

The efficiency of bitcoin: a strongly typed genetic programming approach to smart electronic bitcoin markets

Article

Accepted Version

Creative Commons: Attribution-Noncommercial-No Derivative Works 4.0

Manahov, V. and Urquhart, A. ORCID: <https://orcid.org/0000-0001-8834-4243> (2021) The efficiency of bitcoin: a strongly typed genetic programming approach to smart electronic bitcoin markets. *International Review of Financial Analysis*, 73. 101629. ISSN 1057-5219 doi: 10.1016/j.irfa.2020.101629 Available at <https://centaur.reading.ac.uk/93822/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1016/j.irfa.2020.101629>

Publisher: Elsevier

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

The Efficiency of Bitcoin: A Strongly Typed Genetic Programming Approach to Smart Electronic Bitcoin Markets

Abstract

Cryptocurrencies have gained a lot of attention since Bitcoin was first proposed by Satoshi Nakamoto in 2008. To extend the current literature in this area, we develop four smart electronic Bitcoin markets populated with different types of traders using a special adaptive form of the Strongly Typed Genetic Programming (STGP)-based learning algorithm. We apply the STGP technique to historical data of Bitcoin at the one-minute and five-minute frequencies to investigate the formation of Bitcoin market dynamics and market efficiency. We find that both Bitcoin markets populated by high-frequency traders (HFTs) are efficient at the one-minute frequency but inefficient at the five-minute frequency. This finding supports the argument that at the one-minute frequency investors are able to incorporate new information in a fast and rationale manner and not suffer from the noise associated with the five-minute frequency. We also contribute to the growing volume of cryptocurrency literature by demonstrating that zero-intelligence traders cannot reach market efficiency, therefore providing evidence against the hypothesis of Hayek (1945; 1968). One practical implication of this study is that we demonstrate that cryptocurrency market participants can apply artificial intelligence tools such as STGP to conduct behaviour-based market profiling.

Keywords: Artificial Intelligence, Smart Electronic Markets, Bitcoin Trading, Cryptocurrencies, Evolutionary Computation, Market Efficiency.

1. Introduction

Cryptocurrencies have attracted a lot of attention from the media, regulators and investors since the introduction of Bitcoin by Nakamoto in 2008. The academic literature on Bitcoin is growing,¹ with many papers documenting the hedging the diversification benefits of Bitcoin (Bouri et al 2017; Urquhart and Zhang 2019; Borri 2019; Kajatazi and Moro 2019; Platanakis and Urquhart 2020), the existence of bubbles (Cheah and Fry 2015; Corbet et al 2018), investor sentiment and attention (Urquhart 2018; Shen et al 2019; Ibikunle et al 2020; Rognone et al 2020), the volatility of Bitcoin (Katsiampa 2017; Chaim and Luarini 2018; Catanaia et al 2019; Katsiampa et al 2019; Shen et al 2020), the behaviour of Bitcoin returns (Urquhart 2017; Corbet and Katsiampa 2018; Phillip et al 2018; Katsiampa 2018, Atsalakis et al (2019)).² Given the surge in the number of high frequency traders in financial markets since the turn of the millennium, the studies of Chu et al. (2019), Katsiampa et al. (2019), Zargar and Kumar (2019), Zhang et al. (2019), Ahmed (2020), Aslan and Sensoy (2020), Ma et al. (2020), Manahov (2020), Zhang et al. (2020) investigate cryptocurrency markets using high frequency data. However, the largest area of research on Bitcoin is whether the Bitcoin market conforms to the Efficient Market Hypothesis (EMH)³. That is, do Bitcoin prices reflect all current information thereby making Bitcoin returns random and unpredictable? The first paper to examine the efficiency of Bitcoin markets is Urquhart (2016), who through a battery of tests, finds that Bitcoin returns are not random and therefore predictable over time. This finding is supported by a number of follow-up studies by Nadarajah and Chu (2017), Bariviera (2017), Jiang et al. (2018), Urquhart (2017), Kristoufek (2018) and Tiwari et al. (2018), Hu et al. (2019), Kristoufek and Vosvrda (2019), Noda (2020), Wu and Chen (2020), which all employ alternative testing procedures and sample periods.

All of the above-mentioned studies, however, employed different econometric tests and mathematical models that represent rational agents, which does not allow full investigation of market efficiency and

¹ As Lucey et al. (2018) note that the academic literature needs to address Bitcoin in much more detail in order to gauge the potential benefits and risks of it.

² See Corbet et al (2019) for a recent review fo the empirical literature on cryptocurrencies.

³ The EMH was proposed by Fama (1970), who suggests that assets always trade at their fair value, making it difficult for market participants to either purchase undervalued stocks or sell stocks for inflated prices. Therefore, it should be difficult for investors to outperform the overall market through expert stock selection and that the only way an investor can possibly obtain abnormal returns is by chance or by investing in riskier assets.

