Revisiting the Risks of Bitcoin Currency Exchange Closure

TYLER MOORE, The University of Tulsa
NICOLAS CHRISTIN, Carnegie Mellon University
JANOS SZURDI, Carnegie Mellon University

Bitcoin has enjoyed wider adoption than any previous cryptocurrency; yet its success has also attracted the attention of fraudsters who have taken advantage of operational insecurity and transaction irreversibility. We study the risk investors face from the closure of Bitcoin exchanges, which convert between Bitcoins and hard currency. We examine the track record of 80 Bitcoin exchanges established between 2010 and 2015. We find that nearly half (38) have since closed, with customer account balances sometimes wiped out. Fraudsters are sometimes to blame, but not always. 25 exchanges suffered security breaches, 15 of which subsequently closed. We present logistic regressions using using longitudinal data on Bitcoin exchanges aggregated quarterly. We find that experiencing a breach is correlated with a 13-times greater odds that an exchange will close in that same quarter. We find that higher-volume exchanges are less likely to close (each doubling in trade volume corresponds to a 12 percent decrease in the odds of closure). We also find that exchanges who derive most of their business from trading less popular (fiat) currencies, which are offered by at most one competitor, are less likely to close.

CCS Concepts: •**Security and privacy** → **Economics of security and privacy;** •**Applied computing** → **Digital cash;**

Additional Key Words and Phrases: Bitcoin, currency exchanges, security economics, cybercrime

ACM Reference Format:
Tyler Moore, Nicolas Christin, Janos Szurdi, 2016. Revisiting the Risks of Bitcoin Currency Exchanges *ACM Trans. Internet Technol.* V, N, Article A (January YYYY), 16 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Despite added benefits such as enhanced revenue [Birch and McEvoy 1997] or anonymity [Chaum 1992], and often elegant designs, digital currencies have until recently failed to gain widespread adoption. As such, the success of Bitcoin [Nakamoto 2009] came as a surprise. Bitcoin's key comparative advantages over existing currencies lie in its entirely decentralized nature and in the use of proof-of-work mechanisms to constrain the money supply. Bitcoin also benefited from strongly negative reactions against the banking system, following the 2008 financial crisis: Similar in spirit to hard commodities such as gold, Bitcoin offers an alternative to those who fear that "quantitative easing" policies might trigger runaway inflation.

While Bitcoin's design principles espouse decentralization, an extensive ecosystem of third-party intermediaries supporting Bitcoin transactions has emerged. Intermediaries include currency exchanges used to convert between hard currency and Bitcoin; marketplace escrow services [Christin 2013]; online wallets; mixing services; mining pools; or even investment services, be they legitimate or Ponzi schemes [Jeffries 2012]. Ironically, most of the risk Bitcoin holders face stems from interacting with these intermediaries, which operate as de facto centralized authorities. For instance, one Bitcoin feature prone to abuse is that transactions are irrevocable, unlike most payment mechanisms such as credit cards and electronic fund transfers. Fraudsters prefer irrevocable payments,

This paper is a thoroughly expanded and revised version of a short paper which appeared at the 2013 International Conference on Financial Cryptography and Data Security [Moore and Christin 2013]. This work was partially supported by the National Science Foundation under ITR award CCF-0424422 (TRUST), and 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.

DOI: http://dx.doi.org/10.1145/0000000.0000000

since victims usually only identify fraud after transactions take place [Anderson 2007; Moore et al. 2012]. Irrevocability makes any Bitcoin transaction involving one or more intermediaries subject to added risk, such as if the intermediary becomes insolvent or absconds with customer deposits.

In this paper, we focus on one type of intermediary, currency exchanges, and empirically examine the risk Bitcoin holders face from exchange failures. Since bitcoin mining is now carried out by professional actors, users who wish to acquire bitcoins normally interact with currency exchanges to do so. They pay via bank transfer or credit card in a fiat currency and are credited with the corresponding amount of bitcoin. According to a 2017 survey, 73% of exchanges maintain control of the private keys for the bitcoin purchased by customers [Hileman and Rauchs 2017]. In this sense, the exchanges operate like a bank in the traditional financial system, in that the customers do not actually hold onto the cash but instead maintain an account with a balance that they can withdraw from by request.

As of November 2017, Bitcoin's market capitalization is approximately US$118 billion [Cryptocurrency Market Capitalizations 2017]. With success comes scrutiny, and Bitcoin has been repeatedly targeted by fraudsters. For instance, over 43,000 Bitcoins were stolen from the Bitcoinica trading platform in March 2012 [Leyden 2012]; in September 2012, $250,000 worth of Bitcoins were pilfered from the Bitfloor currency exchange [Lee 2012]. The prevalence of such attacks inspired an earlier version of this paper [Moore and Christin 2013], which found that 45% of Bitcoin exchanges established prior to January 2013 subsequently closed. Shortly after this paper was published, interest in Bitcoin exploded, along with its exchange rate. It is worth revisiting the question to determine whether or not the Bitcoin ecosystem has matured since its early days.

An anecdotal examination of the news suggests that the problems have not gone away with time. Mt. Gox, which had been the leading Bitcoin currency exchange through mid-2013, collapsed spectacularly in early 2014, leaving many of its customers in the lurch [Adelstein and Stucky 2016]. In August 2016, leading exchange Bitfinex was hacked, suffering a $68 million loss and socializing it amongst all users [Chen and Nakamura 2016].

Indeed, upon closer examination we find that the closure rate amongst Bitcoin exchanges remains very high. Of 80 exchanges operational through March 2015, 38 have subsequently closed. 26 have experienced at least one security breach. Section 2 explains our data collection and measurement methodology. In contrast to [Moore and Christin 2013], which computed survival and regression analysis that incorporated all activity across time, in this paper we construct longitudinal (i.e., panel) data calculated quarterly. This is designed to deal with the explosive transformation Bitcoin has experienced since its founding. Section 3 presents summary statistics for the data collected. Section 4 presents a series of logistic regressions to identify factors that contribute to whether an exchange will close. Section 5 reviews related work and Section 6 discusses follow-up research.

2. DATA COLLECTION METHODOLOGY

We collected various indicators from multiple sources: *trade data* from bitcoincharts.com [Bitcoin Charts 2015]; *breach data* from bitcointalk.org [Bitcoin Talk 2015] and from the Bitcoin wiki [Bitcoin Wiki 2015]; *security measures* from from the Bitcoin wiki, bitcointalk.org, and exchange websites; and *compliance data* from the World Bank's Anti-Money Laundering Index [Yepes 2011]. In this section, we describe our data collection methodology for all of these indicators.

In [Moore and Christin 2013], data were collected as attributes that affect an exchange for its entire duration, e.g., the overall trading volume, whether a breach had ever occurred, etc. While appropriate for the time period studied (2010–early 2013), Bitcoin shot to prominence shortly thereafter. By contrast, this paper covers transactions between 2010 and March 2015, and it could be argued that the environment in which Bitcoin currency exchanges operate is dramatically different now compared to its early days. Hence, in this paper, we set out to collect attributes that are time-dependent so that we can perform a longitudinal analysis. Some characteristics do not change over time (e.g., compliance data), but others do. Consequently, we compute indicators that are aggregated quarterly, which is long enough to capture stable measurements but short enough to reflect the dynamic nature of the Bitcoin ecosystem. When we describe the indicators below, we will distinguish whether the measure is computed quarterly or does not vary with time.

Trade data. Bitcoincharts.com provides historical trade data for a large number of Bitcoin exchanges (including all major exchanges), reporting the timestamp, exchange rate, and bitcoin amount for all trades that take place on participating exchanges.[1]. We considered historical trade data for all participating exchanges through March 3, 2015. We note that not all currency exchanges provide data to bitcoincharts.com. We exclude from the analysis any currency exchange that does not.

