Bitcoin Policy Institute

Block Fullness, Security, and the Price of Bitcoin

Douglas Cumming
Joshua Hendrickson
Zachary Glatzer
William J. Luther

Bitcoin Policy Institute

Block Fullness, Security, and the Price of Bitcoin

Douglas Cumming
Department of Finance, Florida Atlantic University

Zachary Glatzer
Department of Finance, Florida Atlantic University

Joshua Hendrickson
Department of Economics, The University of Mississippi

William J. Luther
Department of Economics, Florida Atlantic University

Bitcoin Policy Institute

Block Fullness, Security, and the Price of Bitcoin

About the Bitcoin Policy Institute

The Bitcoin Policy Institute (BPI) is a non-partisan, non-profit think tank. It is dedicated to educating policymakers and the public on Bitcoin and disruptive digital technologies, providing research-based insights to inform sound policy in the United States.

The BPI team comprises experts in economics, law, philosophy, energy, and environmental science, working together to explore the impacts of new technology on existing US public policy interests. The views expressed in this publication do not necessarily reflect the views of all Bitcoin Policy Institute management or its affiliated scholars.

Block Fullness, Security, and the Price of Bitcoin

Abstract

Do cybercriminals time the execution of cryptoexchange breaches and frauds? We argue that, the ability to execute transactions quickly improves the odds of successfully laundering bitcoin, and therefore cybercriminals will be less likely to attempt breaches and frauds when they expect bitcoin blocks will be congested. Furthermore, since those accepting bitcoin when blocks are full can be more confident that it was not acquired by breach or fraud, they will require a smaller discount. As anticipated, we find that block fullness is negatively associated with breaches and frauds and positively associated with the price of bitcoin. Finally, we exploit the segregated witness software upgrade to address endogeneity concerns. The results suggest the aforementioned relationship between block fullness and breach or fraud incidents are causal.

JEL Codes: K42, E42
Keywords: blockchain, blocksize, bitcoin, crime, cybercrime, payments

The first major centralized crypto-asset exchange started with magic and ended in calamity. Mt.Gox, which is short for "Magic: The Gathering Online eXchange" and references the website's earlier use as a platform for trading Magic cards, lost roughly $460 million worth of bitcoin following a security breach in 2014 (McMillan 2014). The failure of Mt.Gox to secure customer funds marked an early blow to bitcoin adoption, and continues to serve as a warning of the vulnerabilities associated with centralized exchanges.

Centralized exchanges are an attractive target for cybercriminals. Shumba (2021) estimates that 10 percent of all bitcoin, worth around $100 billion, is held in digital wallets controlled by a centralized exchange. These wallets are akin to bank vaults: they contain much more value than the typical non-custodial wallet. For this reason, the digital wallets controlled by a centralized exchange are arguably the largest security vulnerability in the bitcoin ecosystem.

Cybercriminals can exploit the existence of centralized exchanges in at least two ways: breaches and frauds. A breach occurs when the cybercriminal hacks—i.e., gains unauthorized access to—an exchange's wallet. If the cybercriminal gains control of an exchange's wallet, he can transfer the crypto assets held by the exchange to a wallet the exchange does not control. Alternatively, the cybercriminal can engage in fraud by posing as a legitimate exchange. After a customer deposits currency or crypto assets at the fraudulent exchange, the cybercriminal can use those funds as he sees fit.[1] The Mt.Gox case, discussed above, is an example of a successful breach. The case of QuadrigaCX, which was an exchange based in Canada, is an example of a successful fraud.[2]

Pulling off a successful breach or fraud is only the first step in the cybercriminal's grift. He must also launder the proceeds. That is especially challenging since cryptoassets are typically tracked on public ledgers. If a cybercriminal fails to launder the proceeds effectively, he may

[1]The fraud need not be intended at the outset. For example, a legitimate exchange may later become fraudulent by misappropriating customer funds.

[2]QuadrigaCX's founder, Gerald William Cotten, reportedly died in India in 2019. At the time, the exchange owed 215 million CAD (roughly $165 million)) to more than 76,000 clients. However, recovered assets totaled just 46 million CAD (roughly $35 million)OSC (2020).

be identified when attempting to spend or cash out his ill-gotten gains at some later date.

Consider three ways a cybercriminal might attempt to launder the proceeds of a breach or fraud. He may transfer the proceeds to a complicit exchange located in a jurisdiction that will not share his information with the relevant authorities in order to purchase some national currency. He may employ mixing technologies[3] and decentralized exchanges to obscure a series of transactions and crypto-asset exchanges on the shared ledgers before ultimately spending the proceeds. Or, he may use mixing technologies and decentralized exchanges before sending the proceeds to a centralized exchange, which may or may not be complicit, in order to acquire some national currency. With each of these three laundering approaches, the cybercriminal is more likely to succeed if he can execute the desired transactions quickly. A complicit exchange operating today may not be operating or may no longer be willing to comply at a later date; and a jurisdiction with strong financial privacy protections might weaken those protections in the future. Likewise, a complex mix routed through a decentralized exchange might obscure the transactions path for awhile, but may ultimately be mapped out through a careful analysis of the public ledgers. A cybercriminal caught attempting to transfer ill-gotten gains may see those transactions denied by the exchange or recipient—and may be apprehended by authorities if his real-world identity is revealed in the process. Hence, fast action boosts the odds of a successful getaway.

It is not always possible for a cybercriminal to launder the proceeds of a breach or fraud quickly without incurring additional costs. For crypto assets like bitcoin, which has a block size limit and adjusts the difficulty of mining to ensure a block is processed every ten minutes, the network can become congested. And, if the network is congested, the cybercriminal must either pay a higher fee to have his transactions included in the next available block or else wait for it to be included in some subsequent block.