Bitcoin's empirical characteristics. In financial and cryptocurrency markets the prices of all financial instruments are determined by a large variety of boundedly rational traders with various decision-making rules and risk attitudes. There is a great degree of heterogeneity between cryptocurrency market participants in terms of their rationality, tech-savviness and motivations. Therefore, a complete market efficiency investigation can only be performed in a model where market participants and their strategies are able to adapt and change over time. The complicated dynamics of these market participants demand a simulation technique that includes multiple heterogeneous market participants and a smart electronic Bitcoin market. Tay (2013) suggests that a natural method to analyse complex systems such as the financial markets is to implement an agent-based modelling technique which entails simulating the markets from the bottom up with a huge number of interacting heterogeneous boundedly rational traders that are designed to copy the behaviour of real-life market participants. In a similar vein, Ehrentreich (2008) argues that the ability to cope with heterogeneous and boundedly rational individuals makes agent-based modelling a perfect tool to study decentralised markets such as the cryptocurrency markets.

We implement a Strongly Typed Genetic Programming (STGP)-based learning algorithm to design four smart electronic Bitcoin markets and obtain historical one- and five-minute HFT data to simulate the interactions of multiple heterogeneous traders. Bichler et al. (2010) argue that smart markets represent theoretically supported computational methods that are useful in the interpretation of the attributes of complex trading environments and help decision makers to make real-time decisions. The aim of this study is to investigate market efficiency and test the Hayek hypothesis within artificial Bitcoin market settings. The Hayek hypothesis postulates that markets can operate in an efficient manner even when market participants possess a limited knowledge of the surrounding environment. We also observe the interactions between HFTs and zero-intelligence (ZI) traders to investigate whether trader intelligence or the market structure dominates Bitcoin markets. In other words, we simulate real-life Bitcoin trading in contrast to the studies discussed earlier. The STGP represents a very suitable tool for examining the Bitcoin market mechanism in isolation from the traders who populate the artificial cryptocurrency market. An important addition is that we are able to examine the relationship between the market dynamics and trading activities and therefore analyse the efficiency of the Bitcoin market in terms of

the EMH. The environment where Bitcoin returns and heterogeneous traders' beliefs co-evolve, adapt and try to survive over time provides an appropriate laboratory platform to investigate market efficiency and test the Hayek hypothesis.

Moreover, the heterogeneity of expectations among artificial traders provides important nonlinear conditions for the Bitcoin market. The other advantage of using the STGP is that trading orders are executed rapidly. Goldstein et al. (2014) argue that HFT involves the application of sophisticated trading algorithms for frequent and rapid execution of trading orders. The STGP technique allows forward testing, where submitted and executed orders change the depth of the cryptocurrency market, and these significant changes can be observed under laboratory conditions. The same process cannot be performed in back-testing because no statistical technique is able to regenerate the reaction of the market to changes in market depth. Furthermore, LeBaron (2006) demonstrates that on the descriptive level, specific complexity analysis – such as agent-based models – better represent the actual behaviour of the system in comparison with traditional economic models. We build our models incrementally, with no repeat visits in the same one- and five-minute datasets (data over fitting), and therefore every Bitcoin quote is used for the evaluation of a trading rule once, as in real-world Bitcoin trading. The STGP gradually develops a systematic model with parameter values that represent the best fit for the two datasets, rather than controlling the datasets to fit 'predefined' models as most statistical packages tend to do.

This is the first study to simulate real-life Bitcoin trading at the high-frequency timeframe and its contributions are threefold. First, we perform twelve different econometric tests and observe that the higher the frequency of trading the more efficient the market. This might indicate that more information is reflected in prices instantaneously in the Bitcoin market at higher frequencies when HFTs are present. It seems that HFTs operating at higher frequencies react to Bitcoin price changes in a timely manner, making the entire market more efficient. However, we find that this is not the case for the ZI traders, who document strong and significant evidence against market efficiency for all tests at the one and five-minute frequency.

Second, our experimental results within the laboratory Bitcoin market settings suggest that ZI traders operating in one- and five-minute frequency Bitcoin markets are unable to achieve market efficiency in any of the two investigated markets. In other words, traders who behave completely randomly cannot strive for market efficiency, which implies that trader cognitive abilities play a significant role in Bitcoin market efficiency. Hence, we can conclude that Bitcoin market efficiency dynamics are mainly influenced by the cognitive abilities of traders.

Third, we observe that the Hayek hypothesis does not hold in Bitcoin markets populated with ZI traders. We demonstrate that market participants equipped with limited information about the environment are unable to strive for market efficiency. Zero intelligence traders in our experiment behave entirely randomly and therefore they are unable to experience learning, observation and adaptation processes over time.

The paper is organised as follows. Section 2 provides an overview of Bitcoin and previous papers that have employed the special adaptive form of the STGP-based learning algorithm. Section 3 provides the experimental design and Section 4 presents the simulation results. Finally, Section 5 concludes the paper and provides a summary.

