ERC-20 Cryptocurrency Market Efficiency: Investigating High Frequency Efficiency Across Systemic Events - The Ethereum Merge and The Collapse of FTX

Author: Dominic Cischke

Advisor: Professor Geoffrey Sanzenbacher

Senior Honors Thesis

Department of Economics

Boston College

May 2023

Abstract:

Recent investigations into the weak form market efficiency of cryptocurrencies find compelling evidence for market efficiency in leading cryptocurrencies such as bitcoin and ether with respect to their own cross-time returns, even at higher frequencies of one hour returns. However, little is known about the market efficiency performance of other cryptocurrencies at high frequencies. Additionally, there is little research into how market efficiency levels are sustained across systemic events in the cryptocurrency ecosystem.

Traditional tests for market efficiency are utilized to expand the high frequency literature to include other established Ethereum-based cryptocurrencies (ERC-20 tokens) centered around systemic events such as the Ethereum Merge and the collapse of the FTX cryptocurrency exchange. The findings indicate mixed levels of efficiency during periods of non-systemic events, in line with the existing literature. However, market efficiency significantly declines during periods of systemic events, indicating a sensitivity to exogenous forces. In particular, there appears to be evidence of stronger inefficiency during unexpected systemic events as exemplified by the collapse of FTX. Further investigations are encouraged to expand and supplement the presented findings.

Acknowledgements:

First I would like to thank Professor Geoffrey Sanzenbacher for his immense support and guidance throughout the thesis process. I am also thankful for his energetic instruction of my econometrics courses which greatly contributed to my quantitative inclination and approach to my thesis.

I'd like to thank the Economics Department of Boston College and its professors for their part in my educational journey at Boston College. I'd also like to thank the Carroll School of Management and in particular its Finance Department for providing me the fundamental financial understanding that allowed me to pursue this topic.

Finally, I'd like to thank the projects of the blockchain technology industry for keeping accessibility as a top priority. In particular, DeFi Llama, CoinMarketCap, and CoinGecko enabled me to smoothly obtain and use market information, allowing me to dedicate more of my attention to investigating the topic.

Table of Contents

1 Introduction

Cryptocurrencies are one of the most recent additions to the ranks of the alternative asset class. Since its introduction with the Bitcoin blockchain in 2009, the cryptocurrency market has grown to include thousands of different cryptocurrencies with a combined market capitalization that has, at times in recent years, eclipsed over one trillion dollars (CoinMarketCap). While the market has grown exponentially, two cryptocurrencies are still responsible for a majority of the cryptocurrency market capitalization: bitcoin[1] and ether[2].

Market efficiency is a fundamental component of contemporary financial economic theory and a cornerstone of traditional financial markets. With the presence of market efficiency, market participants are presented with the fair value of assets. Market participants pay the fair value of an asset when purchasing and receive the fair value of an asset when selling. Market efficiency is of particular importance in the cryptocurrency market due to its highly unregulated nature that leaves it susceptible to market manipulation that does not receive regulatory enforcement action. The prevalence of fair-value exchange in a market helps to support the entrance of new market participants.

Research into the market efficiency of cryptocurrencies has primarily focused on the weak form[3] market efficiency of bitcoin and ether. The existing research has concluded that cryptocurrencies, particularly more established ones such as bitcoin and

---

[1] The native cryptocurrency of the Bitcoin blockchain for transactions and transaction fee payment is bitcoin.

[2] The native cryptocurrency of the Ethereum blockchain for transactions, transaction fee payment, and code deployment fees is ether.

[3] Weak form market efficiency holds that all information held in historical price data cannot be used to predict future returns.

ether, demonstrate increasing levels of market efficiency as they mature that ultimately results in a generally efficient market.

However, as the market continues to mature, a marked growth in the value of other cryptocurrencies has also been observed. Turing-complete blockchains such as the Ethereum blockchain enable this growth by substantially decreasing the barrier to entry in creating a new cryptocurrency[4]. This trend has caused the market share of bitcoin and ether to steadily decline over recent years (CoinMarketCap). Given this trend, recent literature has begun to more commonly investigate the market efficiency of other cryptocurrencies in addition to bitcoin and ether. These other cryptocurrencies largely exist as Ethereum-based tokens which are cryptocurrencies transacted on and maintained by the technological infrastructure of the Ethereum blockchain. The literature contains investigations into the weak form market efficiency of many established cryptocurrencies, commonly at the daily level. However, only bitcoin and ether have been investigated at "high frequencies," with the lowest observation period of one hour between returns. Additionally, the literature lacks ample investigations into the impact of systemic events on cryptocurrency market efficiency.

The forthcoming analysis will attempt to answer the question of: *do Ethereum-based token returns satisfy high frequency market efficiency across systemic events?* The intended contributions to the existing literature include a step taken towards increasing the highest studied frequency for cryptocurrency market efficiency beyond the one hour mark to 30 minute observations, an expansion of the set of investigated cryptocurrencies for market efficiency at the highest studied frequency, and

[4] Cryptocurrencies built on the Ethereum blockchain commonly follow the Ethereum ERC-20 token standard which creates a fungible and transactable asset whose ledger is maintained by the Ethereum blockchain.

the investigation of the impact on market efficiency by two systemic events: the Ethereum Merge and the collapse of FTX.

1.1 The Efficient Market Hypothesis

The Efficient Market Hypothesis is a core theory of modern financial economics. Popularized by Eugene Fama's (1970) investigation and furthering of the Efficient Market Hypothesis, the theory states that efficient markets feature prices which fully reflect all available information. The Efficient Market Hypothesis holds that any new information is instantaneously accounted for in market prices and cannot be used to predict future returns after becoming available.

Forms of market efficiency fall into three described categories: weak form, semi-strong form, and strong form. Weak form market efficiency holds that all information held in historical price data cannot be used to predict future returns. If weak form market efficiency holds, then future prices follow a "random walk" in which future returns are randomly determined and prior trends in the data do not inform future outcomes. Many common statistical tests of market efficiency test for weak form market efficiency, such as those used in the forthcoming analysis. Therefore much of the existing literature on market efficiency in traditional and alternative markets, including the cryptocurrency market, tests for weak form market efficiency.

Semi-strong form market efficiency includes the consideration of prior price information and extends the level of priced-in information to all publicly available information such as new information contained in an earnings report. A semi-strong form efficient market will instantaneously react to new public information and this information will not be useful in predicting the price movement beyond the moment it occurred. Traditional markets such as equity and bond markets are commonly considered to exhibit semi-strong form market efficiency. The theory of strong form market efficiency holds that current prices reflect all private information, in addition to all public information and prior price data. Under strong form market efficiency, private (insider) information such as an undisclosed pending acquisition would already be priced into the asset's value and cannot be used to extract further value. Strong form efficiency is not commonly considered to be exhibited by contemporary markets due to its strict criteria.

1.2 Blockchain Technology and Cryptocurrency

Bitcoin, as set initially forth by its creator Satoshi Nakamoto (2008), is a peer-to-peer electronic payment network that does not require a third, authoritative party to process transactions. Bitcoin was the implementation of Nakamoto's newly invented "blockchain technology." Vitalik Buterin (2014) created Ethereum to expand the use cases of Namakoto's blockchain technology. Ethereum is a Turing-complete blockchain[5], on which code can be executed and applications can be built.

Ethereum enables anyone to create a token (cryptocurrency) of their own by using the Ethereum blockchain to process all transactions. The most common form of tokens built on the Ethereum blockchain are ERC-20 tokens, which are fungible assets that can be transacted within the Ethereum ecosystem. Since the creation of Ethereum, thousands of tokens have been created for varying purposes, with some garnering individual market capitalizations in the billions of dollars. Ethereum also features a

[5] A Turing-complete blockchain is able to accept outside algorithms containing conditional statements and loops and execute the algorithm on its own. The first blockchain, Bitcoin, is not Turing-complete as its transaction scripts cannot perform complete computations of external algorithms.

native cryptocurrency, ether (ETH), which users can send to one another and is needed to pay fees for transactions on the Ethereum blockchain, including all token transactions.

1.3 Systemic Events in the Cryptocurrency Market

Two recent major events that affected the cryptocurrency ecosystem are chosen to investigate the impact of systemic events on market efficiency.

On September 15, 2022 the Ethereum blockchain underwent a long-anticipated, protocol-wide upgrade called “The Merge”. The Merge officially transitioned the Ethereum blockchain to a proof-of-stake consensus algorithm which significantly decreased the network’s operating energy consumption by approximately 99.95% and enabled further scalability upgrades and potential (Ethereum.org, 2023). These protocol changes affect both ether and the cryptocurrencies built on top of the Ethereum blockchain. Because a large share of cryptocurrencies exist as ERC-20 tokens built on the Ethereum blockchain, it is important to investigate the impact that events in the Etherum ecosystem have on the efficiency of these assets. Cryptocurrencies such as bitcoin do not use the Ethereum blockchain, so any impact on bitcoin’s efficiency can likely be attributed to market-wide sentiments that are changing market efficiency levels. Due to the fact that this event was known to occur far in advance, there is no expectation that this event should change market efficiency in a largely efficient market.

