Cryptomarket Discounts*

Nicola Borri
LUISS

Kirill Shakhnov
EIEF

This version: February 20, 2018

Abstract

This paper studies fiat-crypto currency investment strategies across exchanges around the globe from the perspective of US investors. We take Bitcoin as representative cryptocurrency and consider exchanges where investors can trade different fiat and crypto currency pairs (i.e., US dollar for Bitcoin). We treat each currency pairs as a different asset. First, we document large and persistent deviations in the bitcoin prices, converted in U.S. dollars, across the different exchanges. Second, we show that an investment strategy based on information on past cross exchanges price deviations generate large excess returns. We provide evidence that portfolios with the largest price deviations invest in exchanges with a higher probability of temporary shut downs; the smallest bitcoin supplies; the larger volume of transactions; and higher return volatility. These facts are consistent with the convenience yield hypothesis of cryptomarket discounts.

Keywords: bitcoin; cryptomarkets; cryptocurrencies; market anomalies; convenience yield
JEL Classification: G12, G14, G15, F31

*Borri: Department of Economics and Finance, LUISS University, Viale Romania 32, 00197 Rome, Italy; nborri@luiss.it; Tel: +39 06 85225959; http://docenti.luiss.it/borri. Shakhnov: EIEF, Via Sallustiana 62, 00187 Rome, Italy ; kirill.shakhnov@eief.it; Tel: +39 06 47924872; https://sites.google.com/site/kshakhnov/.
We thank seminar participants at LUISS University.

1 Introduction

In this paper we study fiat-crypto currency investment strategies across exchanges in the globe from the perspective of a US investor. We take Bitcoin as the representative crytocurrency and consider exchanges where investors can trade different fiat and crypto currencies (i.e., US dollar for Bitcoin). We treat each currency pairs traded at different exchanges as a different asset. First, we document large and persistent deviations in the price of Bitcoins converted in U.S. dollar across the different exchanges. Second, we show that an investment strategy based on information on past price deviations across exchanges generate large excess returns. Specifically, we allocate assets to different portfolios on the basis of their past price deviations with respect to a benchmark exchange and find a large and monotonic cross section of dollar excess returns even after controlling for transaction costs. We find that these portfolio returns are not explained by standard market risk factors. We provide evidence that portfolios with the largest price deviations contain assets traded in exchanges that have a higher probability of temporarily being shut down, and the smallest bitcoin supply; and have larger volume and return volatility.

Cryptocurrencies are a growing asset class, with a total market capitalization of $US 550 billions at the end of January, 2018. We focus on Bitcoin because it was the first cryptocurrency, created in 2009 using a scheme proposed by Nakamoto (2008), and it currently accounts for 34% of the total market capitalization and trading volume. Bitcoins started trading in 2010 on the Mt. Gox exchange, now defunct, and are now traded 24/7 every day in several exchanges in the world. Most of these exchanges opened in 2013, and this is when our data start.

Bitcoins traded in different exchanges are a homogenous asset. Therefore, in the absence of transaction costs and trade restrictions, the law of one price (LOP) should hold and bitcoin prices should be equal across different exchanges when expressed in the same currency. However, in practice, the dollar price of one bitcoin can differ greatly between different exchanges. We refer to these differences as *discounts* (or *premium*, when the discount is positive). Investors are aware of these price differences. For example, at the end of 2017, the financial press and investors' online forums went to a great length to discuss and analyze the so called "Kimchi Premium", referencing the country's popular fermented cabbage dish,describing the fact that buying Bitcoin on South Korean exchanges in Korean Won was much more expensive than in other exchanges across the globe, after accounting for the currency conversion (Kwon, 2017). This is evident in figure 1 which plots the U.S. dollar price of 1 Bitcoin on three Korean exchanges (Bithumb, Korbit, and Coinone) and on Bitfinex, the largest U.S. dollar based exchange by trading volume. Starting approximately in December 2017, the U.S. dollar prices of Bitcoins started to diverge between the Korean exchanges and Bitfinex and, for example, in January 2018 buying Bitcoins on Korbit was more than 60 percent more expensive than on Bitfinex. Note that these large price differences are very persistent. We document the existence of large and persistent price differences also in other exchanges. In addition, we find that these price differences are time-varying, and can be positive and negative within the same exchange.

Figure 1. Kimchi Premium

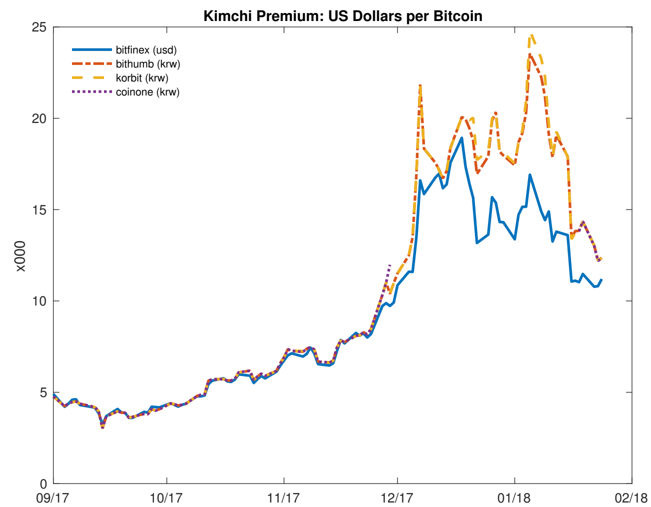

Notes: Daily U.S. dollars price on Bitfinex, Bithumb, Korbit, and Coinone. Data are daily from https://cryptocompare.com/ for the period 2/10/2015–1/24/2018.

If there are no transaction costs, investors should follow arbitrage strategies in the presence of violations of the LOP. However, in practice, these strategies might be difficult or costly to implement for several reasons. First, investors trading must pay a number of fees: transaction and withdrawal fees to the exchanges, currency conversion fees on the spot market, and fees to miners when moving bitcoins through exchanges. Second, constraints with respect to the speed of execution limit substantially the possibility of pure arbitrage strategies. For example, the settlement for fiat currencies conversion typically requires two business days, and the proof of work required by the blockchain technology implies that transferring Bitcoins through exchanges could take about one hour. Third, for investors to trade on a given exchange a registration process is required and several exchanges have restricted the number of new registrations, or introduced a minimum initial mandatory deposit. Fourth, only recently some exchanges have introduced the possibility of short-selling. Finally, restrictions to international capital flows could limit investors' ability to transfer money in and out of some of the countries in which the exchanges in our sample operates.

Rather building arbitrage strategies, we build an investment strategy, for a U.S. dollar based investor, that exploits the large and persistent price differences we observe across the different exchanges. Specifically, at the end of each day t, we sort all assets on the basis of their price discounts and then allocate them to five portfolios. Each asset corresponds to a crypto-fiat currency pair on each exchange and investors use dollars to first buy bitcoins through each asset and then sell bitcoins on our benchmark exchange, which we set to Bitfinex. The first portfolio always corresponds to a trade that goes long bitcoins on exchanges where the discounts are larger, i.e., in exchanges where bitcoins are particularly expensive. The last portfolio corresponds, instead, to a trade that goes long bitcoins on exchanges where the discounts are smaller, i.e., in exchanges where bitcoins are particularly cheap. This strategy produces a large and monotonic cross-section of dollar returns. The average return is approximately equal to -420 basis points on the first portfolio and 69 basis points on the last portfolio. Portfolios with larger returns have also larger return volatility. However, the Sharpe ratio increases monotonically from -71% on the first portfolio to 34 percent on the last. The spread between high and low discount portfolios is, then, 486 basis points per day with a Sharpe ratio of 85%.

We find that two common factors explain most of the variation in portfolio returns. However, we find that these factors, as well as the portfolio returns, are mostly uncorrelated with a large set of standard risk factors. In addition, we find that portfolios with largest discounts are different from the other portfolios with respect to various dimensions. First, they are characterized by a larger volume of transactions. Second, they contain, on average, assets from exchanges that have been shut down more times, and for a larger number of days, because of hacker attacks or software maintenance and crash due to high demand. Third, they contain, on average, assets traded in exchanges with the smallest shared wallets, where the latter are the equivalent of the assets for a financial institution, and capture the supply of bitcoin.

This paper contributes first to the small but fast growing economic literature on bitcoins and cryptocurrencies. Yermack (2013) and Velde et al. (2013) are two excellent primers that describe the functioning of the blockchain and cryptocurrencies[1]. Catalini and Gans (2016), Biais et al. (2018), and Ma et al. (2018) analyze from the prospective of economic theory how blockchain technology and cryptocurrencies will influence the rate and direction of innovation and the incentives and equilibria behind the "proof of work" protocols. Gandal et al. (2017) use a unique dataset to investigate suspicious trading activity on the Mt. Gox exchange in 2013 that appears to have inflated bitcoin prices. Second, this paper contributes to the

[1]There exists also a large literature on blockchain technology with a focus on security, anonymity, scalability, and data integrity from researchers in computer science that is outside the scope of our analysis.

large finance literature on market efficiency and anomalies as well as limits to arbitrage. Lee et al. (1991); Chen et al. (1993) relate discounts on closed-end funds, i.e., the difference between their market prices and the market values of the assets they own, and investors' sentiment. Borri and Verdelhan (2011) provide evidence that closed-end fund discounts are a measure of aggregate risk. Lamont and Thaler (2003) and Cochrane (2002) analyze possible mispricing in tech stock carve-outs. Cochrane argues that a mechanism much like the transaction demand for money drove stock prices above the "fundamental value" they would have had in a frictionless market. High prices are associated with high volume, high volatility, low supply of shares, wide dispersion of opinion, and restrictions on long-term short selling. The latter are also conditions that we document in cryptomarkets. We find that large discounts are associated with markets with the smallest bitcoin supplies, larger volume and price volatility, and where short-selling is not allowed. Shleifer and Vishny (1997); Liu and Longstaff (2003); Mitchell et al. (2002); Scheinkman et al. (2003); Gromb and Vayanos (2010) are empirical and theoretical papers that analyze both markets with arbitrage opportunities and limits to arbitrage.

The rest of the paper is organized as follows: section 2 begins by describing the data, the method used to build the bitcoin portfolios, and the main characteristics of these portfolios. Section 4 considers several extensions. Section 5 presents our conclusions.

2 Bitcoin Portfolios

We take the perspective of U.S. investors buying bitcoin in different markets. We uncover a profitable investment strategy based on information regarding price differences across the different markets.