2. Literature Review

We use a STGP-based technique to obtain new evidence about the Bitcoin market efficiency, and whether the Hayek hypothesis holds in cryptocurrency markets. No other study has investigated either of these through the application of artificial intelligence modelling.

2.1. Bitcoin efficiency

Market efficiency is one of the most important and respected theories in finance, which has been shown the amount of papers examining the market efficiency of various markets. Since the literature on Bitcoin has exploded in recent times, there are countless papers studying the efficiency of this digital currency.

Urquhart (2016) was the first study the efficiency of Bitcoin and over the sample period 2010-2016 through a battery of statistical tests, they find that Bitcoin is inefficiency but through a subsample analysis, show that it may be moving towards an efficiency market. Bariviera (2017) study the efficiency of Bitcoin through the Hurst exponent and show that Bitcoin is inefficiency from 2011 to 2014 but from 2014 to 2017, is comparable to white noise. Since then, many papers have studied the efficiency of Bitcoin with Tiwari et al (2018) documenting the efficiency of Bitcoin, while Caporale et al. (2018), Jiang et al. (2018) and Al-Yahyaee et al. (2018) all reporting the inefficiency of Bitcoin due to its persistence. Conversely, Alvarez-Ramirez et al. (2018) report the anti-persistence and cyclical nature of Bitcoin returns, which is support by Takaishi (2018) using one-minutely data. Sensoy (2018), Zargar and Kumar (2019) and Zhang et al. (2019) also use intraday data to find the inefficiency of Bitcoin while Wu and Chen (2020) show that Bitcoin was inefficient before 2014 and after mid-2017 and that it has become more inefficient since mid-2017 than ever before. Therefore, there is a large literature examining Bitcoin efficiency with mixed findings.

2.2. Early smart electronic markets

Modelling financial markets from the bottom up with a number of interacting artificial agents began in the early 1980s'. Attempts in the literature to simulate real-life financial markets and model the behaviour of market participants are increasing rapidly. In one of the very early smart electronic market simulations, Cohen et al. (1983) investigated the impact of several random behaving market participants on different market structures.

While Frankel and Froot (1988) simulated the dynamics of the foreign exchange, Kim and Markowitz (1989) examined the interactions of specific trading strategies. A few years later, Rieck (1994) adopted actual trading rules in evolutionary settings and observe that trading rules possess very interesting dynamics as the market evolves. Lettau (1997) used a genetic algorithm (GA) to design a very simple setting and replicate the evolution and learning in a small population of traders. The GA approach, introduced by Holland (1975), seeks to replicate the approach of nature in the evolution of species; i.e., is based on the notion of population genetics to facilitate the searching process that results in the 'survival of the fittest' (Balakrishnan and Jacob, 1996). The GA process has been allowed to evolve for

many generations, while the best encountered outcomes are recorded (Wolfe and Sorensen, 2000). However, LeBaron (2000) argues that the processes of evolution and learning in Lettau's experiment takes place in a space disconnected from the real one. Furthermore, a standard GA can handle only a single evaluation criterion, which represents a significant limitation when decision makers need to take into account multiple objectives simultaneously (Kim et al., 2005).

2.3. More advanced smart electronic markets

The following studies are more extensive than the ones discussed above in that they are attempting to replicate more complex financial market structures.

The experiment conducted by Beltratti and Margarita (1992) is different from the earlier studies in that trade is performed in a random environment, where traders forecast future prices using an artificial neural network (ANN). The population of neural nets is divided into 'smart' and 'naïve' with the smart being more expensive to use. However, it appears that smart traders dominate in trading when prices are very volatile in the early stages of the market. Routledge (1994) implemented the framework of Grossman and Stiglitz (1980) and a model with GA-based traders to suggest that traders' forecast parameters can be supported through a dynamic learning tool.

To test learning within artificial settings, Arifovic (1996) used the GA to design a general equilibrium foreign exchange market based on the Kareken and Wallace (1981) principle (monetary equilibrium where demand for savings and money supplies are equal). The author observes that market participants' consumption and portfolio decisions were able to endogenously determine the price levels.

The studies of Brock and LeBaron (1996), as well as Brock and Hommes (1998), utilise different measures of past performance and discrete choice mechanisms in order to simulate the decisions-making process of individual market participants determining whether to buy costly information signals or to adopt more complex but costly forecasting tools. The authors reported that very simple forecasting rules are optimal under conditions of zero cost equilibrium. However, the model became unstable because all traders opted for the forecasting option making the whole market unpredictable.

Midgley et al. (1997) implement a GA to breed artificial agents that represent the actions of brand managers. Each brand's strategies evolve through simulations of repeated interactions in a virtual market, using the estimated weekly profits of each brand as a measure of its fitness for the genetic algorithm. In the tests they conduct, artificial agents outperform the historical actions of brand managers in this market. However, the authors limited the number of previous states of the market that can be held in the memory of the artificial agents. Moreover, the artificial agents faced another limitation – the significant shift in perspective- from week-by-week actions to a pattern of actions across thirteen weeks.

