

Analyzing the Bitcoin Ponzi Scheme Ecosystem

Marie Vasek and Tyler Moore

¹ Computer Science, University of New Mexico, `lastname@cs.unm.edu`

² Tandy School of Computer Science, The University of Tulsa,
`firstname-lastname@utulsa.edu`

Abstract. This paper analyzes the supply and demand for Bitcoin-based Ponzi schemes. There are a variety of these types of scams: from long cons such as Bitcoin Savings & Trust to overnight doubling schemes that do not take off. We investigate what makes some Ponzi schemes successful and others less so. By scouring 11 424 threads on `bitcointalk.org`, we identify 1 780 distinct scams. Of these, half lasted a week or less. Using survival analysis, we identify factors that affect scam persistence. One approach that appears to elongate the life of the scam is when the scammer interacts a lot with their victims, such as by posting more than a quarter of the comments in the related thread. By contrast, we also find that scams are shorter-lived when the scammers register their account on the same day that they post about their scam. Surprisingly, more daily posts by victims is associated with the scam ending sooner.

Keywords: Bitcoin, cybercrime measurement

1 Introduction

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 pseudoanonymous and distributed. Consequently, it is perhaps unsurprising that many Bitcoin users have taken to Ponzi schemes (and Ponzi scheme runners to Bitcoin).

In this paper, we look at the ecosystem around Ponzi schemes advertised to Bitcoin users. Others have 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 skills 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 previous work, 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 paper are both in the data collection methodology and in the analysis of the gathered data.

- Section 2 outlines our data collection contributions: 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 (found in Section 3) leverage this novel dataset to describe supply-side characteristics of scams and scammers as well as describe demand-side characteristics of victims.

2 Methodology

We aim to measure scams by collecting data from the places they are advertised. This helps us generate a comprehensive list of advertised scams. For the purposes of this study, we elect to focus on Ponzi schemes exclusively. Of course, there are many different types of scams affecting Bitcoin, as shown by Vasek and Moore [12]. 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 confident in the comprehensiveness of our approach.

We collect our data from the forum `bitcointalk.org`. This forum was created in the same year as Bitcoin and by the same pseudonymous entity³. This forum is the top place where Bitcoin community members go to discuss Bitcoin and is currently one of the 1 000 most popular websites in the world, according to Alexa. We chose this as our sole source for this work based both on its popularity in the Bitcoin ecosystem and its popularity within the subsection of Ponzi scheme investors within the Bitcoin ecosystem, as we found in our previous work in this space [12].

In order to collect information about the scams, we crawl 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. We can find a number of Ponzi schemes advertised in other subforums of `bitcointalk`. However, we choose the two most popular subforums for Ponzi schemes that had the highest signal to noise ratio: scam accusations and games and rounds. In total, we crawl 11 424 threads on these three subforums from June 2011 through November 2016. We consider all the subforums of `bitcointalk`

³ <https://bitcointalk.org/index.php?topic=5>

where we found any posts advertising a Ponzi scheme. We then look at the forums, particularly for Ponzi schemes. We omit subforums like the gambling subforum which predominantly contained posts about online card games and other non-Ponzi scheme activity.



Fig. 1: Screenshots of the initial posting for the Ponzi scheme and an example victim response.

Since threads on these forums cover other topics than just promoting Ponzi schemes, we refine this further to threads that referenced “ponzi” or “hyip” in the first 10 comments. We then process 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 merge threads containing the same domain or Bitcoin address, since many scams were advertised multiple times or in different places. Note that we throw out threads containing a whitelist of legitimate gambling domains⁴. We also do 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 closely analyzed 10 different scams for which we had ground truth on

⁴ This list was curated by bitcointalk user **mem** here: <https://bitcointalk.org/index.php?topic=75883.0>.

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 analyze their most recent posting history (maximum 20 posts). We parse out the other subforums they posted in as well as the number of times they posted on any given Ponzi-related thread. For scammers, we identify their public interaction with victims; similarly, for victims, we identify their public interactions with scammers. We also find evidence of every user’s public history on the forums.