The second systemic event that will be investigated is the collapse of FTX. FTX, a once popular cryptocurrency exchange, collapsed in November 2022 after it was discovered to be insolvent. It is reported that over $8 billion of customer deposits are missing from the firm’s assets (Chow, 2022). The market is expected to respond to new information and, if efficient, instantaneously price in the new information, therefore maintaining market efficiency. Weak form efficiency will remain intact if previous price information cannot be used to predict returns for future periods beyond the event. However, forces such as poor information flow in the market and a delayed response by some market participants to information can lead to an efficient market becoming inefficient.

FTX allowed users to purchase and sell a variety of cryptocurrencies, so the impact of this event is levied on the entire cryptocurrency market rather than focusing on a specific ecosystem within the cryptocurrency space as was the case with the Ethereum Merge. Additionally, the collapse of FTX is different from the Ethereum Merge as a systemic event in the fact that it was unknown to occur by market participants while the Ethereum Merge had a predetermined date of occurrence. Finally, the collapse of FTX is a systemic event with a negative immediate impact on the ecosystem, while the Ethereum Merge is viewed as a positive technological development for the cryptocurrency ecosystem, particularly for ether and Ethereum-based cryptocurrencies. It is worth noting that the magnitude of impact these events had on the cryptocurrency ecosystem is difficult to compare given the stated differences between the two events.

1.4 Motivation

Efficient markets are an important feature of the global financial system as they provide a more equitable playing field for both institutional and individual investors. Efficient markets enable investors to purchase assets at their fair market value and achieve risk-adjusted returns in line with the market.

One guiding motivation of cryptocurrency market efficiency research is to demonstrate the efficiency of blockchain-based assets to outside capital. The inefficiency of asset prices is one example of an imperfection in capital markets. Capital market imperfections lead to under-investment in innovation which harms the future prospects of the market (Peneder, 2008, 520). While the blockchain technology industry has grown exponentially since its inception in 2008, it is still miniscule in comparison to other industries with efficiently-traded assets. Further innovation and subsequent growth of the blockchain technology industry can be fueled by additional investment into the space. Cryptocurrencies could present an opportunity to further diversify an investor's portfolio. However, outside investors might be dissuaded by the inefficiency of blockchain-based assets when other opportunities exist in efficient markets where the outside investor is confident in receiving a fair price.

Modern-day market movement and the trading of these markets occur at incredibly high frequencies. High frequency trading firms focused on traditional markets can execute orders in fractions of a second. Consequently, traditional markets must exhibit overall efficiency at these high frequencies or else high frequency trading firms could reliably trade on these inefficiencies and extract guaranteed long-run, market-beating returns. As cryptocurrency markets and the technology that participants use to trade these markets continue to develop, so must the efficiency of these markets at higher frequencies to maintain overall market efficiency. The current literature, discussed below, investigates market efficiency at the highest frequency of one hour. The forthcoming investigation will seek to further the literature's understanding of cryptocurrency market efficiency at high frequencies by examining market efficiency at 30 minute observation periods.

The recent occurrence of The Merge on the Ethereum blockchain has not been a time period yet studied by the existing literature. The significance of such an event and the future upgrades to the network that The Merge enables supports a need to study the state and behavior of cryptocurrency market efficiency surrounding the times of such upgrades. Due to the unregulated nature of the cryptocurrency industry, the collapse of FTX is also an important event to investigate as other unknown systemic events arising from institutions within the cryptocurrency industry have a fair chance of repeatedly occurring while the industry continues to develop.

Another motivating factor of researching cryptocurrency market efficiency, and the efficiency of markets in general, is the loss of human capital allocation and societal contribution for those that commit resources towards trading on the market's inefficiencies. Furthermore, this gain by inefficiency traders comes from the unnecessary losses of other, often less sophisticated, investors who must buy and sell at prices that do not reflect available information nor the asset's fair market value. If market-beating risk-adjusted returns are diminished through increased market efficiency, all investors will purchase assets at its fair market value. Additionally, those previously extracting value from the inefficiencies at the expense of other investors will be forced to reallocate their resources to other means of income.

Finally, the implications of inefficient markets are magnified in the case of the cryptocurrency market as, compared to traditional financial markets, the exchanges on which cryptocurrencies are traded are highly unregulated. Unregulated exchanges have a greater ability to abuse their power and engage in dubious behavior such as market manipulation and the front-running of trades. The negative impact of these activities can be amplified when sophisticated exchanges hold an advantage in a state of information asymmetry and are able to act without government enforcement.

The “Wild West” nature of the cryptocurrency market creates an environment where powerful market participants are increasingly able to extract value from other market participants. In markets where there is great opportunity for manipulation, regulation is needed to protect the interests of all market participants (Levmore, 2002, 604). In the past few months alone, a number of prominent, largely unregulated cryptocurrency platforms providing critical financial services such as exchanges, lenders, and asset custodians have collapsed due to fraud, negligence, and the misuse of customer funds, harming both the market and the users that trusted these institutions. The collapse of one of these entities, FTX, will be explored for its systemic effect on cryptocurrency market efficiency. If the presence of regulation in the cryptocurrency market remains at a low level, then efficiency created by well-intentioned market participants through increased levels of liquidity and public information could be left to combat the ability of bad actors to manipulate the market.

2 Related Literature

Prior investigation into the market efficiency of cryptocurrencies have primarily focused on the weak form efficiency of bitcoin and ether. Urquhart (2016) conducted the first investigation into the market efficiency of cryptocurrency with his analysis of daily bitcoin returns from 2010 to 2016. Urquhart examined market efficiency for the entire period and the subperiods of 2010-2013 and 2013-2016. As there was no prior literature on the market efficiency of cryptocurrencies, Urquhart applied a number of market efficiency tests that are commonly used with traditional financial assets such as stocks and bonds.

Urquhart's applied tests include the Ljung-Box and Brock-Dechert-Scheinkman tests for serial correlation, the Runs and Bartels tests for randomness, the Automatic Variance test for a random walk with drift, and the rescaled Hurst exponent for long memory of returns. A number of these tests are common throughout the literature and therefore many will be applied in my analysis. The Methodology section will further explain the implementation and results of the adopted tests.

Urquhart discovered strong evidence of *market inefficiency* with the first period of data ranging from 2010 to 2013. However, Urquhart found increased levels of *efficiency* for the second period of 2013 to 2016 with some tests providing results of market efficiency that previously displayed inefficiency for the first period.

Wei (2018) extends the research of Urquhart by applying the same methods of market efficiency testing, with the addition of specialized liquidity tests to determine the impact of liquidity on market efficiency. 2017 was a time of exponential growth in the number of existing cryptocurrencies due to a time of market jubilance. Capitalizing on the existence of many more cryptocurrencies, Wei obtained daily data during 2017 for 456 different cryptocurrencies from CoinMarketCap, a popular cryptocurrency price aggregation website.

Wei grouped the cryptocurrencies of the sample into five groups ranked by market liquidity defined by their calculated Amihud illiquidity measure and then performed the tests of market efficiency. The results show high levels of inefficiency for the cryptocurrencies of the low liquidity groups. The highest liquidity group made up of more established cryptocurrencies such as bitcoin and ether display positive results that support the presence of an efficient market with price determination being close to a random walk. Wei's investigated impact of cryptocurrency liquidity and market prominence leads me to conduct my analysis of Ethereum-based token market efficiency with a sample of tokens that have intentionally varying levels of market capitalization and liquidity.

The 2017 cryptocurrency boom also brought rise to an increase in the active trading of cryptocurrencies. Chu, Zhang, and Chan (2019) investigate cryptocurrency market efficiency with an improvement in frequency from the daily level, as was common in the existing literature, to the hourly level. They investigate the efficiency of hourly prices from July 2017 to September 2018 for bitcoin and ether under the Alternative Market Hypothesis. The Alternative Market Hypothesis relaxes the constraints of the Efficiency Market Hypothesis, which requires no cross-time correlation of returns at any time, by allowing for the market to grow and result in efficiency over time and still be deemed efficient. To this aim, the analysis uses rolling windows of cross-time observations to evaluate the change in efficiency over time. They find results that support efficiency under the Alternative Market Hypothesis, with increased levels of efficiency over time and ultimately efficient markets for bitcoin and ether.

Noda (2021) reviews the evolution of cryptocurrency market efficiency. Through a robust evaluation of existing literature, Noda compares the findings of various investigations of cryptocurrency market efficiency. Many investigations concerned with the Adaptive Market Hypothesis find general market efficiency for established cryptocurrencies, but the results are not as consistent with tests under the strict rules of the Efficient Market Hypothesis. Although Efficient Market Hypothesis tests have trended towards increased efficiency over time, some investigations still find results that do not support an efficient cryptocurrency market.

Yaya, Ogbonna, Mudida, and Abu (2021) investigate market efficiency alongside cryptocurrency volatility persistence. They argue that market efficiency and volatility are fundamentally connected. The high, sporadic levels of volatility in the cryptocurrency market could be a reason for varying conclusions reached in the literature. They apply standard market efficiency tests while holding to the strict assumptions of the Efficient Market Hypothesis. Using daily prices for prominent cryptocurrencies from August 2015 to November 2018 they find general levels of market efficiency for most studied cryptocurrencies, with the most efficient asset being bitcoin. However, periods of high volatility were still persistent throughout the data set which can lead to varying levels of efficiency across the multi-year time series.