LeBaron et al. (1999) created another GA-based smart electronic market populated with 25 artificial traders who were able to replicate several well-known empirical features of real-world financial markets, such as leptokurtosis, volatility clustering and fundamental and technical predictability. The major disadvantage of this market is the homogeneity of traders because they are all identical except some minor differences in their rule sets. The Santa Fe Stock Market – outlined in great detail in Arthur et al. (1997) and LeBaron et al. (1999) – is one of the most popular smart electronic market projects. The project attempts to combine a well-defined economic structure in the trading mechanism with inductive learning based on a classifier GA system. Although the market represents one of the most complex smart electronic markets structure in existence in that time, the market was relatively difficult to track in terms of computer study due to complicated causalities inside the market.

2.4. Sophisticated smart electronic markets

More recently, there has been a great effort from management science (MS), information systems (IS), and computer science (CS) to attempt to develop new electronic markets that provide better economic outcomes through centralised computational decision making (Anandalingam et al., 2005). More specifically, the very rapid developments in computational technology and the evolution of computer networks allowed scientists to create complex financial market structures where artificial intelligence is able to simulate human decision making.

Chen and Yeh (2001) used genetic programming (GP) to design an artificial stock market and observe that the return series under investigation is independently and identically distributed (IID) indicating support for the EMH.⁴ Interestingly, this random nature of the return series was generated by agents who did not believe in the EMH. However, the authors failed to find a solution to Harrrald's criticism, who highlighted the difference between phenotype and genotype and doubted whether the adaptation of the two can be directly implemented in social processes. In a similar experiment, Chen and Yeh (2002) applied GP to evolving a population of market participants who were learning over time. The authors show that the main principles of the EMH can be satisfied with a fraction of the artificial time series. The main limitation of this study is the presence of a very large search space, which is difficult to cover during the process of modelling traders' learning. A year later, Cincotti et al. (2003) introduced a multi-asset artificial market with a finite number of stocks and amount of cash. Their smart electronic market exhibited a reversion to the mean, volatility clustering and a fat-tailed distribution.

The best performing trading strategy in their market was the one that utilises the mean reversion phenomenon of the asset prices. This finding, however, creates the main issue of this study – only a strategy that exploits the mean reversion provides satisfactory results in all cases.

More recently, Guo et al. (2012) used an adapted iterative market algorithm to examine how different market design factors affect computational market efficiency and liquidity. This study uses 480 random market settings and 160 combinations of market treatments and found that both asynchronous communication and asymmetric market information have negative effects on the speed of market convergence resulting in more traders' welfare losses. Contrary to traditional market assumptions, their results indicate that high trading volume does not correlate with low price volatility and rapid price discovery. Moreover, the authors suggest that a mixture of call and continuous market design is required to avoid premature market closure when agents hold forecasting learning ability. As outlined by the

⁴ Genetic Programming (GP) is a technique for enabling computer programs to develop their own solutions to problems and, consequently, GP is allied to the work that computer scientists have carried out in the related fields of artificial intelligence and machine learning (Poli et al., 2008). GP draws on the idea of adaptive behaviour to make it possible for an initial group of programs to create a next generation of programs, which then create a next generation of programs and so on. Only those programs that produce the best solutions are allowed to 'breed' the next generation.

authors there are several limitations to this study – predetermined, random and static inventory policies for the dealer are used and there are no transaction costs taken into account. Also, the study did not explain how market makers should be compensated for providing liquidity.

In a rather different experiment, Manahov et al. (2014) deploy a STGP trading algorithm and one-minute high frequency data of the most traded currency pairs worldwide in order to suggest that the STGP forecasting technique significantly outperforms the traditional econometric models⁵. They find evidence that the excess returns are both statistically and economically significant, even when appropriate transaction costs are taken into account. Manahov (2016) apply the STGP trading algorithm to real-time millisecond data of Apple, Exxon Mobil and Google to observe that HFTs equipped with scalping strategies generate cascades of cancelled orders within 100 milliseconds of order submission. Their empirical results imply that intense order cancellation activity is associated with a distortion of market quality and has a harmful effect on price discovery. The alleged purpose of order cancellation is to get faster access to the order flow by anticipating and front-running the other market participants in order to generate abnormal profits. Manahov and Zhang (2019) apply STGP to millisecond data of E-mini S&P 500 to investigate the behaviour of high-frequency traders in financial futures markets. The authors contribute to the high-frequency literature by showing that minimum resting trading order period of less than 50 milliseconds could lead to enhanced market efficiency.