In what follows, we consider the timing of breaches and frauds associated with the Bitcoin network and their effect on bitcoin's price. Our intuition is straightforward: since the ability

[3]Involves sending bitcoins to an external protocol that combines your coins with other users' coins and returns coins to a new address you control that cannot be easily linked to your original coins.

to execute transactions quickly improves the odds of successfully laundering bitcoin acquired by breach or fraud, cybercriminals will be less likely to attempt breaches and frauds when they expect bitcoin blocks will be full (i.e., when the block weight is at or near the limit). Furthermore, those accepting bitcoin when blocks are full can be more confident that the bitcoin they receive were not acquired by breach or fraud and, hence, will not be seized by the authorities or refused by others in the future. Consequently, these users will require a smaller discount for bitcoin received when the block weight is at or near the limit and the price of bitcoin will be positively related to the fullness of blocks on the bitcoin blockchain.

1 Background

Bitcoin is a digital asset that can be used to make payments on the Internet.[4] Bitcoin consists of a shared ledger—or, blockchain—and a protocol for updating that ledger.[5] When a user initiates a bitcoin transaction, that transaction is placed in the mempool with other unconfirmed transactions waiting to be processed. Each miner running the Bitcoin protocol decides whether to include the transaction in the block they are attempting to add to the existing blockchain. Miners then race to produce an appropriate SHA-256 hash of their block's header.[6] The first miner to produce an appropriate hash attaches their block of transactions to the prevailing blockchain, effectively updating the shared ledger to reflect the transactions in the miner's block. The successful miner is awarded a quantity of newly created bitcoin, known as the block reward, as well as any transaction fees submitted by

[4]Selgin (2015) contrasts bitcoin with traditional commodity and fiat monies. On its ability to serve as a medium of exchange, see Luther (2016, 2018, 2019), Hendrickson et al. (2016), Hendrickson and Luther (2017, 2022). Hazlett and Luther (2020) consider whether bitcoin is money.

[5]Luther and Smith (2020) explain that bitcoin is a distributed payment system, with important differences from traditional centralized and decentralized payment systems.

[6]A hash function scrambles the inputted data (e.g., the block's header) and returns a unique fixed-length result. The output of a SHA-256 hash is essentially a random number between 0 and the maximum value of a 256-bit number. Requiring the output be less than or equal to the existing difficulty target makes producing an appropriate hash computationally difficult and effectively turns the race to confirm a block of transactions into a lottery, where the odds of winning are approximately equal to the ratio of the miner's computing power to the computing power of all miners.

users whose transactions were included in the miner's block.[7]

To regulate the speed at which new bitcoin is created, the bitcoin system periodically adjusts the difficulty of producing an appropriate hash. These difficulty adjustments ensure that one block is added to the blockchain approximately every ten minutes. If the number of miners increases or their ability to hash improves, this increases the likelihood that blocks will be added to the blockchain more often than once every ten minutes on average. To slow the speed of block creation, the bitcoin system will increase the difficulty of producing an appropriate hash. Likewise, if blocks are being added to the blockchain less often than once every ten minutes on average, the system will decrease the difficulty of producing an appropriate hash. Consequently, a block of transactions is processed roughly once every ten minutes.

The number of transactions that can be included in a block is limited by the block limit. Initially, the bitcoin protocol lacked an explicit block size or weight limit.[8] That changed in 2010, when the psuedonomous Satoshi Nakamoto introduced a 1 MB block size limit to prevent SPAM. Since then, the amount of data—and, hence, the number of transactions—that can be included in a block has been limited.

Since the number of transactions per block is constrained by the block limit and the number of blocks per hour is constrained by the difficulty adjustment, the number of transactions that the bitcoin protocol can process per hour is limited. This can lead to congestion, where transactions are initiated faster than they can be processed.[9] As transactions build up in the mempool, the average time it takes to process a transaction increases. Moreover, since miners are more likely to include a transaction that has a higher fee attached, the average value of transaction fees increases as well.

[7]The block reward was initially set to 50 bitcoin and is cut in half every $210,000$ blocks until the total supply of bitcoin equals 21 million in approximately the year 2140.

[8]The size of a block was initially limited by the number of database locks required to process it. However, this limit varied across nodes and was not widely known until an update of the protocol in early 2013 failed to enforce it. The initial limit was removed by a hardfork in May 2013, at which point the higher 1 MB block size became the universal limit.

[9]Koops (2018) estimates the time it will take for a bitcoin transaction to be processed given the number of transactions in the mempool. See also: Gundlach et al. (2021), Stoepker et al. (2021)

The congestion problem sparked a big debate in the bitcoin community, known as The Blocksize War.[10] On one side, *large blockers* argued that the block size limit should be increased, thereby enabling the bitcoin protocol to process more transactions per hour. On the other, *small blockers* warned that increasing the block size limit would consolidate control of the bitcoin consensus mechanism since the higher costs associated with operating a full node with larger blocks would discourage some from doing so. Small blockers eventually won the war, and the 1 MB limit remains in effect to this day.[11]

In an effort to deal with the congestion problem without increasing the 1 MB block size limit, the Bitcoin protocol was upgraded to segregate the witness data associated with each block. Prior to the activation of SegWit (as the protocol upgrade is known) in August 2017, a transaction included (1) sender and receiver data and (2) witness data, with both counting toward the block size limit in full.[12] Following the activation of SegWit, the witness data was removed and appended as a separate structure and the 1 MB block size limit was replaced with a 4 million weight unit limit. Whereas the now-segregated witness data's weight was set equal to its size, the weight of the sender and receiver data was set equal to four times its size. Hence, the relative contribution of the sender and receiver data's size to the limit was left unchanged while the relative contribution of the witness data's size to the limit was reduced. Correspondingly, the immediate effect of the activation of SegWit was the ability to process more transactions per hour.

