Title

Scalability Controversy: Understanding Past Cryptocurrency Returns through Segregated Witness

Author

Levine, Jeremy

Publication Date

2019

Undergraduate

Scalability Controversy:
Understanding Past Cryptocurrency Returns
through Segregated Witness

Jeremy Levine
March 2019

Department of Economics
University of California, Santa Barbara
jeremylevine@ucsb.edu

Advisor: Travis Cyronek

Abstract

In an attempt to increase transaction speeds, numerous cryptocurrencies installed the software update Segregated Witness (SegWit), causing significant controversy among cryptocurrency users. With this thesis, I use daily data for four cryptocurrencies—Bitcoin, Litecoin, Vertcoin, and Digibyte—to investigate the impact of the introduction of SegWit on a coin's predicted returns. While SegWit does not appear to directly impact returns, fluctuations in trading volume appears to have a significant effect. My model predicts that, upon SegWit's planned introduction to a coin, the impact of fluctuations in trading volume growth on expected returns doubles or triples, depending on the coin. Using OLS regression with lagged controls, this paper explores how technological factors inherent to cryptocurrencies impact them as a financial asset for the first time.

Introduction

Bitcoin supporters initially predicted the new digital currency would rapidly revolutionize international finance, simplify day-to-day transactions, and free humans from the chains of the nation-state and nationally issued currencies.[1] Excitement around Bitcoin derived from the blockchain technology on which the digital "coin" is based, a decentralized network formed by several computers that maintains a record of all transactions permanently and untamperably. As enthusiasm for Bitcoin grew after it was first theorized in 2008 and created in 2009, thousands of other cryptocurrencies sprouted up, some of them serious projects, a few of them practical jokes, and many of them outright scams. After languishing in relative obscurity for most of a decade, Bitcoin—and the entire cryptocurrency market—underwent an enormous boom; the price of one Bitcoin skyrocketed from less than $1,000 to over $20,000 over the course of 2017 before crashing down to approximately $10,000 through January 2018 and gradually declining since. Cryptocurrency fanatics and technologists argue the market simply needs to develop, as the blockchain technology improves and awareness of the benefits of crypto spread.

Yet cryptocurrencies face several major technological hurdles to becoming the panacea many proponents believe possible: Most significantly, the blockchain networks on which most cryptocurrencies rely have a scalability issue, operating too slowly to be useful for international finance or trade. Originally, Bitcoin's network could handle seven transactions per second at maximum capacity. For comparison, consider that credit card company networks typically run

[1] It's worth reading this fun list of material items cryptocurrency can buy: https://cointelegraph.com/news/lambos-bling-and-mansions-what-purchases-do-crypto-millionaire-make

in excess of 25,000 transactions a second, and their networks are capable of handling much more at peak volume. Of the hundreds of digital coins created after Bitcoin, many included significant technological improvements. However, no base code for a new cryptocurrency sufficiently addresses the scalability issue to allow for international finance or commercial usage, since none can handle more than a few hundred transactions per second.

In 2015, Bitcoin developer Pieter Wuille introduced the concept of Segregated Witness, or SegWit for short, a step towards increasing a blockchain network's ability to handle higher transaction volumes by changing the way the network stores information. Supporters of SegWit pushed for the software update through a soft fork in the Bitcoin code, a process by which a software update is introduced that users can opt into. However, SegWit is only one of several possible ways to improve a blockchain network's transaction processing speed; other methods still in development as of 2017, such as increasing the overall block size to let the blockchain store more information, do not mesh well with SegWit, and so to pursue SegWit meant to forego alternative methods of increasing network speed. The cryptocurrency community fractured over the introduction of SegWit; even before the software update was officially locked in to activate for Bitcoin on August 8, 2017, a group of users hard forked, splitting off from Bitcoin's base code into a new blockchain, on August 1, 2018 to create a new currency, Bitcoin Cash. Another group of Bitcoin users followed on January 1, 2018, creating Bitcoin Gold.

Both Bitcoin Gold and Bitcoin Cash operate completely independently of Bitcoin, despite their names; still, they represent community disenchantment with the SegWit update's installation on Bitcoin. Significantly, several cryptocurrencies installed SegWit over the course of 2017, most of them receiving dramatically less press coverage than Bitcoin's adoption of the software. The controversy around Bitcoin's SegWit introduction caused me wonder: How do controversial software updates, namely SegWit in this case, affect average daily returns in cryptocurrencies? And do users across different coins respond differently to SegWit? I compare the effects of SegWit's locking in—the date a crypto network votes that it will activate a soft fork, which typically occurs two weeks after the vote—on dollar returns across four different coins. My research will help software developers understand the net impact of using SegWit on cryptocurrency user confidence. Even though SegWit improves a coin's network speed, I hypothesized the controversy the update created had such a negative effect on a cryptocurrency's growth upon SegWit locking in that cryptocurrencies that installed the update underwent a decrease in overall user confidence, and average returns temporarily fell. I also suspected that SegWit's effect will be fairly consistent across coins, since the cryptocurrency community is fairly insular, with many crypto fanatics investing in dozens or hundreds of coins.

While looking at four years of daily data did not produce evidence to conclusively test my hypothesis over SegWit's direct effects on returns, my data analysis suggests there is an unexpected indirect relationship between SegWit and predicted returns. According to my regression results, a 100 percentage point fluctuation in trade volume—rather than the introduction of SegWit—is associated with a statistically significant effect on returns between 1.5% to 2% across coins analyzed; upon SegWit's introduction, that predicted effect jumps to an average of 4.5% across coins. No other factors besides trade volume and the interaction between trade volume and SegWit had statistically significant correlations with returns; this contradicts previous papers, which found that a coin's own lagged returns may have predictive power. My paper adds to the economic literature on cryptocurrencies by exploring how technological factors unique to cryptocurrencies affect their operation as a financial asset as well as impact cryptocurrency user behavior.

