J. Int. Financ. Markets Inst. Money 77 (2022) 101534

Contents lists available at ScienceDirect

Journal of International Financial Markets, Institutions & Money

journal homepage: www.elsevier.com/locate/intfin

What's the expected loss when Bitcoin is under cyberattack? A fractal process analysis

Klaus Grobys [a,b,*,1], Josephine Dufitinema [c,1], Niranjan Sapkota [d,1], James W. Kolari [e,1]

[a] *Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland*
[b] *Innovation and Entrepreneurship (InnoLab), University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland*
[c] *Mathematics and Statistics Research Group, School of Technology, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland*
[d] *Finance Research Group, School of Accounting and Finance, University of Vaasa, Wolffintie 34, 65200 Vaasa, Finland*
[e] *JP Morgan Chase Professor of Finance, Department of Finance, Texas A&M University, College Station, TX 77843-4218, Finland*

ARTICLE INFO

JEL Classification:
C22
G12
G13
G14
O10

Keywords:
Bitcoin
Cryptocurrency
Cyberattacks
Financial technology
Hacking incidents

ABSTRACT

In the era of digitalization, cryptocurrencies have become an alternative asset for both retail and institutional investors. While the emerging digital ecosystem based on blockchain technology offers numerous advantages, it is important to be aware of potential risks such as hacking incidents. In the 2011–2021 period, approximately 1.7 million units of Bitcoin were stolen due to criminal activity with losses exceeding $700 million. This paper models the distribution of stolen coins as a fractal process using power laws to estimate the expected losses from Bitcoin cyberattacks. Our results show that naïve statistics dramatically underestimate the expected loss by more than 70 percent. Our findings have important policy implications with respect to the urgent need for cryptocurrency market oversight by governments and regulatory agencies.

1. Introduction

Since the launch of Bitcoin in 2009, cryptocurrencies - a finance-related application of blockchain technology - have rapidly expanded. Today, distributed ledger technology and digital currencies are reshaping the way businesses operate. Benefits include decentralization, discretion, increased efficiency in terms of faster settlements, among others. Because technological developments improve the functioning of the financial system, it is not surprising that cryptocurrencies and blockchain systems have become popular topics. However, new digital financial markets are not without risks. Extraordinarily high levels of price volatility as well as

[*] Corresponding author.

E-mail addresses: klaus.grobys@uva.fi (K. Grobys), josephine.dufitinema@uwasa.fi (J. Dufitinema), niranjan.sapkota@uva.fi (N. Sapkota), jkolari@mays.tamu.edu (J.W. Kolari).

[1] The main author was invited by David Yermack (NY Stern University) to present an earlier draft of this study entitled "On the Tail Risk of Cyberattacks in the Bitcoin Market" at the Crypto and Blockchain Economics Research Forum (CBER) Symposium on March 18th 2021. The authors are thankful to David Yermack and the participants of the CBER Symposium for their useful comments. The authors are also thankful for helpful comments from the participants of the World Finance Conference 2021, the Cryptocurrency Research Conference 2021 and the 14th International Risk Management Conference 2021. Finally, we are grateful for having received competent and valuable feedback from two reviewers.

https://doi.org/10.1016/j.intfin.2022.101534
Received 8 November 2021; Accepted 20 February 2022
Available online 22 February 2022

cybersecurity are major obstacles to more widespread adoption. Also, governments and regulatory agencies are struggling to keep pace with rapid changes in these newly emerging digital ecosystems.

On August 10, 2021, a severe cyberattack took place in cryptocurrency markets. Hackers stole approximately $600 million in cryptocurrency from a protocol known as PolyNetwork wherein users can swap tokens across multiple blockchains.[2] Due to its economic magnitude, this cyberattack received widespread media attention with coverage on news broadcasts around the world. The natural question that arises is: What is the expected loss in terms of stolen coins from cyberattacks?

In view of increasing risks and losses of cryptocurrency theft, this paper investigates 58 hacking incidents in the Bitcoin market from 2011 to 2021. We find that 1.7 million units of Bitcoin were stolen, or about 10% of overall Bitcoin supply. Losses due to Bitcoin hacking incidents accounted for more than $700 million, which highlights the economic magnitude of this criminal activity. To deepen our understanding of cyberattacks, we explore the statistical distribution of Bitcoin hacking incidents. Subsequently, we compute the expected loss in terms of stolen coins using a fractal process.[3] The usage of a fractal process is motivated by Taleb (2020), who argued that the tail exponent of a power law function captures (by extrapolation) the low-probability deviation not seen in the data and plays an important role in determining the mean. As shown in Cirillo and Taleb (2020), using naïve statistics may dramatically underestimate risk.[4] We test the power law null hypothesis using a *goodness-of-fit* test based on Kolmogorov–Smirnov (KS) distance. Moreover, we test the plausibility of our model by means of Bayes' rule. Finally, we employ tools borrowed from extreme value theory as a robustness check to determine if the parent distribution is consistent with the assumed Pareto-type distribution.

This study contributes to extant literature in a number of ways. For example, a recent study by Grobys (2021a) explored the effects of cyberattacks due to hackings on uncertainty in the Bitcoin market. Using EGARCH models and daily data from 2013 to 2017 related to 29 Bitcoin hackings, he found that Bitcoin volatility dramatically increased on the day of hacking incidents in addition to the fifth post-trading day. The current research extends the Grobys study in two ways. First, we identify the underlying distribution of cyberattacks, which is crucial to making accurate statistical inferences. To do this, we apply Bayes' rule to explore which distribution is most likely to have generated extreme events in the data. Second, we compute an estimate of expected loss given the occurrence of a Bitcoin cyberattack. As of October 21, 2021, the market capitalization of Bitcoin exceeded $1.2 trillion; hence, the size of the Bitcoin market is substantial (Fry and Cheah, 2016). As losses due to Bitcoin hacking incidents exceed $700 million, the societal impact of this criminal activity is considerable.

