

The Bitcoin Brain Drain: Examining the Use and Abuse of Bitcoin Brain Wallets

Marie Vasek¹, Joseph Bonneau², Ryan Castellucci³, Cameron Keith⁴, and Tyler Moore¹

¹ Tandy School of Computer Science, The University of Tulsa,
`firstname-lastname@utulsa.edu`

² Applied Crypto Group, Stanford University, `jbonneau@cs.stanford.edu`

³ White Ops, `pubs@ryanc.org`

⁴ Computer Science and Engineering Department, Southern Methodist University
`ckeith@smu.edu`

Abstract. In the cryptocurrency Bitcoin, users can deterministically derive the private keys used for transmitting money from a password. Such “brain wallets” are appealing because they free users from storing their private keys on untrusted computers. Unfortunately, they also enable attackers to conduct unlimited offline password guessing. In this paper, we report on the first large-scale measurement of the use of brain wallets in Bitcoin. Using a wide range of word lists, we evaluated around 300 billion passwords. Surprisingly, after excluding activities by researchers, we identified just 884 brain wallets worth around \$100K in use from September 2011 to August 2015. We find that all but 21 wallets were drained, usually within 24 hours but often within minutes. We find that around a dozen “drainers” are competing to liquidate brain wallets as soon as they are funded. We find no evidence that users of brain wallets loaded with more bitcoin select stronger passwords, but we do find that brain wallets with weaker passwords are cracked more quickly.

Keywords: Bitcoin, brain wallets, passwords, cybercrime measurement

1 Introduction

Bitcoin, launched in 2009, is the most successful cryptographic currency to date and has recently attracted considerable research [2][5]. Similar to many other designs for cryptographic currencies, transactions which transfer control of bitcoins are authorized by ECDSA digital signatures. The popularity of Bitcoin, particularly with populations who had not previously used cryptographic software [7], has resulted in a large number of users attempting to manage private keys for the first time.

In this paper we study the use of *brain wallets*, or private keys which are deterministically derived from passwords. 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 [1].

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 on the Bitcoin block chain, 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 [10]. BIP 38 [6] specifies a format for password-protected private key encryption as a second factor. 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 [4], Herley argues that users rationally choose weak passwords for online accounts [13] 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 1 000 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 300 billion passwords 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 \$100K 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 while those storing larger values are emptied faster, nearly all wallets are drained within 24 hours.

2 Data Collection Methodology

We first review how the candidate passwords⁵ were constructed and then explain how we checked for their usage in brain wallets.

Password Corpora We have constructed an extensive set of passwords derived from publicly available sources. This includes prior password leaks (e.g., Rock-

⁵ Technically these are passwords and passphrases. We use password for simplicity of presentation.

you, Yahoo!, LinkedIn) word and derived phrase lists (e.g., English Wikipedia, Wikiquote), and information gleaned from Bitcoin discussion forums. In total, we tested approximately 300 billion passwords for usage in brain wallets. Testing was carried out using the open-source project Brainflayer⁶.

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

1. **English:** English word list packaged with Ubuntu 12.04.
2. **Urban Dictionary:** Terms and phrases from the crowd-sourced slang dictionary⁷. The “combinator” tool was used to check all pairs of terms[12].
3. **Two Words:** English pairs of words using the combinator tool.
4. **English/Slang Urban Dictionary:** Single word entries from Urban Dictionary are combined with English words. Additionally, the results are run through the combinator tool for all phrases up to 20 characters long.
5. **English Wikipedia**
6. **WikiQuotes:** English, Spanish, Russian and German quotes from 3/2013.
7. **Phrases:** Permutations of WikiQuote, wikipedia and Naxxatoe phrases.
8. **xkcd:** Lists obtained on July 10th, 2014 from three sources⁸. Combinations up to three words with and without spaces. All words used for 2 word combinations; words common to all three lists used for 3 word combinations.
9. **Lyrics:** lyrics and song titles purchased from <https://andymoore.info/mysql-lyrics-database/>.
10. **Blockchain.info tags:** All public bitcoin address tags obtained from <https://blockchain.info/tags>.
11. **Password dumps** LinkedIn, MySpace, RockYou, Rootkit.com
12. **Leet MRL:** De-duplicated merge of MySpace, Rockyou and LinkedIn (hence MRL) dumps, with leet-speak substitutions.
13. **Prince MRL:** MRL list applying the Prince attack [14].
14. **Security industry lists:** CrackStation, Naxxatoe, Uniqpass (combination of 2012-01-01 and 2012-04-01 lists), Skull Security⁹ (RockYou list excluded).