To characterize exchange trade data, we focus on three measures. First is whether the exchange remains operational and for how long. The *exchange lifetime* is the number of days an exchange was/has been operational, that is, the number days to have elapsed between the dates of the first and last observed trades. We also calculate a Boolean value whether the exchange has closed in the present quarter.

Second is the *exchange average daily trade volume*, calculated quarterly. We compute the average by dividing the total number of bitcoins transacted by the number of days between the first and last observation during the quarter. For active and high-volume exchanges, this should be approximately 90. However, some low-volume exchanges did not always report trading activity each day, so activity on those days would be counted as zero. By contrast, when an exchange opens or closes, we exclude from the average any adjacent days when not operational.

The third measure centers on the competition level of the currency offered by the exchange. While 41 exchanges have traded USD and 29 have traded EUR, many currencies have only ever been offered at one or two providers. 14 currencies were only ever supported by one exchange, while six more were only supported by two. It is possible that these less competitive currencies may offer greater stability to exchanges. We measure the fraction of an exchange's quarterly trading volume in which trades were conducted in currencies where only one or two exchanges traded that currency in the quarter.

Figure 1 graphically summarizes our trade data. Each row corresponds to one specific exchange. The x-axis represents time. Blue lines indicate times at which each exchange is open; red stars correspond to documented breaches. Two anomalies appear quite clearly in the figure. First, we found trade data for a few exchanges long after those exchanges had appeared to close (see orange and green circles). Such outliers can potentially lead to overestimating the exchange lifetime, and thus overreporting the periods in which the exchange is operational and misreporting when the exchange closes. Consequently, we need to determine whether they are valid. We detect outliers as points above the "median of all absolute deviations from the median" (MAD, [Rousseeuw and Hubert 2011]). The MAD method identifies possible outliers, graphically illustrated in Figure 1, for five exchanges: World Bitcoin Exchange, bitme, bitcoin-24.com, Global Bitcoin Exchange, and Ruxum.

We investigate whether these outliers should be ignored by searching for their probable causes. In the case of World Bitcoin Exchange, bitme and bitcoin-24.com, the outliers indicate an attempt to reopen the exchange. The attempts to reopen World Bitcoin Exchange and bitme failed almost immediately, thus we excluded those outliers from our analysis. On the other hand bitcoin-24.com managed to re-establish trade for a significant period of time. Thus, we included it twice in our data. To account for the gap between intermediate closure and reopening, we consider them as two separate exchanges. For Global Bitcoin Exchange, we removed the outlier as the exchange was reportedly closed, and we could not find any information about an attempt to reopen the exchange. Conversely, we kept the outliers pertaining to Ruxum, since we did not find any corroborating information about a possible exchange closure at the time of the outliers.[2]

[1] Compressed files are available for download from http://api.bitcoincharts.com/v1/csv/

[2] We reran the regressions presented in Section 4 including the outliers and the results did not change in any significant way.

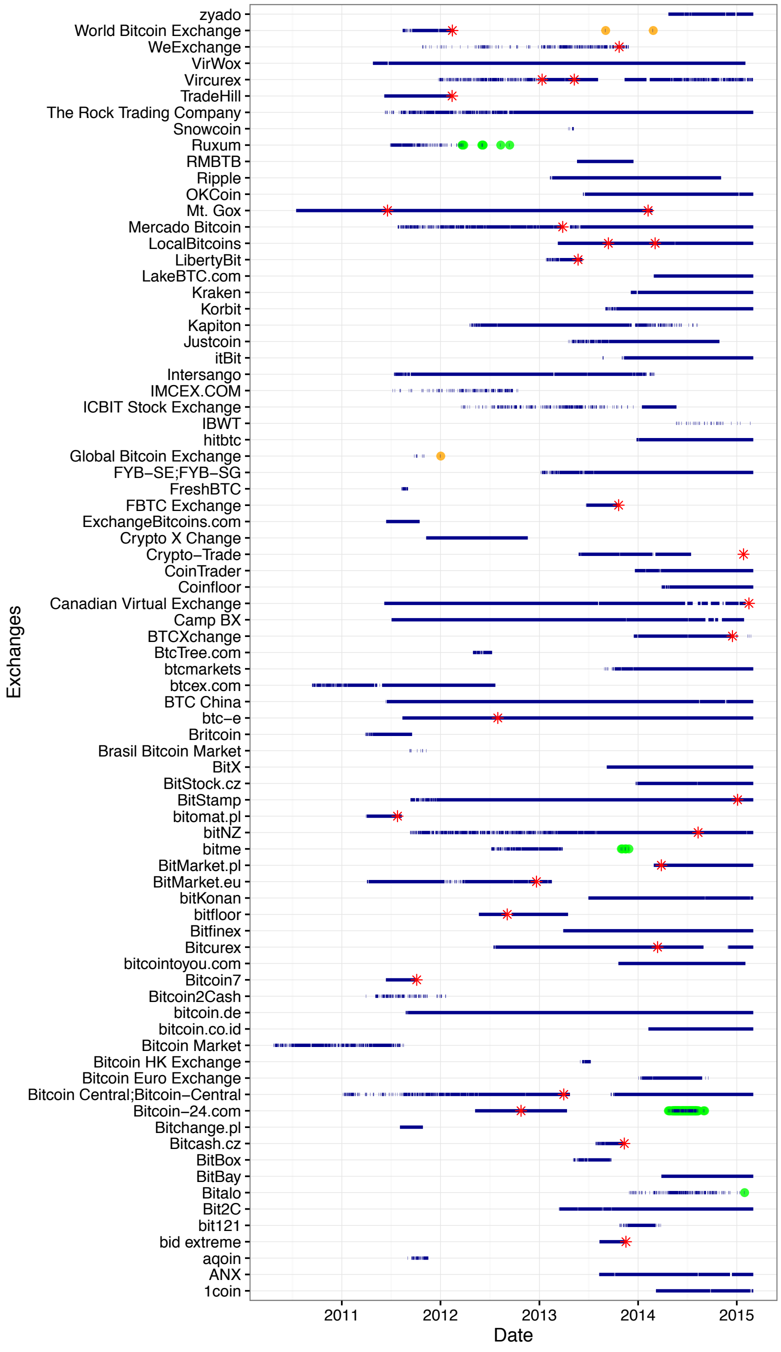

Fig. 1. **Exchange trade activity.** Blue dashed lines show when a given exchange was active. Red stars show when an exchange was breached. Green circles indicate outliers and orange circles indicate extreme outliers.

Second, as also evidenced in Figure 1, there are frequent gaps in trade data over the lifetime of an exchange. These gaps can be explained either by collection issues on bitcoincharts.com, or by real lack of activity on the exchange itself. We consider particularly long gaps (12 weeks or more), which we observe for Vircurex, Bitcurex, and Bitcoin Central. Looking for information on bitcointalk.org, we discovered that both Vircurex and Bitcurex upgraded their software at the beginning of the gap—which could have accounted for bitcoincharts.org not obtaining any data for a while—and we even saw informal evidence of Vircurex being active at the time of the gap in forum discussions. Conversely, Bitcoin Central was reportedly closed after a breach, during the corresponding gap. Nonetheless, in all three cases, the missing days are not included in the quarterly averages.

We deemed an exchange to be closed if there was no trading on the exchange for at least two weeks after the last observed trade day. To make sure that the exchange had truly closed, as opposed to being momentarily offline, we *additionally* confirmed that at least one of these criteria held: 1) the exchange website was consistently down, or 2) there was no trade data after March 27, 2015 (the end of our collection interval).

