssrn.com/abstract=1872065`. Berkman Center Research Publication, Dec. 2010.