Cryptocurrency value and 51% attacks: evidence from event studies

Savva Shanaev

University of Northumbria at Newcastle; Osiris Science; s.shanaev@northumbria.ac.uk

Arina Shuraeva

University of London; Osiris Finance; a.v.shuraeva@gmail.com

Mikhail Vasenin

Higher School of Economics; Osiris Finance; mmvasenin@gmail.com

Maksim Kuznetsov

Finance University under the Government of Russian Federation; Osiris Finance; makskuznetsov19@gmail.com

ABSTRACT

In this article, an event studies approach is utilised to assess the influence of 51% attacks on proof-of-work cryptocurrency prices. The study uses an exhaustive sample of 14 individual attacks on 13 cryptocurrencies. Across multiple event studies techniques, majority attacks on blockchains are consistently shown to immediately decrease corresponding coin prices by 12 to 15 percent. Significantly negative price response is robust in various event windows. Coin prices do not recover to pre-attack levels one week after the event. There is evidence of pump-and-dump schemes prior to the 51% attack, however the market demonstrates high efficiency after the attacks. 51% attacks are suggested to be a fundamental risk factor for cryptocurrency investments, primarily characteristic of small proof-of-work coins with low hash rates.

Keywords: cryptocurrency; alternative investment; 51% attack; event studies

JEL codes: C33, G14, G23

Electronic copy available at: https://ssrn.com/abstract=3290016

Cryptocurrencies have recently been becoming a popular alternative investment product, with at least 2071 coins and tokens being traded on multiple exchanges of total market capitalisation over $130 billion as of 22 November, 2018. As early as 2015, cryptocurrencies have been acknowledged as attractive investments and viable diversification options even in the academic sphere. For example, Briere, Oosterlinck and Szafarz (2015) estimate that a diversified portfolio including 3% or 6% of bitcoin can significantly outperform analogous portfolios that do not include cryptocurrencies both in terms of annualised return and annualised semivariance while offering the same level of return volatility due to spectacular historical performance of cryptocurrencies and low correlation of their returns with other financial markets (a simulated portfolio with 6% of capital allocated to bitcoin enjoys 32.51% return per annum, 12% annualised volatility and 5.38% annualised volatility). Similar findings have been recently presented by Chuen, Guo and Wang (2018), who construct a cryptocurrency index (CRIX) and study its performance. It is shown that while uncorrelated with traditional asset markets, CRIX delivers an outstanding 30.52% annualised return from August, 2014 to March, 2017.

However, the research has acknowledged that emerging asset classes such as cryptocurrencies are prone to bubble-like behaviour (Briere et al., 2015; Fry and Cheah, 2016) and demonstrate co-explosive interdependent price dynamics (Bouri, Shahzad and Roubaud, in press), therefore even seemingly “diversified” portfolios of cryptocurrencies such as CRIX exhibit high levels of tail risk exposure, with -22.64% daily drawdown and -1.04 daily return skewness (Chuen et al., 2018) on a predominantly bullish market sample. This ultimately suggests that cryptocurrencies are heavily exposed to similar risk events, and the bearish trend of 2018 has obviously established the need to investigate the fundamental sources of risk on the cryptocurrency market, filling the respective gaps in the literature.

As cryptocurrencies lack traditional fundamental characteristics typical of stocks, the analysis of their risk exposure has been proven difficult for both investors and researchers.

Electronic copy available at: https://ssrn.com/abstract=3290016

Bouri et al. (in press) assert that this has led investors to over-rely on price data, generating significant co-movement in coin prices. First ambitious attempts to derive risk factors for cryptocurrencies have been undertaken recently, evidencing that coin returns are not exposed to macroeconomic factors (Liu and Tsyvinski, 2018). Among the tested factors, the only ones reliably explaining cryptocurrency price dynamics are coin-specific momentum and investor attention (Liu and Tsyvinski, 2018; Chuen et al., 2018).

However, theory-grounded studies on cryptocurrency risk are virtually non-existent, probably reflecting the complexity of its conceptualisation and quantification. Nevertheless, there exists an easily identifiable event class that is characteristic of risks inherent to a significant amount of blockchains – the so-called 51% attacks. In essence, 51% attacks are short-term "hostile takeovers" of proof-of-work blockchains performed for financial or non-financial gains. This study therefore analyses the impact of 51% attacks on coin prices. It uses an exhaustive sample of 14 attacks on 13 cryptocurrencies and utilises a variety of event studies techniques to determine their impact. It shows that 51% attacks robustly decrease coin prices independent on the event window selected or estimation method chosen.

The rest of the article is organised as follows. First, an extensive literature review on the concept and history of 51% attacks and on the use of event studies in cryptocurrency-related research is elaborated. Second, the data collection techniques and event study methodology employed are discussed. Consequently, the empirical findings are presented and explored in the context of cryptocurrency risk pricing. The last section concludes.

LITERATURE REVIEW

51% attacks: theory and motivation

Since the very elaboration of the proof-of-work (PoW) concept, computer science academics, blockchain enthusiasts and cryptocurrency investors have been aware of the 51% attack threats.