The COVID-19 pandemic was a time of turbulence in traditional markets and this effect was also experienced by the cryptocurrency market. Using one hour returns for bitcoin and ether from a pre-pandemic period of 2019 and a pandemic period of 2020, Kakinaka and Umeno (2022) investigated short-term and long-term weak form market efficiency. They found an increased prevalence of herding behavior during the pandemic which contributed to higher levels of short-term inefficiency, but no lasting impact that contributed to long-term inefficiency. This research into the significant effect that systemic events can have on market efficiency further motivates the need for additional research into other systemic events, such as the forthcoming analysis of the Ethereum Merge and the collapse of FTX.

The existing literature has largely agreed upon the presence of weak form market efficiency for the market leaders of bitcoin and ether. However, there is not a current, definitive determination of market efficiency in line with traditional markets for other cryptocurrencies which primarily exist as Ethereum-based tokens. Market efficiency research in the cryptocurrency market is far behind traditional markets in terms of the frequency of studied observations. The existing literature has investigated macro-economic systemic events such as the COVID-19 pandemic, but there has not been research performed on the impact of systemic events within the cryptocurrency market.

3 Data Discussion

The primary focus of this research is to observe how market efficiency is impacted by the systemic event of the Ethereum Merge. Data will be collected for bitcoin, ether, and a sample of 12 Ethereum-based cryptocurrencies. Price data of each investigated asset will be obtained[6] for every approximate 30 minute period from August 2022 through October 2022. 30 minutes is found to be the highest frequency for which existing price data can be reliably obtained for an array of cryptocurrency assets. This data period is selected to center the research around investigating the state of the Ethereum-based cryptocurrency market's efficiency during the time of the Ethereum Merge which occurred on September 15, 2022. The collected data will be separated into

[6] Data is collected from DeFi Llama's historical API (https://defillama.com/docs/api) which aggregates from a number of price sources.

three periods: the pre-Merge period of August, the Merge period of September, and the post-Merge period of October. Bitcoin price and return data is collected for the purpose of comparison as bitcoin holds a dominant share of the total market capitalization of all cryptocurrencies.

The failure of FTX, a popular cryptocurrency exchange, in early November 2022 resulted in widespread market inefficiencies across the cryptocurrency market as demonstrated in the forthcoming analysis. For this reason, the pre-merge and post-merge period are restricted to one month periods to prevent confounding effects on the investigation of market efficiency around the Ethereum Merge. The effect of FTX's collapse on market efficiency will be performed on a separate data set for the same assets during the period of November 2022. These findings will be presented after the results of the Ethereum Merge and will be compared to highlight the impact of unknown systemic risk events on market efficiency.

The 12 cryptocurrencies are selected based upon a set of evaluative criteria. Firstly, the cryptocurrencies must be primarily focused on the Ethereum blockchain and be tradable as an ERC-20 token that services an application built on top of the Ethereum blockchain. Secondly, the cryptocurrency cannot be of an intended fixed value such as a token representing the US dollar.[7]

Finally, three groups of total market capitalization will be considered: small-cap with a total approximate market capitalization between \$200 million and \$500 million, mid-cap with a total approximate market capitalization between \$500 million and \$1 billion, and large-cap with a total approximate market capitalization above \$1 billion.

[7] "Stablecoins" representing the value of one US dollar do not power an Ethereum-based application and do not follow a true market determined pricing mechanism like non-stablecoins exhibit due to a constantly offered redemption value of 1 US dollar by the issuers of the various prominent stablecoins.

Cryptocurrencies will be assigned the market group for which the asset most commonly resides during the data collection period. Market capitalization groups will be implemented to ensure a representative sample of tokens ranging in prominence and available liquidity which Wei (2018) discovered to have an impact on the outcome of efficiency testing. Four cryptocurrencies from each of the three groups will be investigated for market efficiency with respect to the cross-time returns of itself through the use of multiple market efficiency tests.

3.1 Small-Cap Ethereum-Based Tokens

The small-cap group considers ERC-20 tokens with a market capitalization between $200 million and $500 million. The four small-cap cryptocurrencies to be investigated are Basic Attention Token (BAT), Compound (COMP), Ethereum Name Service (ENS), and yearn.finance (YFI). While denoted as “small-cap” cryptocurrencies, these assets and the protocols they power are well established within the Ethereum ecosystem as all four protocols have existed for multiple years and have maintained their market capitalization up to the date of publication for this investigation (CoinMarketCap).

3.1.1 Small-Cap Ethereum-Based Token Background

Basic Attention Token is the native token of the Brave internet browser. Brave is a privacy-focused internet browser that compensates users for browser-based advertisements with a portion of the Basic Attention Tokens paid by the advertiser.

Compound is a decentralized lending protocol that enables Ethereum-based lending and borrowing services. Users are able to borrow against their over-collateralized funds and are charged an algorithmically-determined interest rate until the loan is repaid or the collateral is liquidated. The Compound token grants voting rights in the governance of the protocol concerning decisions such as interest rate changes and new assets to support.

Ethereum Name Service is an Ethereum-based naming protocol that allows users to obtain unique strings that represent the user's public address. Rather than having a randomly-generated Ethereum address of a 42 character hexadecimal string such as "0xd8dA6BF26964aF9D7eEd9e03E53415D37aA96045", a chosen string such as "thesis.eth" can be used to transact on the network. The Ethereum Name Service token is used to vote in the governance of the protocol and its decentralized autonomous organization.

Yearn.finance is an Ethereum-based protocol which autonomously reallocates user deposits to the highest-earning interest rates available by various Ethereum-based lending protocols. The yearn.finance token is a governance token that enables owners to propose and vote on decisions regarding changes to the protocol.

3.1.2 Small-Cap Ethereum-Based Token Statistics

Figure 1: Small-Cap Cryptocurrency Summary Statistics (Half-Hour) 8/1/22 - 10/31/22

		ETH	BTC	BAT	COMP	ENS	YFI
	Observations	4,243	4,372	3,777	3,977	4,150	4,116
	Mean	1520.13	20664.85	0.33	55.40	15.45	9102.57
	St. Dev.	196.99	1687.85	0.05	5.65	1.96	1278.09
Price	Coeff. of Var.	0.13	0.08	0.16	0.10	0.13	0.14
	Minimum	1212.44	18301.49	0.27	43.93	11.28	7275.98
	Maximum	2016.50	25093.41	0.47	68.22	20.14	12209.67
	Mean	0.000	-0.002	-0.006	-0.001	0.007	-0.005
	St. Dev.	0.52	0.37	0.56	0.65	0.79	0.64
Return (%)	Minimum	-4.42	-3.94	-4.92	-6.20	-10.32	-7.08
	Maximum	4.56	3.14	5.33	9.03	6.81	11.38

Figure 1 shown above provides insight into the small-cap token data set, portraying summary statistic values for both the prices and returns calculated from the approximate 30 minute observations of the four “small” market capitalization tokens along with ether and bitcoin for comparison during the full Ethereum Merge sample period from August 2022 through October 2022. All six assets have a near-zero mean return, but a noticeable difference in variance of returns is present with all four small-cap tokens having a greater standard deviation of returns than that of ether and bitcoin. Bitcoin, the market leader in terms of total market capitalization, is the least volatile across the period, with the smallest coefficient of variation for its price and the smallest standard deviation of returns.

Figure 2: Small-Cap Cryptocurrency Prices 8/1/22 - 10/31/22, Indexed to 100

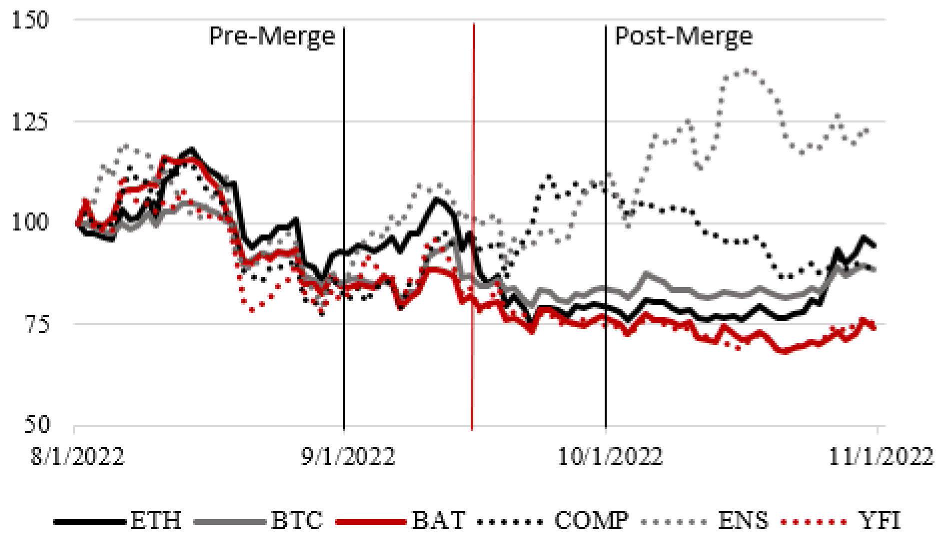

Figure 2 shown above displays the price movement of ether, bitcoin, and the four small-cap tokens across the studied Ethereum Merge period from August 2022 through October 2022. Prices are indexed to 100 starting at the beginning of the period. The black vertical lines differentiate between the pre-merge, merge, and post-merge periods. The red vertical line indicates the date of the Ethereum Merge: September 15. The assets are shown to move closely together during the pre-merge period, with a slight divergence happening during the merge and post-merge periods, before moving together again at the end of the post-merge period.