To execute a bitcoin transaction, one must first acquire some bitcoin. In brief, there are three ways to acquire bitcoin. One can receive newly created bitcoin (and, potentially, fees) as a reward for mining; receive previously-existing bitcoin as a gift or in exchange for some good or service; or, purchase bitcoin on an exchange. In the latter case, one might use a centralized or decentralized exchange. A decentralized exchange matches buyers and sellers

[10] Bier (2021) offers a comprehensive account.

[11] Some large blockers supported a hardfork of the bitcoin blockchain, which resulted in the creation of a distinct cryptocurrency known as bitcoin cash in August 2017. Bitcoin cash would split again in November 2018, with the creation of bitcoin SV. Neither bitcoin cash nor bitcoin SV are as popular as bitcoin is today.

[12] Although the SegWit upgrade began in November 2016, it remained dormant until 95 percent of miners running the bitcoin protocol signaled support.

of bitcoin, but is not a party to the transaction. A centralized exchange, in contrast, acts as a market maker: it stands at the ready to buy or sell bitcoin at the prevailing price. Since purchasing bitcoin on a centralized exchange is the most commonly-used method and also the method most relevant to the study at hand, we limit our attention accordingly.

In order to acquire bitcoin on a centralized exchange, one must initially transfer some asset to the exchange with which to make the purchase. Typically, one will transfer some national currency (or, claim to national currency) from a bank account. However, many exchanges accept cryptocurrency deposits, as well. The exchange will then credit the user with a balance denominated in the national currency or cryptocurrency. The user may then use that balance to purchase bitcoin at the prevailing price.

Once a balance of bitcoin is acquired, the user must decide whether to hold the balance in a custodial or non-custodial wallet. With a non-custodial wallet, the user holds bitcoin and can initiate a transaction with his private key. With a custodial wallet, some other entity—typically, the centralized exchange used to purchase bitcoin—holds bitcoin on the user's behalf, while the user holds a claim to bitcoin; and, since the user does not hold the corresponding private key, he must instruct that entity to initiate a transaction on his behalf. Most centralized exchanges offer custodial wallets.

There are benefits to holding bitcoin (or, more precisely, claims to bitcoin) in a custodial wallet, which is similar in many respects to holding a claim to dollars in a checking account. For one, the user might worry about losing access to his wallet. With a non-custodial wallet, the user holds his own private key; if he loses it, he will not be able to transfer his bitcoin. In contrast, a user who loses the password to his custodial wallet can request that access be restored by the custodian. A custodial wallet may also permit the user to transfer his balance at a lower cost, so long as the transfer is to another wallet maintained by the custodian. Since the custodian can settle such transactions with a mere book entry, it need not incur the cost of settling the transaction on the blockchain and may pass the savings on to its users. A custodial wallet may also permit the user to acquire bitcoin more quickly. Many centralized

exchanges permit the user to purchase bitcoin before the exchange actually receives the funds from the user's bank or wallet, with a restriction that the user holds the balance in a custodial wallet maintained by the exchange until the funds are actually received. Since the user holds a claim to bitcoin, and not the corresponding private key, the centralized exchange can simply reverse the book entry transaction should the funds fail to be delivered. Finally, for those purchasing bitcoin on a centralized exchange, a custodial wallet may be more convenient simply because it is the default option.

Each user must weigh the benefits of holding bitcoin in a custodial wallet against the corresponding risks. In particular, using a custodial wallet leaves one exposed to breaches and frauds. The term breach denotes the unauthorized access of one's wallet (Charoenwong and Bernardi 2021, Moore and Christin 2013; Nguyen and Putniņš 2023). The term fraud denotes a situation where one is deceived into authorizing access (Vasek 2015; Xia 2020). Neither breaches nor frauds are unique to custodial wallets. Likewise, criminal activity may occur between two parties with an exchange merely functioning as an intermediary rather than as the victim of a breach or perpetrator of fraud (Cong et al. 2023; Dhawan and Putniņš 2023). Nonetheless, we maintain that the value of assets held by custodial wallet services, particularly those associated with an exchange, relative to the value of assets held by a typical user makes them a more attractive target for a breach; and the structure of custodial wallets, particularly those maintained by an exchange, make them an especially effective tool for fraud.

Breaches and frauds involving bitcoin are not common. But they do occur. Crystal Intelligence (2021) identifies 70 breaches and 20 frauds[13] occurring between January 1, 2011 and December 31, 2021. Of these, 15 breaches and 7 frauds involved other cryptocurrencies, as well. The average value lost in a breach during this period was $22.3 million, while the average value lost in a fraud was $34.9 million. The largest breach occurred at Mt.Gox in 2014, discussed briefly above. The largest fraud occurred in 2019, when the TokenStore

[13]A breach and a fraud occurred on on March 1, 2012. Consequently, there are only 89 unique days on which an incident occurred.

Table 1: Breaches and frauds involving bitcoin, January 1, 2011 - December 31, 2021

	Only Bitcoin	Bitcoin and Other Cryptocurrencies
Breaches		
Number	55	15
Value	$1,060 million	$500 million
Frauds		
Number	13	7
Value	$184 million	$513 million

absconded with bitcoin and other crypto-assets worth roughly $160 million.