Electronic copy available at: https://ssrn.com/abstract=3290016

In an influential paper, Kroll, Davey and Felten (2013) argue that a selfish or selfless 51% attack on a chain is one of the apparent risks proof-of-work cryptocurrencies are exposed to. If any individual coin miner or a group of miners controls over 50% of the network's mining capacity, they can essentially take over the chain. Most common "selfish" incentive provided in the literature for such a venture is so-called "double-spending", with the attacker being able to obtain coins via manipulating the blockchain consensus and thus securing a financial gain. However, as Kroll et al. (2013) argue, the payoff of such a scheme is significantly limited for two reasons: first, the miner performing an attack has to purchase specialised hardware, requiring a substantial initial investment; second, a 51% attack drastically reduces user and investor confidence in the network, dropping the coin price and thus decreasing the payoff an attacker has just generated via the double-spending procedure. These two factors combined make a 51% attack seem something reasonably unlikely: an individual or a pool already specialising in mining a particular cryptocurrency must forgo the prospective gains from continuous mining (as confidence issues post-attack make future cash flows highly uncertain) to obtain short-term profit from double-spending. This reasoning led researchers to believe that the utility agents derives from the attack must be non-monetary or at least must be generated outside of the blockchain, bringing forward the notion of "selfless attacks" and "Goldfinger attacks" that are expected to be performed by governments (to outright destroy blockchains that are used for purchasing illegal goods and services or for tax evasion), organised crime groups (to generally disrupt network functioning) or large investment funds with short positions in cryptocurrencies (to decrease coin value) (Kroll et al., 2013). This idea has been incorporated into the mathematical modelling of attacks on proof-of-work and proof-of-stake blockchains with an assumption that a successful attack grants some types of agents an exogenous amount of utility (Houy, 2014).

Electronic copy available at: https://ssrn.com/abstract=3290016

51% attacks: a brief history

In fact, 51% attacks used to be exceptionally rare prior to 2018: CoiledCoin was attacked in 2012, Terracoin and Feathercoin in 2013, while small projects Shift and Krypton were targeted in 2016. There were no attacks recorded in 2017 (Canellis, 2018). This is largely consistent with Kroll et al.’s (2013) assertion: among these five, only attacks on Terracoin and Feathercoin have had a double-spending motivation. CoiledCoin is believed to be attacked for personal motives (thus, the “selfless” attacker derived exogenous utility), while Shift and Krypton were threatened with a 51% attack to extract ransom from the projects’ lead developers and consequently attacked after a refusal to pay (“out-of-chain” financial motivation).

However, in 2018, the number of attempted and successful attacks on blockchains has skyrocketed: Group-IB, an international security think-tank, reports at least five major incidents that have resulted in $20 million earned by attackers via double-spending (Canellis, 2018). The majority of recent attacks also has been financially motivated (Cannelis, 2018), which arises a question why such attacks have now been enabled in comparison to not-so-distant past. First, coin prices have significantly increased relative to 2016 and especially 2013, making double-spending more lucrative. Second, a large number of proof-of-work altcoins has emerged, including those with quite small communities and low hash rates, that are, consequently, much more vulnerable to 51% attacks. Finally, and arguably most importantly, the mining industry has developed to offer computational power for rent (the most well-known platform for such services being NiceHash), therefore drastically lowering the initial investment required. Crypto51 (2018) project, for example, regularly estimates costs of 51% attacks on a wide selection of proof-of-work cryptocurrencies for educational purposes, also providing information on how much of the needed computational power can be rented on NiceHash. According to Crypto51 (2018), at least 13 coins can be 51%-attacked solely with

Electronic copy available at: https://ssrn.com/abstract=3290016

rentable power and 9 coins can be attacked for less than $100 an hour (see Exhibit 1 below). Notably, among these “vulnerable” coins, Horizen (more specifically, a previous version of the coin called ZenCash), Bitcoin Private, Litecoin Cash and Feathercoin has already been attacked.

Exhibit 1. Costs of a 51% attack on selected blockchains as of 22 November, 2018

Name	Symbol	Hash Rate	1-hour attack cost	NiceHash-able
Bitcoin	BTC	39,855 PH/s	$278,513	1%
Ethereum	ETH	195 TH/s	$100,555	5%
Bitcoin Cash	BCH	3,565 PH/s	$24,912	13%
Litecoin	LTC	205 TH/s	$23,437	8%
Dash	DASH	2 PH/s	$7,962	24%
Ethereum Classic	ETC	11 TH/s	$5,710	80%
Zcash	ZEC	2 GH/s	$19,416	9%
Dogecoin	DOGE	132 TH/s	$15,110	12%
Bytecoin	BCN	720 MH/s	$653	86%
Electroneum	ETN	3 GH/s	$3,000	19%
Metaverse ETP	ETP	719 GH/s	$371	**1235%**
MonaCoin	MONA	15 TH/s	$824	99%
Horizen	ZEN	161 MH/s	$1,329	**126%**
Bitcoin Private	BTCP	10 MH/s	*$79*	**2111%**
Vertcoin	VTC	4 TH/s	$227	**358%**
Einsteinium	EMC2	129 GH/s	*$15*	**12163%**
Quantum Resistant Ledger	QRL	5 MH/s	*$32*	**695%**
Ubiq	UBQ	198 GH/s	$102	**4485%**
Unobtanium	UNO	4,331 PH/s	$30,263	11%
Viacoin	VIA	58 TH/s	$6,639	27%
ZClassic	ZCL	70 MH/s	$579	**289%**
Litecoin Cash	LCC	14 PH/s	*$98*	**3308%**
Bulwark	BWK	2 TH/s	*$5*	18%
LBRY Credits	LBC	167 TH/s	$150	55%
PACcoin	$PAC	2 TH/s	*$5*	**36484%**
GameCredits	GAME	376 GH/s	*$43*	**4160%**
Feathercoin	FTC	6 GH/s	*$57*	**122%**
FlorinCoin	FLO	118 GH/s	*$14*	**13212%**