Breach data. We define an *exchange breach* as an event, during the life of an exchange, which result in the loss of users' funds due to negligence or misconduct *by the operators of the exchange*. This definition excludes, for instance, phishing attacks against the users of an exchange. Four different scenarios can lead to an exchange breach. In a *security breach*, a malicious entity exploits vulnerabilities in the exchange's software, hardware or system configuration to steal funds. As an example, the Bitfloor exchange suffered a security breach when thieves managed to gain access to backups of the private keys controlling cash flow accounts on the exchange, and used this access to steal an estimated 24 086 bitcoins [Bitcoin Talk 2014]. *Data loss*, e.g. due to hardware problems, can lead to unrecoverable loss of funds. For instance, *Bitomat.pl* reportedly lost all of their users funds, an estimated 17 000 bitcoins, in a data loss caused by an improper server restart [Bitcoin Talk 2014]. In an *insider scam*, unscrupulous exchange operators steal user funds themselves. *Legal action* can also lead to confiscation, and thus loss, of funds. Because it is often unclear which of the scenarios is the root cause of a breach—e.g., is a data loss truly due to incompetence, or malice?—our analysis does not distinguish between the various types of breaches.

The Bitcoin Talk forum[Bitcoin Talk 2015] and the Bitcoin Wiki [Bitcoin Wiki 2015] have dedicated pages to breaches [Bitcoin Talk 2014; Bitcoin Wiki 2014], which, unfortunately, are incomplete. From time to time, breaches are discussed in other areas of the site. To obtain better coverage, we ran customized Google queries on Bitcoin Talk. We generated the queries by combining keywords (theft, hack, scam, breach, loss, incident, stolen, victim) that had the highest frequency of occurence (measured by term frequency–inverse document frequency, or TF-IDF) in the dedicated breach pages with variations of the exchange name. This resulted in several queries to test each exchange, such as *"site:bitcointalk.org theft or hack or scam or breach or loss or incident or stolen or victim Mt. Gox,"* or *"site:bitcointalk.org theft or hack or scam or breach or loss or incident or stolen or victim mtgox.com."*

Overall, we ran 370 such queries during the time interval September 14–19, 2014. For each query we received between zero and ten results from the Google API, which we manually investigated to find breach events. We then complemented this data on March 3, 2015 with a manual investigation of news articles for reports of additional exchange breaches in the period 9/19/14–3/3/15.

Because we use a slightly different breach definition compared to our prior work [Moore and Christin 2013], we obtained a couple inconsistencies in what constitutes a breach. Previously, we considered Bitcoin Market to have been breached and, conversely, Bitcoin-24.com or BitMarket.eu to not have suffered breaches. When applying our revised breach definition more consistently, we reach the opposite conclusion. Bitcoin Market lost funds due to PayPal reportedly freezing their accounts, which we do not consider a breach. On the other hand, Bitcoin-24.com was breached on October 25, 2012, but this was not revealed until March 4, 2013, after [Moore and Christin 2013] was written. Meanwhile, Bitmarket.eu suffered collateral damage from hosting part of their infrastructure or Bitcoinica, which was breached. Since the losses resulted from poor judgement by the operators, we now categorize the event as a breach.

Security-related exchange properties. Because of the value of the resources they host, Bitcoin exchanges are expected to adopt good security hygiene. We conjecture that those who do not practice good security face a greater likelihood of eventual failure. To help us verify this conjecture, we collected the following indicators from each exchange, through manual analysis of their websites: 1) availability of two-factor authentication; 2) use of cold storage, that is, whether the exchange stores most of its bitcoins offline, and minimizes the amount of bitcoins kept online as cash flow for transaction operations; 3) presence of bug bounty programs; and 4) proclamations that the service undergoes routine security audits.

We looked for information about these security indicators on the websites of the exchanges, on the Bitcoin Talk forum and the Bitcoin Wiki. If an exchange was closed, we looked up its webpage on the Internet Archive Wayback Machine [The Internet Archive 2015]. The idea is that exchange operators have a strong incentive to advertise the security features they implement, and thus, evidence should be relatively easy to find.

For cold storage, bug bounty, and security audits, we identified simply whether or not the exchange ever reported these features. We decided to dig a bit deeper for the presence of two-factor authentication (2FA) in order to identify when the feature was added. To do this, we checked Internet Archive's pages for the first mention of supporting 2FA. Of the 58 exchanges found to have supported 2FA, 30 supported it all the way back to the first observation in the Internet Archive. For these exchanges, we label them as having supported 2FA in all quarters that the exchange was open. For the remaining 28 exchanges, we know that 2FA was added at some point between the first observation and the prior cache. The median gap between such observations is 112 days. We approximate when 2FA support is added by taking it to be the quarter in which it is first observed on Internet Archive.

Given the high rate of breaches, a natural question arises: do the exchanges adopt security precautions before or after a breach occurs? If it is the former, then it suggests that the security measures did not help stop a breach. If it is the latter, then the it suggests that the exchanges beefed up their security following a breach. We investigate this for the adoption of 2FA. Of 20 cases where 2FA is adopted and we have Internet Archive data, we can confirm that the breach occurred after 2FA adoption in 15 cases. In four of the remaining five cases the breach occurred before the Internet Archive's first cache, which showed 2FA support. Only in one case (Bitfloor) could we confirm that the exchange did not support 2FA before the breach but did afterwards. Hence, we conclude that for 2FA at least, the security measures were not adopted in response to experiencing a breach.

Compliance properties. Finally, to assess regulatory impact, we attempted to identify the country where each exchange is based. We then used an index (ranging between 0 and 49) computed by World Bank economists [Yepes 2011] to identify that country's compliance with "Anti-Money-Laundering and Combating the Financing of Terrorism" (AML-CFT) regulations [Yepes 2011].

3. ANALYSIS OVERVIEW

We start our analysis by presenting descriptive statistics and graphs that summarize the collected data. Table I lists all 80 known Bitcoin currency exchanges,[3] along with relevant characteristics such as whether the exchange experienced a security breach, subsequently closed, and observed security features. In total, 25 exchanges experienced security breaches, caused either by hackers or other criminal activity. 15 of these exchanges subsequently closed, but 11 have survived so far. Another 23 closed without experiencing a publicly-announced breach.

One key factor affecting the risk posed by exchanges is whether or not its customers are reimbursed following closure. We must usually rely on claims by the operator and investors if they are

[3] As explained in Section 2, the bitcoin-24.com exchange restarted about a year after it first closed. We treat these as two distinct exchanges in the subsequent analysis, which is why the total number of exchanges is 80 rather than 79.

Table I. Bitcoin exchange indicators.