Once a breach or fraud is executed, the perpetrator must launder the ill-gotten bitcoin. This typically involves moving the bitcoin through hundreds of wallets, often with the assistance of a mixer, before eventually exchanging it for some national currency. Since a batch of bitcoin transactions are processed roughly every ten minutes, the process of laundering bitcoin takes time. However, it can take *even more* time when blocks are full, unless the perpetrator pays sufficient fees to ensure the transactions are included in the next possible block. Given the additional cost of laundering bitcoin when blocks are full, a profit-maximizing criminal will try to time their breach or fraud so that they are not laundering bitcoin when blocks are full. This insight leads to our first hypothesis.

H1. *Breaches and frauds are less likely to occur when bitcoin blocks are full.*

Next, consider the effect breaches and frauds have on the price of bitcoin. If it is possible that the bitcoin one receives were previously acquired by breach or fraud, there is some risk that those funds will be confiscated by the relevant law enforcement authorities and returned to their rightful owner or effectively frozen by the bitcoin community, which might refuse to accept bitcoin associated with a breach or fraud. Correspondingly, we should expect those receiving bitcoin to discount it in accordance with that risk. As the risk that bitcoin is acquired by breach or fraud increases, the price the marginal buyer will be willing to pay for

bitcoin decreases (and vice versa). This insight, coupled with the belief that breaches and frauds are less likely to occur when blocks are full, leads to our second hypothesis.

H2. The price of bitcoin is positively related to bitcoin block fullness.

We analyze the available data to assess the empirical validity of these two hypotheses below.

2 Empirical Approach and Data

As indicated in **H1**, we expect breaches and frauds are less likely to occur when bitcoin blocks are full. To test this hypothesis, we estimate a logistic regression model that controls for potentially-confounding factors. Specifically, we estimate the general equation:

$$\log\left(\frac{\Pr(Y_t = 1)}{1 - \Pr(Y_t = 1)}\right) = \beta_0 + \beta_1 F_t + \beta_2 P_t + \beta_3 N_t + \varepsilon_t, \tag{1}$$

where $Y_t \in \{\text{Incident}, \text{Breach}, \text{Fraud}\}$, F_t is block fullness, P_t is the log price of bitcoin, N_t is the log network size, and ε_t is an error term, all at time t. We use Huber-White standard errors in all regressions to mitigate heteroskedasticity concerns.

As indicated in equation 1, we estimate the model using three distinct dependent variables. *Incident* is a dummy variable equal to 1 on each day where a breach or fraud involving bitcoin occurred. *Breach* is a dummy variable equal to 1 on each day where a breach involving bitcoin occurred. *Fraud* is a dummy variable equal to 1 on each day where a fraud involving bitcoin occurred. All of the dependent variables are constructed using the breach and fraud data provided by Crystal Intelligence (2021) and discussed above. The first incident occurs on January 7, 2011. The last incident occurs on August 20, 2021.[14] Correspondingly, we

[14]Crystal Intelligence (2021) records incidents through the end of 2021, but none involving bitcoin occur between August 21, 2021 and December 31, 2021.

limit the analysis window to January 1, 2011 to December 31, 2021.

Block *fullness* is the variable of interest. It is measured as the average daily weight of bitcoin blocks divided by the block limit, 4 million weight units. We construct this measure by collecting the entire population of Bitcoin block data from the Bitcoin Core Client software.[15]

The decision to commit a breach or fraud depends, in part, on the value of the bitcoin that might be acquired by breach or fraud. As Luther (2016) explains, the value of bitcoin depends on the non-network and network benefits of bitcoin. We include the daily *price* of bitcoin reported by CoinCodex and proxies for network size, all in log form.[16] Following Liu and Tsyvinski (2021), we employ four proxies for network size:

- total number of unique addresses used on the bitcoin blockchain that have a non-zero balance, which we denote *Unique*;
- total number of transactions, excluding those involving the networks 100 most popular addresses, which we denote *Unpopular*;
- median time for a transaction with miner fees to be included in a mined block and added to the public ledger, which we denote *Confirmed*; and
- total value of all transaction outputs per day, including coins returned to the sender as change, which we denote *Volume*.

All of the proxies for network size are collected from Blockchain.com. Three of these variables—Unique, Unpopular, and Confirmed—are only available once every three days. We interpolate backward and by average. The method of interpolation does not change the results. Average interpolation results are reported.

[15]We used Google Cloud Services BigQuery program to extract the data, which contains the hash code, size, stripped size, weight, number, version, merkle root, timestamp, nonce, bits, and transaction count for each block.

[16]CoinCodex calculates the daily close price as a volume-weighted average of all close prices reported for each market, where volume is measured as the value of all bitcoin trades occurring on a given market during a 24-hour period.

Table 2: Summary Statistics

	Mean	Median	SD	Min	Max	N
Incident	0.0222	0.0000	0.1472	0.0000	1.0000	4018
Breach	0.0174	0.0000	0.1309	0.0000	1.0000	4018
Fraud	0.0050	0.0000	0.0704	0.0000	1.0000	4018
Fullness	0.5540	0.6581	0.3641	0.0010	0.9989	4018
Price	6.5118	6.4900	2.9460	-1.2208	11.1199	4018
Unique	12.1266	12.8669	1.4860	6.7358	13.7935	4015
Unpopular	11.5225	12.1689	1.3540	6.3439	13.0619	4018
Volume	14.1043	14.1072	0.7425	10.9208	18.0255	4018
Confirmed	2.2174	2.1558	0.2953	1.2139	3.3656	3684

As indicated in **H2**, we expect the price of bitcoin is positively related to block fullness. To test this hypothesis, we estimate an ordinary least squares regression that controls for network size. Specifically, we estimate the general equation:

$$P_t = \beta_0 + \beta_1 F_t + \beta_2 N_t + \mu_t, \tag{2}$$

where μ_t is an error term and P_t, F_t, and N_t are measured as described above. As noted previously, we use Huber-White standard errors in all regressions to mitigate heteroskedasticity concerns.