Notes: blockchains that can be attacked solely with rented hash power are presented **in bold**. Hourly attack costs lower than $100 are presented in *italics.* Source: Crypto51 project (2018).

The awareness of the blockchain community regarding 51% attacks has been growing rapidly. Several targets of past attacks have forked to improve their security protocols or to abandon the proof-of-work model altogether. Various safeguards against 51% attacks have been recently proposed by blockchain researchers and leading developers, including random miner selection (Bae and Jim, 2018) and penalties for suspicious mining (Biscotti, 2018), that have received mixed feedback from the community.

Electronic copy available at: https://ssrn.com/abstract=3290016

Overall, 51% attacks are remaining an acute threat to proof-of-work cryptocurrencies, especially for small-cap coins with low hash rates. Therefore, it would be beneficial to investigate coin price responses to these adverse events to better understand the risk-return characteristics of cryptocurrencies.

Event studies in the crypto-world

Event studies are a spectrum of predominantly non-parametric methods that are typically used to assess the response of a financial instrument's price to a particular event. Historically, this methodology has been extensively utilised in stock market research to investigate stock price reactions to important news, such as publications of disclosure or earnings announcements (MacKinlay, 1997). For these tests to have sufficient power and adequate generalisability, it is commonly preferred to identify a sample of instruments exposed to the same class of event and perform the analysis on an aggregated basis (Brown and Warner, 1985). For cryptocurrencies, event studies are extremely scarce, probably reflecting the complexity and non-triviality of event class conceptualisation for coins and the overall irregularity of this emerging asset class. Nevertheless, recently it has been evidenced that event studies might be useful to identify abnormal returns coin gain around blockchain-specific technological news. For example, Civitarese (2018) shows that altcoins NEO and IOTA have generated significantly negative abnormal returns in the 21-day period after respective news' releases questioning the efficiency of their protocols. Carlsson, Danielsson and Svensson (2018) utilise event studies to determine the stock price reactions of blockchain-related listed companies after corporate name changes and find significantly positive abnormal returns on the event date. However, these remain the only pieces of research to date applying event studies to cryptocurrencies. This study therefore seeks to eliminate the gap in the literature and to apply the most contemporary event studies

Electronic copy available at: https://ssrn.com/abstract=3290016

techniques to one of the most prominent event classes characteristic of cryptocurrencies – 51% attacks.

DATA AND METHODOLOGY

Data collection

Data collection process for this study consists of two distinct stages. First, all historical 51% attacks on proof-of-work blockchains must be identified. Second, financial data on the respective coins must be gathered for the estimation period, event window and post-event period.

On the first stage, the study utilises a variety of sources to create an exhaustive list of 51% attacks. Qualitative data is manually collected from corresponding blockchains' websites, news platforms, specialised blockchain-related social networks, interviews with lead project developers, and existing reports (Canellis, 2018). As no academic sources on 51% attacks are currently available, the study is extremely cautious on this step of the data collection process. Only these attacks that have been successfully implemented and have been confirmed by multiple sources are included into the sample. Therefore, rumoured attacks on Syscoin and AurumCoin and planned but not executed attack on Einsteinium are excluded. If different sources on the same attack have varying dates, the earliest date is taken to reflect the semi-strong market efficiency presumption event studies typically utilise. Overall, 15 confirmed and successfully executed attacks on proof-of-work blockchains have been identified: Bitcoin Gold, Bitcoin Private, CoiledCoin (excluded from the sample since coin data from 2012 is unavailable), Electroneum, Feathercoin, Karbo, Krypton, Litecoin Cash, MonaCoin, Pigeoncoin, Shift, Terracoin, Verge (attacked twice) and ZenCash (recently renamed to Horizen).

Electronic copy available at: https://ssrn.com/abstract=3290016

On the second stage, daily coin price data is extracted for 13 coins from Coinmarketcap for 30-day periods prior to the attack, for the attack date and for 10-day post-attack period. A 30-day pre-event period is selected to remain consistent with the existing event studies literature (MacKinlay, 1997) and the post-event period is slightly shorter both to prevent sample overlaps for Verge (which was attacked twice with a 48-day interval between attacks) and for data availability reasons, as Karbo was attacked on 10 November, 2018, only two weeks prior to this article being written. Raw sample data on 51% attacks can be consulted in Exhibit 2 below. As can be already seen from the preliminary data, all coins' prices respond negatively to 51% attacks, demonstrating an extremely consistent pattern. However, to derive more substantive and accurate estimations of abnormal returns, event study methodology is utilised, discussed in more detail in the next subsection.