From a broader perspective, this study contributes to the emerging literature on new risks associated with digital ecosystems. In this respect, Foley, Karlsen, and Putniņš (2019) proposed a model to identify illegal activities in Bitcoin and found that about one-quarter of all users (26%) and close to one-half of Bitcoin transactions (46%) are associated with illegal activity. Grobys and Sapkota (2020) explored the default risks of cryptocurrencies and found that 60% of cryptocurrencies eventually end up in default. Hileman and Rausch (2017) documented that 73% of exchanges control customers' private keys to funds denominated in cryptocurrency, an attractive target for cyber criminals. These studies showed that, unlike traditional asset markets, new digital financial markets involve different types of risk, including fraud, money laundering, credit, and hacking risks. Extending this literature, our study provides a novel perspective by estimating expected losses given the occurrence of hacking incidents. Estimating cyberattack losses is a crucial part of potential risk assessment, especially in view of the fact that institutional investors store considerably more funds in their digital wallets than retail investors. Relatedly, PricewaterhouseCoopers recently documented that the total assets under management (AuMs) of crypto hedge funds globally increased from about $1 billion in 2018 to more than $2 billion in 2019, whereas average AuMs increased from about $22 million to $44 million.[5]

Finally, our study contributes to studies that explore the degree to which human-engineered systems are exposed to tail risks. Because fat-tailed distributions are often modeled using fractal distributions, studies by Clauset, Shalizi, and Newman (2009) and Gabaix (2009) examined whether real world data sets from a range of different disciplines are governed by fractal processes. Notably, for the majority of the phenomena investigated, the power law null hypothesis could not be rejected. Recent finance studies have used power law functions for modeling cross-sectional stock returns (Warusawitharana, 2018), the realized variance of asset markets (Grobys, 2021b), and the volatility processes of cryptocurrencies (Grobys et al., 2021). Lux and Alfarano (2016) have provided an interesting overview of power laws in financial economics. Our study adds to this body of literature by taking a fractal perspective of Bitcoin hackings. As shown in forthcoming empirical findings, estimated losses from hackings are very different from naïve statistics.

Our results indicate that the share of the cumulative 20% of hacking incidents with the largest amounts of Bitcoin stolen on the cumulative total of the distribution exceeds 80%. We interpret this as strong evidence for a non-Gaussian process. Using a recently proposed test based on Bayes' rule, we find that other often-used distributions can be ruled out as possible candidates governing the distribution of cyberattacks in the Bitcoin market. Because the distribution closely resembles the Pareto 80/20 distribution, we model cyberattacks as a fractal process. Maximum likelihood estimation (MLE) suggests a power law process that is close to a fractal process with no defined theoretical mean, thereby implying an infinite loss (i.e., *t*-statistic 7.96). Because Bitcoin supply is limited, we hypothesize that the mean of lost coins is finite. Our statistical tests strongly support our hypothesis. Estimating the exponent of a power law process via MLE allows us to compute the shadow mean. By combining the shadow mean with the sample mean for the data not

[2] See https://www.bloomberg.com/news/articles/2021–08-10/hackers-steal-600-million-in-likely-largest-defi-crypto-theft.

[3] Fractality describes statistical self-similarity. As an example, the well-known Pareto 80/20 distribution implies that the largest 20% of the observations correspond to 80% of the cumulative total. Fractality implies that 4% (e.g., the largest 20% of the largest 20%) correspond to 64% (e.g., 80% of 80%). Since fractal processes are often termed power laws, we use these terms interchangeably.

[4] A detailed discussion on this issue and comparisons of various methodologies are provided in Taleb (2020).

[5] See https://www.pwc.com/gx/en/financial-services/pdf/pwc-elwood-annual-crypto-hedge-fund-report-may-2020.pdf.

Table 1
Hacking incidents in the Bitcoin market.