Summary statistics are presented in Table 2. An incident occurs on 2 percent of the days in our sample. The vast majority of these incidents are breaches, whereas instances are fraud are very rare. On an average day, blocks are 55 percent full. The standard deviation of block fullness is 36 percent and, on some days, blocks are completely full.

3 Results

The results of the baseline logistic model used to test **H1** are presented in Table 3. Incident is the dependent variable in columns (1) - (5). Breach is the dependent variable in columns (6) - (10). Fraud is the dependent variable in columns (11) - (15). In general, we find that

Table 3: Logistic Regression Results for Block Fullness on Bitcoin Exchange Incidents

	Incident					Breach					Fraud				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
Fullness	-1.317**	-1.609***	-1.862***	-1.195*	-1.602***	-1.326**	-1.959***	-2.208***	-1.336*	-1.653***	-1.222	-0.596	-0.743	-0.409	-1.282
	(0.529)	(0.586)	(0.594)	(0.658)	(0.554)	(0.611)	(0.690)	(0.694)	(0.769)	(0.631)	(0.989)	(1.116)	(1.156)	(1.143)	(1.099)
Price	0.077	-0.006	-0.025	0.040	0.086	0.074	-0.095	-0.080	0.033	0.085	0.059	0.302**	0.186	0.014	0.060
	(0.054)	(0.083)	(0.073)	(0.080)	(0.054)	(0.064)	(0.097)	(0.085)	(0.094)	(0.063)	(0.093)	(0.131)	(0.117)	(0.141)	(0.095)
Unique		0.235					0.490**					-0.620***			
		(0.172)					(0.209)					(0.204)			
Unpopular			0.385**					0.606***					-0.401		
			(0.176)					(0.209)					(0.285)		
Confirmed				0.255					0.414					-0.597	
				(0.507)					(0.538)					(1.332)	
Volume					0.269**					0.307**					0.060
					(0.130)					(0.148)					(0.217)
Pseudo R-squared	0.009	0.011	0.014	0.011	0.013	0.010	0.014	0.019	0.015	0.014	0.007	0.015	0.011	0.004	0.007
Observations	4018	4015	4018	3684	4018	4018	4015	4018	3684	4018	4018	4015	4018	3684	4018

Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

fuller blocks are associated with a smaller likelihood of an incident, breach, and fraud. The coefficient on block fullness is consistently negative. It is statistically significant for regressions in which incidents or breaches are considered. The lack of statistical significance with respect to fraud might be explained by their rare occurrence: there are only 20 incidences of fraud in the entire sample. Hence, standard errors are large when fraud serves as the dependent variable.

To get a sense of the magnitude of theses estimates, consider the marginal effects from the estimation of equation (1) in column (2) of Table 3 (Petersen 1985, Hoetker 2007). The predicted probability that an incident occurs on a day where the average block is 55 percent full is 2.2 percent. A one standard deviation increase in block fullness from the mean (i.e., when the block is 91 percent full) reduces the predicted probability that an incident occurs by 41 percent, from 2.2 percent to 1.3 percent. A one standard deviation decrease in block fullness from the mean (i.e., when the block is 19 percent full) increases the predicted probability that an incident occurs by 77 percent, from 2.2 percent to 3.9 percent. Hence, the magnitude of the effect is far from trivial.

The results of the baseline OLS model used to test **H2** are presented in Table 4. In general, we find that fuller blocks are associated with a higher price. The coefficient on block fullness is consistently positive and statistically significant. However, the magnitude

Table 4: OLS Regression Results for Block Fullness on Price

	(1)	(2)	(3)	(4)	(5)
Fullness	6.901***	0.862***	2.398***	6.573***	6.986***
	(0.066)	(0.099)	(0.103)	(0.056)	(0.057)
Unique		1.661***			
		(0.025)			
Unpopular			1.384***		
			(0.022)		
Confirmed				-2.290***	
				(0.070)	
Volume					-0.095**
					(0.047)
Constant	2.688***	-14.109***	-10.769***	8.114***	3.979***
	(0.047)	(0.253)	(0.199)	(0.156)	(0.666)
R-squared	0.728	0.874	0.823	0.740	0.728
Observations	4018	4015	4018	3684	4018

Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

of the coefficient on block fullness varies considerably—from 0.862 to 6.986—depending on the proxy for network size employed.

4 Instrumental Variables Analysis

Our baseline model is subject to at least three criticisms. First, it is possible that block fullness is related to some omitted variable, such as prevailing cybersecurity, that affects criminal incidents. Second, it is possible that we have the causation reversed (i.e. reverse causality or simultaniety). For example, it could be the case that breaches and scams lead to emptier blocks. Third, the measure of block fullness is potentially measured with error, leading to attenuation bias. To deal with these issues, we exploit the activation of SegWit in August 2017 to conduct an instrumental variables analysis.

As noted above, the SegWit upgrade segregated the witness data from the sender and

receiver data in a block and replaced the 1 MB block size limit with a 4 million weight unit limit. To make SegWit backwards compatible, the weight of all blocks was set equal to three times the stripped size of the block plus the total size of the block. Prior to activation of SegWit, both witness data and sender and receiver data contributed to the stripped size. Hence, block weight is equal to four times block size prior to block number 478,559, when SegWit activated. Beginning with block number 478,559, only sender and receiver data contributed to stripped size. Hence, SegWit reduced the contribution of witness data to the block limit.