3.2 Mid-Cap Ethereum-Based Tokens

The mid-cap group considers ERC-20 tokens with a market capitalization between $500 million and $1 billion. The four mid-cap cryptocurrencies to be investigated are Aave (AAVE), Maker (MKR), Synthetix (SNX), and The Graph (GRT).

3.2.1 Mid-Cap Ethereum-Based Token Background

Aave is a decentralized Ethereum ecosystem-based liquidity protocol that enables collateralized lending and pioneered no-collateral “flash loans” which are blockchain-based uncollateralized loans that are instantaneously borrowed and repaid. A common flash loan use case is the arbitraging of a cryptocurrency between two decentralized cryptocurrency exchanges. The Aave token is a governance token that grants voting rights in the governance of the protocol.

Maker is an Ethereum-based protocol that primarily oversees the DAI stablecoin. Users can borrow the DAI stablecoin against cryptocurrency collateral such as ether. The DAI stablecoin is the most popular crypto-collateralized stablecoin which maintains a “soft-peg” to the US dollar. Collateral must be worth at least 150% of the borrowed DAI at all times or the collateral will be liquidated to maintain the integrity of the protocol and DAI stablecoin. The Maker token is a governance token that grants voting rights in the governance of the protocol.

Synthetix is an Ethereum-based decentralized protocol that enables on-chain ownership and trading of synthetic assets, commonly known as derivatives in traditional finance. Potential on-chain derivatives range from gold to cryptocurrency indexes to inverse cryptocurrencies. Synthetix uses price oracles to constantly update the protocol with real-time, accurate asset prices. The Synthetix token entitles owners to a share of trading fees collected on the protocol.

The Graph is an Ethereum-based indexing protocol for data found on blockchain-based applications. New information is stored in Subgraphs which can be queried by users in exchange for the protocol's native token, which is paid to those that compiled the information.

3.2.2 Mid-Cap Ethereum-Based Token Statistics

Figure 3: Mid-Cap Cryptocurrency Summary Statistics (Half-Hour) 8/1/22 - 10/31/22

		ETH	BTC	AAVE	MKR	SNX	GRT
	Observations	4,243	4,372	4,267	4,068	4,121	4,091
	Mean	1520.13	20664.85	85.82	871.20	2.83	0.11
	St. Dev.	196.99	1687.85	10.66	147.03	0.61	0.02
Price	Coeff. of Var.	0.13	0.08	0.12	0.17	0.22	0.17
	Minimum	1212.44	18301.49	65.52	588.65	1.89	0.08
	Maximum	2016.50	25093.41	115.51	1182.08	4.37	0.16
	Mean	0.000	-0.002	-0.002	-0.003	-0.006	-0.008
	St. Dev.	0.52	0.37	0.60	0.58	0.73	0.64
Return (%)	Minimum	-4.42	-3.94	-4.53	-8.09	-7.09	-4.70
	Maximum	4.56	3.14	5.04	3.95	7.80	6.41

Figure 3 provides summary statistic values for the mid-cap token data set across the Ethereum Merge sample period. All six assets have a near-zero mean return, but the noticeable difference in variance of returns is also present for the mid-cap tokens with all four mid-cap tokens having a greater standard deviation of returns than that of ether and bitcoin. Bitcoin remains the least volatile across the period, with the smallest coefficient of variation for its price and the smallest standard deviation of returns.

Figure 4: Mid-Cap Cryptocurrency Prices 8/1/22 - 10/31/22, Indexed to 100

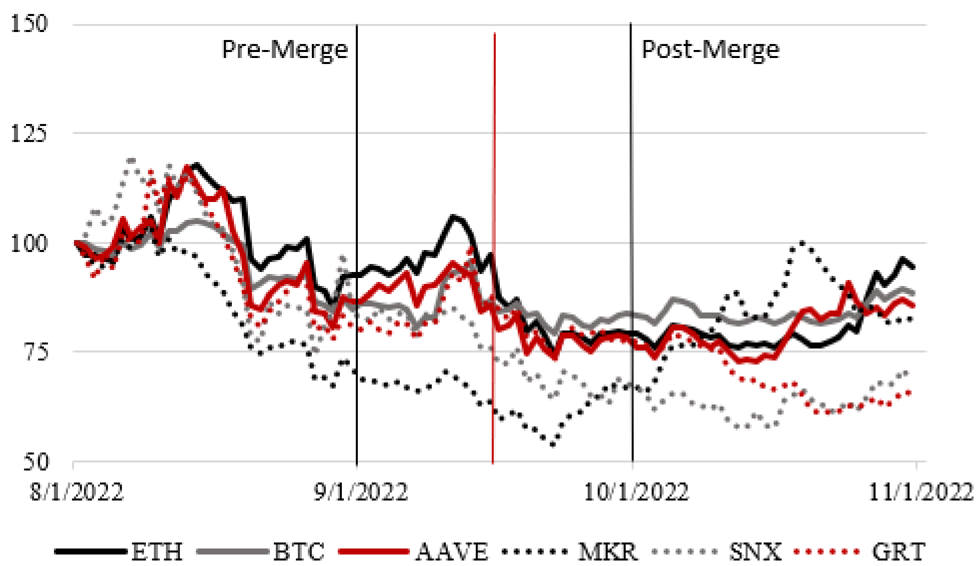

Figure 4 displays the price movement of ether, bitcoin, and the four mid-cap tokens across the studied Ethereum Merge period from August 2022 through October 2022. The assets move more closely in line than the small-cap tokens, but the assets do diverge slightly after the pre-merge period.

3.3 Large-Cap Ethereum-Based Tokens

The large-cap group considers ERC-20 tokens with a market capitalization above $1 billion. The four large-cap cryptocurrencies to be investigated are Polygon (MATIC), Uniswap (UNI), Chainlink (LINK), and Filecoin (FIL).

3.3.1 Large-Cap Ethereum-Based Token Background

Polygon is a blockchain built to serve as a “layer two blockchain” for Ethereum. Because high levels of transaction demand can cause delayed transactions and wildly increased fees for transacting on the Ethereum blockchain, other blockchains such as Polygon serve as a secondary settlement layer for Ethereum-based transactions. The technology is largely the same and users are able to move funds between Ethereum and Polygon with ease. Polygon's native cryptocurrency is used to pay for transaction fees on the Polygon blockchain.

Uniswap is the largest decentralized cryptocurrency exchange on the Ethereum blockchain. When transacting on a decentralized exchange users interact with automatic market makers that "swap" assets held in pooled deposits of other users. Users conducting an exchange pay a small fee which is used to compensate those depositing funds into the trading pools. Users maintain custody of their assets throughout the exchange process. The Uniswap token is a governance token that grants voting rights in the governance of the protocol.

Chainlink is a decentralized oracle network built on the Ethereum blockchain. An oracle brings off-chain data onto the blockchain for use in various blockchain-based applications. Users in need of a data feed from the Chainlink network pay the oracle servicer with the Chainlink cryptocurrency.

Filecoin is a blockchain-based decentralized peer-to-peer file storage network. Those in need of storage space can go to Filecoin network and pay with the Filecoin cryptocurrency. Those with extra storage space can provide storage space to the network in exchange for payment of Filecoin cryptocurrency.

3.3.2 Large-Cap Ethereum-Based Token Statistics

Figure 5: Large-Cap Cryptocurrency Summary Statistics (Half-Hour) 8/1/22 - 10/31/22

		ETH	BTC	MATIC	UNI	LINK	FIL
	Observations	4,243	4,372	4,328	4,205	4,102	4,126
	Mean	1520.13	20664.85	0.85	6.82	7.48	6.24
	St. Dev.	196.99	1687.85	0.07	1.03	0.65	1.25
Price	Coeff. of Var.	0.13	0.08	0.08	0.15	0.09	0.20
	Minimum	1212.44	18301.49	0.70	5.20	6.23	4.76
	Maximum	2016.50	25093.41	1.05	9.54	9.43	11.19
	Mean	0.000	-0.002	0.001	-0.002	0.002	-0.012
	St. Dev.	0.52	0.37	0.52	0.65	0.59	0.68
Return (%)	Minimum	-4.42	-3.94	-4.20	-3.97	-8.73	-7.06
	Maximum	4.56	3.14	3.12	3.93	4.06	6.24

Figure 5 provides summary statistic values for the large-cap token data set across the Ethereum Merge sample period. The difference between the standard deviation of returns for the large-cap tokens and that of the market leaders of bitcoin and ether is the tightest out of the three market capitalization groups.