S. No.	Date	BTC	BTC Price	Lost Market Cap	Event	Circulating Supply BTC
1	2011-06-13	25000.01	19.80	502,750.20	User Allinvain hacked	6,528,850
2	2011-06-19	2000.00	17.50	46,970.91	MtGox theft	6,594,300
3	2011-06-25	4019.00	17.50	72,000.00	MyBitcoin theft	6,662,400
4	2011-07-26	17000.00	17.50	236,000.00	Bitomat loss	6,908,450
5	2011-07-29	78739.58	13.50	1,110,544.00	MyBitcoin theft	6,933,850
6	2011-10-06	5000.00	4.70	50,000.00	Bitcoin7 hack	7,416,650
7	2011-10-28	2609.00	3.20	8,115.12	MtGox loss	7,546,750
8	2012-03-01	46653.00	4.90	228,000.00	Linode hacks	8,462,900
9	Jan-Apr 2012	1000.00	5.00	5,000.00	Bitscalper Scam	8,898,050
10	Feb2012	2211.07	4.97	10,978.00	Andrew Nollan Scam	8,455,100
11	2012-04-13	3171.00	4.90	15,537.90	Betcoin hack	8,776,350
12	2012-04-27	30000.00	5.10	153,000.00	Tony76 Silk Road scam	8,873,500
13	2012-05-12	19980.00	4.57	91,306.00	Bitcoinica hack	8,993,400
14	2012-07-04	1853.00	5.30	9,820.90	MtGox hack	9,376,750
15	2012-07-13	40000.00	5.90	305,200.00	Bitcoinica theft	9,445,700
16	2012-07-17	180819.00	6.20	1,121,077.80	BST Ponzi scheme	9,476,850
17	2012-07-31	4500.00	6.70	42,000.00	BTC-e hack	9,585,150
18	2012-09-04	24086.00	10.40	248,088.00	Bitfoor theft	9,860,150
19	2012-09-28	9222.00	12.40	113,894.00	User Cdecker hacked	10,047,450
20	2012-10-17	3457.00	11.80	38,000.00	Trojan horse	10,186,150
21	2012-12-21	18787.00	13.50	253,624.50	Bitmarket.eu hack	10,576,300
22	2013-02-13	2000.00	25.50	51,000.00	Bit LC Theft	10,774,975
23	2013-05-10	1454.00	117.20	170,408.80	Vircurex hack	11,137,850
24	2013-06-10	1300.00	106.35	138,255.00	PicoStocks hack	11,269,775
25	2013-10-02	29655.00	114.13	3,384,525.15	FBI seizes Silk Road funds	11,782,875
26	2013-10-25	144336.00	186.69	26,946,087.84	FBI seizes Silk Road funds	11,900,650
27	2013-10-26	22000.00	177.32	3,901,040.00	GBL scam	11,905,625
28	2013-11-07	4100.00	296.41	1,215,281.00	Inputs.io hack	11,960,800
29	2013-11-12	484.00	360.33	174,399.72	Bitcash.cz hack	11,981,475
30	2013-11-29	5896.00	1131.97	6,674,095.12	PicoStocks hack	12,054,375
S. No.	Date	BTC	BTC Price	Lost Market Cap	Event	Circulating Supply BTC
31	2013-11-29	5400.00	1131.97	6,112,638.00	Sheep Marketplace closes	12,054,375
32	2014-02-13	4400.00	605.24	2,663,056.00	Silk Road 2 hacked	12,391,375
33	2014-02-25	744408.00	538.71	401,020,033.68	MtGox collapse	12,444,500
34	2014-03-04	896.00	666.78	597,434.88	Flexcoin hack	12,472,725
35	2014-03-04	97.00	666.78	64,677.66	Poloniex hack	12,472,725
36	2014-03-25	950.00	583.92	554,724.00	CryptoRush hacked	12,560,925
37	2014-10-14	3894.00	400.87	1,560,987.78	Mintpal hack	13,383,050
38	2015-01-05	18886.00	274.47	5,183,640.42	Bitstamp hack	13,691,175
39	2015-01-28	1000.00	233.91	233,910.00	796 Exchange hack	13,772,100
40	2015-02-15	7170.00	234.82	1,683,659.40	BTER hack	13,840,225
41	2015-02-17	3000.00	243.61	730,830.00	KipCoin hack	13,848,050
42	2015-05-22	1581.00	240.35	379,993.35	Bit?niex hack	14,189,850
43	2015-09-15	5000.00	230.30	1,151,500.00	Bitpay ?shing scam	14,616,850
44	2016-01-15	11325.00	364.33	4,126,037.25	Cryptsy hack	15,086,025
45	2016-04-07	315.00	422.74	133,163.10	ShapeShift hack	15,404,775
46	2016-04-13	154.00	423.73	65,254.42	ShapeShift hack	15,428,200
47	2016-05-14	250.00	455.67	113,917.50	Gatecoin hack	15,544,975
48	2016-08-02	119756.00	547.47	65,562,817.32	Bit?nex hack	15,792,425
49	2016-10-13	2300.00	636.79	1,464,617.00	Bitcurex hack	15,927,812
50	2017-04-22	3816.00	1231.71	4,700,205.36	Yapizon hack	16,288,550
51	2017-12-06	4736.00	14291.50	67,684,544.00	NiceHash hacked	16,724,975
52	2018-04-09	438.00	6770.73	3,300,000.00	Coin Secure hack	16,986,000
53	2018-09-14	5966.00	6512.71	38,854,827.86	Zaif hack	17,266,000
54	2019-05-07	7000.00	5829.50	40,806,500.00	Binance Phishing attack	17,684,537
55	2020-07-11	336	9240.35	6,658,499.36	Cashaa hack	18,431,243
56	2020-12-21	292	22803.08	2,516,010.70	EXMO exchange hack	18,576,643
57	2020-12-24	106	23735.95	2,516,010.70	Livecoin hack	18,582,818
58	2021-08-19	107	46717.58	4,998,781.06	Liquid Global hack	18,794,212

governed by a fractal process, we compute the overall expected loss at 106,171.49 coins, which is almost four times higher than the naïve sample average equal to 29,050.18 coins. We find that the vast majority of observations do not matter for computing the expected loss as the most salient statistical information resides in the tail of the distribution. Since naïve risk management dramatically underestimates the expected loss due to cyberattacks in the Bitcoin market, our findings have significant policy implications with respect to the urgent need for cryptocurrency market regulations by governments and regulatory agencies to protect investors from potentially severe losses.

Table 2
Descriptive statistics for Bitcoin hackings.

Metric	Sample estimate
Mean	29,050.18
Median	4,059.50
Std. Dev.	101,615.65
Minimum	97.00
Maximum	744,408.00
Skewness	6.43
Kurtosis	44.86

The study is organized as follows. The next section describes the data. The third section provides the empirical framework, and the last section concludes the study.

2. Data

We investigate the distribution of stolen coins due to hacking incidents in the Bitcoin market covering the 2011–2021 period. The data are reported in Table 1.[6] For example, in the well-known MtGox collapse on February 25, 2014, 744,408 coins were stolen amounting to $401.02 million, or 5.98% of the coeval circulating supply of Bitcoin at 12,444,500 coins. From Table 2 we observe that the kurtosis of both distributions exceeds 40, which is a strong statistical indication of fat tails. The sample average is 29,050.18 stolen coins, whereas the median is considerably lower and corresponds to 4,059.00. The sample standard deviation is as high as 101,615.65. It is important to note that the share of the cumulative 20% of the economically largest cyberattacks on the cumulative total of the distribution is 88%, which is very close to the well-known Pareto 80/20 distribution. Because there is no other distribution class apart from power laws that allows for modeling the observed fat tails, we model the underlying process governing stolen units of Bitcoin using a fractal process.