2.1 Building bitcoin portfolios

Bitcoin discounts. We take the perspective of U.S. investors who can trade bitcoins in a set of $m = 1, \ldots, M$ markets. We denote with $P^{\star}_{m,j}$ the units of currency $j = 1, \ldots, J$ required to buy one bitcoin in market m. We also denote with S^j the spot exchange rate expressed in units of currency j per US dollar and with

$$P_{m,j} = \frac{S^j}{P^{\star}_{m,j}}$$

the units of bitcoin that one U.S. dollar can buy in market m. In the absence of any frictions, by the law of one price, U.S. investors should get the same units of bitcoin per dollar in each market m. However, in reality, there exist price differences across markets and these differences are persistent and time-varying. We take the price in market $m = 1$, which we take as Bitfinex, and currency $j = 1$, where $j = 1$ corresponds to the U.S. dollar, as the numeraire and denote with

$$D_{m,j} = \frac{P_{m,j}}{P_{1,1}} - 1 \tag{1}$$

the discount in market m and currency j. Note that $D_{m,j}$ can be positive or negative. Specifically, if $D_{m,j} < 0$, then U.S. investors get a smaller number of bitcoins in market m and currency j than in the reference market. On the contrary, if $D_{m,j} > 0$, U.S. investors get a larger number of bitcoins in market m and currency j than in the reference market. When $D_{m,j} = 0$, then the law of one price holds and investors get the same number of bitcoins in all markets.

Table 1 and figure 2 summarize the characteristics of the discounts that we document. Specifically, for each of the 32 exchanges in our sample, the table reports the mean, standard deviation, maximum, minimum, and first order autocorrelation of the discounts defined as in equation (1). The column "Pairs" reports the total number of currency pairs traded on the exchange and, for exchanges with more than one currency pair, all moments corresponds to the average across the pairs. For exchanges with more than one pair, we also report the tracking error volatility (Tev) defined as the mean standard deviation with respect to the discounts of the first of the currency pairs We find that discounts are, on average, different from zero and usually negative and volatile. In addition, within the same market, discounts are time-varying and can be positive and negative. Finally, we find that discounts are very persistent and mean-reverting. In addition, we find that discounts are different within the same exchange, at the same point in time, for different currency pairs.

Bitcoin excess returns. We assume U.S. investors can borrow at the dollar risk-free rate R^f and denote with lower case letters the log of any variable (i.e., $x = \log(X)$). We denote with

$$\begin{aligned} rx_{m,j,t} &= p^{\star}_{1,1,t} - p^{\star}_{m,j,t} + s^{j}_{t} - r^{f}_{t} \\ &= p_{m,j,t} - p_{1,1,t} - r^{f}_{t} \end{aligned} \tag{2}$$

the excess returns from the following strategy. First, borrow one U.S. dollar. Second, exchange the dollar for currency j at the spot rate. Third, buy bitcoin in market m and with currency j. Fourth, sell bitcoin for dollars in market $m = 1$. Fifth, pay back the dollar initially borrowed plus any accrued interest. Note that these excess returns imply that all

Table 1. Cryptomarket Discounts

Exchange	Mean (%)	Std (%)	Max (%)	Min (%)	AC(1)	Tev (%)	Pairs	T
Bithumb	-7.08	9.35	7.02	-39.69	0.93		1	177
Bitfinex	0.02	0.13	0.31	-0.31	0.16	0.26	2	404
Kraken	0.19	1.64	18.77	-5.11	0.67	0.69	2	642
GDAX	-0.56	1.77	14.39	-9.49	0.68	1.02	3	501
Bitstamp	0.21	1.77	17.17	-6.21	0.74	0.83	2	586
Coinone	-4.41	7.16	10.37	-35.40	0.90		1	192
bitFlyer	-1.17	2.93	18.45	-16.55	0.68		1	593
Gemini	0.05	1.51	17.11	-6.24	0.69		1	564
Korbit	-4.25	6.98	10.18	-36.69	0.92		1	454
BTCC	-5.33	3.41	-1.27	-12.51	-0.15		1	12
Exmo	-0.23	4.28	12.42	-13.72	0.84	2.69	3	232
LakeBTC	-2.30	1.87	3.44	-8.76	0.74	0.55	8	173
YoBit	-3.64	4.64	3.82	-17.29	0.85		1	82
itBit	-0.54	1.41	1.97	-4.14	-0.17	1.34	3	56
Livecoin	0.22	4.18	15.91	-9.75	0.67		1	175
BitBay	-0.64	2.89	17.39	-12.93	0.73		1	477
Luno	-10.78	9.24	6.58	-33.14	0.88	10.92	4	175
Gatecoin	-1.13	2.41	6.35	-12.04	0.70	2.07	2	240
QuadrigaCX	-2.24	3.21	6.64	-17.09	0.40	2.64	2	194
Bitso	-3.17	4.74	7.64	-25.34	0.85		1	270
Coinroom	-1.05	2.87	3.52	-10.68	0.56		1	62
Coinfloor	-0.14	3.65	11.61	-8.15	-0.12	4.63	3	75
BitMarket	-0.49	3.02	9.51	-8.58	0.76	0.89	2	192
Abucoins	-2.46	2.98	1.90	-11.10	0.47	0.66	3	29
Bit2C	-4.88	4.64	15.47	-15.79	0.59		1	88
OKCoin	0.91	5.69	30.02	-12.58	0.93		1	703
MercadoBitcoin	-6.95	6.68	6.75	-31.61	0.87		1	256
Lykke	-3.05	3.77	2.48	-12.55	-0.26	2.96	3	21
Coinbase	-0.56	1.77	14.39	-9.49	0.68	1.02	3	501
Zaif	-1.99	3.38	13.11	-17.98	0.61		1	165
Jubi	-0.12	7.48	21.55	-13.85	0.85		1	66
Unocoin	-8.29	7.42	7.45	-28.39	0.71		1	177

Notes: This table reports, for each of the 32 exchanges in our sample, the mean, standard deviation, maximum, minimum, and first order autocorrelation of the discounts defined as in equation (1). All moments are in percentages. The column "Pairs" reports the total number of currency pairs traded on the exchange. When the number of pairs is larger than one, we report the average across the different pairs. The column T reports the average number of daily observations. For exchanges with more than one pair, the column Tev reports, in percentage, the average standard deviation with respect to the discounts of the first of the currency pairs. Data are daily, from https://cryptocompare.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

transactions take place at the same time t which, in our sample, correspond to 16:00 GMT, i.e., closing time of the London exchange. We follow this timing as, while bitcoin markets are open on a 24–hour basis, spot rates correspond to this timing.

Figure 2. Bitcoin Discounts

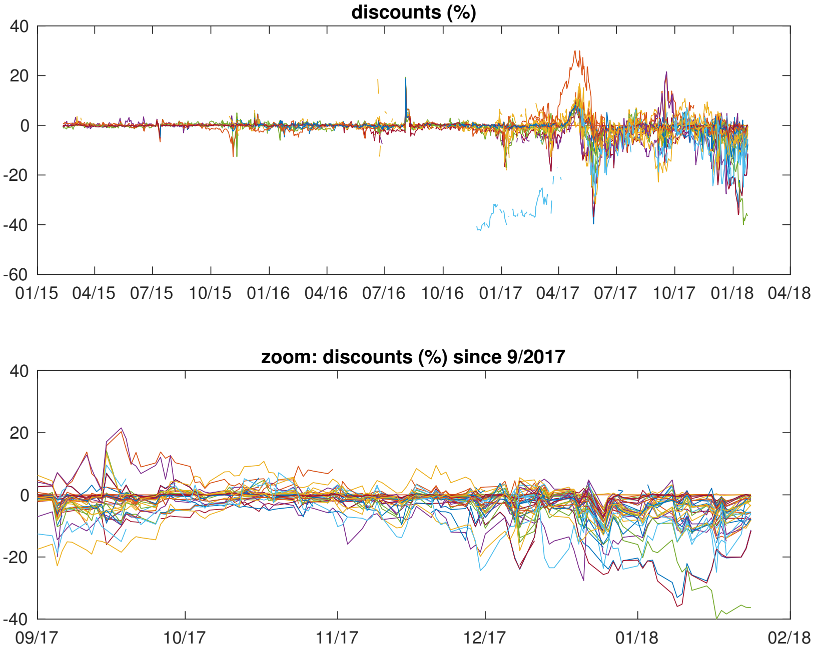

Notes: This figure plots the bitcoin price discounts on different exchanges. Discounts are defined according to equation (1), and denote the percentage difference in the number of bitcoins that one US dollar can buy on different exchanges. The top panel refers to the full sample, and the bottom panel to a shorter sample that starts 9/1/2017. Discounts are in percentages. Data are daily from https://cryptocompare.com/ and Thomson Reuters for the period 2/10/2015–1/24/2018.

Data. We collect hourly bitcoin prices data from 32 exchanges and 64 unique currency-exchange pairs from https://cryptocompare.com/ using a Python script[2]. Note that crypto exchanges operate every day 24/7, including Saturdays, Sundays, and holidays. All exchanges across the globe use the UNIX time-stamp to track time and ensure immediate comparability of market prices. We compute end-of-day prices corresponding to 16:00 GMT to match Bitcoin daily prices in all markets to daily spot rate from WM/Reuters corresponding to 16:00 GMT. The hourly prices of bitcoin on the different exchanges correspond to last transaction of the hour. The advantage of this procedure is that we can exactly match bitcoin prices to spot rates. However, this also implies that we could use, for some markets, bitcoin prices corresponding to hours of the day with smaller volume. For example, 16:00 GMT corresponds to 1AM in Seoul. Note that we treat currency pairs traded on different exchanges as different assets. For example, we treat differently the $US dollar to Bitcoin pair in market A and the $U.S. dollar to Bitcoin pair in market B. We restrict our sample

[2]The exchanges in our sample are: OKCoin, Bitfinex, Bithumb, GDAX, Coinbase, Bitstamp, bitFlyer, Kraken, Gemini, LakeBTC, Coinone, Korbit, Zaif, itBit, Gatecoin, BitBay, Coinfloor, Yunbi, Luno, Exmo, QuadrigaCX, Bitso, Jubi, BitMarket, MercadoBitcoin, Livecoin, Coinroom, Abucoins, YoBit, Unocoin, BTCC, Lykke, and Bit2C.

to exchanges where both fiat and Bitcoin currencies are traded, and exclude peer-to-peer platforms that have typically very low volumes. Table 2 contains descriptive statistics on our sample. The number of exchanges increases over time. Specifically, we start with 5 exchanges and 6 currency pairs in 2013 and end with 32 exchanges and 64 currency pairs in January, 2018. Since spot rates are only available on business days, we drop observations corresponding to non-business days. In addition, we set a liquidity requirement and drop daily observations if the volume of transactions is smaller than $U.S. 0.2 million[3].

There are two types of exchanges: the first, are exchanges on which cryptocurrency pairs are traded (i.e., Bitcoin for Ethereum), and where investors can deposit and withdraw only cryptocurrencies; the second, are instead exchanges on which it is possible to trade fiat currencies for cryptocurrencies (i.e., U.S. dollar for Bitcoin), and where investors can deposit and withdraw both fiat and cryptocurrencies.