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

1. **Reddit User Challenge:** Post about a brain wallet password on Reddit¹⁰.
2. **Brute Force:** All numbers up to 9 digits and printable ASCII up to 5 characters.
3. **Modified BW Passwords:** Appended and prepended one and two ASCII characters to passwords of previously cracked brain wallets using combinator.

Table 1 in Section 3 details the number of brain wallet passwords obtained from each source, along with the total amount drained.

⁶ <https://github.com/ryancdotorg/brainflayer>

⁷ List was sourced from <https://github.com/inieves/urban-dictionary-scraper/blob/4a86fd9ef4c2f8812dc78f5862c327912213436a/dict/UrbanDictionary.txt>.

⁸ <https://xkpasswd.net/s/>, <http://correcthorsebatteryastaple.net/>, and <http://preshing.com/20110811/xkcd-password-generator/>.

⁹ <https://wiki.skullsecurity.org/Passwords>

¹⁰ <https://www.reddit.com/r/Bitcoin/comments/3gycp1/-/cu3316a>

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 library¹¹ [8]. We download the Bitcoin blockchain using Bitcoin core software¹² and extract all the unique Bitcoin addresses using znort987’s block parser¹³. 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.

3 Results

We investigate brain wallet usage by examining all block chain transactions through 8/2015.¹⁴ We report on their prevalence, draining, and password strength.

How Prevalent are Brain Wallets? We have found 884 distinct brain wallets using 845 different passwords. The slight difference is from the small number of instances where compressed and uncompressed wallets were used for the same password. In total, these brain wallets received 1 806 BTC (approx. \$103K¹⁵).

Table 1 reports the brain wallets identified, broken down according to the password sources. The single most popular source is the security word list CrackStation, which included 640 of the 884 brain wallet passwords identified. Notably, 37 of these passwords were only found by CrackStation, also the highest figure for any list. By contrast, the list with the second highest number of matches, Uniqpass, only reported passwords that were also found by at least one other source. Notably, the second-largest source of unique brain wallets, the combinations of English and slang words, only identified 63 wallet passwords.

The password sources used for our study can of course also be used by attackers. One way to estimate the popularity of password sources among attackers is to compare how often repeated drains occur. The fifth column shows the 90th 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.

¹¹ <https://github.com/bitcoin-core/secp256k1>

¹² <https://github.com/bitcoin/bitcoin>

¹³ <https://github.com/znort987/blockparser>

¹⁴ 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 15 brain wallets used in over 20 000 transactions between June and August 2015 as part of a network “stress test”.

¹⁵ All USD calculations presented here are normalized by the corresponding day’s exchange rate on Bitstamp, as reported by bitcoincharts.com.

Source	# Wallets (non-empty)	Unique	90% # drains	Total BTC	Total USD
<i>Word lists</i>					
Urban Dictionary	296	3	2	3.00	561.95 43 120.77
Two Words	13	3	0	4.00	0.79 92.65
Eng/Slang Urban Dict.	63	14	28	2.00	0.90 124.96
Eng. Wikipedia	250	0	0	2.00	505.77 38 833.16
WikiQuotes	35	0	0	12.00	60.96 17 620.50
Phrases	283	0	0	3.00	578.69 57 376.80
xkcd	90	3	3	13.00	97.66 29 140.44
Lyrics	329	4	16	3.00	230.45 26 788.97
Blockchain.info tags	112	0	10	7.00	577.93 31 683.29
Rootkit	123	2	0	6.00	4.50 570.78
MySpace	59	0	0	3.00	1.14 210.44
RockYou	415	3	2	3.00	113.82 33 807.17
LinkedIn	213	0	0	2.00	10.11 738.52
LEET MRL	3	0	0	1.00	0.01 1.49
Prince MRL	295	4	7	3.00	88.93 21 028.02
CrackStation	640	3	37	2.00	396.09 41 326.80
Naxxatoe	388	0	2	2.00	41.56 3 389.31
Skull Security	414	3	3	2.00	71.73 20 756.32
Uniqpass	490	3	0	2.00	134.95 35 266.27
<i>Non-word lists</i>					
Reddit User Challenge	1	0	1	1.00	0.01 2.62
Brute Force	200	3	3	3.00	22.47 3 895.99
Modified BW Passwords	74	1	9	2.00	2.25 209.98
Overall	884	21	139	2.00	1 806.22 103 472.13

Table 1. Brain wallets and values associated with different password sources.