We classify shills as victims who post only about a single scam and nowhere else on the forum. We devise 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 attempt to identify these posters automatically, and the most straightforward way is 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 concluded from manual inspection that this simple approach provides an effective approach to identify many shills without miscategorizing legitimate users.

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 that we use 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.

3 Results

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

Figure 2 shows the lifetimes of the scams, where lifetime is measured as the length of time between the first post about a given scam and the last. 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.




Fig. 2: Survival analysis of the lifetime of scams.

3.1 Scammer Interaction and Scam Lifetime

Figure 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 find 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 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.

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 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 notice that scammers that only post on one scam have a lower lifetime compared to scammers that post outside of just one scam. The




Fig. 3: Lifetime of the scam based on the fraction of the comments about the scam from the scammer.

difference in these survival lifetimes are significant at the $p=0.01$ level. Figure 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 discover 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.

3.2 Victim Behavior

We measure 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 4 we separate 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 1 shows how Ponzi scheme victims' post history compares to that of other users active on bitcointalk. For this, we scrape bitcointalk's aggregated posts statistics for ground truth and categorize each post using bitcointalk's categories. The Ponzi victims' post history is 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




Fig. 4: Lifetime of the scam by interaction by “shill” commenters.

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

We can also observe 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 6 analyzes this effect further. While most victims post within a week, there is quite a long tail of victim posts. We discover new victims posting over half a year after the start of the initial scam posting.

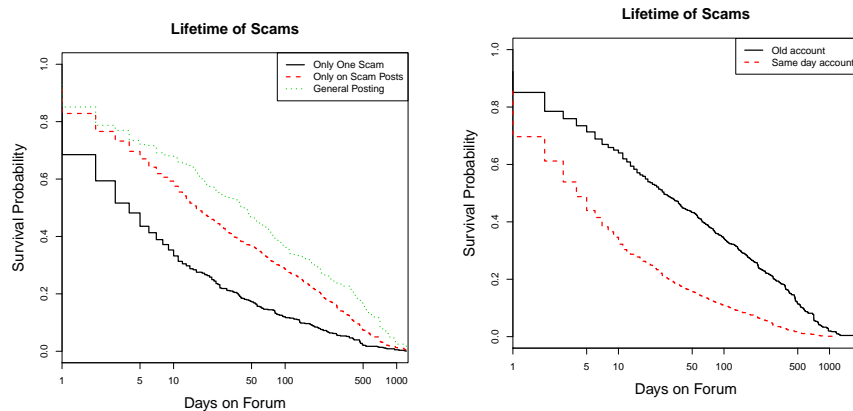
3.3 Proportional Hazards Model

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 use the following:

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 3.1) has posted anywhere in the thread. This accounts for their presence, since the number of comments by these users is so low.



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

Fig. 5: Measuring lifetimes of scams based on attacker accounts.




Fig. 6: Number of victim posts after a thread starts.

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

Table 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 $\exp(\text{coef})$ 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

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 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.

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

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 shorter scam lifetime. This confirms the intuition from Figure 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.

4 Related Work

This paper fits into the greater literature of reputation mechanisms. Resnick et al. provide a general overview for reputation systems as well as drawbacks in them [9]. Shen et al. provide analysis of reviewers posting about products on online retailers [10]. 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.

This paper also fits into greater literature about Bitcoin. Bitcoin has a small community of actors [2]. Maurer et al. associated the distributed network of Bitcoin nodes with the distributed network of conversations, like those found on the Bitcoin forums [6]. We agree 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.

To this end, other researchers have mined Bitcoin forum posts to infer activity in the Bitcoin ecosystem. Vasek et al. searched for reports of DDoS attacks to infer after the fact when they occurred [13]. Fleder et al. searched for Bitcoin addresses to categorize them [4]. Using this information, they were able to tie bitcointalk users to Silk Road transactions. Similarly, Vasek and Moore use bitcointalk to identify addresses for potential Bitcoin scams [12] and Liao et al. use the Bitcoin subreddit to seed their ransomware address finder [5]. Most similarly to this paper’s methodology, Xie et al. analyze how, among other things, the social network in the Bitcoin forums leads to price swings in Bitcoin [14]. They found that bitcointalk users that invite long discussions are more likely to share relevant information. When looking at the connectedness of bitcointalk users, they found that the more connected the users are at a given time, the more intense the trading frequency is.