Table 2. Our Sample

	Price	Trading	Annual volume in thous. BTC			Number	
year	BTC in $	volume in bln $ (24 h)	median	min	max	exchanges	currencies
2013	518.8547	0.003444	0.035534	0.005842	6.510924	5	6
2014	412.8576	0.033955	0.066003	5.77E-05	76.94517	11	14
2015	287.5319	0.091435	1.525172	0.014019	251.5832	15	25
2016	578.2071	0.554458	1.129173	0.013044	891.4508	20	35
2017	5335.826	1.274039	0.997887	0.013871	75.47238	31	63
2018	13906.31	2.842804	0.906161	0	57.10817	32	64

Notes: The table reports descriptive statistics of the sample. Note that annual volume data are computed before imposing the daily liquidity requirement of $U.S. 0.2 million. Data are daily from the crypto–exchange aggregator https://cryptocompare.com/.

Portfolios. At the end of each day t, we allocate the assets corresponding to investing in market m and currency j to five portfolios on the basis of their discount $D_{m,j,t}$ described in equation 1. Note that in the construction of the portfolios we only use information available up to time t. Portfolios are rebalanced at the end of every day. They are ranked from the lowest negative to highest positive discounts; portfolio 1 groups the returns from investing in the markets with the lowest discounts and portfolio 5 groups bitcoin returns from investing in the markets with the highest discounts. We compute the log bitcoin excess returns $rx_{m,j,t+1}$ for portfolio j by taking the average of the log bitcoin excess returns in each portfolio j

$$rx_{m,j,t+1} = \frac{1}{N_j} \sum_{i \in P_j} rx_{m,i,t+1}$$

[3]Our results do not change if we remove the liquidity requirement. Setting a tighter liquidity requirement excludes a large number of observations, especially in the early part of the sample.

The total number of assets in our portfolios varies over time. We have a total of 14 assets at the beginning of the sample in February, 2015 and 39 at the end in January, 2018. The maximum number of bitcoin returns attained during the sample is 50. Note we start building portfolios in February, 2015 because of data availability of Bitfinex, our benchmark exchange.

2.2 Returns to bitcoin speculation for US investor

Figure 3 offers a quick snapshot of our five portfolios obtained by sorting assets on their discounts. Excess returns increase monotonically across the five portfolios, from -429 basis point *per day* to 68 basis points. Even though the standard deviation of excess returns decreases across the portfolios, Sharpe ratios increase monotonically from -0.72 on the first portfolio to 0.33 on the last. Table 3 provides a detailed overview of the properties of the five bitcoin portfolios from the perspective of a US investor. The first panel reports the average discounts across the five portfolios. The first portfolio contains assets from exchanges that are "expensive" with respect to Bitfinex: the average discount is -4.25%. On the contrary, the last portfolio contains assets from exchanges that are "cheap" with respect to Bitfinex: the average discount is 1.06%. The mean discount across the five portfolios is approximately -1%. The second panel reports the dollar excess returns. Note that excess returns on the corner portfolios are significantly different from zero. As back of the envelope computation, if returns were i.i.d., we could simply compute the confidence intervals by dividing the portfolio volatilities by the square root of the number of observations. Note that $(28)^2 \approx 772$, i.e., the number of portfolio daily observations. Therefore, returns on both portfolios are highly significant at standard confidence levels. The third panel reports the high-minus-low excess returns from a strategy that goes long portfolio $j = 2, \ldots, 5$ and short the first portfolio. The spread returns between the last and first portfolio are approximately 500 basis points per day, with a Sharpe ratio of 0.86. For comparison, over the same sample, the mean daily returns on a buy-and-hold strategy that goes long bitcoins on Bitfinex was 51 basis points with a Sharpe ratio of approximately 0.12. The fourth and fifth panels reports information on the average daily volume of transactions, measured both in units of bitcoin and U.S. dollars. The point estimates indicate that portfolios with larger negative discounts contains, on average, assets with larger volume, eve though standard standard deviations are large. Finally, the last portfolio reports the average portfolio turnover as the time-average of the following ratio: the number of portfolio switches divided by the total number of assets in each date. The average frequency is 55 percent, and equal to 38 and 41 percent in the first and last portfolio respectively. Since this investment strategy requires large rebalancing of the bitcoin portfolios, it is important to account for transaction costs. In section 4 we show that accounting for transaction costs does not eliminate the large cross-section of bitcoin portfolio returns.

Figure 3. Five Bitcoin Portfolios

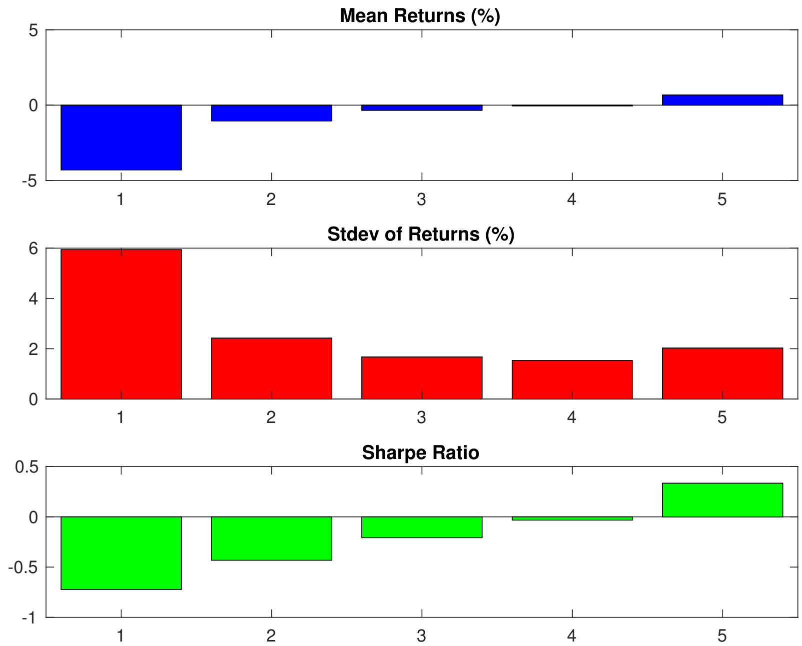

Notes: This figure plots means, standard deviations, and Sharpe ratios for the daily returns on the five bitcoin portfolios sorted on the basis of bitcoin discounts. Discounts are defined according to equation (1), and denote the percentage difference in the number of bitcoins that one US dollar can buy on different exchanges. All returns are expressed in US dollars and in percentages. Data are daily from https://cryptocompare.com/ and Thomson Reuters for the period 2/10/2015–1/24/2018.

3 Common Factors in Bitcoin Returns

This section shows that the large dollar excess returns on the bitcoin portfolios described in the previous section are explained by just two common components. Specifically, a principal component analysis on our bitcoin portfolios reveals that two factors explain approximately 96% of the variation in returns on these five portfolios, and that the first principal component accounts for approximately 80 percent of the total variance. Table 4 and figure 4 report the loadings of our bond portfolios on each of the principal components as well as the fraction of the total variance of portfolio returns attributed to each principal component. The loadings on the first principal component decrease monotonically across the five portfolios. Note how the first component is basically explaining the variation on the first portfolio, i.e., the portfolio that contains assets with the lowest negative discounts. The second component is

Table 3. Bitcoin portfolios: US investor

Portfolio	1	2	3	4	5
			Discounts: D^j		
Mean	-4.25	-1.17	-0.39	0.02	1.06
Std	5.28	2.30	1.68	1.52	2.16
			Bitcoin excess returns: rx^j_{t+1}		
Mean	-4.29	-1.05	-0.35	-0.05	0.68
Std	5.94	2.42	1.67	1.53	2.03
SR	-0.72	-0.43	-0.21	-0.03	0.33
			High minus low: $rx^j_{t+1} - rx^1_{t+1}$		
Mean		3.25	3.95	4.24	4.97
Std		4.58	5.20	5.41	5.80
SR		1.34	0.76	0.78	0.86
			Volume (in btc thousands): V^j_{btc}		
Mean	97.84	40.51	24.44	15.97	19.39
Std	205.72	136.35	88.53	55.96	43.57
			Volume (in US dollar millions): V^j		
Mean	53.43	27.92	19.38	17.81	20.35
Std	112.46	83.06	53.00	42.38	43.86
			Frequency		
Turnover	38.30	64.15	69.53	71.37	41.24

Notes: This table reports, for each portfolio j, the mean and standard deviation for the average discount D^j, the average log excess return rx^j, the average high–low spread hl^j, and the average spread return between portfolios $j = 2, \dots, 5$ and portfolio 1, the average volume V^j_{btc} expressed in thousands of bitcoins, and the average volume V^j expressed in US dollar millions. All moments, with the exception of those for volume, are reported in percentage points. For excess returns, the table also reports Sharpe ratios, computed as ratios of means to standard deviations. Portfolios are constructed by sorting assets into five groups at time t based on the their discounts D^j defined by equation (1). The first portfolio contains assets with the lowest negative discounts. The last portfolio contains assets with the highest positive discounts. The last panel reports the turnover, expressed as average number of trades per asset in each portfolio. Data are daily, from https://cryptocompare.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

a slope factor, and the loadings of the corner portfolios have opposite signs.

However, these two principal components are uncorrelated with standard risk factors. Table 5 reports the results of separate linear regressions, at daily frequency, of the first two principal components and the portfolio returns on a large set of candidate risk factors. Specifically, the Fama and French (1993) three factors, Carhart (1997)'s momentum factor, the log change in price of the Gold Bullion, the log change in the CBOE Vix index, and the change in the OIS spread. Most coefficients are not significantly different from zero, with the exception of the constant; the adjusted R-square are close to zero and the p-values of F-test on all coefficients jointly equal to zero are larger than standard significance levels in all cases. Note that these results are robust to excluding the first portfolio when extracting the common components. In this case, the first three components correspond to "level", "slope", and "curvature", but are still orthogonal to our risk factors.