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 Phrases word list is the most valuable at \$57K, followed by English Wikipedia, CrackStation, and Urban Dictionary. By contrast, the relatively unique English and slang combination passwords are not worth much – all 63 collectively stored just 0.90 BTC.

Figure 1 plots when wallets were first used over time, beginning with the first brain wallet established in September 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.

Relatively speaking, uncompressed wallets are more prevalent. We found 798 uncompressed wallets compared to 71 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). Compressed keys are only supported in versions of Bitcoin clients released after March 30, 2012; we observed 20 brain wallets before then, the first being “one two three four five six seven” seen in September 2011.

Also plotted in Figure 1 is 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




Fig. 1. New brain wallet usage per month (compressed and uncompressed, left); CDF and rank-order plot of total value stored in brain wallets (right).

few thousand dollars, but frequently the amount stored spikes greatly, including to a peak of over \$40K in March 2013. Notably, there is no discernible relationship between the number of new wallets created and the value stored.

The top plot in Figure 1 (right) gives the CDF of brain wallet value in USD. While most brain wallets store little money (just 6% 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 bottom plot of Figure 1 (right) 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.

Draining Brain Wallets As explained in Section 1, 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 2 (left) 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 21 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 wallets loaded with at least \$100. These wallets are consistently drained faster than other wallets, while those loaded with 10 cents or less (indicated by the dotted blue line) are drained more slowly. From this, we can




Fig. 2. CDF of the # of hours to drain brain wallets for wallets by value stored (left); how time-to-drain changes over time (median time-to-drain reported per month, right).

Rank (USD)	Drained # pwd	Drained (USD)	Drained (BTC)	Description
1	1	22 466	250.01	1 woodchuck drain (unintentionally done by researcher Castellucci, https://rya.nc/dc23)
2	1	15 267	250.00	1 woodchuck drain (done by owner)
3	10	14 554	50.02	19 drainer https://bitcointalk.org/index.php?topic=878639.460
4	2	11 528	18.25	2 drainer https://redd.it/2c5jot
5	29	6 784	12.15	49 drainer https://bitcointalk.org/index.php?topic=347828.0
6	1	5 800	500.00	1 “bitcoin is awesome” drain
7	100	3 219	9.96	155 drainer https://bitcointalk.org/index.php?topic=817294.10
8	1	1 863	38.69	1 owner of 1N8gLjZEhRxLRRjg8ymS6Zez8KVegEKtb1
9	1	1 429	14.29	1 “deadsheep” drain
10	1	1 322	97.66	59 “thequickbrownfoxjumpedoverthelazydog” drain

Table 2. Top 10 drain addresses from brain wallets, sorted by amount drained in USD.

conclude that time-to-drain is influenced by the stored value, but that in any case the wallet will almost certainly be drained within one day of funding.

How often are brain wallets drained? 98% of the brain wallets have been drained at least once. We observed 1 895 distinct draining events on 884 brain wallets. 69% of wallets are drained exactly once, while 19% are drained twice, and 1.9% are drained more than ten times. Figure 2 (right) plots the median time-to-drain by month. While this is always brief (less than one day), by September 2013 it becomes measured in minutes and seconds rather than hours.

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 2 (right) also plots in red the number of drainers actively receiving money from brain wallets. Overall, their numbers are increasing – unsurprising given the reduction in time-to-drain.

Digging deeper, we manually inspected all 48 addresses that received at least 100 USD from brain wallets, as well as the 13 addresses receiving payment from at least 20 distinct brain wallets. The top results are presented in Table 2, sorted by the total amount drained in USD. The table indicates how many distinct brain wallets were drained, the associated value in BTC and USD, and the number of drain events that occurred. 34 addresses were associated with a single password drain, suggesting these could be the owner. In a few cases, this is explicitly confirmed by online postings. Nonetheless, we confirmed at least 14 drainers targeting multiple brain wallets, corroborated by reports on discussion forums.

A few drainers are very successful while the rest do not make very much. The top 4 drainers have netted the equivalent of \$35 000 between them. The drainer who has emptied the most brain wallets – 100 in all – has earned \$3 219 for the effort. But other drainers have stolen very little money. For example, one drainer stole from 78 different brain wallets but netted only \$62 worth of bitcoin. Why is this? Looking back at Figure 2 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 [11]:

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).