Exhibit 2. 51% attacks: the sample

Blockchain	51% attack date	Amount stolen, $	Market cap before attack, $	Price response to attack
Verge (2nd attack)	May 22, 2018	35,000,000	809,524,978	-12.91%
Verge (1st attack)	Apr 04, 2018	277,466	841,569,688	-20.07%
Litecoin Cash	May 30, 2018	No data	42,216,948	-17.97%
ZenCash	Jun 03, 2018	550,000	119,264,946	-4.47%
Bitcoin Gold	May 23, 2018	18,000,000	916,739,710	-10.39%
Krypton	Aug 26, 2016	3,642	492,251	-29.84%
Shift	Aug 25, 2016	No data	212,747	-10.67%
MonaCoin	May 17, 2018	90,000	217,351,831	-10.50%
Bitcoin Private	Oct 19, 2018	0	52,036,971	-5.91%
Electroneum	Apr 04, 2018	No data	153,652,602	-13.65%
Karbo	Nov 10, 2018	0	862,004	-6.90%
Terracoin	Jul 25, 2013	No data	487,274	-6.67%
Feathercoin	Jun 10, 2013	2,016	852,774	-11.37%
Pigeoncoin	Sep 26, 2018	30,080	118,500	-61.72%

Notes: An exhaustive lists of confirmed 51% attacks on blockchains as of 22 November, 2018. The CoiledCoin attack is not represented as no market data is available for as early as 2012. Source: authors' calculations based on respective blockchains, interviews with lead developers, specialised news platforms, Canellis (2018) and Coinmarketcap.

Event study methodology

Interestingly, 51% attacks on blockchains have a somewhat direct analogy in the corporate world: hostile takeovers. Historically, hostile takeovers have been analysed with an event

Electronic copy available at: https://ssrn.com/abstract=3290016

studies approach (Casey, Dodd and Dolan, 1987), however stock price reactions to these seemingly adverse risk events have been found predominantly positive. That reflects the underlying motivation differences of corporate raiders (common executors of hostile takeovers) and that of miners performing 51% attacks. Contrary to the latter, seeking one-time financial or non-financial reward from double-spending or from the disruption of the blockchain altogether (Kroll et al., 2018), the former have incentives that may be well-aligned with long-term value creation and therefore can institute better governance practices or increase the overall company efficiency post-takeover (Casey et al., 1987). Perhaps, a more technically accurate analogy to hostile takeovers in the realm of blockchain would be hard forks executed by parts of the community to implement more advanced technological solutions, similar to the recent Bitcoin Cash hard fork. However, this is perhaps a subject for a different paper.

This study therefore primarily consults the literature on stock price reactions to broadly defined cyberattacks, arguably the event class most conceptually close to 51% attacks among those gaining sufficient attention in the recent research. As such, Campbell et al. (2003), Arcuri, Brogi and Gandolfi (2017) and Ganiaridis (2018) study the abnormal returns generated by stocks from various industries around security breaches and cyberattacks. Following this literature as well as more general methodology-related papers (Brown and Warner, 1985; MacKinlay, 1997), this study opts to utilise multiple techniques at once to ensure the robustness of the results. First, the cumulative (CAR) and buy-and-hold (BHAR) abnormal returns for a variety of event windows are calculated as in Brown and Warner (1985). The estimation period is selected as 30 days pre-event, consistent with the majority of academic sources (Brown and Warner, 1985; MacKinlay, 1997). To ensure an adequate level of statistical power and avoid multiple testing concerns, the significance tests are performed on an aggregate level for an equal-weighted pseudo-portfolio of coins. Following MacKinlay (1997), the market model is used, with the return of bitcoin as a proxy of the cryptocurrency market return. The study

Electronic copy available at: https://ssrn.com/abstract=3290016

acknowledges potential issues with such a methodological choice, however alternative ways of constructing a cryptocurrency market proxy could be even more assumption-sensitive[1]. Furthermore, to ensure consistency, the same tests are also performed for a constant returns model with the same estimation window. CAPM is not used, since, first, its benefits over the market model are questionable as it is more restrictive (Brown and Warner, 1985; MacKinlay, 1997) and, second, since there is no obvious risk-free instrument on the cryptocurrency market analogous to treasury bills on a traditional financial market. As in Arcuri et al. (2017), the significance of abnormal returns is tested for a variety of event windows (namely, [0;0], [-3;0], [-1;0], [0;1], [0;3], [0;6]) both to determine if coin prices regain lost value in a reasonably long period post-attack and if the attackers perform insider trading. The logic and incentives of pre-attack insider trading of miners and cyberattackers is conceptually different, however. While cyberattackers holding a planned target's shares are clearly incentivised to sell them before they execute the attack, using their private knowledge of a future price drop (Arcuri et al., 2017), 51% attackers get their reward via double-spending in the form of coins of the very blockchain they attack, therefore it is in their self-interest to push the price upward prior to the event, essentially accompanying the 51% attack with a standard pump-and-dump scheme typical for many cryptocurrency markets.