Table 4. Bitcoin portfolios: Principal Components

Portfolio	1	2	3	4	5
1	0.92	-0.35	-0.18	0.01	0.00
2	0.30	0.40	0.72	-0.47	-0.06
3	0.18	0.40	0.18	0.65	0.59
4	0.14	0.42	-0.08	0.44	-0.78
5	0.12	0.62	-0.64	-0.41	0.19
% Var.	80.79	14.90	2.95	0.92	0.44

Notes: Principal component coefficients of the 5 Bitcoin portfolios. The last row reports (in %) the share of the total variance explained by each common factor. Data are daily, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

Figure 4. Principal Components

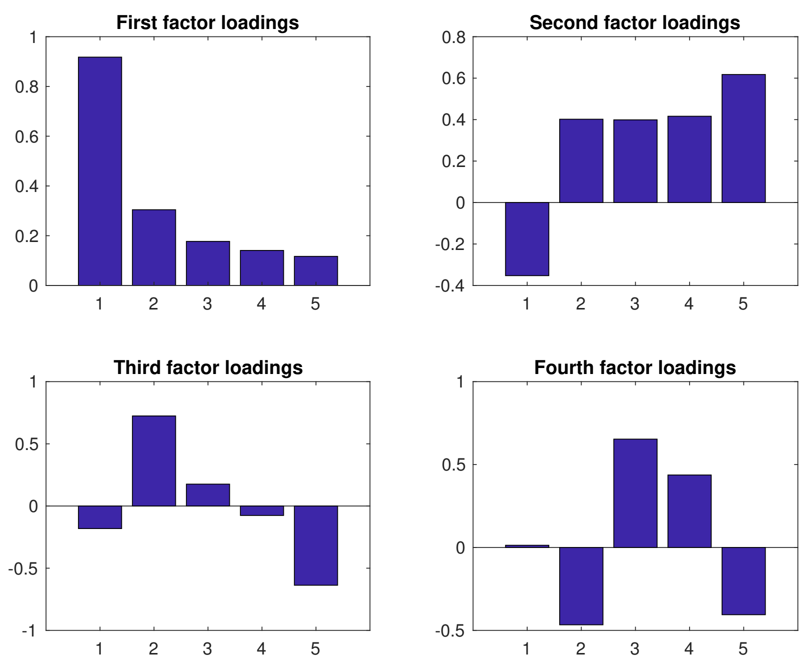

Notes: This figure plots the portfolio loadings on the first four principal components from the five bitcoin portfolios sorted on the basis of bitcoin discounts. Discounts are defined according to equation (1), and denote the percentage difference in the number of bitcoins that one US dollar can buy on different exchanges. All returns are expressed in US dollars and in percentages. Data are daily from https://cryptocompare.com/ and Thomson Reuters for the period 2/10/2015–1/24/2018.

4 Robustness

In this section, we consider several extensions. First, we look at execution risk, i.e., the risk that a transaction won't be executed within the range of recent market prices observed the investor. Second, we consider transaction costs. Third, we consider several critical issues

Table 5. Explaining Bitcoin Portfolio Returns

risk factor	PC1	PC2	P_1	P_2	P_3	P_4	P_5
constant	-0.04	0.01	-0.04	-0.01	-0.00	0.00	0.01
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Mkt	-0.12	0.01	-0.10	-0.10	0.01	0.01	0.00
	[0.21]	[0.10]	[0.21]	[0.07]	[0.05]	[0.05]	[0.07]
SMB	0.33	0.03	0.29	0.14	0.09	0.05	0.02
	[0.46]	[0.21]	[0.45]	[0.15]	[0.11]	[0.11]	[0.16]
HML	0.51	-0.02	0.45	0.24	0.15	0.07	-0.08
	[0.52]	[0.22]	[0.50]	[0.18]	[0.12]	[0.11]	[0.14]
MOM	0.15	-0.03	0.15	0.04	0.02	0.01	-0.02
	[0.25]	[0.11]	[0.24]	[0.08]	[0.06]	[0.06]	[0.08]
Gold	-0.27	-0.05	-0.25	-0.03	-0.03	-0.08	-0.14
	[0.23]	[0.11]	[0.22]	[0.08]	[0.06]	[0.06]	[0.09]
ΔVix	0.01	-0.00	0.01	-0.00	-0.00	0.00	0.00
	[0.02]	[0.01]	[0.02]	[0.01]	[0.01]	[0.01]	[0.01]
ΔOis	-69.84	13.95	-70.69	-9.65	-3.52	-5.07	-5.82
	[25.92]	[9.77]	[25.47]	[6.82]	[4.59]	[4.67]	[7.09]
R2 adj. (%)	0.78	-0.65	0.97	-0.33	-0.54	-0.60	-0.56
F(%)	23.27	88.51	19.78	46.82	85.22	69.53	63.64

Notes: Regressions of the first principal components and the returns on the five Bitcoin portfolios on different risk factors. Standard errors are Newey and West (1986). The last row reports the p-values (in percentages) of a F-test on all coefficients equal to zero. Data are daily, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

that could affect the operations of exchanges across the globe. Fourth, we look at bitcoin supplies in different markets.

Execution risk. U.S. investors buying bitcoins across exchanges around the globe are exposed to different forms of execution risk, i.e., the risk that a transaction is not executed within the range of recent market prices observed by investors. On the contrary, the definition of excess returns from equation (2) is based on the assumption of near-instant speed of execution and daily closing prices. In practice, completing the trade requires some time both because of the time required to transfer bitcoins across exchanges, and because of the time required for the foreign currency transfer[4]. It is hard to exactly quantify these amounts of time, as they depend both on the "type" of investor and the state of the network. First, while retail investors would typically need two business days for the foreign currency transfer, large investors could in principle have agreements with foreign financial intermediaries to reduce this time. Second, the proof-of-work, required by the blockchain to transfer bitcoins

[4] We report in table A3 in the appendix information on the approximate execution times for different cryptocurrencies.

across exchanges, depends on the solution of a computationally challenging problem which takes more time depending on the traffic on the network. As a first step to understand whether execution risk could explain the spread in the excess returns across the five bitcoin portfolios sorted on asset discounts, we consider two aspects of the execution risk. The first is related to possible changes of the exchange rate, and the second of the intra-day bitcoin prices. Since we do not have access to intraday exchange rate quotes, in order to evaluate the possible impact of exchange rate risk, we simply look at the log changes in the spot exchange rate with respect to the U.S. dollar for the assets in each portfolio. The first panel of table 6 reports the means and standard deviations of the log changes in the exchange rate. Intuitively, investors would suffer from an appreciation of the U.S. dollar (i.e., $\Delta s > 0$), as they would get less dollars from their foreign currency balance. On average, exchange rate growth is equal to zero and we do not find significant differences across the portfolios. The standard deviation of the exchange rate growth are slightly higher for the portfolios with the larger negative discounts, indicating that by investing in assets sold at discounts investors are marginally more exposed to exchange rate variability. In order to evaluate the impact of intra-day bitcoin price volatility, we first report, in the second panel of table 6, the spread between the high and low price of the day, measured in units of currency of market j per bitcoin, as a fraction of the average between the high and low price. The high-low spread captures the magnitude of intra-day price oscillations across markets. Also with respect to this measure, we do not find significant differences across the five bitcoin portfolios. In addition, we compute returns assuming that investors always buy bitcoins at the highest price of the day and sell bitcoins at the lowest price of the day[5]

$$\hat{rx}_{m,j,t} = p^{\star,low}_{1,1,t} - p^{\star,high}_{m,j,t} + s^j_t - r^f_t \,.$$

We denote the above as "lower bound" excess returns. The third panel of table 6 reports the mean, standard deviation and sharpe ratios of the portfolio excess returns under the "lower bound" assumption. Excess returns are now negative on all portfolios, ranging from -985 basis points per day on the first portfolio and -449 basis points per day on the first portfolio. This result is not surprising as this scenario is very adverse for the investor because of the extreme daily volatility of bitcoin prices. If one hand execution risk and, specifically, intra-day bitcoin price volatility could explain why discounts persist, on the other it cannot alone explain the cross-section of bitcoin excess returns as it produces a parallel downward shift in excess returns. In fact, the spread in returns between the first and last portfolio is large and significantly different from zero.

[5]Note that in the computation of $\hat{rx}_{m,j,t}$ we cannot match the timing of bitcoin prices and spot exchange rate quotes. This is because our data on spot rates correspond to closing prices at 4PM London time.

Table 6. Bitcoin Portfolios: Additional Characteristics

Portfolio	1	2	3	4	5
	Change in exchange rate: Δs^j				
Mean	0.01	0.02	0.01	-0.01	-0.02
Std	0.41	0.30	0.29	0.25	0.27
	High-Low spread: hl^j				
Mean	5.53	5.21	5.14	5.11	5.15
Std	5.87	5.63	5.00	4.90	4.73
	excess returns (lower bound): $\hat{rx}^j_{t+1}$				
Mean	-9.85	-6.38	-5.62	-5.25	-4.49
Std	9.59	6.56	5.87	5.64	5.44
SR	-1.03	-0.97	-0.96	-0.93	-0.83
	net excess returns: $\tilde{rx}^j_{t+1}$				
Mean	-4.37	-1.18	-0.49	-0.20	0.59
Std	5.93	2.41	1.67	1.53	2.04
SR	-0.74	-0.49	-0.29	-0.13	0.29
	Attacks: episodes				
Mean	2.20	1.09	0.94	0.86	1.09
Std	3.62	1.62	1.26	1.03	1.31
	Attacks: days not active				
Mean	6.22	3.29	1.55	1.75	3.22
Std	13.94	10.15	3.42	5.28	7.60
	Wallets (in btc thousands): W^j				
Mean	0.49	0.79	2.08	3.80	3.82
Std	3.13	4.08	6.74	9.41	9.04

Notes: This table reports additional characteristics of the five bitcoin portfolios obtained by sorting assets with respect to their discounts. The first panel reports the average and standard deviation of the log change in exchange rate with respect to the U.S. dollar for the assets in each portfolio. The second panel reports the average and standard deviation of the high–low spread constructed as the difference between the high and low daily bitcoin prices as a fraction of their average. The third panel reports the average, standard deviation and the Sharpe ratios of the excess returns computed under the assumption that investors always buy bitcoins at the highest price of the day and sell bitcoins at the lowest price of the day. The fourth panel reports the mean, standard deviation and Sharpe ratios of the excess returns net of transaction costs assuming a constant trading fee of 0.2% and abstracting from deposit and withdrawal fees. The fifth and sixth panels report the average and the standard deviation of our measure of market attacks measured, respectively, as total number of episodes and total number of days not active. The last panel reports the mean and standard deviation of the bitcoin supplies measured in thousands of bitcoins. Details on the definition of market attacks and bitcoin supplies are in section 4. Data are daily, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

Trading costs. In section 2 we have showed that the bitcoin investment strategy based on differences in the bitcoin discounts across markets produces a large cross-section of excess returns. However, the strategy requires a large turnover and, therefore, could be expensive in the presence of transaction costs. Most exchanges charge investors trading fees that are proportional to the size of the transactions. Typically, exchanges do not charge investors when they deposit crypto or fiat currencies, but they do charge lumps sum fees in case of withdrawals. It is convenient to consider a simple example to understand all the fees that a U.S. investor could in theory face. Figure 5 considers the case of a U.S. investor that first trade on Bithumb, one of the Korean exchanges, and then on Bitfinex, our benchmark exchange. First, the investor could have to pay a deposit fee to deposit Korean Won on Bithumb. In practice, exchanges do not charge deposit fees. Second, the investor must pay a trading fee when she exchanges Korean Won for Bitcoin. Third, the investor must transfer Bitcoins from Bithumb to Bitfinex. At this stage, she must pay a miner's fee to the node in the network that first validates the transaction on the blockchain. Fourth, she could have to pay a new deposit fee when depositing Bitcoins on Bitfinex. Fifth, she must pay a trading fee to exchange Bitcoins for U.S. dollars. Sixth, she must pay a withdrawal fee to take her U.S. dollar balance outside Bitfinex. This complex sequence of fees clearly reduces the net returns to the investor.