3. Empirical framework

Because many human-engineered systems are typically governed by Pareto distributions (and related power laws), Taleb (2020, p. 91) posited: "There are a lot of theories on why things should be power laws, as sort of exceptions to the way things work probabilistically. But it seems that the opposite idea is never presented: power laws should be the norm, and the Gaussian a special case." Given descriptive statistics of our data discussed in the previous section, we model the distribution stolen coins as a fractal process and test the following power law null hypothesis:

$$P(X > x) = p(x) = Cx^{-\alpha}, \tag{1}$$

where $C = (\alpha - 1)x_{MIN}^{\alpha-1}$ with $\alpha \in \{\mathbb{R}_{+}|\alpha > 1\}$, $x \in \{\mathbb{R}_{+}|x_{MIN} \leq x < \infty\}$, x_{MIN} is the minimum number of coins stolen, and α is the magnitude of tail exponent.[7] It can be shown that the expectation, or $E[X]$, is given by

$$E[X] = \int_{x_{MIN}}^{\infty} xp(x)dx = \frac{(\alpha - 1)}{(\alpha - 2)}x_{MIN}, \tag{2}$$

whereas the second moment, or $E[X^2]$, is defined as

$$E[X^2] = \int_{x_{MIN}}^{\infty} x^2 p(x)dx = \frac{(\alpha - 1)}{(\alpha - 3)}x_{MIN}^2, \tag{3}$$

and higher moments of order k are analogously defined as

$$E[X^k] = \frac{(\alpha - 1)}{(\alpha - 1 - k)}x_{MIN}^k. \tag{4}$$

From Eq. (2), we see that the mean only exists for $\alpha > 2$, whereas the variance only exists for $\alpha > 3$. Following White et al. (2008) and Clauset et al., we employ MLE and estimate the tail exponent as

$$\widehat{\alpha} = 1 + N\left(\sum_{i=1}^{N}\ln\left(\frac{x_i}{x_{MIN}}\right)\right)^{-1}, \tag{5}$$

[6] We collected the data on hackings by merging multiple sources, such as Table 1 and Table A.3 of Grobys (2021a), Table 3 of Biais, Bisiere, Bouvard, Casamatta, and Menkveld (2019), zdnet.com, coindesk.com, bbc.com, and bitcointalk.org.

[7] We follow notation in Clauset, et al. (2009).

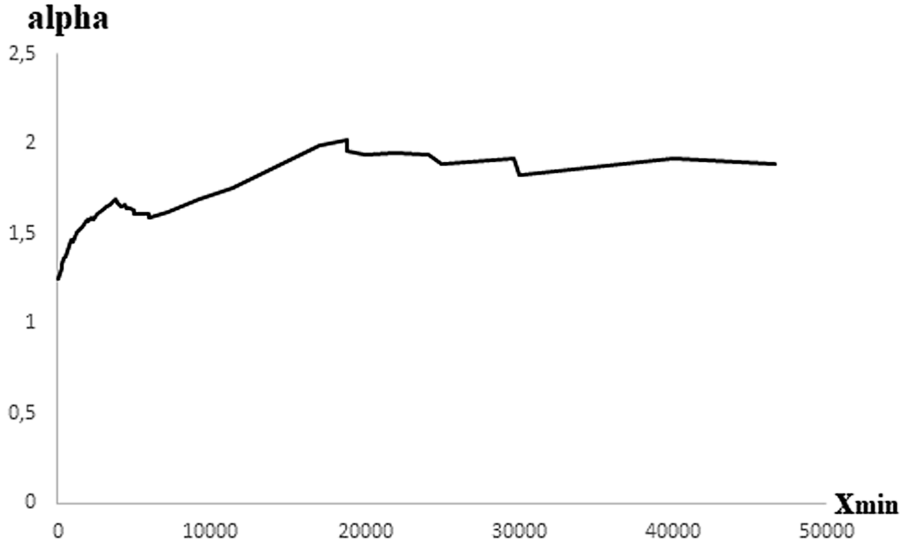

Fig. 1. Hill plot. This figure shows the Hill plot for our data on hacking incidents. Following common practice, the x-axis is truncated so that only $x_{min} < 50,000$ is taken into account.

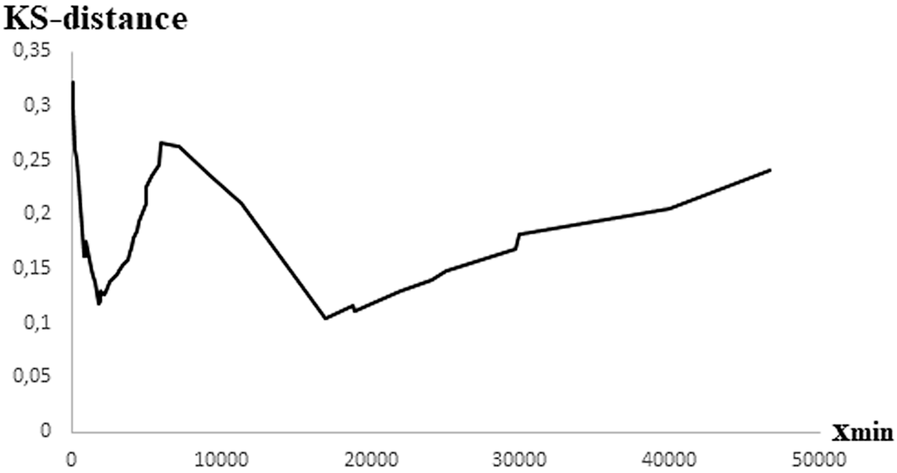

Fig. 2. KS distance. This figure plots the KS distance depending on x_{min}. Following common practice, the x-axis is truncated so that only $x_{min} < 50,000$ is taken into account.

where $\widehat{\alpha}$ denotes the MLE estimator, and N denotes the number of sample observations exceeding x_{MIN}, that is, $x_i \geq x_{MIN}$. As seen from Eqs. (2)–(4), minimum value x_{MIN} is essential for the calculation of the power law exponent. Fig. 1 plots MLE estimators $\widehat{\alpha}$ depending on x_{MIN}. This plot is often referred to as a Hill plot.[8] A question concerns which MLE estimator $\widehat{\alpha}$ in association with x_{MIN} is most accurate in describing the data-generating process and in line with the technical constraints of Bitcoin. Following Clauset et al., we estimate lower threshold x_{MIN} by making use of the Kolmogorov-Smirnov (KS) approach. This statistic is simply the maximum distance D between the data and fitted cumulative density functions (CDFs) given by

$$D = MAX_{x \geq x_{MIN}} |S(x) - P(x)|, \tag{6}$$

where $S(x)$ is the CDF of the data for the observation with a value of at least x_{MIN}, and $P(x)$ is the CDF for the power law model that best fits the data in the region of $x \geq x_{MIN}$. The estimate of x_{MIN} is the value of $\widehat{x}_{MIN}$ that minimizes D. Fig. 2 plots D depending on x_{MIN}. We observe there that D reaches a minimum at 0.1081 corresponding to $\widehat{x}_{MIN} = 17,000$. Hence, according to this approach, the parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN}) = (1.99, 17,000)$ is optimal.