Exchange	Origin	Start	End	Closed	Breach	Repaid	2FA	Bounty	Audit	Cold S.	AML
bitomat.pl	PL	4/11	8/11	yes	yes	yes	no	no	no	no	21.7
Bitcoin Market	US	4/10	8/11	yes	no	–	–	–	–	–	34.3
FreshBTC	PL	8/11	9/11	yes	no	–	–	–	–	–	21.7
Britcoin	GB	3/11	9/11	yes	no	–	no	no	no	yes	35.3
Bitcoin7	US/BG	6/11	10/11	yes	yes	par.	–	no	no	–	33.3
ExchangeBitcoins.com	US	6/11	10/11	yes	no	yes	no	no	no	no	34.3
Bitchange.pl	PL	8/11	10/11	yes	no	–	yes	yes	no	no	21.7
Brasil Bitcoin Market	BR	9/11	11/11	yes	no	–	–	–	–	–	24.3
aqoin	ES	9/11	11/11	yes	no	–	–	no	no	–	30.7
Global Bitcoin Exchange	GB	9/11	1/12	yes	no	par.	no	no	–	–	35.3
Bitcoin2Cash	US	4/11	1/12	yes	no	–	yes	no	no	no	34.3
TradeHill	US	6/11	2/12	yes	yes	no	no	no	no	no	34.3
BtcTree.com	US/CN	5/12	7/12	yes	no	yes	–	–	–	–	29.2
btcex.com	RU	9/10	7/12	yes	no	–	no	no	no	no	27.7
Ruxum	US	6/11	9/12	yes	no	–	–	–	–	–	34.3
IMCEX.COM	SC	7/11	10/12	yes	no	–	yes	no	no	no	11.9
Crypto X Change	AU	11/11	11/12	yes	no	no	yes	no	no	no	25.7
BitMarket.eu	PL	4/11	2/13	yes	yes	–	–	–	–	–	21.7
bitfloor	US	5/12	4/13	yes	yes	par.	–	no	no	yes	34.3
Snowcoin	IN	4/13	5/13	yes	no	–	yes	no	yes	yes	26.7
LibertyBit	CA	1/13	6/13	yes	yes	yes	–	–	–	–	25.0
Bitcoin HK Exchange	HK	6/13	7/13	yes	no	–	no	no	no	no	28.3
BitBox	US	5/13	9/13	yes	no	–	no	no	no	no	34.3
FBTC Exchange	NL	6/13	10/13	yes	yes	no	yes	no	yes	yes	27.3
Bitcash.cz	CZ	7/13	11/13	yes	yes	–	–	no	no	no	24.8
bid extreme	PL	8/13	11/13	yes	yes	–	–	–	–	–	21.7
WeExchange	US/AU	10/11	11/13	yes	yes	no	–	–	–	–	30.0
bitme	US	7/12	11/13	yes	no	–	–	–	–	–	34.3
RMBTB	CN	5/13	12/13	yes	no	–	yes	–	–	yes	24.0
Mt. Gox	JP	7/10	2/14	yes	yes	no	yes	no	–	–	22.7
World Bitcoin Exchange	AU	8/11	2/14	yes	yes	par.	yes	no	no	no	25.7
Intersango	GB	7/11	3/14	yes	no	no	no	no	no	yes	35.3
bit121	GB	10/13	3/14	yes	no	yes	–	–	–	–	35.3
ICBIT Stock Exchange	SE	3/12	5/14	yes	no	–	yes	no	no	no	27.0
Crypto-Trade	HK	5/13	7/14	yes	yes	–	no	no	no	no	28.3
Kapiton	SE	4/12	8/14	yes	yes	–	no	no	no	no	27.0
Bitcoin-24.com	DE	5/12	9/14	yes	yes	par.	–	no	no	no	26.0
Bitcoin Euro Exchange	CZ	1/14	9/14	yes	no	–	–	–	–	–	24.8
Justcoin	NO/HK	4/13	10/14	no	no	–	yes	yes	no	yes	29.7
Ripple	US	2/13	11/14	no	no	–	no	yes	no	no	34.3
Bitalo	FI	11/13	12/14	no	no	–	yes	no	no	no	24.3
OKCoin	SG	6/13	1/15	no	no	–	yes	yes	no	yes	33.7
Camp BX	US	7/11	1/15	no	no	–	yes	no	yes	no	34.3
LakeBTC.com	CN	3/14	1/15	no	no	–	yes	no	no	yes	24.0
Mercado Bitcoin	BR	7/11	2/15	no	yes	no	yes	no	no	yes	24.3
bitcointoyou.com	BR	10/13	2/15	no	no	–	yes	no	no	yes	24.3
VirWox	AT	4/11	2/15	no	no	–	yes	no	no	no	26.5
Canadian Virtual Exchange	CA	6/11	2/15	no	yes	yes	yes	no	yes	yes	25.0
IBWT	GB	5/14	2/15	no	no	–	yes	no	no	yes	35.3
BitX	SG	9/13	2/15	no	no	–	yes	no	yes	yes	33.7
BTCXchange	RO	12/13	2/15	no	yes	yes	yes	yes	no	no	26.3
BitMarket.pl	PL	3/14	2/15	no	yes	–	yes	no	no	yes	21.7
1coin	CN	3/14	3/15	no	no	–	yes	no	no	no	24.0
ANX	HK	8/13	3/15	no	no	–	yes	yes	no	yes	28.3
Bit2C	IL	3/13	3/15	no	no	–	yes	no	no	yes	29.3
BitBay	PL	3/14	3/15	no	no	–	yes	no	no	yes	21.7
Bitcoin Central	FR	1/11	3/15	no	yes	yes	yes	no	no	yes	31.7
bitcoin.co.id	ID	2/14	3/15	no	no	–	yes	no	no	yes	17.7
bitcoin.de	DE	8/11	3/15	no	no	–	yes	yes	yes	no	26.0
Bitcurex	PL	7/12	3/15	no	yes	yes	yes	yes	no	no	21.7
Bitfinex	HK	3/13	3/15	no	no	–	yes	no	yes	yes	28.3
bitKonan	HR	7/13	3/15	no	no	–	yes	no	no	no	19.0
bitNZ	NZ	9/11	3/15	no	yes	yes	no	no	no	no	21.3
BitStamp	GB	9/11	3/15	no	yes	yes	yes	no	yes	yes	35.3
BitStock.cz	CZ	12/13	3/15	no	no	–	no	no	yes	no	24.8
BTC China	CN	6/11	3/15	no	no	–	yes	no	yes	yes	24.0
btc-e	BG/CY	8/11	3/15	no	yes	yes	yes	no	no	no	33.7
btcmarkets	AU	8/13	3/15	no	no	–	yes	no	no	yes	25.7
Coinfloor	GB	3/14	3/15	no	no	–	yes	no	yes	yes	35.3
CoinTrader	CA	12/13	3/15	no	no	–	yes	yes	no	yes	25.0
FYB-SE;FYB-SG	SG/SE	1/13	3/15	no	no	–	yes	no	no	yes	30.3
hitbtc	DK	12/13	3/15	no	no	–	yes	yes	no	yes	24.3
itBit	SG/US	8/13	3/15	no	no	–	yes	no	yes	yes	34.0
Korbit	KR	9/13	3/15	no	no	–	–	–	–	–	20.0
Kraken	US	12/13	3/15	no	no	–	yes	yes	yes	yes	34.3
LocalBitcoins	FI	3/13	3/15	no	yes	yes	yes	yes	no	no	24.3
The Rock Trading Company	MT	6/11	3/15	no	no	–	yes	no	no	no	33.7
Vircurex	CN	12/11	3/15	no	yes	par.	yes	no	no	yes	24.0
zyado	PT/DE	4/14	3/15	no	no	–	yes	no	no	no	29.3

Table II. Summary statistics for categorical variables (overall and by exchange-quarter).

		Breached?	Closed?	2FA?	$\geq$ 90% Duopoly?
Overall	Yes	25	38	58	18
	No	55	42	16	62
	Unknown	0	0	6	0
Exchange-quarter	Yes	27	38	380	107
	No	520	509	151	440
	Unknown	0	0	16	0

Table III. Contingency and correlation tables for observed security characteristics. Significant Spearman correlations are indicated by p values $0 < .001 : ***, 0.001 < 0.01 : **, 0.01 < 0.05 : *, 0.05 \leq 0.10 : \cdot$.

	Yes	No	?	2-Factor Auth.	Bug Bounty	Security Audit	Cold Storage
2-Factor Auth.	58	16	6	1.00	0.02	0.15	0.40 **
Bug Bounty	12	54	14	0.02	1.00	-0.04	0.02
Security Audit	13	51	16	0.15	-0.04	1.00	0.29 *
Cold Storage	31	32	17	0.40 **	0.02	0.29 *	1.00

made public. Of the 38 exchanges that closed, we have found evidence on whether customers were reimbursed in 16 cases. Six exchanges have not reimbursed affected customers, while five have fully refunded customers and five more have partially done so. Thus, the risk of losing funds stored at exchanges after closing is real but uncertain.