Figure 5. Trasaction Costs for U.S. Investor

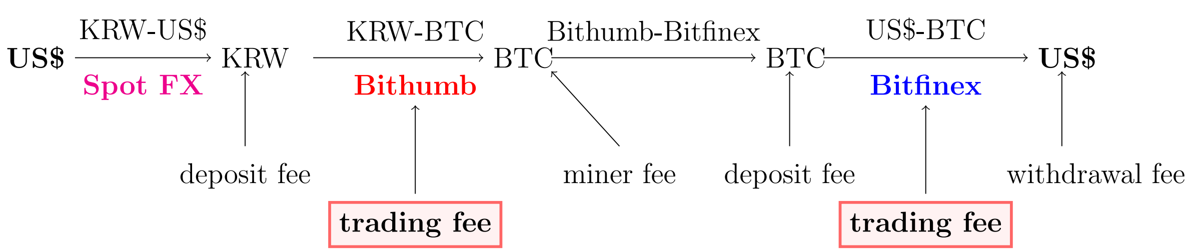

Notes: This figure describes the transaction costs that a U.S. investor face to conclude a bitcoin transaction that goes through Kraken and Bitfinex. We assume that the U.S. investor first convert dollars for Korean Won (KRW) on the spot market. Then, she buys bitcoins on Bithumb and transfer them immediately to Bitfinex. Finally, she convert bitcoins to dollars on Bitfinex and then withdraw her dollar balance.

In order to evaluate the likely impact of transaction costs on the cross section of portfolio returns, we assume a constant trading fee equal to 0.2%. This number is the average "taker" trading fee, applied to retail investors, across all the exchanges in our sample at the end of January, 2018. Note that for each completed trade, investors must pay the fee twice: i.e., at the time they start the investment, and at the time they conclude their investment. We abstract from deposit and withdrawal fees, miner's fees, and for additional fees on the foreign exchange spot market that are usually lumps sums and report in table A4 details on the fees applied by all the exchanges. In the fourth panel of table 6 we report, for each of the five bitcoin portfolios, the excess returns net of transaction costs. Excess returns decrease by about 20bp, an order of magnitude less than the cross section of portfolio returns. The net returns on the first and last portfolio are, respectively, -417bp and 69bp at daily frequency. Note that the returns of the corner portfolios are the least affected by the introduction of transaction costs as they require a smaller turnover with respect to the middle portfolios.

Cryptomarket critical issues Cryptomarkets are often subject to critical issues. First, cybersecurity, i.e., thefts and distributed denial-of-service attacks (DDoS). Second, jam due to high user traffic. Third, software maintenance updates and crashes. Fourth, uncertain regulatory environment. In the data, we define a DDoS attack as a day in which daily volume is equal to zero. Moore and Christin (2013) find that by early 2013, 45% of Bitcoin exchanges had closed, and many of the remaining markets were subject to frequent outages and security breaches. Vasek and Moore (2015) investigate denial-of-service attacks against cryptocurrency exchanges and document 58 such attacks. Note that our identification strategy is by construction imprecise. For example, if the attack starts in the middle of the day we could observe non-zero volume for that day. Also, we label as "DDoS attacks", events that are not necessarily linked to cybersecurity attacks (e.g., exchange inactive because of software maintenance). In order to evaluate the effectiveness of this strategy, we manually verify that our identification exactly captures a set of main critical events associated with the exchanges in the sample (see table A5 in the appendix). Figure 6 reports the daily number of total assets in our sample (blue line), together with the number of active assets (red line), where we define as "asset" a fiat-crypto currency pair for each exchange. Two facts stand out: first, the steep increase in the number of assets that goes from 14 to 62 in our sample; second, the fact that only about two-third of the assets are "active" in any given day. In the fifth and sixth panels of table 6 we report, for each portfolio, the average number of attacks, measured both in terms of total number of episodes and average number of days of inactivity, that characterizes the assets in the portfolio. Interestingly, we find that portfolios with the lowest negative discounts contain, on average, assets from markets that face a higher risk of being shut-down because of an attack. Specifically, the first portfolio, on average, contains assets that have, up to that day, been inactive for 6.22 day and object of 2.2 attacks. On the contrary, the last portfolio contains assets that have been inactive for 3.22 days and object of only 1.1 attacks.

Bitcoin supplies Cryptoexchanges function in many ways like brokers, or banks. Customers buy and sell bitcoins (or other cryptocurrencies), but typically maintain balances

Figure 6. Total and Active Number of Exchanges

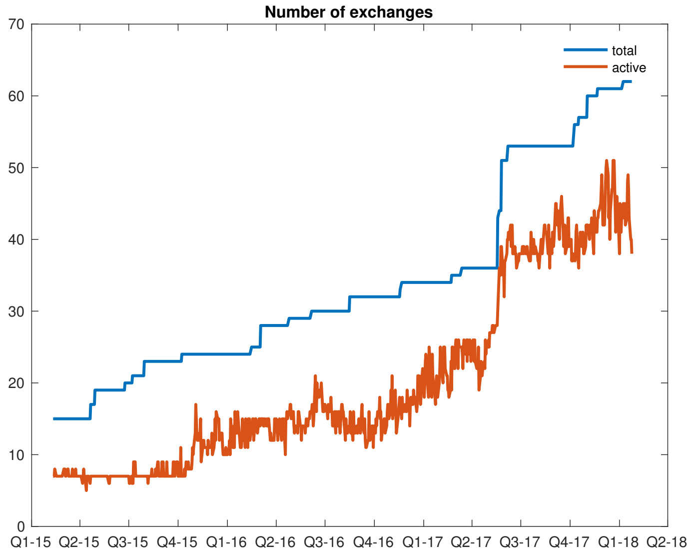

Notes: This figure plots the total (blue line) and active (red line) number of assets in our sample. An asset is denoted as "active" when the daily volume of transactions is non-zero. Data are daily, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

of both fiat currencies and bitcoins on the exchange without retaining direct access to the currency (Gandal et al., 2017). Investors's trades on an exchange are done off the blockchain. When investors deposit bitcoins on an exchange, these are put in a shared wallet that the exchange controls (i.e., these are like assets for a bank). The exchange keeps track of investors' balances, and of all the transactions. The blockchain only knows that investors send coins to the exchange, and consider these coins as owned by the exchange. When investors withdraw coins, then the blockchain is informed and bitcoins are transferred to the investors' personal wallets. In part because of the uncertain regulatory environment around cryptoexchanges, the size of the shared wallets is one of the indicators that investors use to evaluate the reliability of different markets and the risk of not being able to withdraw their balances. We collect daily data on exchange wallets from walletexplorer.com and bitinfocharts.com. We cannot cover all the exchanges in our sample, and in particular we do not have data from the Korean exchanges. Also note that wallets data are self-reported, and exchanges started to identify their wallets only after the Mt. Gox bankruptcy in 2013 to improve their transparency. The last panel of table 6 reports, for each portfolio, the average wallets size, in units of bitcoins. We find that the first portfolio, which contains on average assets from markets with the lowest negative discounts, has a smaller wallet than the last portfolio, which, instead, contains, on average, assets from markets with the largest positive discounts. The point estimates are large: the size of the wallets is, approximately, equal to 0.5 billions of bitcoins for the first portfolio and to almost 4 billions of bitcoins for the last portfolio. Therefore, we find evidence that exchanges with the lowest negative discounts are also those with the smallest supply of bitcoins. This evidence supports results in Cochrane (2002) according to which high prices with respect to "fundamentals" are typically associated with low share supplies.

5 Conclusions

In this paper we uncover a novel strategy based on persistent deviations of bitcoin price discounts across different exchanges around the globe. We build bitcoin portfolios using available information on discounts and find that the variation in returns is explained by just two common factors, which are mostly uncorrelated with a large set of standard risk factors. We document that portfolios containing assets from exchanges with the largest negative discounts tend to be different along several dimensions. Specifically, they tend to invest in assets from exchanges with a higher probability of temporary shut downs and the smallest bitcoin supplies; the larger transactions' volume; and higher return volatility. However, more work is required to understand the source of the price deviations. In particular, we plan to explore both currency and geographical risk factors. First, (Lustig et al., 2011) show that at least two risk factors, labeled "Dollar" and "Carry", respectively a "level" and "slope" factor, are required to explain aggregate returns in the currency market. Since currency returns are likely to play an important role for bitcoin returns across all the exchanges, we expect the "Dollar" and "Carry" factor to play a role. Second, we need to further investigate the distinct features of the bitcoin markets in terms of the allocation of buyers (i.e., retail investors) and sellers (i.e., miners). For example, the Korean and Japanese exchanges generate 60% of all trading volume in bitcoins. On the contrary, miners are mostly concentrated in China, the U.S., Europe, and Russia.

References

Biais, B., C. Bisiere, M. Bouvard, and C. Casamatta. 2018. The blockchain folk theorem .

Borri, N., and A. Verdelhan. 2011. Closed-end Fund Discounts and Aggregate Risk. Mimeo.

Carhart, M. M. 1997. On persistence in mutual fund performance. *The Journal of finance* 52:57–82.

Catalini, C., and J. S. Gans. 2016. Some simple economics of the blockchain. Tech. rep., National Bureau of Economic Research.

Chen, N.-f., R. Kan, and M. H. Miller. 1993. Are the Discounts on Closed-End Funds a Sentiment Index? *The Journal of Finance* 48:795–800.

Cochrane, J. H. 2002. Stocks as money: convenience yield and the tech-stock bubble. National Bureau of Economic Research Working Paper 8987.

Fama, E. F., and K. R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of financial economics* 33:3–56.

Gandal, N., J. Hamrick, T. Moore, and T. Oberman. 2017. Price Manipulation in the Bitcoin Ecosystem. *Journal of Monetary Economics* Forthcoming.

Gromb, D., and D. Vayanos. 2010. Limits of arbitrage. *Annu. Rev. Financ. Econ.* 2:251–275.

Kwon, J. 2017. South Korea is Going Bitcoin Crazy. CNN Tech.

Lamont, O. A., and R. H. Thaler. 2003. Can the market add and subtract? Mispricing in tech stock carve-outs. *Journal of Political Economy* 111:227–268.