[8] For implementing the MLE, we sort our observations in increasing order and use observations $i = 1, \cdots, 49$ as potential candidates for x_{MIN}.

Next, to explore whether this estimate is a reasonable choice from a statistical point of view, we simulate 10 million drawings from a power law distribution with $\alpha = 1.99$ and calculate the share of the cumulative 20% of largest observations on the cumulative total of the distribution. We find that the latter cumulative total equals 91%, which is larger than our empirical observation of 88%. In a simulation experiment searching for α that generates a share of the cumulative 20% of largest observations on the cumulative total equal to 88%, we find that the exponent is $\alpha = 2.03$. Because we know from the Hill plot that the parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$ is another MLE estimator, we test the following hypothesis:

$H_0 : \alpha = 2.03$ versus $H_1 : \alpha = 1.99$.

Given that $\widehat{\alpha}$ is normally distributed with a standard deviation of $\widehat{\sigma} = (\widehat{\alpha} - 1)/\sqrt{N}$, the corresponding *t*-statistic is estimated at 0.16 (*p*-value = 0.5636). Hence, we cannot reject the null hypothesis. Note also from Fig. 2 that the KS distances are virtually the same. One may argue that the parameters $\alpha = 2.03$ and $\alpha = 1.99$ are quite close to each other, such that setting $\alpha = 1.99$ under the null hypothesis might not reject it. While this possibility is a valid concern, Taleb (2010, p.265) has observed that: “Just a 0.2 difference in the exponent changes the results dramatically - such a difference can come from a single measurement error. This difference is not trivial: just consider that we have no precise idea what the exponent is because we cannot measure it directly.” Moreover, there is an economic argument for why exponent α should be $\alpha > 2$. From Eq. (2), we see that $\alpha < 2$ suggests an infinite or undefined theoretical mean. However, because Bitcoin supply is limited, an infinite theoretical mean does not make economic sense.

To support our hypothesis that $\alpha = 2.03$, following Clauset et al., we employ the parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$ in a *goodness-of-fit* test, thereby generating a *p*-value that quantifies the plausibility of the power law null hypothesis. Specifically, this test compares *D* from Eq. (6) with distance measurements for comparable synthetic data sets drawn from the hypothesized model. The *p*-value is defined to be the fraction of the synthetic distances that are larger than the empirical distance. Given a significance level of 5%, the power law null hypothesis is not rejected, as the difference between the empirical data and the model can be attributed to statistical fluctuations alone. Using the corresponding $D = 0.1190$, which is associated with $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$ and 1,000 synthetic data series, we find that the estimated *p*-value is 0.4070. Hence, we cannot reject our null hypotheses $\alpha = 2.03$. For these reasons, we utilize the empirically and economically optimal MLE estimator $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$ [9]

Using Eq. (2), we compute the shadow mean as $E[x] = 645,020.33$ coins, which is considerably larger than the sample tail mean estimated at $\frac{\#x_i \geq \widehat{x}_{MIN}}{\#x_i} \sum_{x_i \geq \widehat{x}_{MIN}} x_i = 102,873.70$ coins. In this notation, $\#x_i \geq \widehat{x}_{MIN}$ means that we count the number of observations for which $x_i \geq \widehat{x}_{MIN}$ is satisfied. The difference is enormous with the shadow mean exceeding the sample tail mean by a factor of 6.27. As discerned by Taleb (2020) and mentioned earlier, the tail exponent of a power law function captures (by extrapolation) low-probability deviations which largely determine the mean. Interestingly, only 25.86% of the distribution is governed by a power law process. Because the vast majority of the distribution of stolen coins is not governed by a power law process, following Cirillo and Taleb (2020), we can compute the corresponding expectation simply as the sample average, which is $\frac{\#x_i < \widehat{x}_{MIN}}{\#x_i} \sum_{x_i < \widehat{x}_{MIN}} x_i = 3,297.79$. Combining the shadow mean with the latter, we compute the overall expected loss as 106,171.49 coins, which is 3.65 times higher than the naïve sample average equal 29,050.18 coins.[10] In this respect, we should note that the vast majority of observations (74.14%) do not matter for computing the expected loss, as most relevant statistical information resides in the tail of the distribution. Naïve risk management may dramatically underestimate the expected loss in terms of stolen coins in Bitcoin cyberattacks.

To test the reasonability of our power law hypothesis, we apply Bayes’ rule as proposed in Taleb (2020). As an example, it is worthwhile to consider the MtGox cyberattack. In terms of stolen coins, this well-known attack on February 25, 2014 represents the largest theft with 744,408.00 coins stolen. Given our power law model with $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$, we can compute $P(X > 744, 408.00) = 2.81\%$. Assuming a normal distribution, such an event corresponds to a 6.72-sigma event with a corresponding probability of 9.09E–12. According to Benoit Mandelbrot, events of this magnitude never happen.[11] In this regard, Taleb’s statistical test can be used to explore how likely it is that a data-generating process is thin-tailed as opposed to fat-tailed. Specifically, using Bayes’ rule, the conditional probability that the underlying distribution governing this event is normally distributed, given that a 6.72-sigma event occurred on February 25, 2014, is defined as

$$P(ND|E) = \frac{P(ND)P(E|ND)}{(1 - P(ND)\ )P(E|PL) + P(ND)P(E|ND)}$$

where $P(ND|E)$ is the probability that the distribution is Gaussian given that the event occurred, $P(E|ND)$ is the probability of the event given that the distribution is normal, and $P(E|PL)$ is the probability of the event given that the distribution is a power law process with $\Phi = (\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03,18,787)$. Assuming various probabilities for $P(ND)$, we can compute the likelihood of $P(ND)|E)$. Table 3 reports

[9] It is interesting to note that $(\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18, 787)$ is the only MLE estimator satisfying $\widehat{\alpha} > 2.\ 00$. We used the codes *plfit* and *plpva* written by Aaron Clauset to estimate the $\widehat{\alpha}$ and the *goodness-of-fit* test. Since the code *plfit* does not provide the corresponding $\widehat{x}_{MIN}$ as additional output, we assess the corresponding $\widehat{x}_{MIN}$ directly from the Hill plot. The codes are available at http://www.santafe.edu/∼aaronc/powerlaws/. We would like to thank Aaron Clauset for making these codes available.