Our work also falls in the literature on online Ponzi schemes, also known as high yield investment programs (HYIPs). Moore et al. overviewed the ecosystem using HYIP aggregator websites [7]. They found that the lifetime of any given HYIP could be predicted by interest payments and the mandatory investment term. Neisius and Clayton followed up on this work, concentrating on the incentives promoting this criminal behaviors [8]. They found that HYIP operators paid to be listed on aggregator websites and also received a referral bonus for users directed to HYIPs. They also crawled the criminal forums behind people that run HYIPs and HYIP aggregator websites, and found that the majority of these criminals are based in the US. Drew and Moore found clusters of replicated HYIP websites, pointing to the high use of HYIP kits in creating Ponzi scheme websites [3]. Vasek and Moore carried out the first analysis of Bitcoin-based Ponzi schemes [12]. They directly measured the profits of 32 Ponzi schemes and found that these scammers were bringing in over \$7 million. Bartoletti et al. analyzed Ponzi schemes using the cryptocurrency Ethereum and found similar results as Vasek and Moore found with Bitcoin-based Ponzi scams [1][12]. 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 [11]. 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.

5 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 analyze that shill interaction, or users that only post in one thread, and discover

that it lengthens the life of the scam. We demonstrate 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.

References

1. Massimo Bartoletti, Salvatore Carta, Tiziana Cimoli, and Roberto Saia. Dissecting ponzi schemes on Ethereum: identification, analysis, and impact. *arXiv preprint arXiv:1703.03779*, 2017.
2. Jeremiah Bohr and Masooda Bashir. Who uses bitcoin? an exploration of the bitcoin community. In *Twelfth Annual International Conference on Privacy, Security and Trust*, pages 94–101. IEEE, 2014.
3. Jake Drew and Tyler Moore. Automatic identification of replicated criminal websites using combined clustering. In *International Workshop on Cyber Crime*, pages 116–123. IEEE, 2014.
4. Michael Fleder, Michael S Kester, and Sudeep Pillai. Bitcoin transaction graph analysis. *arXiv preprint arXiv:1502.01657*, 2015.
5. Kevin Liao, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. Behind closed doors: Measurement and analysis of a CryptoLocker ransoms in Bitcoin. In *Eleventh APWG eCrime Researcher’s Summit*, June 2016.
6. Bill Maurer, Taylor C Nelms, and Lana Swartz. “When perhaps the real problem is money itself!”: the practical materiality of Bitcoin. *Social Semiotics*, 23(2):261–277, 2013.
7. Tyler Moore, Jie Han, and Richard Clayton. The postmodern Ponzi scheme: Empirical analysis of high-yield investment programs. In *Financial Cryptography and Data Security*, volume 7397 of *Lecture Notes in Computer Science*, pages 41–56. Springer, 2012.
8. Jens Neisius and Richard Clayton. Orchestrated crime: The high yield investment fraud ecosystem. In *Proceedings of the Eighth APWG eCrime Researcher’s Summit*, Birmingham, AL, September 2014.
9. Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.
10. Wenqi Shen, Yu Jeffrey Hu, and Jackie Rees Ulmer. Competing for attention: An empirical study of online reviewers’ strategic behavior. *Mis Quarterly*, 39(3):683–696, 2015.
11. Kyle Soska and Nicolas Christin. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *USENIX Security Symposium*, pages 33–48, 2015.
12. Marie Vasek and Tyler Moore. There’s no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency scams. In *Financial Cryptography and Data Security*, volume 8975 of *Lecture Notes in Computer Science*, pages 44–61. Springer, January 2015.

13. Marie Vasek, Micah Thornton, and Tyler Moore. Empirical analysis of denial-of-service attacks in the Bitcoin ecosystem. In *1st Workshop on Bitcoin Research*, volume 8438 of *Lecture Notes in Computer Science*, pages 57–71. Springer, March 2014.
14. Peng Xie, Hailiang Chen, and Yu Jeffrey Hu. Network structure and predictive power of social media for the Bitcoin market. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2894089, 2017.