Lee, C., A. Shleifer, and R. H. Thaler. 1991. Investor sentiment and the closed-end fund puzzle. *The Journal of Finance* 46:75–109.

Liu, J., and F. A. Longstaff. 2003. Losing money on arbitrage: Optimal dynamic portfolio choice in markets with arbitrage opportunities. *The Review of Financial Studies* 17:611–641.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *The Review of Financial Studies* 24:3731–3777.

Ma, J., J. S. Gans, and R. Tourky. 2018. Market Structure in Bitcoin Mining. Tech. rep., National Bureau of Economic Research.

Mitchell, M., T. Pulvino, and E. Stafford. 2002. Limited arbitrage in equity markets. *The Journal of Finance* 57:551–584.

Moore, T., and N. Christin. 2013. Beware the middleman: Empirical analysis of Bitcoin-exchange risk. In *International Conference on Financial Cryptography and Data Security*, pp. 25–33. Springer.

Nakamoto, S. 2008. Bitcoin: A peer-to-peer electronic cash system.

Newey, W. K., and K. D. West. 1986. A simple, positive semi-definite, heteroskedasticity and autocorrelationconsistent covariance matrix.

Scheinkman, J., W. Xiong, et al. 2003. Overconfidence, short-sale constraints, and bubbles. *Journal of Political Economy* 111:1183–1219.

Shleifer, A., and R. W. Vishny. 1997. The limits of arbitrage. *The Journal of Finance* 52:35–55.

Vasek, M., and T. Moore. 2015. There's no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency scams. In *International Conference on Financial Cryptography and Data Security*, pp. 44–61. Springer.

Velde, F., et al. 2013. Bitcoin: A primer. *Chicago Fed Letter* .

Yermack, D. 2013. Is Bitcoin a real currency? An economic appraisal. National Bureau of Economic Research Working Paper 19747.

Appendix

A Data

In this section we report additional details on the bitcoin data. Specifically, first we provide details on the major cryptomarkets across the globe. Second, we consider bitcoin prices and discounts. Third, we provide information on execution times and and transaction costs for different cryptocurrencies. Fourth, we analyze cryptomarket attacks.

A.I Cryptomarkets

Bitcoins are traded 24/7, every day of the week, including holidays, on several exchanges across the globe. There are two types of exchanges:

1. exchanges on which investors can trade cryptocurrency pairs, and where they can deposit and withdraw only cryptocurrencies,

2. exchanges on which investors can trade fiat for cryptocurrencies, and where investors can deposit and withdraw both fiat and cryptocurrencies.

Table A1 reports a list of the top exchanges by trading volume as for January, 2018. The third column reports the total number of currency pairs that can be traded, which is as large as 409 for HitBTC, an exchange on which only cryptocurrencies can be traded. The table also reports the U.S. dollar price of 1 Bitcoin in all exchanges, the daily trading volume (in U.S. dollar billions) and the launch dates. The exchange with the largest trading volume is Binance, an exchange in which only cryptocurrency pairs are traded. The second largest exchange is Bithumb, a Korean exchange, where investors can trade Korean Won (KRW) for several cryptocurrencies. The largest U.S. dollar based exchange is Bitfinex, where investors can use U.S. dollars and Euros to buy several cryptocurrencies.

Table A1. Top Crypto Exchanges

#	Name	Currency type	Number of currency pairs	Price BTC in USD	Trading volume in bln USD (24h)	Launch Date
1	Binance	crypto only	247	$10,979.60	$2.65	
2	Bithumb	crypto, KRW	12	$12,076.70	$2.32	10-Sep-13
3	Upbit	crypto only	214	$12,051.40	$2.07	24-Oct-17
4	Okex	crypto only	410	$10,921.10	$1.99	
5	Bitfinex	crypto, USD, EUR	85	$10,931.00	$1.62	8-Jun-13
6	Huobi	crypto only	145	$10,944.70	$1.27	
7	Bittrex	crypto only	270	$10,937.80	$0.59	
8	Kraken	crypto, EUR, USD, JPY	45	$11,042.20	$0.57	10-Sep-13
9	GDAX	crypto, USD, EUR, GBR	12	$10,916.00	$0.57	25-Jan-15
10	HitBTC	crypto only	409	$10,684.80	$0.48	

Notes: This table reports the list of the top crypto exchanges by trading volume. Data are for January, 24 2018 from https://coinmarketcap.com/. The number of pairs refers to the total number of currency pairs (fiat and crypto) that are traded on each exchange. The trading volume is in U.S. dollar billions and is annual.

The number of cryptocurrencies has increased substantially since the launch of bitcoin in 2009. Table A2 lists the three cryptocurrencies with the largest market capitalization, as of January 2018, together with their U.S. dollar price, daily volume, release date, maximum and currently circulating supply. The maximum supply denotes the largest number of units that can be mined according to the technology of each cryptocurrency. Not all the cryptocurrencies are the same. For example, Ripple uses a centralized clearing system, and allows almost instant transactions with very limited costs. Ethereum, instead, is not just a digital currency. It is a blockchain platform that features smart contracts, the Ethereum Virtual Machine (EVM), and allows users to create digital tokens that can be used to represent virtual shares, assets, proofs of membership, etc.

Table A2. Cryptocurrencies by Market Cap

Name	Market Cap in bln $	Price in $	Volume in bln $ (24h)	Release Date	Supply Max in mln	Supply Circulating in mln
Bitcoin	190	11,291	10	9-Jan-09	21	16.8
Ethereum	101	1,036	4	30-Jul-15	No Cap	97.2
Ripple	52	1	2	26-Sep-13	100,000	38739.1
Total	547		29			

Notes: This table lists the main cryptocurrencies by market capitalization. Data are for January, 24 2018 from https://coinmarketcap.com/. "Total" refers to all the cryptocurrencies tracked by the data aggregator. Market capitalization is computed as market value of circulating supply, as is in U.S. dollar billions. Volume is annual, and in U.S. dollar billions. Supply is in millions of units.

A.II Bitcoin Prices and Discounts

In this paper, we take bitcoin as representative cryptocurrency as it currently accounts for approximately 30% of trading volume and market capitalization. The price and volume of bitcoins has increased at an incredible pace since its launch in 2009. Figure A1 plots the daily U.S. dollar price on Bitfinex, together with the daily volume of transactions on the same exchange. The bitcoin price increased from approximately $US220 at the beginning of the sample to $US11,000 at the end of the sample and reached a maximum value of approximately $US18,900 on December, 18 2017.

U.S. dollar based investors can buy bitcoin on different exchanges. In some exchanges, like Bitfinex, they can directly trade U.S. dollars for bitcoins (and other cryptocurrencies). In others, investors must first convert U.S. dollars in an another fiat currency (e.g., the Korean Won), and then use this second currency to buy bitcoins. For example, figure A2 describes the available trade to a U.S. investor on May 25th, 2017 when we assume no transaction costs and real-time speed of execution. The U.S. investor could first exchange $1000 U.S. dollars for 1,126,000 Korean Won on the spot market. Then, she could trade Korean Won for bitcoins on Bithumb and obtain 0.2659 bitcoins. Third, the investor could transfer her bitcoins to Bitfinex and exchange them for $603.1 dollars. This would not be a very good deal for the investor who would, clearly, prefer to start the trade on Bitfinex and end the trade on Bithumb before converting Korean Won for U.S. dollars. In this paper, we always take Bitfinex as the exchange on which the U.S. investor ends her trade. Note that Bithumb

Figure A1. The U.S. Dollar Price of 1 Bitcoin on Bitfinex

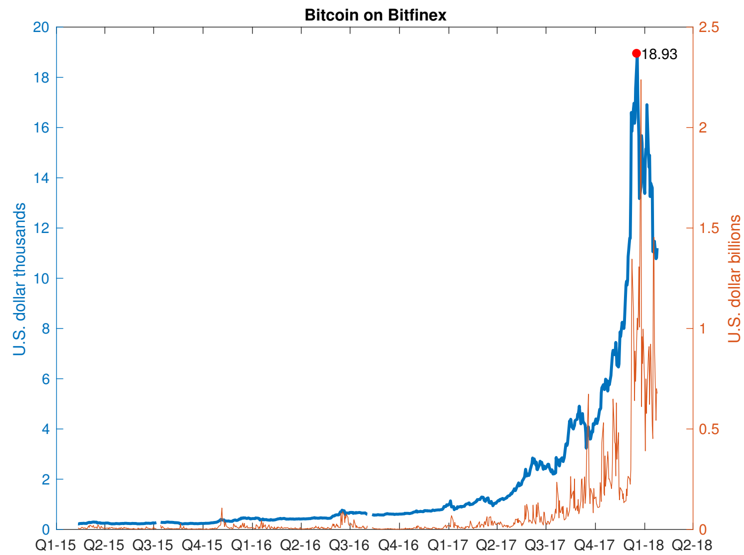

Notes: This figure plots the U.S. dollar price (in thousands) of 1 Bitcoin and the daily volume (in U.S. dollar billions) on Bitfinex. Data are daily from https://cryptocompare.com/ for the period 2/10/2015–1/24/2018.

is not the only exchange to post bitcoin prices, converted in U.S. dollars, that are very different from those of other exchanges. Figure A3 plots the daily U.S. dollar price of bitcoins on all exchanges in our sample and for all the currency pairs. There exist price differences throughout the sample, but the magnitude of these differences increases dramatically at the end of 2017. Figure A4 reports descriptive statistics on all the asset discounts. Specifically, the figure reports the mean, standard deviation, maximum, minimum, and first order autocorrelation of the discounts. We find that discounts are usually different from zero and volatile. In addition, they can can be positive and then negative within the same asset and persistent and mean-reverting.

Figure A2. U.S. investor buying bitcoins through Bithumb

$$\textbf{1,000 US\$} \xrightarrow[\textbf{Spot FX}]{\text{US\$ to KRW}} \text{1,126K KRW} \xrightarrow[\textbf{Bithumb}]{\text{KRW to BTC}} \text{0.2659 BTC} \xrightarrow{\text{Bithumb-Bitfinex}} \text{0.2659 BTC} \xrightarrow[\textbf{Bitfinex}]{\text{BTC to US\$}} \textbf{603.1 US\$}$$

Notes: Available transaction to U.S. investor based on Bitcoin price differences across Bithumb and Bitfinex on May 25th, 2017 assuming no transaction costs and near-instant execution time. Data are from https://cryptocompare.com/ and Thomson Reuters.