The immediate effect of the SegWit upgrade on block fullness can be seen in Figure 1. Recall that block fullness is measured as the average daily weight of bitcoin blocks divided by the block limit, 4 million weight units. For comparison, we also include a legacy measure of block fullness, which is constructed as if the witness data had not been removed from the stripped size. Prior to the activation of SegWit, the two series are identical. After the activation of SegWit, block fullness falls relative to our legacy measure of block fullness. Indeed, the legacy measure typically exceeds 1 following the introduction of SegWit, suggesting a counterfactual where blocks are *more than* completely full. That would not technically be possible. Instead, the failure to activate SegWit would have likely resulted in completely full blocks and a long queue of transactions waiting in the mempool. By reducing the contribution of witness data to the block limit, SegWit permitted more transactions to be processed per day and, correspondingly, reduced block fullness.

We use the activation of SegWit as an instrumental variable. A valid instrument must influence block fullness exogenously and therefore not be related to the error term (mitigating omitted variable bias) or dependent variable (mitigating simultaneity) through any other channel other than the variable of interest, including (and therefore mitigating) measurement error. In order to serve as a valid instrument, SegWit must satisfy two conditions: relevance and exclusion(Roberts and Whited 2013). The relevance condition requires that the partial correlation between the instrument (SegWit) and the endogenous variable (block fullness)

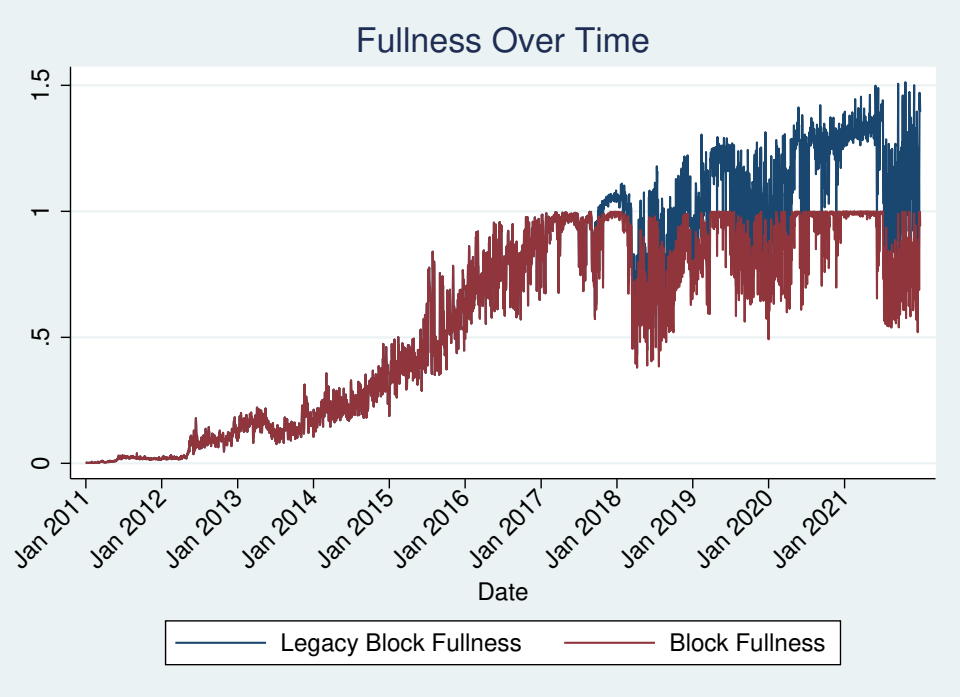

Figure 1: Block Fullness and Segregated Witness Data

not be equal to zero. The exclusion condition requires that the covariance between the instrument and the error term is equal to zero.

To test for relevance, we estimate the following first-stage regression:

$$F_t = \alpha_0 + \alpha_1 S_t + \alpha_2 P_t + \alpha_3 N_t + v_t, \tag{3}$$

where S_t is a dummy variable equal to 0 prior to the activation of SegWit and 1 otherwise and v_t is an error term. The predicted sign, significance, and a model F-statistic greater than 10 should ameliorate concerns of weak instrument bias.

The exclusion restriction is violated when the instrument is correlated with both the causal variable of interest and other determinants of the dependent variable (Angrist and Pischke 2009). The exclusion condition cannot be tested because ε_t is unobservable. However, it is reasonable to believe the activation of SegWit satisfies the exclusion restriction. The activation of SegWit exogenously altered the capacity of the bitcoin protocol. It did not impact any other determinants of breaches and fraud, such as hacking technology or

Table 5: First Stage Results

	(1)	(2)	(3)	(4)	(5)
SegWit	-0.139***	-0.124***	-0.137***	-0.124***	-0.120***
	(0.021)	(0.021)	(0.021)	(0.022)	(0.021)
Price	0.200***	0.120***	0.158***	0.167***	0.251***
	(0.031)	(0.036)	(0.033)	(0.033)	(0.035)
Unique		0.225***			
		(0.059)			
Unpopular			0.148**		
			(0.059)		
Confirmed				0.083***	
				(0.028)	
Volume					0.158***
					(0.050)
Observations	121	121	121	121	121
R-squared	0.25	0.31	0.29	0.29	0.32
F-statistic	25.64	21.58	17.92	18.36	19.91

Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

exchange-related procedures. Hence, SegWit is not related to incidents of breaches or fraud other than through its effect on block fullness.