Data

I tried to address these questions looking at daily coin data over the time period from June 1, 2014 to March 31, 2018 using data from coinmetrics, an aggregator of open source datasets for cryptocurrencies, which the development team updates daily.[2] The coinmetrics team receives much of their data from coinmarketcap.com, the standard data source used for economic research on cryptocurrencies. Information such as circulating supply, total supply, number of transactions, and fees all comes from coinmarketcap. Coinmetrics independently aggregates information such as coins generated and exchange volume (or the amount traded solely on exchanges) from the blockchain ledger of each coin.

Daily closes for price occur at 7 p.m. EST, and the daily price listed is the opening price the next day. Trade volume is the total value of transactions measured in US Dollars on a given blockchain. This variable is independently unreliable for a variety of reasons; the coinmetrics team generates an adjusted trade volume variable for each coin, which offers a better estimate of real trade volume, with all known instances of double counting and internal-crypto-exchange coin transfers removed from the estimate of trade volume. Essentially representing the nominal inflation rate, generated coins describes the number of new coins generated by miners on a given blockchain. Unlike most sources for cryptocurrency data, which use theoretical estimates of coins generated based on the cryptocurrencies' algorithms, coinmetrics aggregates information on actual coins generated from all known existing trading pools. Representing a

[2] See the coinmetrics team's own explanation of their data: https://coinmetrics.io/on-data-and-certainty/

small transaction cost standard to most cryptocurrencies, fees fluctuate according to an algorithm for the coins I look at. Coinmetrics provides numbers for median daily fee, found by aggregating all fees paid in a day and taking the median. Trading exchanges that bypass fees are excluded from the fees data because individuals can trade with themselves if there is no fee, increasing trading volume artificially.

No source of information I have mentioned so far includes data on the introduction of SegWit. That information I had the joy of finding for myself by crawling through internet forums and news articles. I searched for three sets of data: the date of SegWit's introduction to be voted on in a blockchain network, the date of SegWit's approval by consensus model on a blockchain network, and the date of SegWit's actual installation on a blockchain network. I successfully found data for four cryptocurrencies regarding the date of SegWit's approval and activation on the network; information on the introduction of voting for SegWit to each cryptocurrency I analyze remains elusive. I ultimately only used the date for SegWit's approval in my regressions, since the market should respond to the information upon that announcement. From scrolling the internet, I found four coins that installed SegWit throughout 2017 and have reliable data available: Bitcoin (BTC), Litecoin (LTC), Vertcoin (VTC), and Digibyte (DGB).

Each of these coins have been in existence since at least January 2014, with an active user base developed by at least June of 2014. All four coins use a proof of work system, by which users on a cryptocurrency's network (i.e. miners) devote computing power to solve algorithms that simultaneously verify transactions and generate new coins. These miners also determine the rules of the cryptocurrency network, such as which software updates are installed, depending on the proportion of total computing power (i.e. hash power) they contribute to a network, in a system called a consensus model.

The four coins I analyze operate using different algorithms, designed to influence the dynamic of the consensus model. Miners on Bitcoin and Litecoin can use expensive specialized Bitcoin mining computers (see glossary: ASIC), allowing for one miner or a team of miners to conceivably take over the network by contributing a majority of computing power. Vertcoin and Digibyte, on the other hand, only allow mining with certain equipment (see glossary: graphics cards), which is intended to ensure no individual or oligopoly dominates the consensus model to rewrite the rules in their own favor. Beyond their variations in valid mining equipment, the coins have naturally different networks speeds. Upon its creation, Bitcoin's network could

Table 1: Price Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Max
Bitcoin	1,457	2,275.441	3,633.646	176.900	19,475.800
Litecoin	1,457	32.395	61.304	1.150	359.130
Vertcoin	1,457	0.757	1.637	0.006	9.460
Digibyte	1,457	0.008	0.016	0.00003	0.127

Table 2: Returns Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Max
Bitcoin	1,457	4.626	268.309	−2,405.000	3,536.800
Litecoin	1,457	0.073	5.898	−49.990	102.900
Vertcoin	1,457	0.001	0.181	−2.500	1.340
Digibyte	1,457	0.00002	0.002	−0.027	0.028

handle 7 transactions per second, making it among the slowest coins in existence; Digibyte's could handle several hundred.

As the differences in price reveal from the summary statistics table, the coins analyzed here have widely varying natural prices. This is both because the coins vary in popularity and also because they have naturally varying total supplies. Bitcoin's algorithm will produce a maximum of 21 million coins; Digibyte's will create 84 billion. So even if Digibyte were to be as popular as Bitcoin, the price per coin would be lower. Differences in supply do not account for all of the differences in price, however. Bitcoin and Litecoin are historically the most popular coins, whereas Digibyte and Vertcoin each went through periods of being among the top five most highly valued coins, each at one point with a market valuation of over $1 billion, but have since faded into realtive obscurity. Over the course of my study period, Bitcoin and Litecoin generate positive daily returns on average, whereas Vertcoin and Digibytes' returns indicate that an investor would have approximately broken even over my study period.

SegWit locked in for Bitcoin on April 4, 2017, for Litecoin on May 5, 2017, for Vertcoin on May 7, 2017, and for Digibyte on May 10, 2017. As can be seen from the graphs on the following page, SegWit's lock-in occurred directly before a huge boom in price growth occurred across the cryptocurrency market in the latter half of 2017. However, readers should keep in mind that this paper analyzes returns rather than price, so the visible price spike that occurs across coins shortly after SegWit's introduction was certainly caused by other factors. Percentage returns, on the other hand, could be equally high at an invisibly small blip in price during 2014 as it could during the price spike in 2017. See appendix D for graphs of returns over time.