A.III Execution Times and Transaction Costs

Cryptocurrencies differ in terms of their execution times and transaction costs. Figure A3 reports the approximate times for transactions in different cryptocurrencies, and for the

Figure A3. The U.S. Dollar Price of 1 Bitcoin on All Exchanges

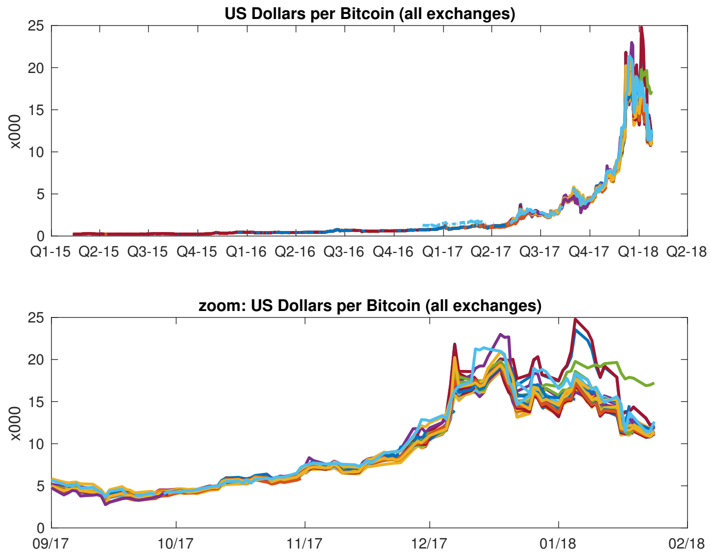

Notes: This figure plots the U.S. dollar price (in thousands) of 1 Bitcoin. Daily price for 64 different currency pairs in 32 different exchanges. Data are daily from https://cryptocompare.com/ for the period 2/10/2015–1/24/2018.

Figure A4. Bitcoin Discounts

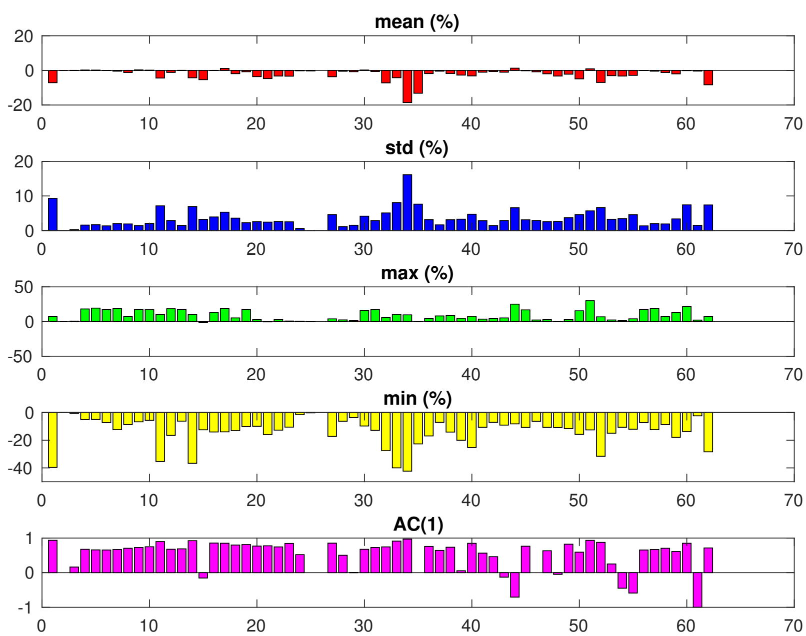

Notes: This figure plots the average, standard deviation, maximum, minimum and first order autocorrelation coefficient for the discounts on each asset. Discounts are defined by equation (1). Assets are all the fiat-crypto currency pairs in our sample. Data are daily from https://cryptocompare.com/ and Thomson Reuters. Samples are different for different assets. The longest sampls are for the period 2/10/2015–1/24/2018.

credit card provider Visa as a reference. The time estimates assume that the transaction is confirmed in the first block after being submitted and are approximate because they depend on the network congestion. For bitcoin, the maximum number of transactions per second is 7, with an estimated execution time of one hour. Both Ethereum and Ripple allow a larger number of transactions per second (respectively, 20 and 1,500) and a shorter execution time (respectively, 6 minutes and near-instant execution). The credit card provider Visa allows for a significant larger number of transactions per second, but also has larger transaction costs than cryptocurrencies.

Table A3. Cryptocurrencies by Execution Time

	Max Transactions per second	Estimated Execution Time	Transaction fee in crypto	Transaction fee in $
Bitcoin (BTC)	7	60 Minutes	0.001	11.29
Etherum (ETH)	20	6 Minutes	0.005	5.18
Ripple (XRP)	1,500	Near-instant	0.02	0.03
Visa	24,000	Near-instant		1.4–2.4%

Notes: We report approximate executions times for transactions in different cryptocurrencies. Execution times can vary depending on the conditions of the network. The time estimates assume that the transaction is confirmed in the first block after being submitted. Dollar fees depend on crypto-US% exchange rate and are for January, 2018. Data are manually collected by authors from different sources.

In table A4 we report, for all of the exchanges in our sample, information on the fees charged to investors. The second and third columns report, in percentage, the "taker" and "maker" trading fees. Taker fees are paid when investors remove liquidity form the order book by placing any order that is executed against an order on the order book. On the contrary, maker fees are paid when investors add liquidity to the order book of the exchange by placing a limit order below the ticker price for buy, and above the ticker price for sell. Typically, taker trading fees are equal or larger than marker trading fees. The fourth and fifth columns report, respectively, the withdrawal fees that investors pay when they withdraw bitcoins and fiat currencies. Most exchanges charge higher withdrawal fees when investors withdraw fiat currencies, rather than bitcoins. Withdrawal fees for fiat currencies are usually lump sums, but some exchanges charge fees proportional to the size of the transaction. Note that the numbers reported on table A4 are for January, 2018 and are manually collected by the authors from the exchange websites. Therefore, it is likely that they represent the fees paid by retail investors. Large investors are, instead, likely to pay smaller fees.

A.IV Cryptomarket attacks

Cryptomarkets are subject to several critical issues. One of the most common is represented by distributed denial-of-service attacks (DDoS attacks). A denial-of-service attack (DoS attack) is a cyber-attack where the perpetrator seeks to make a machine or network resource unavailable to its users by temporarily or indefinitely disrupting services of a host connected to the Internet typically by flooding the service with requests. In a DDoS attack, the traffic

Table A4. Investing in Bitcoin: Transaction Costs

Exchange	Trading fee (taker,%)	trading fee (maker,%)	withdrawal (in btc)	withdrawal (in fiat)	units withdrawal (in fiat)
OKCoin	0.20	0.20	0.00	0.10	%
Bitfinex	0.20	0.10	0.00	0.10	%
Bithumb	0.15	0.15	0.00	1000.00	KRW
GDAX	0.25	0.20	0.00	25.00	USD
Coinbase	1.49	1.00	0.01	–	
Bitstamp	0.25	0.25	0.00	0.90	EUR
bitFlyer	0.15	0.15	0.00	432.00	JPY
Kraken	0.26	0.16	0.00	0.90	EUR
Gemini	0.25	0.25	0.00	0.00	USD
LakeBTC	0.20	0.15	0.00	0.10	%
Coinone	0.10	0.10	0.00	1000.00	KRW
Korbit	0.20	0.08	0.00	1000.00	KRW
Zaif	0.00	0.00	0.00	486.00	JPY
itBit	0.20	0.00	–	–	
Gatecoin	0.35	0.25	–	–	
BitBay	0.43	0.43	–	–	
Coinfloor	0.30	0.30	0.00	1.50	EUR
Yunbi	0.05	0.05	0.00	0.10	%
Luno	0.20	0.00	–	–	
Exmo	0.20	0.20	0.00	2000.00	RUB
QuadrigaCX	0.50	0.50	0.01	1.00	%
Bitso	1.00	0.10	–	–	
Jubi	0.10	0.10	0.00	–	
BitMarket	0.45	0.15	0.00	2.00	EUR
MercadoBitcoin	0.70	0.30	0.01	1.99	%
Livecoin	0.18	0.18	0.00	3.00	%
Coinroom	0.25	0.06	0.00	0.00	PLN
Abucoins	0.25	0.00	0.00	0.00	PLN
YoBit	0.20	0.20	0.00	4.00	%
Unocoin	0.70	0.70	0.01	–	
BTCC	0.10	0.10	0.00	0.30	%
Lykke	0.14	0.09	0.00	50.00	CHF
Bit2C	0.50	0.50	0.01	60.00	NIS

Notes: Manually collected by authors on January, 2018 from exchanges websites. Trading fees are always in percentage. Withdrawal fees are fixed for most exchanges, and last column specifies the units.

flooding originates from many sources. Cryptomarkets are also subject to thefts, jams due to high user traffic, and suspensions of service due to software maintenance. In the data, we denote a DDoS attack *any* event that leads to a suspension of service and we identify these events with observations of zero daily volume of transactions. Figure A5 plots the daily fraction of exchanges that are inactive because of a DDoS attack as defined above. The mean fraction of daily inactive exchanges is approximately equal to 5%, and the maximum number of inactive exchanges is 27%. Table A5 reports a list of the main critical issues that cryptomarkets faced since 2014 that we manually collected using various sources. Most of the events are associated to thefts of bitcoins. For these events, in the third and fourth columns we report the amount of bitcoin stolen both in units of bitcoins and in millions of U.S. dollars. For each event, we also report the date corresponding to the first day of suspension of service as well as, when available, the date in which the exchanges resumed operations. The last columns reports the final outcome of the attack. Most exchanges are still active and have resumed operations, with the exceptions of Mt. Gox and Youbit who had to declare bankruptcy.

Figure A5. Fraction of Inactive Exchanges for DDoS attacks

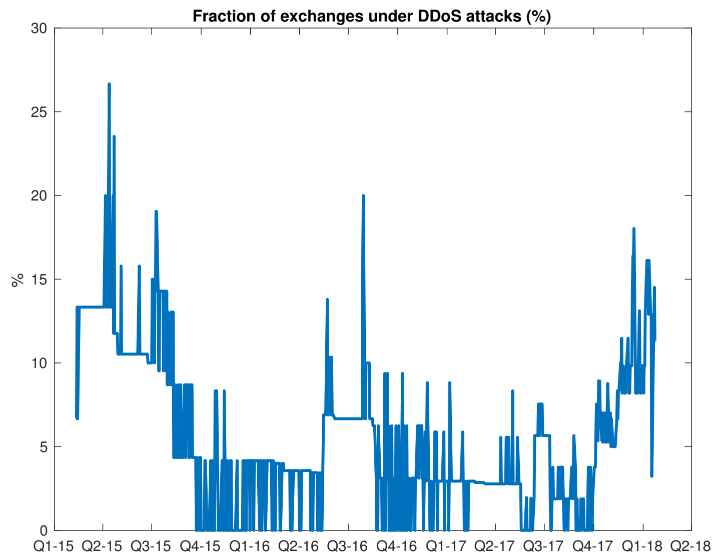

Notes: This figure plots the daily fraction of exchanges that are inactive as a consequence of a DDoS attack. The definition of DDoS attack is specified in section 4. Data are daily from https://cryptocompare.com/ for the period 2/10/2015–1/24/2018.