Figure 6: Large-Cap Cryptocurrency Prices 8/1/22 - 10/31/22, Indexed to 100

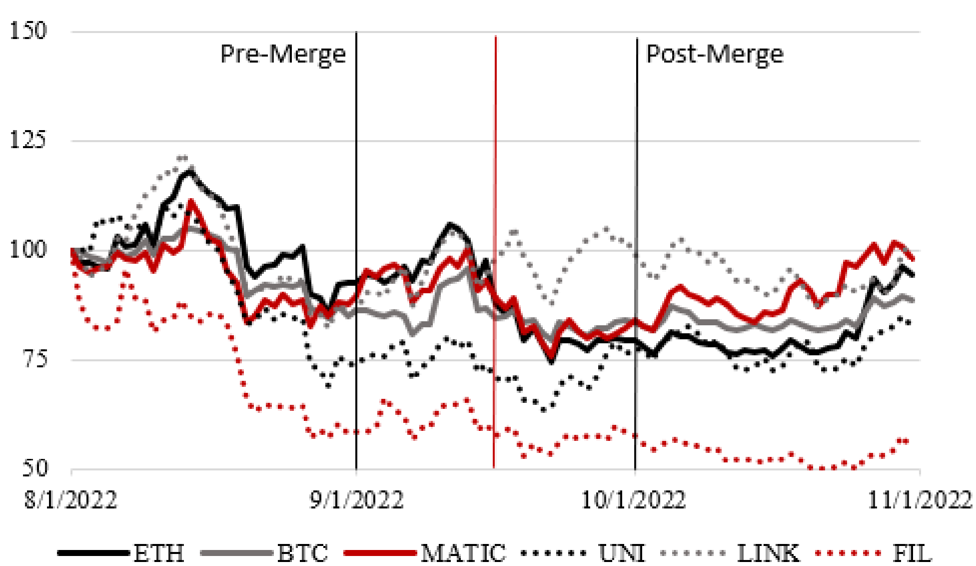

Figure 6 displays the price movement of ether, bitcoin, and the four large-cap tokens across the studied Ethereum Merge period from August 2022 through October 2022. The assets move fairly closely in line with each other, but the large-cap tokens are observed to move lower and higher than bitcoin and ether across most of the period, demonstrating greater volatility and supporting their greater standard deviations of returns compared to bitcoin and ether.

4 Methodology

A number of tests persist throughout the existing literature as researchers are working to reach consensus on the efficiency of the cryptocurrency market. The following tests will primarily be applied to the data set of approximate 30 minute observations of prices and returns for ether, bitcoin, and the three groups of small-cap, mid-cap, and large-cap tokens taken from the Ethereum Merge period from August 2022 to October 2022. The methods will be applied to the full period, the pre-merge period of August, the merge period of September, and the post-merge period of October. Findings will be compared to investigate how market efficiency changed across this period as a systemic event occurred. Then, the methods will be applied to the period of November 2022 to investigate the impact of an unknown systemic event, the collapse of FTX, on market efficiency.

The Ljung-Box (Ljung and Box, 1978) test is a test for serial correlation in a time series data set. The test examines randomness across the time series concerning a user-defined number of lags. In its application to cryptocurrency returns, a lag of one period is implemented. The Ljung-Box test has a null hypothesis of no serial correlation where the time series is independent of cross-time observations.

The Runs (Wald and Wolfowitz, 1940) test is a test for independence in a time series data set. The test has a null hypothesis of a randomly-determined distribution. In its application to asset returns, the Runs test is concerned with runs of observations above or below the median value. The Runs test can uncover dependence across observations and highlight the presence of non-random streaks in the data set.

The Lo-MacKinlay variance ratio test (Lo and MacKinlay, 1988) is used to determine if the time series of returns follows a random walk with allowed drift by conducting an overlapping variance ratio test on a time series data set. The natural log of the price values for assets is used. A failure to reject the null hypothesis of the data exhibiting a random walk with drift results in a conclusion that there is not statistically significant evidence to claim that the marketplace for this asset is inefficient. Values for the overlapping span of differencing in the conducted testing are 2, 12, 24, and 48 observations to observe how market efficiency evolves across additional time from an observation.

5 Results

Simplified results of the performed methods are shown below. The full results of all conducted tests are shown in the Appendix. Results are presented with conditional formatting to portray lower p-values with red coloration and higher p-values with green coloration.

Figure 7: Small-Cap Cryptocurrency Methodology Results 8/1/22 - 10/31/22

		ETH	BTC	BAT	COMP	ENS	YFI
Full Period	Ljung-Box Test	0.00	0.07	0.17	0.09	0.03	0.01
	Runs Test (two-tailed)	0.65	0.15	0.33	0.08	0.23	0.37
	Lo-MacKinlay VRT (q=24)	0.02	0.06	0.48	0.75	0.74	0.79
Pre-Merge	Ljung-Box Test	0.50	0.57	0.05	0.39	0.12	0.20
	Runs Test (two-tailed)	0.94	0.76	0.19	0.56	0.78	0.38
	Lo-MacKinlay VRT (q=24)	0.06	0.49	0.11	0.80	0.37	0.55
Merge	Ljung-Box Test	0.00	0.22	0.83	0.00	0.06	0.00
	Runs Test (two-tailed)	0.65	0.19	0.28	0.63	0.83	0.27
	Lo-MacKinlay VRT (q=24)	0.23	0.01	0.98	0.64	0.46	0.80
Post-Merge	Ljung-Box Test	0.05	0.22	0.86	0.36	0.72	0.21
	Runs Test (two-tailed)	0.60	0.45	0.11	0.04	0.55	0.10
	Lo-MacKinlay VRT (q=24)	0.12	0.66	0.87	0.56	0.29	0.68

Figure 7 features the results of the Ljung-Box test, Runs test, and Lo-Mackinlay Variance Ratio test for ether, bitcoin, and the four small-cap tokens of the data set for the full Ethereum Merge period of August 2022 through October 2022 along with the results for each of the sub-periods: pre-merge, merge, and post-merge. While only one asset demonstrates a Ljung-Box result that supports a significant rejection of the null of no present serial correlation in both the pre-merge and post-merge periods, three assets demonstrate a significant rejection of the null of no serial correlation for the merge period, including ether.

Apart from one result during the post-merge period for COMP (Compound), no other asset from the group of ether, bitcoin, and the small-cap tokens present a significant rejection of the Runs test's null of a randomly determined distribution.

The Lo-Mackinlay Variance Ratio test yields results that support the existing literature. At lower spans of differencing between observations (q=2, as shown in the Appendix) for this data set of high frequency observations at the 30 minute level, we observe mixed results of efficiency for assets across the different periods analyzed. However, as the span of differencing increases to periods more in line with the existing literature such as the daily level, nearly all assets fail to reject the null hypothesis of a random walk with drift. One interesting observation is that for the full period across the Ethereum Merge, ether is the only asset to reject the Lo-Mackinlay Variance Ratio Test at the level of 24 observations apart, but it does not reject the null in any of the sub-periods. Another interesting observation is a fairly pervasive fluctuation in p-values as the span of differencing increases, while one would expect a nearly strict increase in p-value as the span of differencing increases.

These results indicate that the observed small-cap assets, ether, and bitcoin exhibit the most conclusive evidence of market inefficiency during the merge period, compared to the pre-merge and post-period. These observed inefficiencies drive many of the observed results that show inefficiency across the full period.

Figure 8: Mid-Cap Cryptocurrency Methodology Results 8/1/22 - 10/31/22

		ETH	BTC	AAVE	MKR	SNX	GRT
Full Period	Ljung-Box Test	0.00	0.07	0.00	0.20	0.00	0.18
	Runs Test (two-tailed)	0.65	0.15	0.31	0.00	0.00	0.17
	Lo-MacKinlay VRT (q=24)	0.02	0.06	0.43	0.26	0.00	0.74
Pre-Merge	Ljung-Box Test	0.50	0.57	0.36	0.86	0.00	0.71
	Runs Test (two-tailed)	0.94	0.76	0.43	0.01	0.03	0.92
	Lo-MacKinlay VRT (q=24)	0.06	0.49	0.20	0.77	0.00	0.11
Merge	Ljung-Box Test	0.00	0.22	0.07	0.06	0.00	0.01
	Runs Test (two-tailed)	0.65	0.19	0.47	0.41	0.22	0.04
	Lo-MacKinlay VRT (q=24)	0.23	0.01	0.61	0.15	0.97	0.86
Post-Merge	Ljung-Box Test	0.05	0.22	0.06	0.09	0.21	0.37
	Runs Test (two-tailed)	0.60	0.45	0.89	0.02	0.07	0.39
	Lo-MacKinlay VRT (q=24)	0.12	0.66	0.67	0.79	0.06	0.27

Figure 8 features the results of the Ljung-Box test, Runs test, and Lo-Mackinlay variance ratio test for ether, bitcoin, and the four mid-cap tokens of the data set for the full Ethereum Merge period of August 2022 through October 2022 along with the results for each of the sub-periods: pre-merge, merge, and post-merge. Similar to the results for the small-cap tokens, the merge period is the most popular period for assets in this group to reject the null hypothesis of the Ljung-Box test, favoring the greatest prevalence of market inefficiency during this period.

The Runs test provides significant rejections of the null hypothesis of a randomly-determined distribution for the full period and at least one sub-period for two of the four mid-cap tokens: MKR (Maker) and SNX (Synthetix). Interestingly, this is more than observed with the small-cap token group.

The Lo-Mackinlay variance ratio test yields similar results to those found with the small-cap group, further supporting the existing literature's findings of general efficiency at higher frequencies such as the daily level. Additionally, these results also provide support to the determination of mixed efficiency at the studied frequency of 30 minutes.

The sub-period occurrences of market inefficiency results being concentrated in the sub-period of the merge is similar to that found with the small-cap group. However, while the mid-cap assets have greater market capitalizations than the small-cap group and therefore typically greater liquidity, this investigation finds more rejections of market efficiency with the mid-cap group than with the small-cap group. This finding raises questions about the relationship between liquidity and market efficiency discussed in the literature.