3. Experimental Design

Makarov and Schoar (2018) and Hu et al. (2018) suggest that the majority of cryptocurrency exchanges take the form of an open electronic limit order book and operate in the same way as traditional equity markets, where market participants submit buy and sell trading orders and the exchange clears all trades without market limits, order size rules and centralised regulation. Therefore, we use a special adaptive form of the STGP to build four smart electronic markets that operate on the same basis as real-life equity markets. The STGP enables us to select and modify different parameters to suit our specification,

⁵ Strongly Typed Genetic Programming (STGP) is an enhanced version of GP first introduced by Montana (2002). STGP enforces data type constraints and whose use of generic functions and data types makes it a lot more powerful than GP.

such as different trading preferences, the level of transaction costs, the number of market participants and the minimum price increment. A complete list of the exact number of evolutionary parameters that can be specified is listed in Table 4.

Artificial traders represent each market participant who possesses their own trading rule and each trader is given an initial wealth of \$100,000. In Appendix B, we explain in more detail the production of the new genomes, which is conducted through the recombination of the parent genomes by crossover and mutation operations. Over time, each trader improves their trading rule over time through to maximize their gains by a process of survival of the fittest. Therefore, artificial traders are able to identify, learn and utilise profit opportunities thereby continually changing their trading decisions based on past performance and the changing cryptocurrency market conditions. In this way, the STGP trading algorithm evolves with real one- and five-minute quotes of Bitcoin and subsequently the smart electronic markets evolve by referencing to the characteristics of the Bitcoin market. Below, we set out in detail the design, structure and process of the markets.

3.1. The creation of initial trading rules

Initially, each individual trader has one trading rule that is randomly created⁶. We then apply the crossover recombination technique and mutation operation to create future generations of successful rules where the crossover recombination technique randomly chooses parts of two trading rules to exchange in order to create two new trading rules. This means successful trading rules survive and adapt. Then, the mutation operation randomly changes a small part of a trading rule (similar to Dempster and Jones 2001) and this process is repeated until at least one rule achieves the desired level of return over a specified period. As previously stated, each trading rule takes real-time one- and five-minute prices of Bitcoin and generates forecasts consisting of the desired position in Bitcoin, as well as an order limit price for buying and selling Bitcoin.⁷

⁶ There are two reasons for the random nature of the initial trading rules. First, the random initial creation of trading rules allows us to examine the whole range of possible trading rules. Second, the random nature of the initial trading rules allows us to observe how cryptocurrency market participants learn, adapt, and survive in constantly changing market conditions.

⁷ This process is further explained in Sub-sections 3.2 and 3.3.

In the original GP procedure, trading rules were evaluated by the same fitness function in each generation: however, the STGP evaluates the fitness of traders through a dynamic fitness function, which enables the return estimation period to include the most recent quotes in the markets. This is important as noted by Sermpinis et al. (2015), who state that having a novel fitness function is crucial in financial modelling.

Also, in contrast to the GP, the STGP only substitutes a small fraction of the entire population, ensuring a gradual change in population and thus greater model steadiness (Manahov, 2016). This avoids large fluctuations in the population of the model. Each trader generates their own fundamental value of Bitcoin without interaction with other traders, ensuring the individuality of each artificial trader and avoiding issues with herding. We distinguish between two different types of trades, namely: zero-intelligence (ZI) traders and high-frequency traders (HFTs) as discussed in the following subsection.

3.2. The structure of the smart electronic cryptocurrency markets and the difference between ZI traders and HFTs

In this paper, we study the efficiency of the Bitcoin market that is populated with 100,000 boundedly rational traders and develop four smart electronic Bitcoin markets using the STGP outlined in Appendix B. In two of the markets there are 100,000 ZI traders, while the other two markets are populated with 100,000 HFTs where all traders receive the same one- and five-minute Bitcoin data. Below, we spell out the difference between the ZI and HFTs.

Bid and ask prices of the ZI traders are developed randomly by enabling genes ‘RndPos’ and ‘RndLim’. ‘RndPos’ represents a function with a random position value between -100% and 100% obtained from a uniform distribution. ‘RndLim’ represents a function with a random limit price generated by a method based on Raberto et al. (2001), where we use $m = 1$ instead of 1.01. This is because no spread is added or subtracted for increasing the likelihood of an order being executed, because it is unknown whether the trader will place a buy or sell order at the time of ‘RndLim’. Consequently, their position and order limit price are developed randomly.

ZI traders behave entirely randomly and therefore they are unable to experience learning, observation and adaptation processes over time. ZI traders The ZI market is characterised with conditions of closed economy and there is no breeding⁸ or any broker fees. Therefore, there is no capital coming in or out of the market during the process. Consequently, the entire amount of cash within the ZI market will remain constant and the average amount of cash per ZI trader will also remain constant. The total net number of Bitcoins in the model is zero and will stay zero because no Bitcoins will be added or removed through ZI trader replacement. This is because ZI traders are initially only able to deal with cash but can then either establish long or short positions in Bitcoin, facilitating the total net number of Bitcoins to remain at zero. Therefore, any change in the price of Bitcoin has no implication on the total wealth of the market because the market is closed, and the losses of one trader are gained by another trader who holds the opposite position. Overall, ZI traders are fairly unintelligent and do not learn how the market works or the market conditions over time.