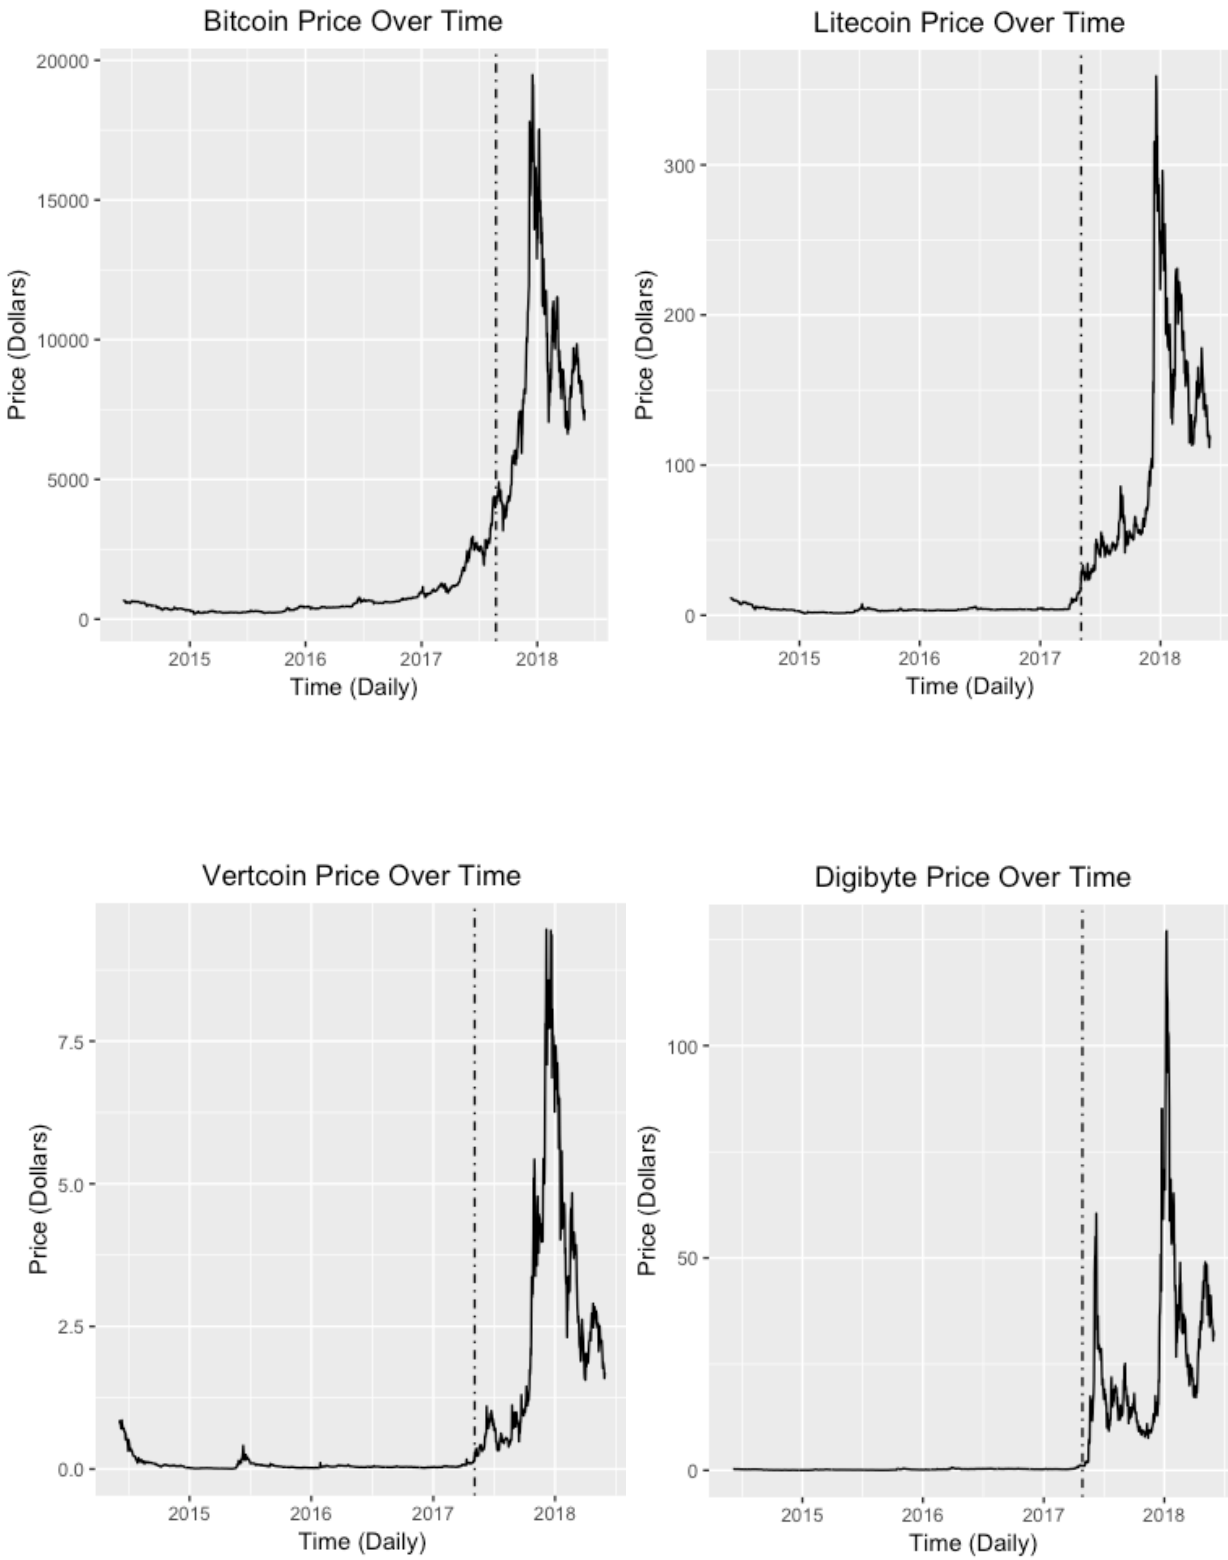

Price Tables: SegWit updates occurred at similar times for all coins.

Literature Review

Research on cryptocurrencies connects closely to traditional financial research, although key differences exist. Generally, papers studying cryptocurrencies (including this one) present statistical models informed by findings from canonical Fama papers, which argue that regressing on returns rather than price estimates the value of a stock more accurately. Fama's efficient market hypothesis further theorizes the market responds to new information, influencing my decision to use the date SegWit was voted to be installed for each cryptocurrency rather than the date it was activated—there is typically a two week gap between the two. However, the academic literature on cryptocurrencies diverges significantly from that of traditional financial assets in other ways.