Figure 9: Large-Cap Cryptocurrency Methodology Results 8/1/22 - 10/31/22

		ETH	BTC	MATIC	UNI	LINK	FIL
Full Period	Ljung-Box Test	0.00	0.07	0.00	0.01	0.00	0.31
	Runs Test (two-tailed)	0.65	0.15	0.84	0.41	0.22	0.00
	Lo-MacKinlay VRT (q=24)	0.02	0.06	0.07	0.10	0.34	0.46
Pre-Merge	Ljung-Box Test	0.50	0.57	0.05	0.56	0.00	0.55
	Runs Test (two-tailed)	0.94	0.76	0.87	0.23	0.07	0.25
	Lo-MacKinlay VRT (q=24)	0.06	0.49	0.23	0.42	0.09	0.16
Merge	Ljung-Box Test	0.00	0.22	0.00	0.17	0.62	0.08
	Runs Test (two-tailed)	0.65	0.19	0.44	0.25	0.70	0.07
	Lo-MacKinlay VRT (q=24)	0.23	0.01	0.29	0.42	0.81	0.76
Post-Merge	Ljung-Box Test	0.05	0.22	0.83	0.04	0.20	0.41
	Runs Test (two-tailed)	0.60	0.45	0.14	0.54	0.91	0.01
	Lo-MacKinlay VRT (q=24)	0.12	0.66	0.16	0.03	0.19	0.42

Figure 9 features the results of the Ljung-Box test, Runs test, and Lo-Mackinlay variance ratio test for ether, bitcoin, and the four large-cap tokens of the data set for the full Ethereum Merge period of August 2022 through October 2022 along with the results for each of the sub-periods: pre-merge, merge, and post-merge. Similar to the results for the small-cap and mid-cap group, the merge period is the most popular period for assets in this group to reject the null of the Ljung-Box test, favoring the greatest prevalence of market inefficiency during this period.

Most assets of this group fail to reject the null hypothesis of a randomly-determined distribution of the Runs test, except for FIL (Filecoin) which rejects the null hypothesis in the post-merge period and the full period.

The findings of the Lo-Mackinlay variance ratio test are in line with those of the other groups, with a general finding of mixed results at the studied high frequency and an overall trend towards increased efficiency as the span of differencing increases. The interesting finding of fluctuating results as the span of differencing increases persists in the large-cap group as well.

5.1 Exploring Unexpected Systemic Events: FTX Collapse

The collapse of FTX in early November 2022 was another systemic event that had a broad impact on the cryptocurrency market. Unlike the Ethereum Merge which had a date of occurrence known by the market well in advance, the collapse of FTX was widely unexpected by the market. It is also important to investigate the impact on market efficiency of unknown systemic events and to compare these findings to those of a known systemic event. Additionally, the discussed unregulated nature of the cryptocurrency industry increases the likelihood of similar events concerning institutions of the cryptocurrency market occurring in the future as the cryptocurrency space continues to develop. Therefore, the presented methods are applied to the period of November 2022 for approximate 30 minute observation periods for the assets previously explored.

Figure 10: Select Cryptocurrency Prices 11/1/22 - 11/31/22, Indexed to 100

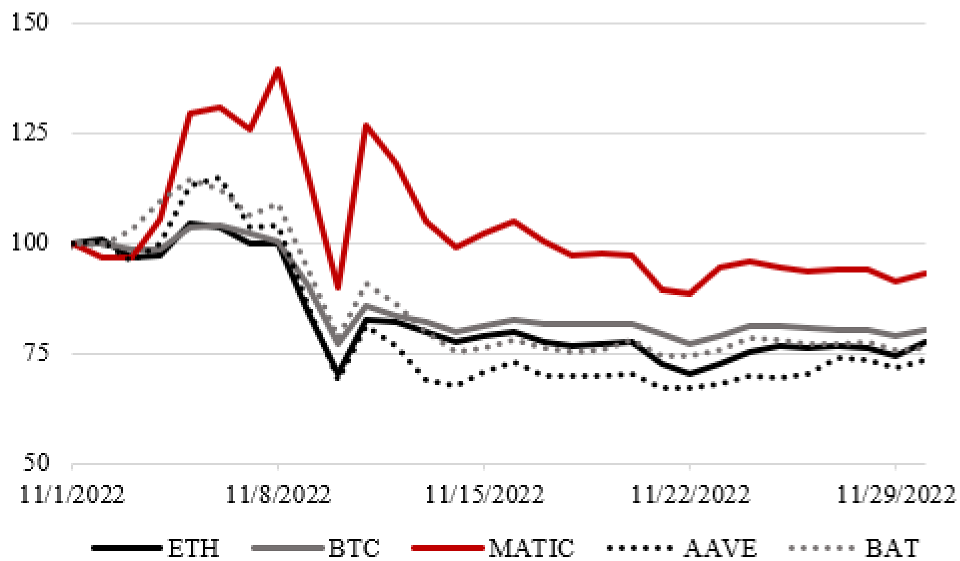

Figure 10 shown above, with indexed prices for a subset of cryptocurrencies of the full sample, illustrates the impact of the news and resulting fallout of the collapse of FTX in early November 2022. The steep, but not instantaneous, drop and following volatility displayed in the figure in response to the news of FTX's collapse creates a data set that is unlikely to demonstrate market efficiency. Under the Efficient Market Hypothesis, the market would instantaneously react to the new information of FTX and retain market efficiency. However, as the following results will demonstrate, the market failed to maintain weak form efficiency in response to the collapse of FTX.

Figure 11: FTX Collapse: Cryptocurrency Methodology Results 11/1/22 - 11/31/22

	ETH	BTC	MATIC	UNI	LINK	FIL
Ljung-Box Test	0.00	0.00	0.00	0.00	0.00	0.00
Runs Test (two-tailed)	0.34	0.06	0.56	0.17	0.13	0.64
Lo-MacKinlay VRT (q=24)	0.81	0.17	0.14	0.93	0.53	0.90

	ETH	BTC	AAVE	MKR	SNX	GRT
Ljung-Box Test	0.00	0.00	0.00	0.00	0.00	0.00
Runs Test (two-tailed)	0.34	0.06	0.04	0.67	0.05	0.56
Lo-MacKinlay VRT (q=24)	0.81	0.17	0.87	0.08	0.94	0.81

	ETH	BTC	BAT	COMP	ENS	YFI
Ljung-Box Test	0.00	0.00	0.00	0.00	0.00	0.00
Runs Test (two-tailed)	0.34	0.06	0.34	0.00	0.04	0.28
Lo-MacKinlay VRT (q=24)	0.81	0.17	0.75	0.37	0.67	0.89

Figure 11 features the results of the Ljung-Box test, Runs test, and Lo-Mackinlay variance ratio test for ether, bitcoin, and the three market capitalization groups of small-cap, mid-cap, and large-cap tokens for the month of November 2022. The results of the Ljung-Box test highlight an immediate difference compared to the previous findings focused on the Ethereum Merge. All 14 assets present a significant rejection of the null hypothesis of no serial correlation. This demonstrates significant market inefficiency as a result of serial correlation between returns.

The Runs Test portrays greater market efficiency resiliency when compared to the drastic findings of the Ljung-Box test, but there are still a number of assets that reject the null hypothesis of a randomly-determined distribution during the FTX collapse observation period.

Nearly all assets reject the null hypothesis of a random walk with drift of the Lo-Mackinlay variance ratio test at low levels of observation differencing. However, all assets recover to efficient levels once the span of differencing reaches 24 observations. Interestingly, this was not the case with the same assets during the three months across the Ethereum Merge.

6 Conclusion

This investigation explores the high frequency market efficiency of established Ethereum-based cryptocurrencies across systemic events. The tests conducted operate with a hypothesis of weak form market efficiency, meaning that prior return information cannot be used to reliably future returns.

The results indicate that observable levels of weak form market efficiency decrease during times of systemic events such as the Ethereum Merge where most results supporting market inefficiency were discovered in the “during merge” period, including an inefficiency in ether itself.

Compared to the results of the Ethereum Merge, the results of tests conducted on the November data set focused on the collapse of FTX yielded more significant results that supported market inefficiencies at the studied high frequency. This finding could be potentially attributed to a few things, including the severity of the FTX collapse and the unknown factor of this systemic event compared to the known occurrence of the Ethereum Merge.

Contributions to the existing literature include a step taken towards increasing the highest studied frequency for cryptocurrency market efficiency beyond the one hour mark to 30 minute observations, an expansion of the set of investigated cryptocurrencies for market efficiency at the highest studied frequency, and the investigation of the impact on market efficiency by two systemic events: the Ethereum Merge and the collapse of FTX.

6.1 Discussion of Limitations

While data can be reliably obtained for approximately every 30 minutes, these observations do not occur at the same exact time for all assets as a result of available data-provider collection methods. Due to the volatility and ever-changing nature of asset prices in contemporary markets, the loss of some price information and fluctuation between observations is a potential limitation to fully understanding market efficiency.

The data set used to achieve the presented results is gathered from historical sources of calculated market price averages. Unlike the stock market where trades flow through a select few entities with reliable data flows, the cryptocurrency market is far more decentralized with trades flowing through many decentralized entities and centralized entities with no obligation to report price and volume information, accurately or at all.