Alternatively, HFTs are free to develop and advance their trading rules through time because they are involved in continual learning, adaptation and evolution processes. The survival-of-the-fittest principle implemented in our experiment implies that the worst performing cryptocurrency market participants are eliminated based on their Breeding Fitness Return performance. The Breeding Fitness Return is a trailing return, where it is the n quotes of data of an exponential moving average of traders' wealth, where n is set to the minimum breeding age with a maximum value of 250.⁹ The Breeding Fitness Return measures the fitness criterion for the selection of HFTs to breed. All HFTs generate their wealth by investing in Bitcoin and the risk-free instrument, which is represented by cash. Given the continuous evolution of the cryptocurrency market, cryptocurrency HFTs who perform well (badly) become wealthier (poorer) which will positively (negatively) influence performance and the accuracy of the model. We follow Manahov (2016) to calculate the wealth of market participants as:

⁸ Breeding represents a process of creating new market participants to replace the underperforming ones.

⁹ In the case where the age is less than n , no value is calculated.

$$W_{i,t} = M_{i,t} + P_t h_{i,t} \quad (1)$$

where $W_{i,t}$ is the wealth generated by cryptocurrency trader i in period t ; $M_{i,t}$ and $h_{i,t}$ represents the risk-free investment and the amount of Bitcoin held by cryptocurrency trader i respectively, in period t and P_t is the price of Bitcoin in period t . Also, HFTs are designed to detect and avoid sweep risk because they are highly risk adverse and therefore will not hold potentially risky positions.

3.3. The clearing mechanism of cryptocurrency markets and order generation process for HFTs

Smart electronic Bitcoin markets are simulated double-auction markets, where all the buy and sell orders are collected. Our HFTs receive real-time quotes of Bitcoin and evaluate their trading rule and subsequent position in Bitcoin. If Bitcoins need to be purchased (sold), an order is generated to buy (sell) the amount determined by the specified limit price. For example, if a cryptocurrency trader holds three Bitcoins priced at \$10,000 each and has \$70,000 in cash, then their wealth is \$100,000 and their position in Bitcoin is 30%. If the trading rule generates a signal of a position of 60% and a limit price of \$10,000, then the limit order will be produced to purchase three¹⁰ additional Bitcoins with a price of \$10,000. The Bitcoin market therefore calculates the clearing price and all trading orders are executed at that price, which is where the highest trading volume from limit orders can be matched. If the clearing price is matched at multiple price levels, then the clearing price is formed on a basis of the average of the lowest and highest of those prices. The number of Bitcoins purchased by traders is always equal to the number sold by other traders. If the number of Bitcoins offered and the number of Bitcoins requested are unequal, then the remaining orders will be partially executed. Consequently, orders at the clearing price will be chosen for execution with priority for market orders over limit orders and then on a first-in-first-out (FIFO) basis. In the unlikely event of no matching limit orders, no market orders are executed (Manahov, 2016).

¹⁰ $60\% * (100,000/10,000) - 3 = 3$ Bitcoins.

3.4. Bitcoin data

Bitcoin is based on a scatter network of participating computers and has no physical equivalent. The Bitcoin operational system has a pre-programmed supply of funds that grows at a declining rate up to a fixed point (semi-fixed supply). Every single Bitcoin is able to generate an address through which to send and receive transactions. The decisive component that makes Bitcoin trading effective is that it answers the double-spending question without using a central mechanism. In economic terms, it is feasible to send a Bitcoin, without sending the same Bitcoin again, without other market participants being able to cast a transaction and also without them charging that Bitcoin back.

All Bitcoin transactions are registered on a decentralised register known as ‘blockchain’ supported by a grid of different computers known as ‘miners’. The main role of the ‘miners’ in the ‘blockchain’ is to retain concurrence by answering difficult mathematical issues. ‘Miners’ are usually remunerated with actual Bitcoins or voluntary transaction fees. Therefore, the rewarded Bitcoins increase the total of Bitcoin supply in the ‘blockchain’.

We obtain one- and five-minute high-frequency¹¹ Bitstamp data from www.bitcoincharts.com for the period of 1st January, 2015, to 3rd August, 2020. After the initial cleaning, the data is organised as comma-separated files with each row listing a time and date stamp, volume and unit price.

3.5. Data availability statement

The data that support the findings of this study are available from the corresponding author, upon reasonable request.

4. Simulation Results

4.1. Bitcoin market efficiency

We extract the processed Bitcoin data from the STGP to perform the following tests and examine the returns for statistical independence. We compute Bitcoin returns as:

¹¹ D’Amico and Petroni (2018) suggest that high-frequency trading is one of the most intensively researched fields in modern finance.

$$R_t = \text{Ln}\left[\left(P_t\right) / \left(P_{t-1}\right)\right] \times 100 \quad (2)$$

where R_t measures the return of Bitcoin and $\text{Ln}(P_t)$ and $\text{Ln}(P_{t-1})$ represent Bitcoin's natural logs at time t and $t-1$, respectively. Table 1 reports the descriptive statistics of Bitcoin for the HFTs and ZI traders at one- and five-minute frequencies. In each frequency, the distribution of returns are negatively skewed with excess kurtosis for the two trader groups at one- and five-minutes. The Jarque-Bera statistics suggest that the null hypothesis that Bitcoin returns generated by both HFTs and ZI traders are normally distributed is rejected in all periods.