We estimate our two-stage least squares model on the sample 60 days before and after the activation of SegWit. Limiting the analysis to a tight window around the exogenous shock reduces the chance that confounding factors are driving the results. The smaller sample size from the narrow window reduces statistical power, increasing the risk of type 2 error (i.e., false negative). It does not lead to an increased risk of type 1 error (i.e., false positive). Therefore, in this context, marginally significant results can be compelling.

The results of our first stage regression are presented in Table 5. The coefficient on the SegWit dummy is statistically significant and negative in the first stage regression, consistent with our contention that SegWit reduced block fullness by increasing the capacity of the bitcoin protocol to process transactions. The estimated coefficient ranges from -13.9 percent

Table 6: Instrumental Variable Regression of Block Fullness on Incidents

	(1)	(2)	(3)	(4)	(5)
Fullness	-0.565	-0.456	-0.575*	-0.510	-0.739*
	(0.344)	(0.346)	(0.349)	(0.335)	(0.409)
Price	0.008	0.070	-0.021	0.015	0.108
	(0.048)	(0.086)	(0.058)	(0.045)	(0.069)
Unique		-0.233			
		(0.239)			
Unpopular			0.109		
			(0.091)		
Confirmed				-0.043	
				(0.061)	
Volume					0.202
					(0.153)
Observations	121	121	121	121	121

Robust standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

to -12.0 percent across regressions. Hence, the magnitude of the effect of SegWit on block fullness is meaningful. Moreover, concerns about weak instrument bias are ameliorated by the high F-statistics, all of which exceed 10.

The results of the second stage regression are shown in Table 6. The estimated coefficient on block fullness is consistently negative, which is consistent with our hypothesis. However, the coefficient is only statistically significant in two of the five specifications. Given the low power of the test, evidence of statistical significance is strong evidence in favor of our hypothesis.

5 Conclusion

The vast sums of wealth held on centralized exchanges make them an attractive target for cybercriminals, who might gain unauthorized access to an exchange's wallet (breach) or pose

as a legitimate exchange (fraud) to steal customer funds. To profit from an attack, however, a cybercriminal must launder the ill-gotten gains quickly and at a sufficiently low cost. Network congestion makes this difficult. Just as robbing a bank during rush hour reduces the odds of a speedy getaway, network congestion undermines a cybercriminal's ability to abscond with stolen bitcoin.

We have explained how block fullness—that is, the weight of bitcoin blocks relative to the block limit—reduces transaction times and, hence, should discourage incidents of breaches and frauds involving bitcoin. Moreover, since those accepting bitcoin when blocks are full can be more confident that the bitcoin received was not obtained by breach or fraud, the price of bitcoin should tend to rise and fall with block fullness.

We have also offered some evidence for our view. In our baseline model, we find that a one standard deviation increase in block fullness reduces the predicted probability of an incident of breach or fraud by roughly 41 percent. Similarly, a one standard deviation decrease in block fullness increases the predicted probability of an incident by roughly 77 percent. Furthermore, changes in block fullness coincide with price changes, as predicted.

Finally, we have exploited the activation of SegWit to show that the aforementioned relationship between block fullness and breach or fraud incidents is causal. Specifically, we have conducted an instrumental variable analysis using SegWit as an instrument. Despite the low statistical power of our test, we find evidence that an increase in block fullness reduced incidents of breaches or frauds.

References

Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, January 2009. ISBN 978-0-691-12035-5. Google-Books-ID: YSAzEAAAQBAJ.

Jonathan Bier. *The Blocksize War: The battle over who controls Bitcoin's protocol rules*. Independently published, March 2021. ISBN 9798721895609.

Ben Charoenwong and Mario Bernardi. A Decade of Cryptocurrency 'Hacks': 2011 – 2021, October 2021. URL https://papers.ssrn.com/abstract=3944435.

Lin William Cong, Xi Li, Ke Tang, and Yang Yang. Crypto Wash Trading. *Management Science*, 69(11):6427–6454, November 2023. ISSN 0025-1909. doi: 10.1287/mnsc.2021.02709. URL https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2021.02709. Publisher: INFORMS.

Anirudh Dhawan and Tālis J. Putniņš. A new wolf in town? Pump-and-dump manipulation in cryptocurrency markets. *Review of Finance*, 27(3):935–975, 2023. URL https://academic.oup.com/rof/article-pdf/27/3/935/50306667/rfac051.pdf. Publisher: Oxford University Press.

Rowel Gundlach, Martijn Gijsbers, David Koops, and Jacques Resing. Predicting confirmation times of Bitcoin transactions. *ACM SIGMETRICS Performance Evaluation Review*, 48(4):16–19, May 2021. ISSN 0163-5999. doi: 10.1145/3466826.3466833. URL https://dl.acm.org/doi/10.1145/3466826.3466833.

Peter K. Hazlett and William J. Luther. Is bitcoin money? And what that means. *The Quarterly Review of Economics and Finance*, 77:144–149, 2020. URL https://www.sciencedirect.com/science/article/pii/S1062976919301528. Publisher: Elsevier.

Joshua R. Hendrickson and William J. Luther. Banning bitcoin. *Journal of Economic Behavior & Organization*, 141:188–195, 2017. URL https://www.sciencedirect.com/science/article/pii/S0167268117301798. Publisher: Elsevier.

Joshua R. Hendrickson and William J. Luther. Cash, crime, and cryptocurrencies. *The Quarterly Review of Economics and Finance*, 85:200–207, August 2022. ISSN 1062-9769. doi: 10.1016/j.qref.2021.01.004. URL https://www.sciencedirect.com/science/article/pii/S1062976921000041.