From this, one potential limitation arises from the calculation of the “market price” by data providers. Market prices for cryptocurrencies are typically calculated with a volume-weighted formula that considers the average price and volume of an asset’s trades on each exchange that provides its data during a given time period. This formula provides a representative price of an asset, but as this is an average it could very likely be the case that the price used in market efficiency testing is not the exact price at which one can actually buy or sell an asset at any marketplace.

Another limitation that arises from the data-reporting systems of the cryptocurrency market comes from the inherent trust placed in marketplaces to accurately report their data and the lack of oversight on the accuracy of data reporting. An exchange might be incentivized to alter its price and/or volume information to appear more prevalent or appealing to market participants in the hopes of gaining more users. This would falsely alter the calculation of an asset's market price and could impact the results of market efficiency testing.

6.2 Discussion of Continued Research

6.2.1 High Frequency

The highest frequency achievable in searching for data to address the goals of this research was found to be the 30 minute level. At frequencies of shorter timing than 30 minutes, the availability of observations particularly for non-market leading cryptocurrencies becomes unreliable in existing data sources.

As the cryptocurrency market continues to develop and data becomes increasingly available and reliable, further research can extend the literature to shorter frequencies to eventually match the high frequencies at which traditional markets are studied. Additionally, more cryptocurrencies can be studied to gain a more holistic understanding of efficiency in the cryptocurrency market.

Alternatively, an approach of collecting real-time price data from one exchange can be implemented to achieve datasets for higher frequencies. However, one potential drawback of this approach is the loss of all other price and volume information that occurs on other marketplaces which could skew results.

6.2.2 Continued Analysis of Systemic Events

The findings of decreased levels of market efficiency during periods of systemic events when compared to other periods of similarly collected data supports the need for additional research into how the cryptocurrency market reacts to systemic events. If the cryptocurrency market is to continue its growth to become a fixture of the world's asset markets and receive widespread institutional interest, it must maintain efficiency, like its traditional counterparts, even in times of events that have a broad impact on the cryptocurrency market or economy as a whole.

6.2.3 Ether Dependence

The presented results highlight that the cryptocurrency market is still growing into becoming broadly efficient across its many assets, particularly at higher frequencies. As the market develops and assets become broadly weak form efficient with respect to their own prior returns, research into cryptocurrency return independence from market leaders such as bitcoin and ether should be investigated.

The complete hosting of all technological functions by the Ethereum blockchain for tokens built on top of it creates a unique relationship. Tokens created on Ethereum are entirely reliant on the functionality of Ethereum. If the Ethereum blockchain were to "disappear", so would all Ethereum-based tokens. If an event occurred that (positively or negatively) influenced the ability to transact ether on the Ethereum blockchain, it would (in all likelihood) similarly impact Ethereum-based tokens. The heavy, intrinsic reliance on the single entity of the Ethereum blockchain by Ethereum-based tokens is not readily observed between assets in traditional markets.

For example, publicly traded companies do not exhibit the described relationship between Ethereum and Ethereum-based tokens. The business models of publicly traded companies are diversified enough to eliminate the complete reliance on another firm. If this relationship were to be observed between the operations of two companies, the diversity of business verticals could dilute and confound the impact of the technology-providing firm's performance on the performance of reliant firm. In the case of Ethereum and the tokens built on top of it, ether most clearly represents the ability to use the technology of Ethereum and the tokens represents the use case of the project that created itself on the Ethereum blockchain.

For the Ethereum-based cryptocurrencies that are found to be efficient with respect to their own returns, the consideration of ether's prior period return can be added to existing market efficiency tests. The popular Ljung-Box and Lo-MacKinlay tests can be modified to use the lagged time series of ether's returns in place of the cryptocurrency's lagged returns. The null hypotheses of these tests remain the same, all with the expectation of market efficiency.

An autoregressive model can also be implemented to test the joint significance of both the cryptocurrency's prior period return and ether's prior period return for an Ethereum-based cryptocurrency that is suspected of being influenced by the returns of ether. This model takes the following general form:

$$\text{TokenReturn}_t = B_0 + B_1\text{TokenReturn}_{t-1} + B_2\text{EtherReturn}_{t-1} + u_t$$

Additional controls could be introduced such as time between observations if perfectly spaced data is unavailable. Under the adapted standard tests for market efficiency, both B_1 and B_2 are expected to be statistically insignificant. Under the theory of the Efficient Market Hypothesis, these coefficients are equal to zero. The Efficient Market Hypothesis also stipulates that a test for joint significance between the cryptocurrency's prior return and ether's prior return yields no ability to predict the cryptocurrency's current period return. An ability to use a combination of prior returns to reliably predict the current period return of an asset violates the theory of weak form market efficiency.

References

Baum, Christopher F. 2006. "LOMACKINLAY: Stata module to perform Lo-MacKinlay variance ratio test," Statistical Software Components S456740, Boston College Department of Economics, revised 11 Mar 2019. https://ideas.repec.org/c/boc/bocode/s456740.html.

Buterin, Vitalik. 2014. “Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform.” https://ethereum.org/669c9e2e2027310b6b3cdce6e1c52962/Ethereum_Whitepaper_-_Buterin_2014.pdf.

Chow, Andrew. 2022. “Where Did FTX's Missing $8 Billion Go? Crypto Investigators Offer New Clues.” *Time*. https://time.com/6243086/ftx-where-did-money-go/.

Chu, Jeffrey, Yuanyuan Zhang, and Stephen Chan. 2019. "The adaptive market hypothesis in the high frequency cryptocurrency market." *International Review of Financial Analysis* 64: 221-231. https://doi.org/10.1016/j.irfa.2019.05.008.

CoinMarketCap. “Global Currency Charts.” https://coinmarketcap.com/charts/.

DeFi Llama. 2022. “DeFiLlama API.” https://defillama.com/docs/api.

Ethereum.org. 2023. “Upgrades: The Merge.” https://ethereum.org/en/upgrades/merge/.

Fama, Eugene F. 1970. "Efficient capital markets: A review of theory and empirical work." *The journal of Finance* 25, no. 2: 383-417. https://doi.org/10.2307/2325486.

Kakinaka, Shinji, and Ken Umeno. 2022. "Cryptocurrency market efficiency in short-and long-term horizons during COVID-19: An asymmetric multifractal analysis approach." *Finance Research Letters* 46: 102319. https://doi.org/10.1016/j.frl.2021.102319

Levmore, Saul. 2002. "Simply Efficient Markets and the Role of Regulation: Lessons from the Iowa Electronic Markets and the Hollywood Stock Exchange." *J. Corp. L,* 28: 589. https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=2604&context=journal_articles.

Ljung, Greta M., and George EP Box. 1978. "On a measure of lack of fit in time series models." *Biometrika* 65, no. 2: 297-303. https://doi.org/10.2307/2335207.

Lo, Andrew W., and A. Craig MacKinlay. 1988. "Stock market prices do not follow random walks: Evidence from a simple specification test." *The review of financial studies* 1, no. 1: 41-66. https://doi.org/10.1093/rfs/1.1.41.

Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System.” https://bitcoin.org/bitcoin.pdf.

Noda, Akihiko. 2021. "On the evolution of cryptocurrency market efficiency." *Applied Economics Letters* 28, no. 6: 433-439. https://doi.org/10.1080/13504851.2020.1758617.

Peneder, Michael. 2008. "The problem of private under-investment in innovation: A policy mind map." *Technovation* 28, no. 8: 518-530. https://doi.org/10.1016/j.technovation.2008.02.006.

Urquhart, Andrew. 2016. "The inefficiency of Bitcoin." *Economics Letters* 148: 80-82. https://doi.org/10.1016/j.econlet.2016.09.019.

Wald, Abraham, and Jacob Wolfowitz. 1940. "On a test whether two samples are from the same population." *The Annals of Mathematical Statistics* 11, no. 2: 147-162. https://doi.org/10.1214/aoms/1177731909.

Wei, Wang Chun. 2018. "Liquidity and market efficiency in cryptocurrencies." *Economics Letters* 168: 21-24. https://doi.org/10.1016/j.econlet.2018.04.003.

Yaya, OlaOluwa S., Ahamuefula E. Ogbonna, Robert Mudida, and Nuruddeen Abu. 2021. "Market efficiency and volatility persistence of cryptocurrency during pre-and post-crash periods of Bitcoin: Evidence based on fractional integration." *International Journal of Finance & Economics* 26, no. 1: 1318-1335. https://doi.org/10.1002/ijfe.1851.