[10] As per the power laws properties specified earlier, the variance does not exist. Cirillo and Taleb (2016) derived variances for fractal processes with seemingly infinite mean by means of an application of extreme value theory. Unlike Cirillo and Taleb (2020), our study deals with a fractal process exhibiting a mean that is not necessarily infinite though close to infinity. Provided $\alpha < 2$ is satisfied, Cirillo and Taleb’s (2016) method can be applied, which is beyond the scope of the present study and therefore left for future research.

[11] See Benoit Mandelbrot’s lecture at MIT is available at https://www.youtube.com/watch?v=ock9Gk_aqw4.

Table 3
Testing the normal distribution assumption against a power law function.

$P(ND)$	$P(ND)\|E)$
0.50	3.25E–10
0.90	2.92E–09
0.99	3.21E–08
0.999	3.24E–07
0.9999	3.25E–06
1	1

We apply Bayes' rule as outlined in detail by Taleb (2020) to investigate how likely the data-generating process of cyberattacks in the Bitcoin market is normally distributed as opposed to a power law process with $\Phi = (\alpha, x_{MIN}) = (2.03, 18,787)$. Assuming that the process follows a normal distribution, the 6.72-sigma event of February 25, 2014 occurs with probability 9.09E–12, whereas the corresponding probability that the data are governed by our estimated power law process is 0.0280. According to Bayes' rule, the conditional probability that the underlying distribution governing this event follows a normal distribution given the occurrence of a 6.72-sigma event is defined as

$$P(ND|E) = \frac{P(ND)P(E|ND)}{(1 - P(ND)\ )P(E|PL) + P(ND)P(E|ND)},$$

where $P(ND|E)$ is the probability that the distribution is Gaussian given that the event occurred, $P(E|ND)$ is the probability of the event given that the distribution is normal, and $P(E|PL)$ is the probability of the event given that the distribution is a power law process with $\Phi = (\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03, 18,787)$. Assuming various probabilities for $P(ND)$, we can compute the likelihood of $P(ND)|E)$. Assuming various probabilities for the normal distribution, this table reports the computed likelihoods, $P(ND|E)$.

the results. Assuming that the normal distribution and the power law process are equally likely, the likelihood that the distribution is normal given a 6.72-sigma event is virtually zero. Even if we assume that the probability that the underlying data-generating process is normally distributed is as high as 99.99%, the likelihood that the distribution is normal given a 6.72-sigma event is near zero. Thus, we can rule out the normal distribution as the underlying data-generating process governing stolen units of Bitcoin.

As an additional robustness check, we use the same approach to test the exponential distribution versus our suggested fractal process. Given that a 6.72-sigma event occurred on February 25, 2014, the conditional probability that the underlying distribution governing this event is exponentially distributed equals 1.21E–03. From Table 4, we observe that assuming the exponential distribution and the power law process are equally probable, the likelihood that the distribution is exponential given a 6.72-sigma event is less than 5%. Even if we assume that the probability that the underlying data-generating process is exponentially distributed is as high as 90%, the likelihood that the distribution is exponential conditional on the occurrence of a 6.72-sigma event is only 28%. While we cannot entirely rule out the exponential distribution, we need to make an unreasonably strong assumption about how likely the exponential distribution would be, given the observed properties of the underlying data-generating process. Specifically, we need to require that the probability of the exponential distribution equals 96% so that the conditional probability of that distribution given the observed extreme event slightly exceeds 50%.[12] Of course, this assumption is highly improbable.

As an additional robustness check, we conduct an implicit test of the power law hypothesis borrowed from extreme value theory (EVT). In EVT the generalized Pareto distribution (GPD) is often used to model the tails of a parent distribution given by:

$$GPD(\xi, \beta) = 1 - \left(1 + \frac{\xi z}{\beta}\right)^{\frac{-1}{\xi}}\ if\ \xi \neq 0 \tag{7.1}$$

[12] We can say that even then the conditional probability only lightly exceeds 50%. The exact probability is 50.91%.

Table 4
Testing the exponential distribution assumption against a power law function.

$P(EXP)$	$P(EXP)\|E)$
0.50	0.0414
0.90	0.2800
0.99	0.8105
0.999	0.9774
0.9999	0.9977
1	1

We apply Bayes’ rule as outlined in detail by Taleb (2020} to investigate how likely it is that the data-generating process of cyberattacks in the Bitcoin market is exponentially distributed as opposed to a power law process with $\Phi = (\alpha, x_{MIN}) = (2.03, 18, 787)$. Assuming that the process follows an exponential distribution, the 6.72-sigma event of February 25, 2014 occurs with probability 1.21E–03, whereas the corresponding probability that the data are governed by our estimated power law process is 0.0280. According to Bayes’ rule, the conditional probability that the underlying distribution governing this event follows an exponential distribution, given that a 6.72-sigma event occurred, is defined as

$$P(EXP|E) = \frac{P(EXP)P(E|EXP)}{(1-P(EXP)\ )P(E|PL)+P(EXP)P(E|EXP)},$$

where, $P(EXP|E)$ is the probability that the distribution is exponential given that the event occurred, $P(E|EXP)$ is the probability of the event given that the distribution is exponential, $P(E|PL)$ is the probability of the event given that the distribution is a power law process with $\Phi = (\widehat{\alpha}, \widehat{x}_{MIN}) = (2.03,\ 18,\ 787)$. Assuming various probabilities for $P(EXP)$, we can compute the likelihood of $P(EXP)|E)$. Assuming various probabilities for the exponential distribution, this table reports the computed likelihoods, $P(EXP|E)$.

$$GPD(\xi,\beta) = 1 - e^{\frac{-z}{\beta}}\ if\ \xi = 0 \tag{7.2}$$

where ξ and β define the shape and scale parameters, respectively, and $z \in \{x|x > u\}$ in which u defines the tail threshold. EVT uses this limiting distribution to model tails of distributions, i.e., for data exceeding a certain threshold (or peaks over threshold abbreviated as POTs). Even if the main advantage of asymptotic laws used to derive the asymptotic extreme value distributions in Eqs. (7.1)–(7.2) is that they do not require knowledge of the parent distribution, it is possible to take into account ξ to draw conclusions concerning the parent distribution.