We then apply the Rissanen's predictive stochastic complexity (PSC) criterion to the Bitcoin return series generated by the HFTs and ZI traders in order to identify any linear autoregressive moving average (ARMA) models by selecting the model with the minimum PSC. If Bitcoin returns satisfy the EMH, both p and q of the ARMA model should be equal to zero at any marker level and trading group. Therefore, there will not be any presence of linear dependence and Bitcoin returns will not be linearly predictable. However, Table 1 illustrates the possible presence of an ARMA process in the return series. The results from the Rissanen's PSC criterion indicate that the market populated with HFTs processing one-minute Bitcoin data is linearly independent ($p = 0; q = 0$). The absence of linearity in this particular market suggests an important initial finding that Bitcoin is so efficient that there is no presence of any linear signals. However, Bitcoin returns generated by ZI traders at one- and five-minutes and HFTs at the five-minute market are linearly dependent and are therefore inefficient.

To perform the Brock, Dechert and Scheinkman (BDS) test (Brock et al., 1996) for nonlinearity, we compute the most appropriate $\text{ARMA}(p; q)$ model and fit it to the two datasets to eliminate all linearity.¹² Hence, any signals left in the two datasets must be nonlinear. The BDS test postulates that if the Bitcoin return series are identically and independently distributed (IID), then the test statistic possesses a limiting standard normal distribution. Under the null hypothesis of the test, Bitcoin returns

¹² Consistent with Lim and Hooy (2013), Urquhart and Hudson (2013) and Urquhart and McGroarty (2016).

are IID. Also, this particular test discovers major deviations in the correlation of integral behaviour from that expected under the IID of the two datasets. The correlation integral can be estimated as:

$$C_{\varepsilon,m} = \frac{\sum_{1 \leq s \leq t \leq n} \sum_{t \leq n} \frac{I_{\varepsilon}(y_t^m, y_s^m)}{\binom{n}{2}} \quad (3)$$

where $y_t^m = (y_t, y_{t+\tau}, \dots, y_{t+(\tau-1)m})$ is an ‘ m -history’ estimated from the underlying uni-variate datasets and $I_{\varepsilon}(\cdot)$ represents an indicator function: $I_{\varepsilon}(y_t^m, y_s^m) = 1$ if $\|y_t^m - y_s^m\| < \varepsilon$ and zero otherwise. Chen et al. (2000) demonstrate that the correlation integrals determine the frequency and connectivity of different points that are within the distance of ε . Here m is the embedding dimension in which lag τ is integrated to compute ‘ m -history’. This computation is needed in order to prevent the accumulation of an extremely high correlation between the components of an m -tuple.

If the Bitcoin returns generated by the two groups of traders for the two frequencies are IID, then the correlation function of Equation 3 should indicate that $\lim_{\varepsilon \rightarrow \infty} C_{\varepsilon,m} = \left(\lim_{\varepsilon \rightarrow \infty} C_{\varepsilon,1} \right)^m$ for all $\varepsilon > 0$ and $m = 2, 3, \dots$. Brock et al. (1996) introduced the test statistic with limiting standard normal distribution:

$$V_{\varepsilon,m} = \sqrt{n} (C_{\varepsilon,m} - C_{\varepsilon,1}^m) / \sigma_{\varepsilon,m} \quad (4)$$

We applied the BDS test to one-minute returns generated by HFTs since we found no presence of linearity in those series. We also consider the BDS test parameters and we report the results with $\varepsilon = 0.6$ and DIM=5 in the last column of Table 1. The null hypothesis of IID is significantly rejected in the five-minute market where HFTs operate and in both markets populated with ZI traders. However, the null hypothesis has not been rejected in the one-minute HFTs’ market, suggesting that the Bitcoin returns in that market are identically and independently distributed. This market is therefore more random and more efficient than the other markets.

It is well-known econometric fact, however, that a large part of data nonlinearity is in their second moment. Thus, we implement the Lagrange multiplier (LM) test with up to 17 lags in order to identify the presence of an autoregressive conditional heteroscedasticity (ARCH) effect of the residuals in the two datasets. If the null hypothesis is rejected, we detect the generalised autoregressive conditional heteroscedasticity (GARCH) order of the return series by using the Schwartz Information Criterion:

$$r_t = \mu + h_t^{1/2} \varepsilon_t; h_t = \alpha_0 + \sum_{i=1}^p \alpha_i x_{t-i}^2 + \sum_{i=1}^q \beta_i h_{t-i} \quad (5)$$

where ε_t represents the IID normal innovations with restrictions $\alpha_0 > 0$, $\alpha_1, \beta_1 \geq 0$, $\sum_i \alpha_i + \sum_i \beta_i < 1$.