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 [16]. In practice, many websites use inconsistent and poorly specified methods for giving users feedback on password strength [9]. The gold-standard of non-parametric statistics requires very large sample sizes and is hence impractical in our setting [3]. Instead, we use the `wheelerzxcvbn` formula as a rough measure of password strength. 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. We did observe a positive correlation of $\rho = 0.54$ ($p = 0.013$) between the estimated strength and the time it took for the wallet to initially be drained of funds.

This result suggests we can reject the null hypothesis with over 95% confidence (applying a Holm-Bonferonni correction for $m = 3$ hypotheses tested).

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. But we do see that addresses protected by weaker passwords are generally attacked quicker than stronger passwords. The cause of this correlation is that attacks have improved over time, so stronger passwords may have survived earlier cracking efforts but fall to later cracking efforts, giving a longer overall survival time. This could partially be due to the rise of ASIC mining [15], leaving Bitcoin enthusiasts with idle GPUs ripe for brain wallet cracking.

4 Conclusion

The idea behind brain wallets is elegant and alluring: remembering a password is surely easier than a private key. Unfortunately, as this paper 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 300 billion candidate passwords, we found 884 brain wallets that were active at some point in time. Unfortunately, we also found that nearly all were drained – usually quickly. 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.

Acknowledgements We thank the anonymous reviewers and paper shepherd Sarah Meiklejohn for their helpful feedback. Some authors are funded by the Department of Homeland Security (DHS) Science and Technology Directorate, Cyber Security Division (DHSS&T/CSD) Broad Agency Announcement 11.02, the Government of Australia and SPAWAR Systems Center Pacific via contract number N66001-13-C-0131. Support from the Oak Ridge Associated Universities Ralph Powe Junior Faculty Enhancement Award is also gratefully acknowledged. This paper represents the position of the authors and not that of the aforementioned agencies.

References

1. Simon Barber, Xavier Boyen, Elaine Shi, and Ersin Uzun. Bitter to better: How to make Bitcoin a better currency. In *Financial Cryptography and Data Security*, pages 399–414. Springer, 2012.
2. Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38, 2015.
3. Joseph Bonneau. Statistical metrics for individual password strength. In *20th International Workshop on Security Protocols*, April 2012.
4. Joseph Bonneau. The science of guessing: analyzing an anonymized corpus of 70 million passwords. In *2012 IEEE Symposium on Security and Privacy*, May 2012.

5. Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *IEEE Symposium on Security and Privacy*, May 2015.
6. Mike Caldwell and Aaron Voisine. BIP 38: Passphrase-protected private key, November 2012.
7. Nicolas Christin. Traveling the silk road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd International World Wide Web Conference*, pages 213–224, 2013.
8. Nicolas Courtois, Guangyan Song, and Ryan Castellucci. Speed optimizations in Bitcoin key recovery attacks. <http://eprint.iacr.org/2016/103.pdf>.
9. Xavier de Carné de Carnavalet and Mohammad Mannan. From very weak to very strong: Analyzing password-strength meters. In *Network and Distributed System Security Symposium (NDSS 2014)*. Internet Society, 2014.
10. Shayan Eskandari, David Barrera, Elizabeth Stobert, and Jeremy Clark. A First Look at the Usability of Bitcoin Key Management. In *Proceedings of the NDSS Workshop on Usable Security (USEC)*, 2015.
11. gmaxwell. #bitcoin-wizards, 2015. <https://botbot.me/freenode/bitcoin-wizards/2015-09-22/>.
12. hashcat. Combinator attack, 2015. https://hashcat.net/wiki/doku.php?id=combinator_attack.
13. Cormac Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 Workshop on New Security Paradigms*, pages 133–144. ACM, 2009.
14. Jens Steube. PRINCE: modern password guessing algorithm. <https://hashcat.net/events/p14-trondheim/prince-attack.pdf>.
15. Michael Bedford Taylor. Bitcoin and The Age of Bespoke Silicon. In *Proceedings of the 2013 International Conference on Compilers, Architectures and Synthesis for Embedded Systems*, page 16. IEEE, 2013.
16. Matt Weir, Sudhir Aggarwal, Michael Collins, and Henry Stern. Testing metrics for password creation policies by attacking large sets of revealed passwords. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 162–175. ACM, 2010.