Since there are notable overlaps in the sample periods for different cryptocurrencies (for example, Verge and Electroneum have been attacked on the same day), variances of abnormal returns cannot be assumed independent and volatility clustering might be present (Brown and Warner, 1985; MacKinlay, 1997). Therefore, to perform an additional robustness check, abnormal returns are also estimated via a panel regression approach with dummy variables, as

[1] The most obvious broad market proxy for cryptocurrencies would be a value-weighted portfolio of all coins and tokens listed on Coinmarketcap. However, the construction of such a portfolio is itself a rather laborious task and it is unclear whether using it instead of simple bitcoin return would be effective in further reducing the abnormal return variance. As such, all sample coins are proof-of-work cryptocurrencies, and bitcoin, itself being a proof-of-work cryptocurrency, might be a benchmark of sufficient quality for the purposes of the study.

Electronic copy available at: https://ssrn.com/abstract=3290016

described in Brown and Warner (1985), with cross-sectional panel-corrected standard errors (PCSE), developed by Beck and Katz (1995). Moreover, following more contemporary approaches (Campbell et al., 2003), the same issue is also addressed via employing a seemingly unrelated regression (SUR) approach (Zellner, 1962). As both methods are essentially panel regressions with varying algorithms for standard error calculation, the estimators themselves will not differ, however this procedure can ensure robustness. All panel regression models are estimated in four forms: with common intercept and common slope (market beta), with fixed effects and common slope, with common intercept and differential slopes and with fixed effects and differential slopes.

As already stated above, the study performs its estimations on an aggregated basis using a pseudo-portfolio of coins. However, to account for possible heterogeneity biases, it also tests two hypotheses on an individual coin level: the null joint hypothesis is that all abnormal returns on the day of a 51% attack are equal to zero and the null average hypothesis is that the mean abnormal return on the day of 51% attack is equal to zero as in Campbell et al. (2003). These hypotheses are tested in a SUR framework using a redundant variables F-test and Wald coefficient restriction F-test, respectively.

In the next section, the results of the statistical procedures described above are presented and discussed.

EMPIRICAL FINDINGS AND DISCUSSION

First, as in MacKinlay (1997), buy-and-hold abnormal return of a pseudo-portfolio of coins around the 51% attack date is calculated (see Exhibit 3 below). Market model and constant return model present consistent results, showing significant pseudo-portfolio value drop on the event date (-15.41% and -12.60% in constant return and market model, respectively, both significant at 1%). Interestingly, there is apparent positive abnormal returns generated by the

Electronic copy available at: https://ssrn.com/abstract=3290016

coins on the two days preceding the attack (9.86% and 5.76% in constant return and market model, significant on 1% and 10%, respectively), possibly evidencing some insider trading or pump-and-dump schemes performed by the attackers prior to the event.