In fact, a 2018 National Bureau of Economic Research working paper finds that cryptocurrencies are virtually unexposed to traditional stock market factors as well as fluctuations in the value of stocks and commodities (Liu and Tsyvinski 2018). The authors found essentially no causality between cryptocurrency returns and gold, silver, and oil commodity prices; 155 common factors found to predict cross-sections of stock returns in finance literature (see summary of known factors in Feng et al 2017); and profit returns in over 400 industries across the U.S. and China. Only investing momentum—i.e. past coin returns—had reliable predictive power for future returns, and google trends data had uncertain predictive power.

Liu and Tsyvinski's findings confirm those by the other researchers, who increasingly try to predict crypto asset prices by analyzing datasets of online commentary through a process called textual sentiment analysis (TSA), which uses key words to determine general trends from the content of large datasets of individual online. Economists have studied cryptocurrency using TSA on Twitter (Kaminski 2014), various cryptocurrency forums (Kim et all 2016), and Google Trends search data (Matt et all 2017) in order to predict cryptocurrency prices. All studies that have used emotional sentiment found online through TSA to predict cryptocurrency prices conclude that online sentiment, although strongly correlated with prices, reflects responses to the market rather than helps predict prices.

Other papers trying to understand the forces influencing cryptocurrency prices and returns find that cryptocurrency markets are subject to manipulation that is currently illegal and heavily regulated in traditional financial markets, and that they follow unpredictable patterns. Pump-and-dump schemes, by which groups of individuals coordinate trading frenzies which rapidly inflate a coin's price before causing a crash with mass sell-offs, occurred at least 500 times (the author's sample size over the period from May 15, 2017 to August 26, 2018), affecting dozens of coins (Li et all 2018). Other researchers have found some evidence that price booms and busts occur even without bubbles forming—i.e. booms and busts are a natural part of cryptocurrency price fluctuations—as well as that cryptocurrencies are liable to face arbitrary price collapse (Fry 2018).

Some of the earliest theoretical models of cryptocurrency prices hypothesize exactly the patterns previously observed in data. One model presented in a different NBER working paper implies that Bitcoin prices are likely to undergo submartingales or supermartingales, depending on preexisting conditions (Schiller and Uhlig 2018). Key variables Schiller and Uhlig use for their theoretical model include trade volume and fees, a small price paid for each cryptocurrency transaction. This model can be generalized to the majority of cryptocurrencies, which run on an inflationary model somewhat similar to Bitcoin's.

Studying the modern cryptocurrency market requires acknowledging Bitcoin's dominance. With over 50% of cryptocurrency market value typically held in Bitcoin, the cryptocurrency reaps the benefits of being the first successful entrant into a new market, an advantage prior literature suggests implies that the entire market lags Bitcoin (Glazer 1985, Lieberman 1988). Despite Bitcoin's technological limitations, enthusiasm for the currency remains high, and the entire cryptocurrency market tends to track Bitcoin. Nonetheless, theoretical papers about information acquisition, which argue that freely offered new information about a new market induces entry by multiple actors, hold implications for Bitcoin and other cryptocurrencies: Initial participants in the cryptocurrency industry usually start investing in Bitcoin—the currency of exchange in the cryptocurrency world—before potentially branching out to new coins (Banerjee 2018). Furthermore, Banerjee's paper supports the hypothesis that freely endowed new information regarding Bitcoin, such as information about software updates like SegWit, may gradually induce abandonment of Bitcoin into new cryptocurrencies.

The boom in new theoretical models and data-drive studies regarding cryptocurrencies, while useful for establishing how cryptocurrencies differ from traditional financial assets, has not yet investigated how technological elements inherent to cryptocurrencies specifically affect their operation as an asset class. I construct a model to estimate SegWit's impact on returns, incorporating lessons from both theory and data analysis to justify my choices of predictors.

Model

I perform various transformations on my data before approaching my analysis. I take log(price) and then difference at one time period in order to get percent returns (returns_t =

$\log(\text{price}_t) - \log(\text{price}_{t-1})$), so I can easily compare the effect of my coefficients across coins. I take the log of trade volume as well as coins generated to further improve interpretability, and difference both variables at one time period to capture the impact of changes in trade volume and coins generated—necessary for model stationarity—on returns. I leave median fee undifferenced because fees are fairly static, and so change in fees likely would have no correlation with returns. Since I am estimating realized returns, I use robust standard errors to control for the inevitable heteroskedasticity present in timeseries data (Merton 1980).

Most importantly for my experiment, I generate a dummy variable for SegWit: 1 after SegWit has locked in to activate for a cryptocurrency and 0 beforehand. I then ran the following regressions:

$$(1) \qquad \text{Ret}_t = \boldsymbol{\beta}_0 + \boldsymbol{\beta}_1\text{Segwit}_t + \boldsymbol{\beta}_2\boldsymbol{\Delta}(\log(\text{TV}_t)) + \boldsymbol{\beta}_3\text{medianFee}_t + \boldsymbol{\beta}_4\boldsymbol{\Delta}(\log(\text{CG}_t)) \ldots$$

Ret_t refers to returns, Segwit_t refers to my dummy for the SegWit update locking in, TV_t refers to daily trade volume, medianFee refers to the median fee paid in a day to make a transaction with a cryptocurrency, and CG_t refers to coins generated in a day on a blockchain. In this regression, I try to isolate the effect of SegWit, using fluctuations in trade volume as a proxy for general popularity of a cryptocurrency; Schiller and Uhlig's theoretical model of bitcoin used trade volume as a key variable. Median fee and the rate of change of coins generated intuitively fit in the model, since fees represent an economic burden to use cryptocurrencies, and coins generated is essentially inflation.

$$(2) \qquad (1) \ + \boldsymbol{\beta}_5\text{Segwit}_t * \boldsymbol{\Delta}(\log(\text{TV}_t)) \ldots$$

I then add an intuitive interaction term, since by technological definition SegWit should decrease financial frictions between changes in trading activity and returns. At this point, my model derives legitimacy primarily from theoretical papers about Bitcoin. Fry also uses trading volume as a predictor in his analysis.