We expect that the observed security characteristics may be correlated with one another – for example, exchanges that support two-factor authentication might be more likely to run a bug-bounty program. The correlation table in Table III shows correlations for the security variables. It also shows the rate of occurrence for each characteristic. Two-factor authentication was most widely supported, while bug bounty programs and security audits were comparatively rare. Note that it was easier to determine whether or not two-factor authentication was offered: we could identify whether it was offered in all but 6 cases. Missing values were more common for the other security characteristics. Consequently, our subsequent analysis focuses on the presence of two-factor authentication.

Table II reports the incidence of several binary variables, notably the frequency of breaches and closure, along with the occurrence of two-factor authentication and low-competition exchanges. The first rows count by the number of exchanges, while the second grouping counts "exchange-quarter" occurrences, which is used in the time-based regressions. An exchange-quarter combines an exchange with the quarter during which it operates. For example, consider an exchange operating for two years that is breached twice, closing after the second occurrence. There are eight exchange-quarters, during two of which a breach is recorded, and during one of which it is closed.

While 25 exchanges are breached, 27 breach incidents are included for different time periods. Three exchanges are breached twice (Mt. Gox, Vircurex and Local Bitcoins), while one exchange (Crypto-Trade) was breached long after it stopped reporting trades to Bitcoin charts and had been reported closed. By comparison, 520 quarters went by without an exchange being breached.

Figure 2 shows the number of currency exchanges in operation each quarter, along with the number that close. By this measure, relatively few exchanges close compared to the number that remain open (roughly 1–5 each quarter close compared to 20–40 that remain open). But this is misleading, since 90 days is a short time window for an exchange to close within (and we have seen from Table II that around half of exchanges eventually close). Consequently, Figure 3 plots an annualized probability that an exchange will close (top) or be breached (bottom). Overall, the annualized probability of closing peaked in 2012 at around 40%, but dropped to around 15% by the end of our

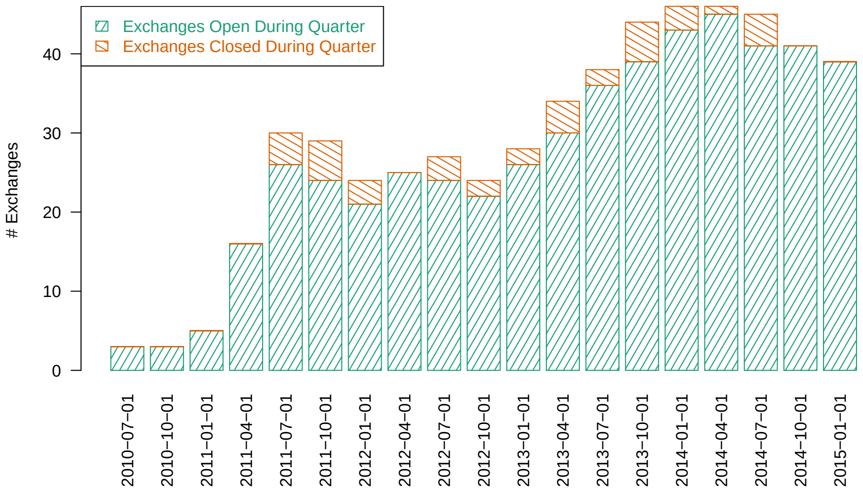

Fig. 2. Number of exchanges open and closed per quarter.

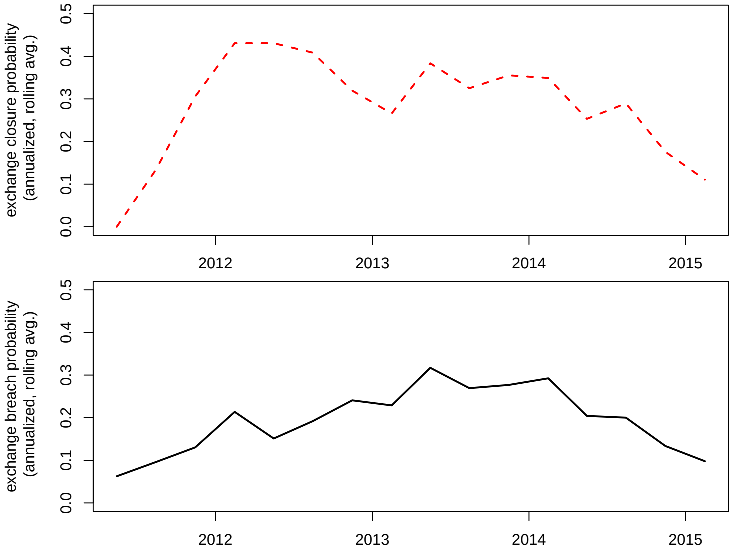

Fig. 3. Annualized probability of exchange closure (top) and breach (bottom).

study. Meanwhile, the annualized probability of a breach peaked at around 30% in late 2013–early 2014.

The overall failure rate of Bitcoin exchanges is 48%, and the median lifetime of exchanges is 451 days. These summary statistics obscure two key facts: exchanges are opened at different times and so their maximum potential lifetimes vary, and a majority of exchanges remain viable at the end of our observation period. Survival analysis can properly account for this.

We first report on the overall survival probability for Bitcoin exchanges. Figure 4 plots the overall survival probability of exchanges in black using data from all exchanges, along with a 95% confidence interval (black dashed lines). 90% of exchanges survive at least 97 days, but only 70% survive more than one year. The median survival for an exchange is estimated to be 796 days.

The popularity of exchanges varied greatly. Figure 5 (left) plots the CDF of the average daily exchange volume (in BTC) for each quarter under study. Some of the variation can be explained by the growth in Bitcoin, as data from earlier time periods would necessarily have lower volume. But

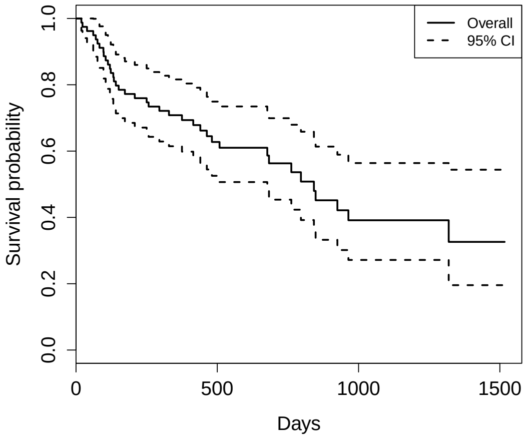

Fig. 4. Survival probability for exchanges.

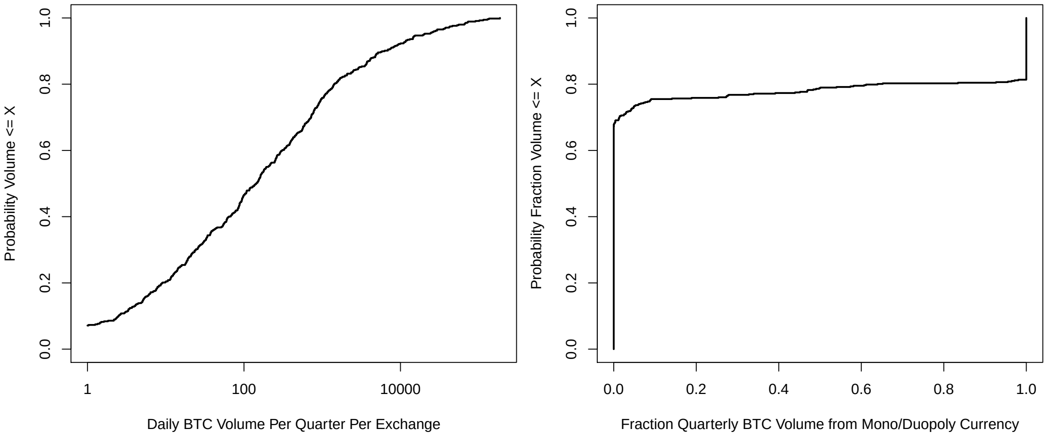

Fig. 5. Cumulative distribution functions for daily exchange volume and fraction of exchange volume derived from currencies with low exchange competition.

much of the extremity can be tied to variation between exchanges – the most successful exchanges consistently trade orders of magnitude more bitcoin than smaller ones. Around 40% of exchange-quarters average 100 BTC in daily trades or less, while the top 10% average more than 10,000 BTC per day. Given the wide disparities in trading volume, we will use a log-transformation for the regressions.