Appendix

Full Methodology Results: Small-Cap Across Ethereum Merge (August-October 2022)

Period	Test	Statistic	ETH	BTC	BAT	COMP	ENS	YFI
Full Period	Ljung-Box Test	coefficient	79.70	53.90	48.29	52.22	57.80	62.30
		p-value	0.000	0.070	0.173	0.093	0.034	0.014
	Runs Test	p-value (two-sided)	0.65	0.15	0.33	0.08	0.23	0.37
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.01	0.05	0.45	0.00
		p-value (q=12)	0.01	0.11	0.21	0.73	0.90	0.95
		p-value (q=24)	0.02	0.06	0.48	0.75	0.74	0.79
		p-value (q=48)	0.05	0.29	0.55	0.42	0.82	0.70
Pre-Merge	Ljung-Box Test	coefficient	39.29	37.71	56.15	41.78	50.58	47.38
		p-value	0.502	0.574	0.046	0.393	0.122	0.197
	Runs Test	p-value (two-sided)	0.94	0.76	0.19	0.56	0.78	0.38
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.00	0.73	0.98	0.02
		p-value (q=12)	0.01	0.45	0.02	0.68	0.84	0.23
		p-value (q=24)	0.06	0.49	0.11	0.80	0.37	0.55
		p-value (q=48)	0.02	0.32	0.63	0.76	0.59	0.37
Merge	Ljung-Box Test	coefficient	75.72	46.42	31.55	72.58	54.50	71.55
		p-value	0.001	0.225	0.828	0.001	0.063	0.002
	Runs Test	p-value (two-sided)	0.65	0.19	0.28	0.63	0.83	0.27
	Lo-MacKinlay VRT	p-value (q=2)	0.19	0.00	0.95	0.02	0.14	0.16
		p-value (q=12)	0.56	0.04	0.72	0.50	0.74	0.38
		p-value (q=24)	0.23	0.01	0.98	0.64	0.46	0.80
		p-value (q=48)	0.59	0.15	0.86	0.34	0.35	0.71
Post-Merge	Ljung-Box Test	coefficient	56.20	46.47	30.58	42.59	34.45	47.02
		p-value	0.046	0.223	0.858	0.360	0.718	0.207
	Runs Test	p-value (two-sided)	0.60	0.45	0.11	0.04	0.55	0.10
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.57	0.78	0.09	0.93	0.26
		p-value (q=12)	0.03	0.55	0.56	0.50	0.81	0.49
		p-value (q=24)	0.12	0.66	0.87	0.56	0.29	0.68
		p-value (q=48)	0.06	0.91	0.50	0.33	0.38	0.38

Full Methodology Results: Mid-Cap Across Ethereum Merge (August-October 2022)

Period	Test	Statistic	ETH	BTC	AAVE	MKR	SNX	GRT
Full Period	Ljung-Box Test	coefficient	79.70	53.90	71.94	47.20	125.50	48.07
		p-value	0.000	0.070	0.001	0.202	0.000	0.179
	Runs Test	p-value (two-sided)	0.65	0.15	0.31	0.00	0.00	0.17
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.00	0.00	0.00	0.79
		p-value (q=12)	0.01	0.11	0.17	0.05	0.00	0.57
		p-value (q=24)	0.02	0.06	0.43	0.26	0.00	0.74
		p-value (q=48)	0.05	0.29	0.82	0.65	0.39	0.78
Pre-Merge	Ljung-Box Test	coefficient	39.29	37.71	42.62	30.65	107.84	34.70
		p-value	0.502	0.574	0.359	0.856	0.000	0.707
	Runs Test	p-value (two-sided)	0.94	0.76	0.43	0.01	0.03	0.92
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.01	0.07	0.00	0.82
		p-value (q=12)	0.01	0.45	0.08	0.46	0.00	0.09
		p-value (q=24)	0.06	0.49	0.20	0.77	0.00	0.11
		p-value (q=48)	0.02	0.32	0.24	0.91	0.06	0.18
Merge	Ljung-Box Test	coefficient	75.72	46.42	53.89	55.00	70.70	63.39
		p-value	0.001	0.225	0.070	0.058	0.002	0.011
	Runs Test	p-value (two-sided)	0.65	0.19	0.47	0.41	0.22	0.04
	Lo-MacKinlay VRT	p-value (q=2)	0.19	0.00	0.20	0.14	0.86	0.59
		p-value (q=12)	0.56	0.04	0.80	0.05	0.64	0.80
		p-value (q=24)	0.23	0.01	0.61	0.15	0.97	0.86
		p-value (q=48)	0.59	0.15	0.88	0.76	0.44	0.42
Post-Merge	Ljung-Box Test	coefficient	56.20	46.47	54.59	52.56	47.09	42.37
		p-value	0.046	0.223	0.062	0.088	0.205	0.369
	Runs Test	p-value (two-sided)	0.60	0.45	0.89	0.02	0.07	0.39
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.57	0.21	0.00	0.01	0.61
		p-value (q=12)	0.03	0.55	0.84	0.60	0.00	0.40
		p-value (q=24)	0.12	0.66	0.67	0.79	0.06	0.27
		p-value (q=48)	0.06	0.91	0.72	0.57	0.30	0.32

Full Methodology Results: Large-Cap Across Ethereum Merge (August-October 2022)

Period	Test	Statistic	ETH	BTC	MATIC	UNI	LINK	FIL
Full Period	Ljung-Box Test	coefficient	79.70	53.90	78.69	63.03	72.87	43.84
		p-value	0.000	0.070	0.000	0.012	0.001	0.312
	Runs Test	p-value (two-sided)	0.65	0.15	0.84	0.41	0.22	0.00
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.00	0.00	0.00	0.14
		p-value (q=12)	0.01	0.11	0.03	0.06	0.03	0.79
		p-value (q=24)	0.02	0.06	0.07	0.10	0.34	0.46
		p-value (q=48)	0.05	0.29	0.34	0.35	0.59	0.26
Pre-Merge	Ljung-Box Test	coefficient	39.29	37.71	55.99	38.12	77.00	38.22
		p-value	0.502	0.574	0.048	0.555	0.000	0.551
	Runs Test	p-value (two-sided)	0.94	0.76	0.87	0.23	0.07	0.25
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.00	0.00	0.02	0.00	0.02
		p-value (q=12)	0.01	0.45	0.14	0.24	0.02	0.29
		p-value (q=24)	0.06	0.49	0.23	0.42	0.09	0.16
		p-value (q=48)	0.02	0.32	0.40	0.38	0.18	0.02
Merge	Ljung-Box Test	coefficient	75.72	46.42	74.44	48.46	36.70	53.04
		p-value	0.001	0.225	0.001	0.169	0.620	0.081
	Runs Test	p-value (two-sided)	0.65	0.19	0.44	0.25	0.70	0.07
	Lo-MacKinlay VRT	p-value (q=2)	0.19	0.00	0.01	0.56	0.02	0.70
		p-value (q=12)	0.56	0.04	0.59	0.67	0.72	0.98
		p-value (q=24)	0.23	0.01	0.29	0.42	0.81	0.76
		p-value (q=48)	0.59	0.15	0.56	0.63	0.91	0.95
Post-Merge	Ljung-Box Test	coefficient	56.20	46.47	31.48	57.46	47.44	41.46
		p-value	0.046	0.223	0.830	0.036	0.195	0.407
	Runs Test	p-value (two-sided)	0.60	0.45	0.14	0.54	0.91	0.01
	Lo-MacKinlay VRT	p-value (q=2)	0.00	0.57	0.00	0.00	0.07	0.98
		p-value (q=12)	0.03	0.55	0.02	0.01	0.05	0.20
		p-value (q=24)	0.12	0.66	0.16	0.03	0.19	0.42
		p-value (q=48)	0.06	0.91	0.23	0.22	0.48	0.40

Full Methodology Results: All Assets During Month of FTX Collapse (November 2022)

			ETH	BTC	MATIC	UNI	LINK	FIL
Full Period	Ljung-Box Test	coefficient	154.53	168.77	134.76	99.52	135.61	88.88
		p-value	0.000	0.000	0.000	0.000	0.000	0.000
	Runs Test	p-value (two-sided)	0.34	0.06	0.56	0.17	0.13	0.64
	Lo-MacKinlay VRT	p-value (q=2)	0.01	0.79	0.00	0.00	0.00	0.03
		p-value (q=12)	0.51	0.04	0.09	0.65	0.71	0.34
		p-value (q=24)	0.81	0.17	0.14	0.93	0.53	0.90
		p-value (q=48)	0.26	0.82	0.04	0.86	0.53	0.40

			ETH	BTC	AAVE	MKR	SNX	GRT
Full Period	Ljung-Box Test	coefficient	154.53	168.77	81.35	130.70	74.56	90.53
		p-value	0.000	0.000	0.000	0.000	0.001	0.000
	Runs Test	p-value (two-sided)	0.34	0.06	0.04	0.67	0.05	0.56
	Lo-MacKinlay VRT	p-value (q=2)	0.01	0.79	0.01	0.06	0.00	0.01
		p-value (q=12)	0.51	0.04	0.96	0.10	0.63	0.83
		p-value (q=24)	0.81	0.17	0.87	0.08	0.94	0.81
		p-value (q=48)	0.26	0.82	0.50	0.00	0.34	0.53

			ETH	BTC	BAT	COMP	ENS	YFI
Full Period	Ljung-Box Test	coefficient	154.53	168.77	70.42	101.72	119.59	67.38
		p-value	0.000	0.000	0.002	0.000	0.000	0.004
	Runs Test	p-value (two-sided)	0.34	0.06	0.34	0.00	0.04	0.28
	Lo-MacKinlay VRT	p-value (q=2)	0.01	0.79	0.00	0.00	0.00	0.02
		p-value (q=12)	0.51	0.04	0.24	0.03	0.20	0.98
		p-value (q=24)	0.81	0.17	0.75	0.37	0.67	0.89
		p-value (q=48)	0.26	0.82	0.73	0.33	0.21	0.77