According to de Zea Bermudez and Kotz (2010), the flexibility of the GPD to assume many different forms enables its application to a variety of practical situations. Referring to Eqs. (7.1) and (7.2), we can infer that: (1) $\xi > 0$ if the GPD reduces to a fat-tailed distribution (e.g., Pareto distribution), (2) $\xi = 0$ if we have an exponential distribution, and (3) $\xi < 0$ if we have a thin-tailed distribution. For example, a recent study by Cirillo and Taleb (2020) argued that α in Eqs. (1)–(5) can be expressed as $\alpha = 1/\xi$. Hence, if $\xi > 0$, we can infer that the parent distribution is governed by a power law process. In empirical finance, researchers often use 5% or 1% of the largest realizations of a parent distribution and then estimate ξ and β. Using the method of moments (MOM) and the largest 5, 6,…, 14, 15 observations of the parent distribution, we iteratively estimate ξ and β. That is, in each iteration, the largest 5, 6,…, 14, 15 observations are allocated to the POT cluster, i.e., $z \in \{x|x > u\}$. We report these estimates in Table A.1 in the Appendix. Strikingly, we observe that irrespective of which cluster of the largest observations is used, $\widehat{\xi} > 0$ for all iterations. Consistent with Cirillo and Taleb, we interpret this evidence to mean that the parent distribution is a Pareto-type distribution. However, as noted by de Zea Bermudez and Kotz (2010), due to squaring of sample observations, MOM estimators in heavy-tailed or outlier situations can increase sampling errors. Thus, the MLE procedure in Eq. (5) could provide a more accurate estimate of the power law exponent.

Finally, we consider whether the estimated power law exponent is stable. As recognized by Taleb (2010), extremely fat-tailed processes are typically subject to a masquerade problem, which means that estimated power law exponents are likely to be overestimated (viz., a larger exponent implies a smaller role for large deviations). One manifestation of this problem is that an event may be less Black Swannish than perceived. It can take a considerable length of time for some fractal processes to reveal their properties, such that a shock's severity is underestimated. To investigate whether our estimate is prone to overestimation, we restrict the sample to the end of 2018. Again employing the KS distance D suggests that the parameter vector $(\widehat{\alpha}, \widehat{x}_{MIN}) = (1.99, 17,000)$ to be optimal. However, testing the null hypothesis $\alpha = 2.03$ using the *goodness-of-fit* test as proposed in Clauset et al. and 1,000 synthetic data series, the estimated p-value of 0.4010 suggests that we cannot reject the null hypothesis. Hence, we infer that our results are not sample specific.

Lastly, a model should reasonably resemble the properties of the underlying data-generating process. Perfect data fitting is not possible, which only the empirical distribution can provide. Interestingly, Taleb (2020) argued that the (nonparametric) empirical distribution is not valid in the sense that it misrepresents the realizations of the distribution in the tails. In fact, future maximums are poorly tracked by past data without some intelligent extrapolation. Power law functions address this issue. The tail exponent of a power law process captures (via extrapolation) low-probability deviations in the data and is a major determinant of mean estimation. Moreover, it is important to note that MLE, which is used in this study, is a valid method even if we are dealing with fat-tailed data, whereas the naïve sample cannot be used. Taleb stressed that the law of large numbers (LLN) either does not work in the presence of fat tails or works too slowly. Because MLE is a valid estimation method for fat-tailed data, he proposed a two-step estimation procedure when deriving moments from power law functions: (1) the power law exponent is estimated via MLE, and (2) the theoretical moments for the corresponding power law distribution are estimated. We adopt this procedure in the present study; as such, we do not deviate from the power law assumption, which Taleb believed should be standard practice.

4. Conclusion

This study investigated the expected loss in terms of stolen coins for cyberattacks in the Bitcoin market. We found that the naïve sample mean underestimates the expected loss by more than 70 percent compared to a model based on a fractal process using power laws. We modeled the distribution of stolen coins as a fractal process and estimated the expected loss given a cyberattack to be 106,171.49 coins. Assuming Bitcoin is traded above $60,000, the expected loss would be $11.3 billion.[13] Also, the well-known MtGox attack on February 25, 2014, in which 744,408.00 coins were stolen, would yield a loss of $44.7 billion in today's market environment. Notably, our model suggests that we have not seen the largest cyberattack yet.

Taleb and Cirillo have argued that it is perilous to employ naïve but reassuring statistics as a motivation for policymaking by governments or regulatory agencies. The more fat-tailed the distribution of losses from cybercrime, the greater the extent to which pertinent statistical information resides in the extremes rather than in the bulk of the distribution. Using our approach to investigate the distribution of stolen coins in cyberattacks, an important implication for policymakers is that more accurate risk assessments are needed to guide government and regulatory oversight of the rapidly evolving digital financial market.

While it is relatively difficult for hackers to gain access to the blockchain, they can easily steal bitcoins by accessing the digital wallets of naïve users through various scamming techniques. Also, hackers manipulate the market by using the trading system of the crypto exchanges. In many cases, the exchanges reimburse the lost coins by distributing their coins or tokens to their clients as a compensation. We have shown that the overall estimated loss is almost four times larger than the sample averages of all previous losses. Thus, we recommend that policymakers impose stricter regulation of crypto exchanges to improve cyberspace security.