Figure 5 (right) plots a CDF of the fraction of quarterly exchange volume that is derived from trading fiat currencies where only one or two exchanges provide the currency. For example, only two exchanges, Korbit and Kraken, trade in KRW. Korbit only trades in KRW, so 100% of its volume is derived from currencies with low competition. By contrast, Kraken trades in many currencies, including USD. Most quarters the proportion derived from low-competition currencies including KRW is between 0.1–5%.

We can see from the CDF there is a clear dichotomy between those exchanges who derive almost no trading volume from low competition currencies (around 2/3 of the total) and those who derive nearly all of their volume from these minor currencies (around 20% of the total). Consequently, we will employ a Boolean variable where we mark any exchange that derives at least 90% of its trading volume for a given quarter from low competition currencies. The right-most columns in Table II indicate that 18 exchanges met this criteria at least once, spread over a total of 107 quarters.

Table IV. Proportion of exchanges that close and associated categorical variables (overall and by exchange-quarter). Differences in proportion that are statistically over- or under-represented are indicated in bold (according to χ^2 test).

		Breached?		Top 10% Vol.?		$\geq$ 90% Duopoly?		2FA?	
Overall	Open	11 of 42	(26%)	6 of 42	(14%)	16 of 42	(38%)	**39 of 42**	**(93%)**
	Closed	13 of 38	(34%)	7 of 38	(13%)	7 of 38	(18%)	18 of 32	(56%)
Exchange-Quarter	Open	17 of 509	(3%)	60 of 509	(12%)	106 of 509	(21%)	362 of 509	(73%)
	Closed	**10 of 38**	**(26%)**	3 of 38	(8%)	**1 of 38**	**(3%)**	18 of 32	(56%)

Table V. Correlation table for candidate predictor variables in logistic regressions. Significant Spearman correlations are indicated by p values $0 < .001 : ***, 0.001 < 0.01 : **, 0.01 < 0.05 : *, 0.05 \leq 0.10 : \cdot$.

	Closed		Breached		lg(Daily Vol.)		Duopoly		2FA		AML		Time	
Closed in Q	1.00		0.27	***	-0.08	·	-0.12	**	-0.09	*	0.03		-0.10	*
Breached in Q	0.27	***	1.00		0.06		-0.07		0.01		-0.09	*	-0.02	
lg(Daily Vol. Q)	-0.08	·	0.06		1.00		-0.12	**	0.10	*	0.07		-0.04	
Duopoly in Q	-0.12	**	-0.07		-0.12	**	1.00		-0.22	***	-0.14	**	-0.12	**
2FA in Q	-0.09	*	0.01		0.10	*	-0.22	***	1.00		0.05		0.35	***
AML	0.03		-0.09	*	0.07		-0.14	**	0.05		1.00		-0.10	*
Time	-0.10	*	-0.02		-0.04		-0.12	**	0.35	***	-0.10	*	1.00	

Table IV examines the occurrence of various measures for exchanges that are open and closed. Again we provide statistics for exchanges overall, plus a breakdown of quarterly activity at exchanges. The table reveals why it is helpful to view these characteristics at multiple points in time, rather than treating them as a single occurrence. For example, while a similar fraction of open and closed exchanges were breached, the breaches occur disproportionately often during the quarter in which an exchange closes. In 26% of the quarters when an exchange closes, the exchange also suffered a breach. This compares to breaches occurring in just 3% of the quarters in which an exchange remains open. This difference in proportion is statistically significant with 95% confidence according to a chi-squared test.

Similarly, we observe that *just one* of 38 exchanges closed during the quarter in which more than 90% of its trading volume was derived from currencies traded on only one or two platforms. By contrast, in 21% of the quarters in which exchanges remained open the exchange operated with such low competition.

We also draw two broader conclusions from the results in Table IV. First, it is useful to study the characteristics of an exchange (e.g., trading volume, experiencing a breach) at different points in time in order to determine what affects the likelihood the exchange will close at that point in time. Second, several of these characteristics are under- or over-represented with exchanges closing. This in turn motivates the regression model described next.

4. REGRESSION ANALYSIS OF EXCHANGE CLOSURE

We hypothesize that five variables affect the probability that a Bitcoin exchange will close in a given quarter:

Experiencing an exchange breach in the quarter: suffering a breach can erase profits, reduce cash flow, and scare away existing and prospective customers. We thus expect breached exchanges to be more likely to subsequently close.

Average daily transaction volume per quarter: an exchange can only continue to operate if it is profitable, and profitability usually requires achieving scale in the number of fee-generating transactions performed. We expect that exchanges with low transaction volume are more likely to shut down. We use a log-transformation of the average daily transaction volume per quarter given how skewed transaction volumes are.

Table VI. Logistic regression model. Significant coefficients are indicated by adjusted p values $0 < .001 : ***, 0.001 < 0.01 : **, 0.01 < 0.05 : *, 0.05 \leq 0.10 : \cdot$. The p values are corrected for multiple hypothesis testing using the method of Benjamini and Hochberg [Benjamini and Hochberg 1995], implemented using R's $p.adjust()$ method.

$\log{(p_c/(1-p_c))}$	Baseline			+Duopoly			+AML			+2FA		
	coef.	OR		coef.	OR		coef.	OR		coef.	OR	
(Intercept)	−1.09	0.33	·	−0.61	0.54		−1.98	0.14	·	−1.49	0.23	*
Breached in Q	2.60	13.45	***	2.49	12.04	***	2.67	14.48	***	2.60	13.46	***
lg(Trans. Vol.)	−0.12	0.88	*	−0.13	0.88	*	−0.13	0.88	*	−0.09	0.92	·
Time	−0.09	0.92	*	−0.10	0.91	*	−0.08	0.92	*	−0.05	0.95	
Duopoly in Q				−2.40	0.09	*						
AML							0.03	1.03				
2FA in Q										−0.58	0.56	
N	547			547			547			531		
Log-likelihood	-121.0			-115.2			-120.57			-106.3		

Most transactions involve from mono- or duopoly currencies: Many exchanges specialize in facilitating trades with less popular fiat currencies. We hypothesize that when there is less competition, the exchanges are less likely to close.

Two-factor authentication offered in the quarter: some exchanges advertise the availability of security features as described previously. We only include two-factor authentication for two reasons. First, we were able to obtain the most comprehensive data on 2FA (definitive answer on all but 6 exchanges). Second, we were able to identify in many cases when support for 2FA was added, unlike for the other security features.

AML/CFT compliance: some Bitcoin exchanges complain of being hassled by financial regulators. Thus, exchanges operating in countries with greater emphasis on anti-money laundering efforts may be pressured into shutting down.

Additionally, we include a time trend (defined as the number of quarters since Q3 2010) to account for changes in the Bitcoin ecosystem as it became more popular.

Table V shows the correlations between these variables. We note that closure in the quarter is positively correlated with experiencing a breach, while negatively correlated with the average daily transaction volume, low currency competition and time. Note that breach and transaction volume are not correlated, while low competition is weakly correlated with transaction volume.

2FA is strongly negatively correlated with low competition and positively correlated with transaction volume. AML is strongly negatively correlated with low competition and weakly negatively correlated with experiencing a breach. Given that both AML and 2FA correlate with low competition, we run regressions with each included independently plus a baseline model excluding them all.