(3) (2) $+ \boldsymbol{\beta}_6 Ret_{t-1} + \boldsymbol{\beta}_7 Ret_{t-2} + \boldsymbol{\beta}_8 Ret_{t-3} \ldots$.

Next, I add regressors for a coin's own lagged returns at one, two, and three time periods (lagged at one, two, and three days), which Liu found to be the only variables with any significant predictive power on Bitcoin and Litecoin. Although Liu finds statistically significant autocorrelation at lags larger than three, such coefficients are typically statistically significant noise, and so I only add lagged returns from the past three periods.

(4) (3) $+ \boldsymbol{\beta}_9 BTCRet_{t-1} + \boldsymbol{\beta}_{10} BTCRet_{t-2} + \boldsymbol{\beta}_{11} BTCRet_{t-3} + \boldsymbol{\varepsilon}_t$

Lastly, I added Bitcoin's lagged returns at one, two, and three days to my regressions for Litecoin, Vertcoin, and Digibyte. Referring to my price graphs from the data section, the entire cryptocurrency market appears to track Bitcoin. Furthermore, with an expected first mover advantage derived from Glazer's paper, Bitcoin should theoretically lead the cryptocurrency market, at least over my study period.

Results

Running my baseline regression across coins (see table 3 on following page), I find that Segwit's lock-in on a cryptocurrency is associated with a non-significant 0.4 percentage point increase in expected daily returns. Only the change in trade volume shows statistically significant coefficients across the table, with a 100 percent point increase in the rate of growth of trade volume being associated with a 2.6 percent point change in expected returns for Bitcoin, a 2.3 percentage point change for Litecoin, a 1.8 percentage point change for Vertcoin, and a 2.0 percentage point change for Digibyte, all significant at the 99% confidence level. For example, this model suggests a 100 percentage point increase in trading volume of bitcoin, from $1.5 billion to $3 billion, would be associated with an average returns of Bitcoin to increase by $39 million across the entire currency (1.026 * 1.5 billion, where 1.5 billion is the change in in the value of total trade volume). Since approximately 17 million bitcoin currently exist, the expected $39 million increase in returns translates to approximately $2 in extra returns per bitcoin.

No other coefficients for other variables are significant in my baseline regression, except for the coefficient on change in coins generated for Vertcoin. Vertcoin's $\Delta(\log(CG_t))$ coefficient, at -0.141, is significant at the 99% confidence level: the negative direction fits basic economic assumptions, since an increase in the rate of change of coins generated is essentially an increase in the rate of inflation growth. This coefficient is very large, implying a 100 percentage point change in the rate of coins generated results in a 14.1 percentage point decrease in predicted returns, this is a somewhat reasonable estimate since fluctuations in

Table 3: Baseline Regressions

	Dependent variable: log(Returns)			
	Bitcoin	Litecoin	Vertcoin	Digibyte
	(1)	(2)	(3)	(4)
Segwit	0.004	0.004	0.004	0.007
	(0.003)	(0.004)	(0.006)	(0.008)
Δlog(TV)	0.026***	0.023***	0.018***	0.020***
	(0.004)	(0.004)	(0.003)	(0.004)
Median Fee	−2.093	0.033	−6.728	0.001
	(6.560)	(0.110)	(6.034)	(0.005)
Δlog(GC)	−0.0001	0.036	−0.141***	−0.009
	(0.008)	(0.023)	(0.033)	(0.012)
Segwit * Δlog(TV)				
Own Lag				
BTC Lag				
Observations	1,457	1,457	1,457	1,457
Adjusted R^2	0.029	0.025	0.038	0.020
SER	0.038	0.059	0.105	0.102
F Statistic	11.698***	10.384***	15.376***	8.414***

*p<0.1; **p<0.05; ***p<0.01

coins generated occur gradually. Furthermore, heteroskedasticity due to my use of timeseries data likely biases my coefficients anyways, so confirming the intuitively negative direction of the coefficient is more important for verifying the validity of my regressions.

I ran various robustness checks to ensure my model meets the baseline regression assumptions, checking in particular for collinearity between SegWit and Trade Volume. Collinearity between those two variables is an intuitive concern, since SegWit's activation is supposed to enable a higher possible trade volume than pre-SegWit. Since I am analyzing the effect of SegWit's lock-in rather than that of SegWit's activation, which are approximately two weeks apart, the technological impact of SegWit on Trade Volume should not appear in my model as collinearity. Furthermore, after running regressions without trade volume variables as collinearity checks, collinearity does appear to be an issue: Regressions run without trade volume do not show significantly different coefficients or standard errors [see appendix B]. The coefficient on SegWit changes from 0.00440 to 0.00438, with the standard error changing from 0.0027 to 0.0033. Similarly negligible changes occur across variables and coins, suggesting that collinearity between SegWit and fluctuations in trade volume does not significantly interfere with my model's accuracy.