Exhibit 3. Buy-and-hold abnormal return around the 51% attack date

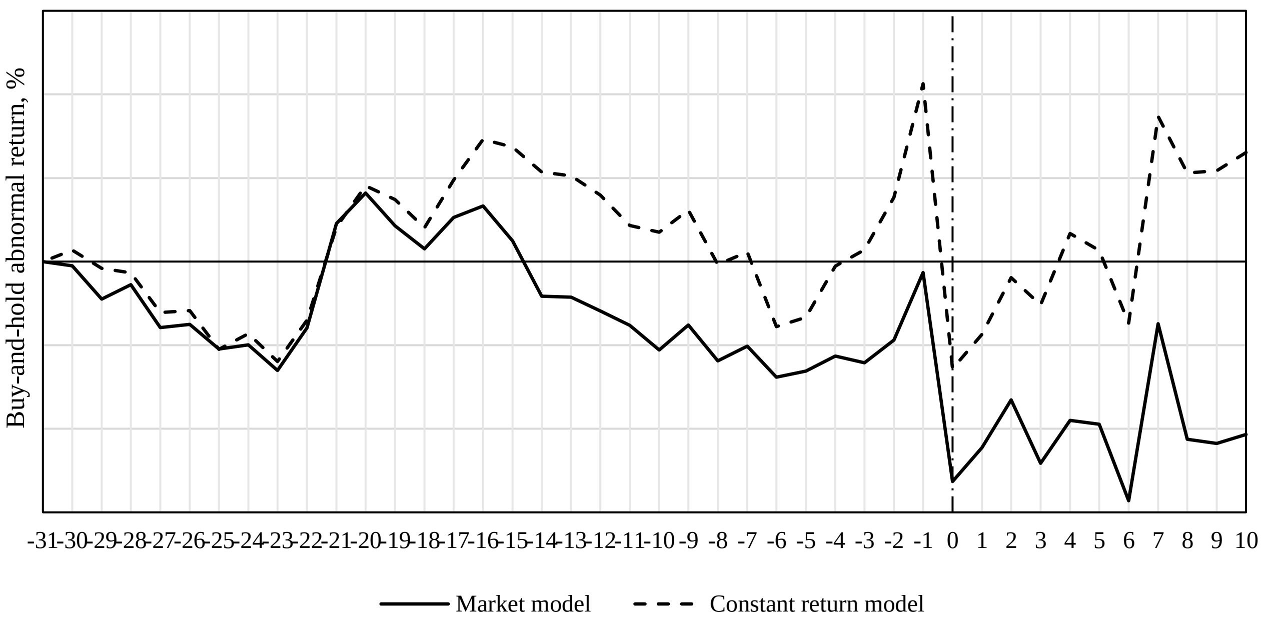

Notes: Buy-and-hold abnormal return (BHAR) estimated with market model and constant return model around the event date. The event date is represented with a vertical dash dot line. Source: authors' calculations.

Next, cumulative abnormal returns and buy-and-hold abnormal returns for multiple event windows are calculated (see Exhibit 4 below). Both the constant return model and the market model demonstrate that 51% attacks cause significant and persistent decline in coin prices. Even a week post-attack (event window [0;6]) cryptocurrencies on average are traded 13% below their pre-attack levels, with only 2 coins (less than 15% of the sample) recovering to their pre-attack value. The response of different cryptocurrencies to 51% attacks is extremely consistent, with all coins generating negative CAR on the event date and over 70% of the sample having negative CARs and BHARs for all event windows. The findings are robust across CAR and BHAR measures for both models. Noticeably, there are no significant gains or losses realised by cryptocurrencies after the attack date, suggesting both a surprising degree of market efficiency and lack of overreaction bias, thus supporting the emergent efficiency

Electronic copy available at: https://ssrn.com/abstract=3290016

claim of Urquhart (2016) and the earlier findings of altcoin market efficiency from event studies (Civitarese, 2018). Therefore, while the cryptocurrency market is seemingly inefficient pre-attack due to possible market manipulation by 51% attack beneficiaries, it rather promptly absorbs new information post-attack.

Exhibit 4. Cumulative and buy-and-hold abnormal returns for 51% attack events

	Event window	[0;0]	[-3;0]	[-1;0]	[0;1]	[0;3]	[0;6]
Market model	**CAR, %**	-12.60***	-7.32*	-8.36***	-10.27***	-11.18**	-13.58**
		0.0000	*0.0964*	*0.0097*	*0.0020*	*0.0138*	*0.0227*
	% negative	100.00	71.43	64.29	85.71	78.57	85.71
	BHAR, %	-12.60***	-7.96*	-8.89***	-10.56***	-11.50**	-13.75**
		0.0000	*0.0722*	*0.0063*	*0.0015*	*0.0116*	*0.0212*
	% negative	100.00	78.57	64.29	85.71	85.71	85.71
Constant return model	**CAR, %**	-15.41***	-4.76	-8.87**	-13.20***	-11.31**	-12.24*
		0.0000	*0.3450*	*0.0172*	*0.0008*	*0.0302*	*0.0725*
	% negative	100.00	71.43	71.43	100.00	85.71	71.43
	BHAR, %	-15.41***	-6.15	-9.88***	-13.54***	-11.96**	-12.94*
		0.0000	*0.2252*	*0.0087*	*0.0006*	*0.0226*	*0.0584*
	% negative	100.00	78.57	71.43	100.00	85.71	71.43

Notes: Cumulative and buy-and-hold abnormal returns (CAR and BHAR) estimated with market model and constant return model for various event windows. Corresponding p-values are reported *in italics.* *, ** and *** denote significance at 1%, 5% and 10%, respectively. Source: authors' calculations.

Then, a robustness check is employed, with average abnormal returns estimated in a panel regression framework using cross-sectional panel-corrected standard errors (PCSE) and a seemingly unrelated regressions (SUR) approach to address possible volatility clustering (see Exhibit 5 below). The findings are consistent with the previous results. Even accounting for variance dependency, all abnormal returns in [0;0] and [0;1] event windows are statistically significant. Notably, both average abnormal returns and cumulative abnormal returns for the coins around 51% attacks are substantially higher than those of stocks around cyberattacks and security breaches: Arcuri et al. (2017) report a [0;5] CAR of -0.54% and Campbell et al. (2003) estimate a [-1;1] CAR at -1.88%. This reflects both the seriousness of a 51% attack as compared to an average cyberattack as well as the high volatility of cryptocurrency markets in general.