4.1. Results

We run multiple logistic regressions with fixed effects for the longitudinal data. Models are fit using maximum likelihood estimation and implemented using R's `pglm` package. The results are presented in Table VI.

We start with a baseline model that includes whether a breach occurred, the log of the daily transaction volume, and a time trend variable. Each of these variables are statistically significant. When a breach occurs, there is a 13.5 times increase in the odds that the exchange will close that same quarter. This is highly significant, and the strong significance remains across all models.

The average daily transaction volume in the quarter is negatively correlated with closing. Each doubling of the daily transaction volume corresponds to a 12% decrease in the odds that the exchange will close that quarter. The time trend is also significant and negatively correlated with a breach, though we have only included the time trend in order to control for its effects. Already, the baseline model has a large log likelihood measure of -121.

The subsequent three models each add the additional variables that exhibit correlation with each other. Adding the low competition or duopoly variable is also significant and corresponds to a massive reduction in the odds that an exchange will close. In other words, exchanges who derive the vast majority of their transaction volume from currencies that few other exchanges also trade are 91% less likely to close than other exchanges who mostly trade fiat currencies with greater competition. We do note that there is a potential for endogeneity between trading volumes and low competition exchanges, as less popular currencies may be traded less often. These values are in fact negatively correlated (R=0.12), but what is striking is that both variables are negatively correlated with closure.

The other two variables, AML and 2FA, are not significant when incorporated into the regression. Note that in both cases, experiencing a breach and the daily transaction volume do remain statistically significant.

5. RELATED WORK

Bitcoin's success has not faded in the recent years and it inspired numerous research papers. Just like our work, a set of papers focuses on better understanding the strength and weaknesses of the Bitcoin protocol and its ecosystem. Another line of research proposes improvements to the current Bitcoin protocol. Finally, there is work that builds on Bitcoin to create new applications.

Barber et al. [2012] discussed the good and bad of Bitcoin and proposed enhancements to make it "a good candidate for a long-lived stable currency". Both Barber et al. [2012] and Böhme et al. [2015] in their survey paper looked at properties that play an important role in Bitcoin's success. Four such important properties of Bitcoin greatly inspired research due to their potential weaknesses: built-in incentives for mining (block rewards and transaction fees), transaction irreversibility, decentralization and pseudonymity.

First, currently Bitcoin mining is mainly fueled by the high block reward and only very low transaction fees are collected. As the block rewards decease overtime, the transaction fees cannot remain so low and new strategies are needed to set the transaction fees as discussed by Kaskaloglu [2014] and Möser and Böhme [2015].

Second, Bitcoin was designed to be a decentralized protocol, but for practical reasons Bitcoin users and miners use centralized services such as exchanges, mixers, online wallets and mining pools [Gervais et al. 2014]. For instance, the majority of mining is in the hands of a few large mining pools due to their ability to specialize in more cost-effective mining. This centralization of mining forced average users to buy from exchanges in order to own bitcoins and to use mixers if they want to enhance their anonymity. This centralization in the Bitcoin ecosystem increases the risk of users as observed by several researchers [Böhme et al. 2015; Kiran and Stanett 2015]. The risk associated with using an exchange is particularly high. To partially mitigate this risk, some researchers have devised methods for an exchange to demonstrate proof of solvency [Dagher et al. 2015; Decker et al. 2015].

Third, pseudonymity and the decentralized nature of Bitcoin gained the attention of criminals and it is used as the main currency at online black markets [Christin 2013; Soska and Christin 2015]. However the gap between pseudonymity and anonymity is not clear. Researchers shown how using the public transaction ledger in combination with personal data can be combined to reduce anonymity [Ron and Shamir 2013; Meiklejohn et al. 2013; Vasek and Moore 2015; Reid and Harrigan 2013; Androulaki et al. 2013].

Fourth, irreversibility of Bitcoin transaction increases the risk of transactions and lead to many scams both at illegal black markets and legitimate uses such as mining pools, digital wallets, exchanges and mixers. Vasek et al. studied Bitcoin scams and their victims in detail [Vasek and Moore 2015] . One (controversial) way to deter such scams and other malefactors is to introduce a scoring system for Bitcoin transactions [Möser et al. 2014].

Bitcoin exchanges necessarily make user participation more centralized while reducing users' anonymity. Compared to the related work, ours is the only effort that quantifies the risks associated with exchange closure.

6. CONCLUDING REMARKS

In this paper, we empirically investigated risks linked to the closure of Bitcoin exchanges. We conducted logistic regressions using data on exchanges aggregated quarterly. Compared to the earlier study that ran through January 2013 [Moore and Christin 2013], we gathered additional explanatory variables, most notably the competition level of the currencies traded and security features such as two-factor authentication that exchanges may support. The longer timespan also enabled us to conduct longitudinal analysis that focused on answering the question of what attributes of an exchange could prompt its immediate closure.

We found that experiencing a breach in a given quarter is strongly correlated with the exchange closing that same quarter. This is in contrast to the original paper, which did not find such an association. We believe that this insight is made possible by studying exchange closure longitudinally rather than over its complete period of operation.

We found that an exchange's trading volume is positively correlated with its continued operation. This is consistent with the previous paper's finding that the log of overall daily transaction volume is positively associated with survival. Given the longer time frame of this study, and the explosion in interest and trading activity that occurred, it is more appropriate to compare trading volumes over time to closure over time, as we now do.

Notably, we did not find any correlation between the presence of security controls (in this case support for two-factor authentication) and exchange lifetime. This does not mean that we believe there is no such relationship in general; instead, we conclude that more direct measures of security investment tied to the prevention of breaches is needed.

Finally, one completely new observation compared to the prior study is that there is an inverse relationship between competition and an exchange's continued operation. Exchanges who derived at least 90% of their activity by trading currencies that at most one other exchange also trades were much less likely to close. According to the regression, these low-competition exchanges experienced a 91% reduction in the odds of closure. One possible explanation for this is that with reduced competition comes reduced pressure on profit margins for trading. Another explanation could be that such currencies may not experience as much volatility in terms of money flows in and out of the exchange compared to those trading more popular fiat currencies.

Despite the advances compared to prior efforts, limitations to the statistical analysis remain. For instance, there is substantial randomness affecting when an exchange closes or is breached that is not captured by our model. Future work might investigate additional explanatory variables, such as the exchange reputation. Moreover, the effects we do capture could be measured more precisely. For example, breaches are treated as a binary variable (either one occurs or it does not). In reality, breaches can have widely varying impact, in terms of the financial loss imposed. In future work one could also examine the cumulative effect of an exchange suffering multiple breaches. Additionally, one could examine whether the minority of exchanges that do not hold customer bitcoins for long periods are more or less likely to close.

Moreover, while we have shed light on some factors that affect the likelihood that an exchange will close, we did not draw any specific conclusions about how closure affects individual investors. We note that in some cases users lost their deposits, but we are unable to precisely quantify the likelihood that this has happened. This point is often disputed between exchange operators and customers. Furthermore, there is no external regulator like the FDIC available to handle exchange failures in an orderly fashion. As the Bitcoin ecosystem matures, the community should consider devising mechanisms to ensure a fair and equitable distribution of available funds when exchanges close.

Finally, we focused on economic considerations, such as closure risks, that a rational actor would want to estimate before investing in a given exchange. However, reducing Bitcoin to a mere speculative instrument misses an important piece of the puzzle. Most Bitcoin users are early adopters, often motivated by non-economic considerations. For instance, online anonymous black market users, who constitute a large share of the Bitcoin economy [Soska and Christin 2015], may shy away from exchanges that require identification, and instead prefer assurances of anonymity. This may in turn lead them to use exchanges posing greater economic risk. Studying the unique characteristics of Bitcoin users and investors, compared to typical foreign exchange traders, is an avenue for future work we think well worth exploring.