Table 4: Regressions with Lagged Returns Bitcoin

	Dependent variable: log(Returns)			
	Bitcoin	Litecoin	Vertcoin	Digibyte
	(1)	(2)	(3)	(4)
Segwit	0.004	0.003	0.004	0.006
	(0.003)	(0.004)	(0.006)	(0.008)
log(Trading Volume)	0.019***	0.018***	0.016***	0.015***
	(0.004)	(0.004)	(0.004)	(0.003)
Median Fee	−2.089	−0.029	−6.534	0.001
	(6.454)	(0.109)	(6.054)	(0.005)
log Generated Coins	−0.001	0.036	−0.136***	−0.010
	(0.008)	(0.022)	(0.034)	(0.012)
Segwit * log(Trading Volume)	0.034***	0.022*	0.021**	0.038**
	(0.013)	(0.013)	(0.010)	(0.018)
Lagged Returns by Coin	✓	✓	✓	✓
BTC Lagged Returns		✓	✓	✓
Observations	1,457	1,457	1,457	1,457
Adjusted R^2	0.036	0.029	0.039	0.029
SER	0.038	0.059	0.105	0.102
F Statistic	7.729***	5.015***	6.303***	5.008***

*p<0.1; **p<0.05; ***p<0.01

I then proceed to add my interaction term and controls, improving my model across coins as measured by R^2, but the coefficients for my SegWit dummy remain nonsignificant (see table 4 above; see appendix C for all regressions on each coin). Again, my model gives no strong indications as to whether or not SegWit's introduction directly has a positive effect on predicted returns. Although my baseline regressions don't enable me to accept or reject my original hypothesis, when I add the interaction term between SegWit and change in trade volume, my regression coefficients remain remarkably similar (see appendix A). However, the coefficients on change in trade volume decrease slightly, by at most 0.007, while the interaction term has a strong significant effect. In Bitcoin's case, SegWit's lock-in for activation nearly triples the suggested impact of a 100 percentage point increase in the change of trade volume on returns, from 1.9 percentage points to 5.3 percentage points; i.e., the interaction term is 0.034.

My regressions with the interaction term and controls return consistent results across coins: SegWit's introduction is associated with an approximate 150% increase of the effect of the change in trade volume on estimated returns. Upon SegWit's introduction, my model estimates the effect of trade volume fluctuations on returns went from 1.8 percentage points to 4.0 for Litecoin, from 1.6 percentage points to 3.7 for Vertcoin, and 1.5 percentage points to 5.3 for Digibyte (See appendix C for all models on individual coins). These interaction terms are significant at the 95% level, except for the coefficient on Litecoin, which is only significant at the 90% level. Other variables, including the controls, remain statistically insignificant; the coefficient for change in coins generated is significant only on Vertcoin, where it again goes the intuitive, negative direction: As expected, my model suggests an increase in the rate of inflation leads to a decrease in predicted returns.

A few interesting results do not relate directly to SegWit, rather indicating where this paper fits in the prior literature. Consistent with Liu's findings in 2018, adding lagged returns by coin to my regression model increases the model's fit as estimated by the adjusted R squared values; however, none of those variables are statistically significant. Contradicting my hypotheses regarding Bitcoin's first-mover advantage drawing from industrial organization literature, adding lagged Bitcoin returns had either no effect or a small negative effect on the model fit for each coin (see appendix C). The entire cryptocurrency market may not track Bitcoin as strongly as I suspected; more likely, a different set of regressions would be appropriate to investigate to what degree the entire crypto market tracks Bitcoin.

Conclusion

My regressions suggest Segregated Witness has no direct impact on predicted returns for cryptocurrencies; the results presented here fail to imply one way or the other whether or not SegWit's lock-in had a direct effect on predicted returns. Rather, fluctuations in trade volume (predictably) seem have the only significant impact on prices. My results also suggest that the dummy for SegWit's lock-in interacts with trade volume, doubling or tripling the predicted effect a fluctuation in the growth of trading volume has on returns.

Supporters of SegWit's introduction on cryptocurrencies that have yet to install the software can celebrate these results: I found no evidence that SegWit's planned installation caused a large enough exodus from any cryptocurrency—or any major decrease in use—to seriously dampen growth. As important, SegWit appears to decrease financial frictions, increasing the responsiveness of predicted returns to changes in growth of trade volume.

Future research to test my findings will focus on getting data for other coins that installed SegWit, such as Syscoin and Navcoin. Since these two coins are more obscure than those I analyze, they do not have reliable trade volume data, with double counting and inter-exchange transactions removed as they are by the coinmetrics team for the adjusted trade volume data I use. Not including these coins in my comparison leads to a huge oversight—Vertcoin is the only coin in my sample with similarly small market capitalization (around $20 million) to Syscoin and Navoin. Regressions on volatility measures would also be interesting to compare with my regressions on returns.

The introduction of second-layer protocols, updates built on top of SegWit's technology that further increase transaction speed, over SegWit will further require new rounds of research. Also, future regressions should compare coins that instituted SegWit with coins that pursued other, currently theoretical methods of improving network speed, such as increasing block size. Since one of the most significant limitations to cryptocurrencies' widespread use in international commerce remains slow blockchain network speed, further understanding how the market responds to controversial technological changes will be necessary for improving cryptocurrency networks in a way acceptable to the crypto community. For now, no evidence indicates that SegWit is so controversial it should prevent crypto communities from trying to use the technology.

Glossary

ASIC (Application Specific Integrated Circuits): specialized mining machines that can be used on certain cryptocurrency networks, enabling individuals or groups to more easily contribute a majority of hash power.

Blockchain: A digital ledger recording transactions chronologically and publicly.

Block: https://www.coindesk.com/information/what-is-segwit

Coin: Cryptocurrency operating independently of any others.

Consensus Models: Most cryptocurrencies' blockchain networks have rules determined by a consensus of users on the network. Consensus model describes this system of governing currencies.

Cryptocurrency: A financial instrument constructed using blockchain technology. Transactions are verified and new coins are generated algorithmically.

ERC-20: Tokens built on the Ethereum network

Forgers: The individuals who generate new coins by solving algorithms in PoS systems.

Graphics Cards: general computer chips that can mine cryptocurrency more efficiently than most PCs. However, graphics cards are far cheaper and more widely available than ASICs.

Hard Fork: A cryptocurrency splits in two, occurring when new code overwriting old code coexists simultaneously with the old code. Hard forks effectively result in two new cryptocurrencies.

Hash Power: The computing power contributed by individual computers to cryptocurrency networks through mining or forging. Hash power determines the sway an individual or group has on a consensus model—essentially, more hash power equals more votes in the consensus model.

Proof of Work (PoW) System: One style of cryptocurrency algorithms in which new coins are generated by miners solving algorithms that verify transactions.

- Mining power determines coins generated, so 51% attacks can occur if an individual (or a small oligopoly) has massive amounts of mining power.