B Additional Robustness Checks

In this section we perform additional robustness checks on the bitcoin portfolios. Specifically, we form portfolios using weekly frequency data. We build end-of-week asset returns and discounts. Returns are computed every Friday by taking the log difference of the asset prices between Friday and Monday (i.e., four days prior):

$$rx_{m,j,t} = p_{m,j,t-4} - p_{1,1,t} - r^f_{t-4}$$

where t denotes the weekly frequency dates. Figure A6 provides a snapshot of the five bitcoin portfolios sorted on assets discounts. As for the daily frequency sample, excess returns, and Sharpe ratios, increase monotonically across the five portfolios. Table A6 provides details on the five portfolios. The mean discounts increase from -4.3% for the first portfolio, to 0.94% for the last portfolio. Excess returns increase, monotonically, from approximately 297 basis points per week to 186 basis points per week. The sample length is equal to 151 weeks. Recall that $12.3^2 \approx 151$. Therefore, assuming that portfolio returns are i.i.d., the excess returns for the corner portfolios are significant at standard confidence levels. As for portfolios build with daily frequency data, also for portfolios at weekly frequency the average volume is

Table A5. Critical Issues

	Name	btc	USD mln	suspended	resume	outcome
1	Mt Gox	850000	500	7-Feb-14	24-Feb-14	bankruptcy
2	NiceHash	4700	75	7-Dec-17	20-Dec-17	active
3	Bitfinex	120000	65	2-Aug-16		active
4	Parity		32	19-Jul-17		active
5	Tether		31	21-Nov-17		active
6	Bitstamp	19000	5	5-Jan-15	9-Jan-15	active
7	Youbit	4000	5	28-Apr-17		active
8	Youbit	1000	2	19-Nov-17		bankruptcy
9	Coincheck	523 mil NEM	534.8	28-Jan-18		active
1	Parity	locked funds, 513,774.16 ETH	160.8	7-Nov-17		active
2	Bitstamp	down		12-Feb-14	14-Feb-14	active
3	Kraken	down		10-Jan-18	12-Jan-18	active

Notes: The table reports a selected sample of critical issues on the major cryptocurrencies exchanges. The second and third columns report, in units of currency and US dollars respectively, either the amount of bitcoins or other cryptocurrencies stolen, locked or unavailable due to software maintenance or a DDoS attack. Data are manually collected by the authors from different sources.

decreasing going from portfolio 1 to 5. Table A7 presents additional characteristics of these portfolios that confirm our results on daily frequency data. First, we find no differences, across portfolios, with respect to the growth rates in the exchange rate and the high-low spread. Second, excess returns computed under the ”lower bound” scenario are all negative, but maintain the cross-section. Third, the spread in excess returns, net of transaction costs, is positive and significant. Fourth, portfolios with the largest negative discounts contain assets characterized by a larger number of attacks and of days being inactive because of attacks. Fifth, portfolios with the largest negative discounts contain assets from exchanges with the smallest shared wallets.

Finally, we report in table A8 the factor loadings and total explained variance on the first five principal components. The first two components explain virtually all the variance of portfolio returns. All portfolios load uniformly on the first component, which is a level factor. On the contrary, loadings on the second component, which explains only approximately 6% of the total variance, increase monotonically from the first to the last portfolio.

Figure A6. Five Bitcoin Portfolios (Weekly Frequency)

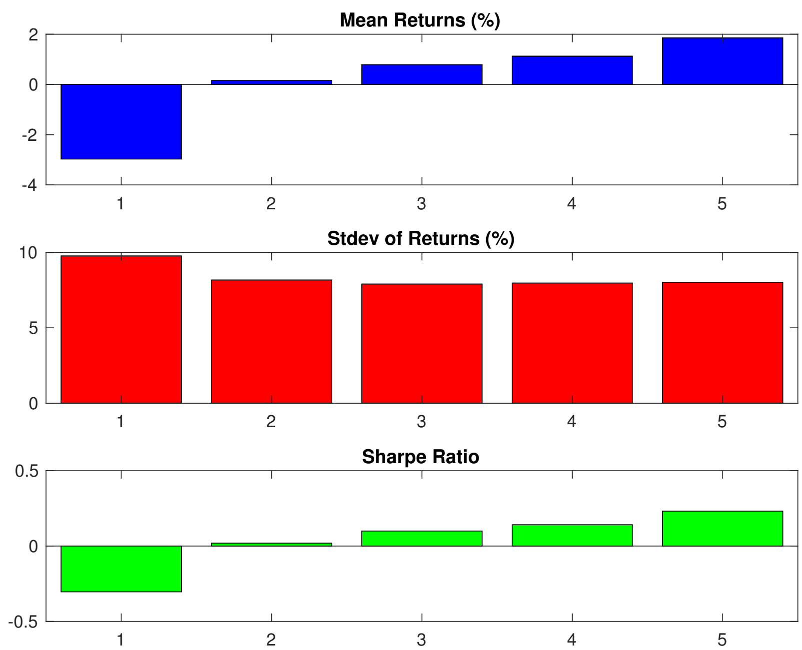

Notes: This figure plots means, standard deviations, and Sharpe ratios for the weekly returns on the five bitcoin portfolios sorted on the basis of bitcoin discounts. Discounts are defined according to equation (1), and denote the percentage difference in the number of bitcoins that one US dollar can buy on different exchanges. All returns are expressed in US dollars and in percentages. Data are weekly from https://cryptocompare.com/ and Thomson Reuters for the period 2/10/2015–1/24/2018.

Table A6. Bitcoin portfolios: US investor (Weekly Frequency)

Portfolio	1	2	3	4	5
			Discounts: D^j		
Mean	-4.30	-1.31	-0.52	-0.11	0.94
Std	5.28	2.35	1.72	1.52	1.93
			Bitcoin excess returns: rx^j_{t+1}		
Mean	-2.97	0.16	0.79	1.13	1.86
Std	9.77	8.18	7.91	7.97	8.02
SR	-0.30	0.02	0.10	0.14	0.23
			High minus low: $rx^j_{t+1} - rx^1_{t+1}$		
Mean		3.13	3.76	4.10	4.83
Std		4.36	5.05	5.31	5.68
SR		0.38	0.74	0.77	0.85
			Volume (in btc thousands): V^j_{btc}		
Mean	88.40	37.93	32.10	17.53	26.70
Std	185.01	131.24	106.17	45.91	76.13
			Volume (in US dollar millions): V^j		
Mean	46.50	26.73	22.75	18.19	24.36
Std	97.95	76.99	61.43	35.57	60.45
			Frequency		
Turnover	47.95	72.09	75.39	75.05	52.37

Notes: This table reports, for each portfolio j, the mean and standard deviation for the average discount D^j, the average log excess return rx^j, the average high–low spread hl^j, and the average spread return between portfolios $j = 2, \ldots, 5$ and portfolio 1, the average volume V^j_{btc} expressed in thousands of bitcoins, and the average volume V^j expressed in US dollar millions. All moments, with the exception of those for volume, are reported in percentage points. For excess returns, the table also reports Sharpe ratios, computed as ratios of means to standard deviations. Portfolios are constructed by sorting assets into five groups at time t based on the their discounts D^j defined by equation (1). The first portfolio contains assets with the lowest negative discounts. The last portfolio contains assets with the highest positive discounts. The last panel reports the turnover, expressed as average number of trades per asset in each portfolio. Data are weekly, from https://cryptocompare.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

Table A7. Bitcoin Portfolios: Additional Characteristics (Weekly Frequency)

Portfolio	1	2	3	4	5
	Change in exchange rate: Δs^j				
Mean	-0.02	-0.08	0.00	-0.04	0.03
Std	0.87	0.64	0.47	0.53	0.49
	High-Low spread: hl^j				
Mean	5.50	5.76	5.64	5.21	5.29
Std	5.80	7.21	7.60	5.32	5.54
	excess returns (lower bound): $\hat{rx}^j_{t+1}$				
Mean	-7.76	-4.56	-4.11	-3.58	-2.87
Std	11.68	9.78	9.48	9.45	9.32
SR	-0.66	-0.47	-0.43	-0.38	-0.31
	net excess returns: $\tilde{rx}^j_{t+1}$				
Mean	-3.10	0.01	0.64	0.99	1.77
Std	9.84	8.24	7.96	8.03	8.07
SR	-0.32	0.00	0.08	0.12	0.22
	Attacks: episodes				
Mean	0.62	0.28	0.32	0.29	0.40
Std	0.92	0.42	0.49	0.40	0.44
	Attacks: days not active				
Mean	8.54	3.06	1.67	2.43	3.01
Std	21.68	11.13	3.40	8.72	7.97
	Wallets (in btc thousands): W^j				
Mean	0.56	0.59	2.11	3.88	3.62
Std	3.89	3.66	6.69	9.85	8.35

Notes: This table reports additional characteristics of the five bitcoin portfolios obtained by sorting assets with respect to their discounts. The first panel reports the average and standard deviation of the log change in exchange rate with respect to the U.S. dollar for the assets in each portfolio. The second panel reports the average and standard deviation of the high–low spread constructed as the difference between the high and low weekly bitcoin prices as a fraction of their average. The third panel reports the average, standard deviation and the Sharpe ratios of the excess returns computed under the assumption that investors always buy bitcoins at the highest price of the day and sell bitcoins at the lowest price of the day. The fourth panel reports the mean, standard deviation and Sharpe ratios of the excess returns net of transaction costs assuming a constant trading fee of 0.2% and abstracting from deposit and withdrawal fees. The fifth and sixth panels report the average and the standard deviation of our measure of market attacks measured, respectively, as total number of episodes and total number of days not active. The last panel reports the mean and standard deviation of the bitcoin supplies measured in thousands of bitcoins. Details on the definition of market attacks and bitcoin supplies are in section 4. Data are weekly, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.

Table A8. Bitcoin portfolios: Principal Components (Weekly Frequency)

Portfolio	1	2	3	4	5
1	0.49	-0.85	-0.19	0.06	0.00
2	0.45	0.07	0.67	-0.57	-0.15
3	0.43	0.23	0.25	0.47	0.69
4	0.43	0.30	-0.05	0.50	-0.69
5	0.43	0.37	-0.67	-0.45	0.15
% Var.	93.60	5.79	0.36	0.17	0.08

Notes: Principal component coefficients of the 5 Bitcoin portfolios. The last row reports (in %) the share of the total variance explained by each common factor. Data are weekly, from https://coinmarketcap.com/ and Thomson Reuters. The sample period is 2/10/2015–1/24/2018.