Some exchanges have been permanently shut down after a large hacking incident. For example, the Russian crypto exchange Livecoin was closed after it lost about $19 million worth of funds due to a cyberattack on their trading system. Clients accused the exchange of a possible exit scam, but no legal action was taken after the exchange distributed the remaining fund to users. One can speculate that both EXMO and Livecoin exchanges were hacked by the same attackers as identical wallet addresses were used by the hacker(s). We believe that, if they had not ignored lower magnitude hacking incidents, the exchanges could have stopped the hacking incident from happening in the first place or at least taken measures to mitigate losses.

Faced with a cyberattack, cryptocurrency exchanges typically notify the security breach to local law enforcement authorities. Unfortunately, no response occurs in most cases. Based on our findings of potential large losses from cyberattacks, we advocate that governments implement stricter assessment tests of cybersecurity systems before issuing a permit or license to cryptocurrency exchanges. Relatedly, the Financial Conduct Authority (FCA) requires that any company carrying out activity related to crypto-assets in the U.K. must register and comply with anti-money laundering and counter-terrorist financing requirements. Since the risks of cyberattacks are considerably larger, we recommend that regulators adopt policies to harden cybersecurity systems and therefore augment current anti-money laundering and counter-terrorist financing requirements. Furthermore, regulators should require exchanges to establish reserve funds as insurance against cybersecurity threats for users.

In terms of future research, because our study provided estimates of Bitcoin losses due to cyberattacks using an unconditional probability function modeling approach, future research is encouraged using a conditional probability function approach.

[13] The closing price of Bitcoin on October 29, 2021 was $62,227.96.

CRediT authorship contribution statement

Klaus Grobys: Conceptualization, Methodology, Software, Writing – original draft, Validation. **Josephine Dufitinema:** Supervision, Methodology, Validation. **Niranjan Sapkota:** Investigation, Writing – review & editing. **James W. Kolari:** Supervision, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

See Table A1

Table A1
Generalized Pareto Distribution.

#obs	ξ	β	$E[u]$
5	0.08	233,205.33	255,099.28
6	0.15	186,393.72	222,103.26
7	0.20	155,566.39	197,678.84
8	0.23	132,702.80	178,262.65
9	0.26	116,302.72	163,122.77
10	0.28	103,379.12	150,606.07
11	0.30	93,263.20	140,291.86
12	0.31	84,996.58	131,546.25
13	0.33	78,090.05	124,012.41
14	0.34	72,281.44	117,487.56
15	0.34	67,386.28	111,826.36
Mean	0.26	120,324.33	162,912.48
t-statistic	10.07	7.59	11.71

We employ the GPD defined as

$$GPD(\xi,\beta) = 1 - \left(1 + \frac{\xi z}{\beta}\right)^{\frac{-1}{\xi}} \text{ if } \xi \neq 0,$$

$$GPD(\xi,\beta) = 1 - e^{\frac{-z}{\beta}} \text{ if } \xi = 0,$$

to model the tails of the parent distribution. In this model, ξ and β define the shape and scale parameters, respectively, and $z \in \{x|x > u\}$ where u defines the tail threshold. Using MOM and the largest 5, 6,…, 14, 15 observations of the parent distribution, we iteratively estimate ξ and β. In each iteration, the largest 5, 6,…, 14, 15 observations are allocated to the POT cluster; that is, $z \in \{x|x > u\}$. This table reports for each iteration the number of largest observations (first column), $\widehat{\xi}$ (second column), $\widehat{\beta}$ (third column), and expected tail mean (fourth column).

References

de Zea Bermudez, P., Kotz, S., 2010. Parameter estimation of the generalized Pareto distribution—Part I. J. Statist. Plann. Inference 140 (6), 1353–1373.

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., Menkveld, A., 2019. Equilibrium Bitcoin pricing. 2019 Meeting Papers of the Society for Economic Dynamics.

Clauset, A., Shalizi, C.R., Newman, M.E.J., 2009. Power-law distributions in empirical data. SIAM Rev. 51 (4), 661–703.

Cirillo, P., Taleb, N.N., 2016. Expected shortfall estimation for apparently infinite-mean models of operational risk. Quant.Finance 16 (10), 1485–1494. https://doi.org/10.1080/14697688.2016.1162908.

Cirillo, P., Taleb, N.N., 2020. Tail risk of contagious diseases. Nat. Phys. 16 (6), 606–613.

Foley, S., Karlsen, J.R., Putniņš, T.J., 2019. Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies? Rev. Financ. Stud. 32 (5), 1798–1853.

Fry, J., Cheah, E.-T., 2016. Negative bubbles and shocks in Cryptocurrency markets. Int. Rev. Financ. Anal. 47, 343–352.

Gabaix, X., 2009. Power Laws in Economics and Finance. Ann. Rev. Econ. Financ. 1, 255–294. https://doi.org/10.1146/annurev.economics.050708.142940.

Grobys, K., 2021a. When the blockchain does not block: on hackings and uncertainty in the cryptocurrency market. Quantit. Finance 21 (8), 1267–1279.

Grobys, K., 2021b. What do we know about the second moment of financial markets? Int. Rev. Financ. Anal. 78, 101891. https://doi.org/10.1016/j.irfa.2021.101891.

Grobys, K., Sapkota, N., 2020. Predicting cryptocurrency defaults. Appl. Econ. 52 (46), 5060–5076.

Grobys, K., Junttila, J., Kolari, J.W., Sapkota, N., 2021. On the stability of stablecoins. J. Empir. Finance 64, 207–223.

Hileman, G., Rausch, M., 2017. Global Cryptocurrency Benchmarking Study. Cambridge University, Cambridge Center for Alternative Finance, Cambridge, UK.

Lux, T., Alfarano, S., 2016. Financial power laws: Empirical evidence, models, and mechanisms. Chaos, Solitons Fractals 88, 3–18.

Taleb, N.N., 2010. The Black Swan. Random House, New York, NY.

Taleb, N.N., 2020. Statistical Consequences of Fat Tails: Real World Preasymptotics, Epistemology, and Applications (STEM Academic Press).
Warusawitharana, M., 2018. Time-varying volatility and the power law distribution of stock returns. J. Empir. Finance 49, 123–141.
White, E.P., Enquist, B.J., Green, J.L., 2008. On estimating the exponent of power-law frequency distributions. Ecology 89 (4), 905–912.