Proof of Service (PoS) System: New coins are generated randomly, with weights given based on stake size and term of holding. Forgers verify transactions for a small fee, with penalties for verifying invalid transactions.

- Only buying a large stake of a cryptocurrency with standard money can enable a 51% attack.

Miners: The individuals who generate new coins by solving algorithms in PoW systems.

SegWit (Segregated Witness): A soft fork software used to change the nature of transactions in Bitcoin, Litecoin, and a host of other currencies. Most significantly, SegWit increased the speed with which blockchain networks could process transactions.

- SegWit increases the amount of information that can be stored in a blockchain block
- Bitcoin cash was created largely to protest SegWit. Bitcoin cash increased block size

Soft Fork: New code is introduced over the old base for a cryptocurrency but only one version of the blockchain is active at a time as users transition. Soft forks are essentially software updates instituted when enough miners or forger signal consensus for the update.

- Soft forks require consensus among users on a crypto network to activate, but then not all users must participate
- Hard forks require a subset of users on a crypto network to agree to a new protocol

Token: A cryptocurrency whose network depends on another cryptocurrency to operate.

51% attack: A maneuver by which an individual or small oligopoly seeks to take control of a network by controlling a majority of the hash power and overwriting the standing rules, usually in the attacker's favor. An individual or group that executes a 51% attack could verify its own transactions, even if they were illegitimate.

Table of Contents

Banerjee, Snehal and Breon-Drish, Bradyn M., “Dynamic Information Acquisition and Entry into New Markets” (September 20, 2018). Available at SSRN: https://ssrn.com/abstract=2846059 or http://dx.doi.org/10.2139/ssrn.2846059

Cutler, David M., James M. Poterba, and Lawrence H. Summers. "What Moves Stock Prices?" Available at NBER (1988): https://www.nber.org/papers/w2538.
Variances in stock prices cannot be fully explained, or even very well partially explained, by events that the market responds to.

Fama, Eugene F. "Random walks in stock market prices." *Financial analysts journal* 51, no. 1 (1995): 75-80.
https://www.chicagobooth.edu/~/media/34F68FFD9CC04EF1A76901F6C61C0A76.PDF.

Fama, Eugene F., and Kenneth R. French. "Dividend yields and expected stock returns." *Journal of financial economics*22, no. 1 (1988): 3-25.
https://www.sciencedirect.com/science/article/pii/0304405X88900207.

Glazer, A. "The Advantages of Being First." *The American Economic Review* 75, no. 3 (1985): 473-80. http://www.jstor.org/stable/1814812.

Feng, Guanhao, Stefano Giglio, and Dacheng Xiu. "Taming the Factor Zoo: A Test of New Factors." *Fama-Miller Working Paper* (2019): 17-04.

Fry, John. "Booms, busts and heavy-tails: The story of Bitcoin and cryptocurrency markets?." *Economics Letters* 171 (2018): 225-229.
https://reader.elsevier.com/reader/sd/pii/S016517651830315X?token=D90E632C94C322C28110FDA374AB479004B5717C57C59AADCC6D4D4439B8D72880C176121BF633B350EA3AFCD11A12FB.

Kaminski, Jermain. “Nowcasting the Bitcoin Market with Twitter Signals.” Available at Cornell: arXiv:1406.7577v3 [cs.SI] or https://arxiv.org/abs/1406.7577.

Kim YB, Kim JG, Kim W, Im JH, Kim TH, Kang SJ, et al. (2016). “Predicting Fluctuations in Cryptocurrency Transactions Based on User Comments and Replies.” PLoS ONE 11(8): e0161197. https://doi.org/10.1371/journal.pone.0161197.

Li, Tao, Donghwa Shin, and Baolian Wang. "Cryptocurrency pump-and-dump schemes." *Available at SSRN* (2018).

Lieberman, Marvin B., and David B. Montgomery. "First-mover advantages." *Strategic management journal* 9, no. S1 (1988): 41-58.
https://onlinelibrary.wiley.com/doi/abs/10.1002/smj.4250090706.

Liu, Yukun, and Tsyvinski, Aleh. “Risks and Returns of Cryptocurrency.” Available at NBER (2018): https://www.nber.org/papers/w24877.pdf. (DOI): 10.3386/w24877.

Matta, Martina, Ilaria Lunesu, and Michele Marchesi. "Bitcoin Spread Prediction Using Social and Web Search Media." In *UMAP Workshops*, pp. 1-10. 2015. https://pdfs.semanticscholar.org/1345/a50edee28418900e2c1a4292ccc51138e1eb.pdf.

Merton, Robert C. "On estimating the expected return on the market: An exploratory investigation." *Journal of financial economics* 8, no. 4 (1980): 323-361. http://www.people.hbs.edu/rmerton/onestimatingtheexpectedreturn.pdf.

Schilling, Linda and Uhlig, Harold, “Some Simple Bitcoin Economics” (April 2018). Available at NBER: https://www.nber.org/papers/w24483.

Shiller, Robert J., Stanley Fischer, and Benjamin M. Friedman. "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity* 1984, no. 2 (1984): 457-510. doi:10.2307/2534436.

A list of nonacademic reading for the interested reader:
https://www.coindesk.com/information/what-is-segwit - explaining Segregated Witness
https://youtu.be/fst1IK_mrng?t=36m – first conference presentation introducing SegWit
https://coincentral.com/making-sense-of-proof-of-work-vs-proof-of-stake/ - explaining the two types of cryptocurrency (only PoW is addressed in my paper)
https://bitcoin.org/bitcoin.pdf - Satoshi’s founding paper for bitcoin