Electronic copy available at: https://ssrn.com/abstract=3290016

Exhibit 5. 51%-attack abnormal returns estimated with panel regressions and SUR

	Event window	**Panel regression approach**				**Seemingly unrelated regressions (SUR) approach**			
		Common slope		Differential slopes		Common slope		Differential slopes	
		Common intercept	Fixed Effects	Common intercept	Fixed Effects	Common intercept	Fixed Effects	Common intercept	Fixed Effects
Average abnormal return	[0;0]	-15.04***	-15.05***	-13.28***	-13.26***	-15.04***	-15.05***	-13.28***	-13.26***
		(2.75)	(2.76)	(2.76)	(2.77)	(3.26)	(3.29)	(3.22)	(3.25)
		0.0000	*0.0000*	*0.0000*	*0.0000*	*0.0000*	*0.0000*	*0.0000*	*0.0001*
	[-3;0]	-1.46	-1.46	-1.92	-1.90	-1.46	-1.46	-1.92	-1.90
		(1.46)	(1.46)	(1.45)	(1.45)	(2.07)	(2.10)	(1.94)	(1.97)
		0.3176	*0.3186*	*0.1842*	*0.1907*	*0.4816*	*0.4860*	*0.3231*	*0.3340*
	[-1;0]	-4.40**	-4.40**	-4.35**	-4.32**	-4.40	-4.40	-4.35*	-4.32
		(2.00)	(2.01)	(2.00)	(2.00)	(2.79)	(2.82)	(2.64)	(2.67)
		0.0288	*0.0291*	*0.0298*	*0.0316*	*0.1158*	*0.1194*	*0.0995*	*0.1059*
	[0;1]	-6.73***	-6.73***	-5.33***	-5.31***	-6.73**	-6.74**	-5.33**	-5.31**
		(2.00)	(2.00)	(1.99)	(2.00)	(2.68)	(2.71)	(2.59)	(2.62)
		0.0008	*0.0008*	*0.0076*	*0.0080*	*0.0122*	*0.0131*	*0.0402*	*0.0430*
	[0;3]	-3.40**	-3.40**	-2.91**	-2.89**	-3.40*	-3.40*	-2.91	-2.89
		(1.45)	(1.46)	(1.44)	(1.45)	(2.02)	(2.04)	(1.92)	(1.94)
		0.0197	*0.0201*	*0.0439*	*0.0460*	*0.0924*	*0.0961*	*0.1289*	*0.1358*
	[0;6]	-2.23*	-2.23*	-2.06*	-2.05*	-2.23	-2.23	-2.06	-2.05
		(1.15)	(1.15)	(1.14)	(1.11)	(1.61)	(1.63)	(1.52)	(1.54)
		0.0533	*0.0538*	*0.0716*	*0.0731*	*0.1669*	*0.1713*	*0.1762*	*0.1818*

Notes: average abnormal returns calculated using Panel and SUR regressions and dummy variable approach for various event windows (as discussed in Brown and Warner, 1985). Panel regressions are estimated with robust cross-sectional panel-corrected standard errors (Beck and Katz, 1995). Cross-sectional SUR regressions (Zellner, 1962) are performed as in Campbell et al. (2003). Standard errors are reported (in parentheses) while corresponding p-values are presented *in italics.* *, ** and *** denote significance at 1%, 5% and 10%, respectively. Source: authors' calculations.

Electronic copy available at: https://ssrn.com/abstract=3290016

Finally, addressing the issue of heterogeneity bias frequently encountered in similar studies (Campbell et al., 2003; Arcuri et al., 2017; Carlsson et al., 2018), joint and average hypotheses are tested in a SUR framework for disaggregated return data as in Campbell et al. (2003). Both the joint hypothesis (that all 51% attack abnormal returns are equal to 0) and the average hypothesis (that a mean 51% attack abnormal return is equal to 0) have to be rejected at the 1% confidence level (see Exhibit 6). Coupled with consistently negative CARs and BHARs for all event windows and for over 70% of the coins, this evidences the absence of heterogeneity bias in the sample.

Exhibit 6. Hypothesis testing

Hypothesis	F-statistic	p-value
Joint (all 51% attack abnormal returns are equal to 0)	4.2883***	*0.0000*
Average (mean 51% attack abnormal return is equal to 0)	15.8698***	*0.0001*

Notes: the hypotheses are tested on a disaggregated basis (each attack on a coin has a unique dummy variable) in a SUR framework (Zellner, 1962) as in Campbell et al. (2003). The joint hypothesis is tested using a redundant variables F-test and the average hypothesis is tested using a Wald coefficient restriction F-test. *, ** and *** denote significance at 1%, 5% and 10%, respectively. Source: authors' calculation.

CONCLUSION

This study has estimated the impact of 51% attacks on blockchains on corresponding proof-of-work cryptocurrency prices applying a spectrum of event study methodologies (CAR, BHAR, panel regressions with dummy variables and seemingly unrelated regressions approach) to an exhaustive sample of 14 attacks on 13 cryptocurrencies. It has found a robust and statistically significant negative price reaction on the attack date, ranging from -12% to -15% depending on the estimation technique. The CARs and BHARs are consistently negative across all event windows for more than 70% of the sample coins. There is some evidence of insider trading prior to 51% attacks, representing a typical pump-and-dump scheme which is well-aligned with the attackers' incentives. However, cryptocurrency markets have been proven to immediately

Electronic copy available at: https://ssrn.com/abstract=3290016

incorporate the new information into the coin prices and to demonstrate no overreaction bias, signalling increasing market efficiency (Urquhart, 2016; Civitarese, 2018).