Joshua R. Hendrickson, Thomas L. Hogan, and William J. Luther. THE POLITICAL ECONOMY OF BITCOIN. *Economic Inquiry*, 54(2):925–939, April 2016. ISSN 0095-2583, 1465-7295. doi: 10.1111/ecin.12291. URL https://onlinelibrary.wiley.com/doi/10.1111/ecin.12291.

Glenn Hoetker. The use of logit and probit models in strategic management research: Critical issues. *Strategic Management Journal*, 28(4):331–343, April 2007. ISSN 01432095,

10970266. doi: 10.1002/smj.582. URL https://onlinelibrary.wiley.com/doi/10.1002/smj.582.

Crystal Intelligence. Crystal Blockchain Report:, December 2021. URL https://crystalintelligence.com/rohirov/2024/06/Adolescent-Anarchy-Thirteen-Years-of-Crypto-Crimes-Unveiled-2024.pdf.

David Koops. Predicting the confirmation time of Bitcoin transactions, September 2018. URL http://arxiv.org/abs/1809.10596. arXiv:1809.10596 [cs, math].

Yukun Liu and Aleh Tsyvinski. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6):2689–2727, 2021. Publisher: Oxford University Press.

William J. Luther. Cryptocurrencies, Network Effects, and Switching Costs. *Contemporary Economic Policy*, 34(3):553–571, 2016. ISSN 1465-7287. doi: 10.1111/coep.12151. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/coep.12151. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/coep.12151.

William J. Luther. Is bitcoin intrinsically worthless? *AIER Sound Money Project Working Paper*, (2018-07), 2018. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3000068.

William J. Luther. Getting off the ground: the case of bitcoin. *Journal of Institutional Economics*, 15(2):189–205, 2019. URL https://www.cambridge.org/core/journals/journal-of-institutional-economics/article/getting-off-the-ground-the-case-of-bitcoin/08F266520BB3C5FDB1C346681550FF1C. Publisher: Cambridge University Press.

William J. Luther and Sean Stein Smith. Is Bitcoin a decentralized payment mechanism? *Journal of Institutional Economics*, 16(4):433–444, 2020. URL https://www.cambridge.org/core/journals/journal-of-institutional-economics/article/is-bitcoin-a-decentralized-payment-mechanism/16CA6A2E4440BB4E715E3E745A34C355. Publisher: Cambridge University Press.

Robert McMillan. The Inside Story of Mt. Gox, Bitcoin's $460 Million Disaster. *Wired*, 2014. ISSN 1059-1028. URL https://www.wired.com/2014/03/bitcoin-exchange/. Section: tags.

Tyler Moore and Nicolas Christin. Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk. In Ahmad-Reza Sadeghi, editor, *Financial Cryptography and Data Security*, pages 25–33, Berlin, Heidelberg, 2013. Springer. ISBN 978-3-642-39884-1. doi: 10.1007/978-3-642-39884-1_3.

Khanh Quoc Nguyen and Tālis J. Putniņš. Stealing Digital Assets: Evidence From Cryptocurrency Exchange Hacks, July 2023. URL https://papers.ssrn.com/abstract=4525029.

OSC. QuadrigaCX: A Review by Staff of the Ontario Securities Commission, 2020. URL https://www.osc.ca/quadrigacxreport/.

Trond Petersen. A Comment on Presenting Results from Logit and Probit Models. *American Sociological Review*, 50(1):130–131, 1985. ISSN 0003-1224. doi: 10.2307/2095348. URL https://www.jstor.org/stable/2095348. Publisher: [American Sociological Association, Sage Publications, Inc.].

Michael R. Roberts and Toni M. Whited. Endogeneity in Empirical Corporate Finance1. In George M. Constantinides, Milton Harris, and Rene M. Stulz, editors, *Handbook of the Economics of Finance*, volume 2, pages 493–572. Elsevier, January 2013. doi: 10.1016/B978-0-44-453594-8.00007-0. URL https://www.sciencedirect.com/science/article/pii/B9780444535948000070.

George Selgin. Synthetic commodity money. *Journal of Financial Stability*, 17:92–99, 2015. URL https://www.sciencedirect.com/science/article/pii/S1572308914000722. Publisher: Elsevier.

Camomile Shumba. The amount of bitcoin held on exchanges is at its lowest in 6 months, while ether reserves are at record lows, indicating any sell-offs could be a lot more modest. URL https://markets.businessinsider.com/news/cryptocurrencies/crypto-price-outlook-bitcoin-ether-supply-exchanges-at-lows-2021-7.

Ivo Stoepker, Rowel Gundlach, and Stella Kapodistria. Robustness analysis of Bitcoin confirmation times. *ACM SIGMETRICS Performance Evaluation Review*, 48(4):20–23, May 2021. ISSN 0163-5999. doi: 10.1145/3466826.3466834. URL https://dl.acm.org/doi/10.1145/3466826.3466834.

Vasek. Vasek: There's no free lunch, even using Bitcoin:... - Google Scholar, 2015. URL https://scholar.google.com/scholar?cluster=886669128961772690&hl=en&as_sdt=0,10.

Pengcheng et al. Xia. Characterizing cryptocurrency exchange scams. *Computers & Security*, 98:101993, November 2020. ISSN 0167-4048. doi: 10.1016/j.cose.2020.101993. URL https://www.sciencedirect.com/science/article/pii/S0167404820302662. Publisher: Elsevier Advanced Technology.

Bitcoin Policy Institute

www.btcpolicy.org