Appendix A

Table 5: Regressions with Lagged Returns All Coins

	Dependent variable: log(Returns)			
	Bitcoin	Litecoin	Vertcoin	Digibyte
	(1)	(2)	(3)	(4)
Segwit	0.004	0.004	0.004	0.006
	(0.003)	(0.004)	(0.006)	(0.008)
log(Trading Volume)	0.019***	0.018***	0.016***	0.015***
	(0.004)	(0.004)	(0.004)	(0.003)
Median Fee	−2.089	0.006	−6.769	0.001
	(6.454)	(0.105)	(6.024)	(0.005)
log(Generated Coins)	−0.001	0.036	−0.136***	−0.010
	(0.008)	(0.022)	(0.034)	(0.011)
Segwit * log(Trading Volume)	0.034***	0.022*	0.020**	0.037**
	(0.013)	(0.013)	(0.010)	(0.017)
Lagged Returns by Coin	✓	✓	✓	✓
Observations	1,457	1,457	1,457	1,457
Adjusted R^2	0.036	0.030	0.039	0.029
SER	0.038	0.059	0.105	0.102
F Statistic	7.729***	6.589***	8.455***	6.348***

*p<0.1; **p<0.05; ***p<0.01

Appendix B

Table 6: Collinearity Checks Using Bitcoin Regressions

	Dependent variable: log(Returns)			
	Bitcoin	Litecoin	Vertcoin	Digibyte
	(1)	(2)	(3)	(4)
Segwit	0.004	0.004	0.004	0.004
	(0.003)	(0.003)	(0.003)	(0.003)
log(TV)	0.026***		0.019***	
	(0.004)		(0.004)	
Median Fee	−2.093	−1.936	−2.089	−1.929
	(6.560)	(6.797)	(6.454)	(6.826)
log(GC)	−0.0001	−0.003	−0.001	−0.003
	(0.008)	(0.008)	(0.008)	(0.008)
Segwit * log(TV)			0.034***	
			(0.013)	
Lagged Returns by Coin			✓	✓
Observations	1,457	1,457	1,457	1,457
Adjusted R^2	0.029	-0.00000	0.036	0.0001
Residual Std. Error	0.038	0.039	0.038	0.039
F Statistic	11.698***	0.998	7.729***	1.019

Note: *p<0.1; **p<0.05; ***p<0.01

Appendix C

Table 7: Regressions on Bitcoin Returns

	Dependent variable:			
	log(Returns)			
	(1)	(2)	(3)	(4)
Segwit	0.004	0.004	0.004	0.004
	(0.003)	(0.003)	(0.003)	(0.003)
log(Trade Volume)		0.026***	0.019***	0.019***
		(0.004)	(0.004)	(0.004)
Median Fee		−2.093	−2.054	−2.089
		(6.560)	(6.399)	(6.454)
log Generated Coins		−0.0001	−0.001	−0.001
		(0.008)	(0.008)	(0.008)
Segwit * log(Trading Volume)			0.034***	0.034***
			(0.013)	(0.013)
Own Lag				✓
Observations	1,457	1,457	1,457	1,457
Adjusted R^2	0.001	0.029	0.036	0.036
SER	0.039	0.038	0.038	0.038

*p<0.1; **p<0.05; ***p<0.01

Table 8: Regressions on Litecoin Returns

	Dependent variable:				
	log(Returns)				
	(1)	(2)	(3)	(4)	(5)
Segwit	0.003	0.004	0.004	0.004	0.003
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
log(TV)		0.023^{***}	0.019^{***}	0.018^{***}	0.018^{***}
		(0.004)	(0.004)	(0.004)	(0.004)
Median Fee		0.033	0.016	0.006	−0.029
		(0.110)	(0.109)	(0.105)	(0.109)
log(GC)		0.036	0.037	0.036	0.036
		(0.023)	(0.023)	(0.022)	(0.022)
Segwit * log(TV)			0.021^{*}	0.022^{*}	0.022^{*}
			(0.013)	(0.013)	(0.013)
Own Lag				✓	✓
BTC Lag					✓
Observations	1,457	1,457	1,457	1,457	1,457
Adjusted R^2	-0.00004	0.025	0.028	0.030	0.029
SER	0.060	0.059	0.059	0.059	0.059

*p<0.1; **p<0.05; ***p<0.01

Table 9: Regressions on Vertcoin Returns

	Dependent variable: log(Returns)				
	(1)	(2)	(3)	(4)	(5)
Segwit	0.006	0.004	0.004	0.004	0.004
	(0.006)	(0.006)	(0.006)	(0.004)	(0.006)
log(TV)		0.018^{***}	0.016^{***}	0.016^{***}	0.016^{***}
		(0.003)	(0.004)	(0.004)	(0.004)
Median Fee		−6.728	−6.854	-6.769^{***}	−6.534
		(6.034)	(6.019)	(0.105)	(6.054)
log(GC)		-0.141^{***}	-0.137^{***}	-0.136^{***}	-0.136^{***}
		(0.033)	(0.034)	(0.022)	(0.034)
Segwit * log(TV)			0.020^{**}	0.020	0.021^{**}
			(0.010)	(0.013)	(0.010)
Own Lag				✓	✓
BTC Lag					✓
Observations	1,457	1,457	1,457	1,457	1,457
Adjusted R^2	-0.0001	0.038	0.040	0.039	0.039
Residual Std. Error	0.107	0.105	0.105	0.105	0.105

$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 10: Regressions on Digibyte Returns

	Dependent variable: log(Returns)				
	(1)	(2)	(3)	(4)	(5)
Segwit	0.006	0.007	0.006	0.006	0.006
	(0.007)	(0.008)	(0.008)	(0.008)	(0.008)
log(TV)		0.020***	0.015***	0.015***	0.015***
		(0.004)	(0.003)	(0.003)	(0.003)
Median Fee		0.001	0.001	0.001	0.001
		(0.005)	(0.005)	(0.005)	(0.005)
log(GC)		−0.009	−0.009	−0.010	−0.010
		(0.012)	(0.011)	(0.011)	(0.012)
Segwit * log(TV)			0.038**	0.037**	0.038**
			(0.018)	(0.017)	(0.018)
Own Lag				✓	✓
BTC Lag					✓
Observations	1,457	1,457	1,457	1,457	1,457
Adjusted R^2	-0.00002	0.020	0.028	0.029	0.029
Residual Std. Error	0.103	0.102	0.102	0.102	0.102

*p<0.1; **p<0.05; ***p<0.01