The findings of the study have broad implications for cryptocurrency investors and are possibly paving the way for further research in the field. This study has been the first to provide an early conceptualisation of fundamental risk factors associated with investing into cryptocurrency markets. The quantitative results of this study can be helpful for cryptocurrency investors to understand possible downside risks and to perform adequate stress-testing of their coin portfolios. As 51% attacks are becoming more frequent and are often synchronised, i.e. executed on multiple blockchains at the same time, it would perhaps be interesting to study whether 51% attack risk is diversifiable or, alternatively, if it presents a unique cryptocurrency-specific source of systematic risk. If the latter is true, then, reflecting the fact that 51% attacks are much easier to execute on proof-of-work blockchains with low hash rates, the first "fundamental" risk factor portfolios can be constructed for the cryptocurrency market. Moreover, the methodology developed by this study can be utilised for other related event classes, such as blockchain hard forks or switches between proof-of-work and proof-of-stake, both topics also being of great importance to cryptocurrency investors.

Reference list

Arcuri, M., Brogi, M., & Gandolfi, G. (2017). How Does Cyber Crime Affect Firms? The Effect of Information Security Breaches on Stock Returns. In *ITASEC* (pp. 175-193).

Bae, J., & Lim, H. (2018, June). Random Mining Group Selection to Prevent 51% Attacks on Bitcoin. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)* (pp. 81-82).

Beck, N., & Katz, J. (1995). What to Do (and Not to Do) with Time-Series Cross-Section Data. *American Political Science Review*, *89*(3), 634-647.

Electronic copy available at: https://ssrn.com/abstract=3290016

Biscotti, C. (2018, November 16). A Solution to Crypto's 51% Attack? Fine Miners Before It Happens. *Hawthorne Caller.* Retrieved from https://hawthorncaller.com

Bouri, E., Shahzad, S., & Roubaud, D. (in press). Co-Explosivity in the Cryptocurrency Market. *Finance Research Letters*. https://doi.org/10.1016/j.frl.2018.07.005

Briere, M., Oosterlinck, K., & Szafarz, A. (2015). Virtual Currency, Tangible Return: Portfolio Diversification with Bitcoin. *Journal of Asset Management, 16*(6), 365-373.

Brown, S., & Warner, J. (1985). Using Daily Stock Returns: The Case of Event Studies. *Journal of Financial Economics*, *14*(1), 3-31.

Campbell, K., Gordon, L, Loeb, M., & Zhou, L. (2003). The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market. *Journal of Computer Security, 11*(3), 431-448.

Canellis, D. (2018, October 24). Cryptocurrency Hackers Earned $20 million in 2018 with So-Called 51% Attacks. *Business Insider*. Retrieved from http://uk.businessinsider.com

Carlsson, C., Danielsson, F., & Svensson, C. (2018). *The Effect of Blockchain Related Corporate Name Changes on Stock Prices: An Investigation into the Creation of Cumulative Abnormal Returns Following a Blockchain Related Corporate Name Change.* http://www.diva-portal.org/smash/get/diva2:1235823/FULLTEXT01.pdf

Casey, R., Dodd, P., & Dolan, P. (1987). Takeovers and Corporate Raiders: Empirical Evidence from Extended Event Studies. *Australian Journal of Management*, *12*(2), 201-220.

Chuen, D., Guo, L., & Wang, Y. (2018). Cryptocurrency: A New Investment Opportunity? *The Journal of Alternative Investments, 20*(3), 16-40.

Civitarese, J. (2018). *Technical Development, Asset Prices and Market Efficiency in Alternative Cryptocurrencies.* http://dx.doi.org/10.2139/ssrn.3154124

Electronic copy available at: https://ssrn.com/abstract=3290016

Crypto51 (2018). *Cost of a 51% Attack for Different Cryptocurrencies*. Retrieved from: https://www.crypto51.app/

Ganiaridis, P. (2018). *Evaluating the Financial Effect from Cyber Attacks on Firms and Analysis of Cyber Risk Management*. https://dspace.lib.uom.gr/handle/2159/21675

Houy, N. (2014). *It Will Cost You Nothing to 'Kill' a Proof-of-Stake Crypto-Currency*. http://dx.doi.org/10.2139/ssrn.2393940

Fry, J., & Cheah, E. (2016). Negative Bubbles and Shocks in Cryptocurrency Markets. *International Review of Financial Analysis*, *47*(1), 343-352.

Liu, Y., & Tsyvinski, A. (2018). *Risks and returns of cryptocurrency* (No. w24877). National Bureau of Economic Research.

MacKinlay, A. (1997). Event Studies in Economics and Finance. *Journal of Economic Literature*, *35*(1), 13-39.

Kroll, J., Davey, I., & Felten, E. (2013, June). The Economics of Bitcoin Mining, or Bitcoin in the Presence of Adversaries. In *Proceedings of WEIS* (Vol. 2013, p. 11).

Urquhart, A. (2016). The Inefficiency of Bitcoin. *Economics Letters*, *148*(1), 80-82.

Zellner, A. (1962). An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias. *Journal of the American Statistical Association*, *57*(2), 348-368.

Electronic copy available at: https://ssrn.com/abstract=3290016