ACKNOWLEDGMENTS

We thank Rainer Böhme and our anonymous reviewers for their extensive feedback on an earlier version of this paper.

REFERENCES

Jake Adelstein and Nathalie-Kyoko Stucky. 2016. Behind the Biggest Bitcoin Heist in History: Inside the Implosion of Mt. Gox. (June 2016). http://www.thedailybeast.com/articles/2016/05/19/behind-the-biggest-bitcoin-heist-in-history-inside-the-implosion-of-mt-gox.html.

Ross Anderson. 2007. Closing the Phishing Hole: Fraud, Risk and Nonbanks. In *Federal Reserve Bank of Kansas City – Payment System Research Conferences*.

Elli Androulaki, Ghassan O Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. 2013. Evaluating user privacy in Bitcoin. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 7859. Springer, 34–51.

Simon Barber, Xavier Boyen, Elaine Shi, and Ersin Uzun. 2012. Bitter to better - how to make Bitcoin a better currency. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 7397. Springer, 399–414.

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57, 1 (1995), 289–300. DOI:http://dx.doi.org/10.2307/2346101

David Birch and Neil McEvoy. 1997. Electronic Cash – Technology Will Denationalise Money. In *Financial Cryptography (Lecture Notes in Computer Science)*, Vol. 1318. Springer, Antigua, B.W.I., 95–108.

Bitcoin Charts. 2015. Markets API. (2015). https://bitcoincharts.com/about/markets-api/. Last accessed March 3, 2015.

Bitcoin Talk. 2014. List of Major Bitcoin Heists, Thefts, Hacks, Scams, and Losses. (2014). https://bitcointalk.org/index.php?topic=576337. Last accessed: November 5, 2016.

Bitcoin Talk. 2015. (2015). https://bitcointalk.org/. Last accessed April 3, 2015.

Bitcoin Wiki. 2014. (2014). https://en.bitcoin.it/wiki/List_of_Major_Bitcoin_Heists,_Thefts,_and_Losses. Last accessed: November 5, 2016.

Bitcoin Wiki. 2015. (2015). https://en.bitcoin.it/. Last accessed April 3, 2015.

Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. 2015. Bitcoin: Economics, technology, and governance. *The Journal of Economic Perspectives* 29, 2 (2015), 213–238.

David Chaum. 1992. Achieving Electronic Privacy. *Scientific American* (Aug. 1992), 96–101.

Lulu Yilun Chen and Yuji Nakamura. 2016. Hacked Bitcoin Exchange Says Users May Share $68 Million Loss. *Bloomberg* (Aug. 2016). https://www.bloomberg.com/news/articles/2016-08-05/hacked-bitcoin-exchange-says-it-will-spread-losses-among-users.

Nicolas Christin. 2013. Traveling the Silk Road: A measurement analysis of a large anonymous online marketplace. In *22nd World Wide Web Conference (WWW)*. Rio de Janeiro, Brazil, 213–224.

Cryptocurrency Market Capitalizations. 2017. Bitcoin Market Capitalization. (2017). https://coinmarketcap.com/currencies/bitcoin/. Last accessed November 2, 2017.

Gaby G Dagher, Benedikt Bünz, Joseph Bonneau, Jeremy Clark, and Dan Boneh. 2015. Provisions: Privacy-preserving proofs of solvency for Bitcoin exchanges. In *ACM Conference on Computer and Communications Security*. ACM, 720–731.

Christian Decker, James Guthrie, Jochen Seidel, and Roger Wattenhofer. 2015. Making Bitcoin exchanges transparent. In *European Symposium on Research in Computer Security (Lecture Notes in Computer Science)*, Vol. 9327. Springer, 561–576.

Arthur Gervais, Ghassan Karame, Srdjan Capkun, and Vedran Capkun. 2014. Is Bitcoin a decentralized currency? *IEEE Security & Privacy* 12, 3 (2014), 54–60.

Garrick Hileman and Michel Rauchs. 2017. 2017 Global Cryptocurrency Benchmarking Study. (2017). http://www.jbs.cam.ac.uk/fileadmin/user_upload/research/centres/alternative-finance/downloads/2017-04-20-global-cryptocurrency-benchmarking-study.pdf.

Adrianne Jeffries. 2012. Suspected multi-million dollar Bitcoin pyramid scheme shuts down, investors revolt. *The Verge* (27 August 2012). http://www.theverge.com/2012/8/27/3271637/bitcoin-savings-trust-pyramid-scheme-shuts-down.

Kerem Kaskaloglu. 2014. Near zero Bitcoin transaction fees Cannot last forever. In *International Conference on Digital Security and Forensics (DigitalSec2014)*. The Society of Digital Information and Wireless Communication, 91–99.

Mariam Kiran and M Stanett. 2015. Bitcoin risk analysis. *NEMODE Policy Paper* (2015).

Timothy B. Lee. 2012. Hacker steals $250k in Bitcoins from online exchange Bitfloor. *Ars Technica* (Sept. 2012). http://arstechnica.com/tech-policy/2012/09/hacker-steals-250k-in-bitcoins-from-online-exchange-bitfloor/.

John Leyden. 2012. Linode hackers escape with $70K in daring bitcoin heist. *The Register* (March 2012). http://www.theregister.co.uk/2012/03/02/linode_bitcoin_heist/.

Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. 2013. A fistful of bitcoins: characterizing payments among men with no names. In *ACM Internet Measurement Conference*. ACM, 127–140.

Tyler Moore and Nicolas Christin. 2013. Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 7859. Springer, 25–33. http://tylermoore.ens.utulsa.edu/fc13.pdf

Tyler Moore, Jie Han, and Richard Clayton. 2012. The Postmodern Ponzi Scheme: Empirical Analysis of High-Yield Investment Programs. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 7397. Bonaire, N.A., 41–56.

Malte Möser and Rainer Böhme. 2015. Trends, tips, tolls: A longitudinal study of Bitcoin transaction fees. In *Financial Cryptography and Data Security Workshops (Lecture Notes in Computer Science)*, Vol. 8976. Springer, 19–33.

Malte Möser, Rainer Böhme, and Dominic Breuker. 2014. Towards risk scoring of Bitcoin transactions. In *Financial Cryptography and Data Security Workshops (Lecture Notes in Computer Science)*, Vol. 8438. Springer, 16–32.

Satoshi Nakamoto. 2009. Bitcoin: A Peer-to-Peer Electronic Cash System. (2009). http://www.bitcoin.org/bitcoin.pdf.

Fergal Reid and Martin Harrigan. 2013. An analysis of anonymity in the Bitcoin system. In *Security and privacy in social networks*. Springer, 197–223.

D. Ron and A. Shamir. 2013. Quantitative Analysis of the Full Bitcoin Transaction Graph. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 7859. Springer, 6–24.

Peter J Rousseeuw and Mia Hubert. 2011. Robust statistics for outlier detection. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 1, 1 (2011), 73–79.

Kyle Soska and Nicolas Christin. 2015. Measuring the longitudinal evolution of the online anonymous marketplace ecosystem. In *24th USENIX Security Symposium (USENIX Security 15)*. 33–48.

The Internet Archive. 2015. Wayback machine. (2015). https://archive.org/web/.

Marie Vasek and Tyler Moore. 2015. There's No Free Lunch, Even Using Bitcoin: Tracking the Popularity and Profits of Virtual Currency Scams. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Vol. 8975. Springer, 44–61.

Concepcion Verdugo Yepes. 2011. *Compliance with the AML/CFT International Standard: Lessons from a Cross-Country Analysis*. IMF Working Papers 11/177. International Monetary Fund. http://ideas.repec.org/p/imf/imfwpa/11-177.html