Table 2 reports the GARCH results where we report the GARCH ordering. We find that ZI traders on both one- and five-minute frequencies and HFTs at five-minute demonstrate ARCH effects, where we report the GARCH order. However, HFTs trading at the one-minute frequency do not exhibit any ARCH effect which supports the findings of the BDS test.

Barnett et al. (1998) argue that the BDS and the Kaplan tests are the two best performing procedures in terms of nonlinearity and therefore we adopt the Kaplan testing procedure in our experiment to add further robustness to our results.

Kaplan (1994) suggests that in any deterministic process the points are nearby under their image in the environment of state space. In econometric terms, if points X_i and Y_j are close to each other, then X_{i+1} and Y_{j+1} are close to each other too. In cases when $X_i = (r_i, r_{i-\tau}, r_{i-2\tau}, \dots, r_{i-(m-1)\tau})$ is fixed in m dimensional phase space, there is the following recursive function:

$$X_{i+\tau} = f(X_i) \quad (6)$$

where τ denotes the positive fixed-integer time decay. Therefore, we are able to compute the final components of the test as:

$$\delta_{i,j} = |X_i - X_j| \text{ and } \epsilon_{i,j} = |X_{i+\tau} - X_{j+\tau}| \quad (7)$$

One of the most significant properties of the Kaplan test is the precise estimation of the piecewise regression line for $(\delta_{i,j}; \epsilon_{i,j})$ and the use of the intercept to estimate the actual value of Kaplan (K).

Under the null hypothesis of the test is the IID condition and it is accepted when K is smaller than the test statistic and rejected when the opposite is valid. Column Three of Table 2, Panels A and B show that the trading activity of HFTs operating at one-minute frequency is consistent with the Kaplan test. So far in our experiment, this particular market does not reject the null hypothesis of IID based on the BDS, ARCH and Kaplan test statistics.

The Hurst exponent, introduced by Hurst in 1951, provides a measure for capturing long-term memory and fractality of the two Bitcoin datasets. The Hurst exponent (H) illustrates the scaling behaviour of the range of cumulative departures of Bitcoin return series from its mean. The test uses the range of the partial sums of deviations of Bitcoin return series from its mean, re-scaled by its standard deviation. When we face a sequence of continuous compounded Bitcoin returns $\{r_1, r_2, \dots, r_\tau\}$, τ measures the length of the estimation interval and \bar{r}_τ represents the sample mean:

$$(R/S)_\tau \equiv \frac{1}{S_\tau} \left[\max_{1 \leq t \leq \tau} \sum_{t=1}^{\tau} (r_t - \bar{r}_\tau) - \min_{1 \leq t \leq \tau} \sum_{t=1}^{\tau} (r_t - \bar{r}_\tau) \right] \quad (8)$$

where S_τ measures the standard deviation,

$$S_\tau = \left[\frac{1}{\tau} \sum_t (r_t - \bar{r}_\tau)^2 \right]^{1/2} \quad (9)$$

Hurst (1951) estimated the following relation:

$$(R/S)_\tau = (\tau/2)^H \quad (10)$$

While a Hurst exponent of 0.5 indicates randomness, $0.5 < H < 1$ indicates a persistent return series and $0 < H < 0.5$ suggests an anti-persistent Bitcoin return series. A persistent return series points towards the direction of the next Bitcoin return value most likely being the same as the current return value (trend reinforcing process). The greatest drawback of the re-scaled method is the lack of a natural significance

test because the value of H will almost always deviate from 0.5. However, Qian and Rasheed (2004) use Monte Carlo simulations to generate a range when there is only weak evidence of persistence ($0.5 < H < 0.65$), whereas a H greater than 0.65 suggests there is strong evidence of persistence. Conversely, a Hurst exponent below 0.45 shows strong evidence for anti-persistence and mean reversion, while a Hurst exponent between 0.45 and 0.5 suggests weak evidence for anti-persistence. Therefore we suggest any Hurst exponent greater than 0.65 and less than 0.45 indicates significant predictability in the precious metals' returns, similar to Urquhart (2017).

Table 2 shows that ZI traders operating one- and five-minutes frequencies and HFTs trading five-minute Bitcoin returns exhibit anti-persistent behaviour, manifested in Hurst exponents of less than 0.45 indicating strong evidence of anti-persistence and mean reversion. However, the Hurst exponent of HFTs trading one-minute Bitcoin returns is 0.588, indicating randomness of returns that is consistent with market efficiency.

Next, we examine the autocorrelation of Bitcoin returns by employing the Ljung-Box (Ljung and Box, 1978) test that has the null hypothesis of no autocorrelation (independently distributed returns):

$$Q = n(n+2) \sum_{k=1}^h \frac{p_k^2}{n-k} \quad (11)$$

where n is the sample size, p_k measures the sample autocorrelation at lag k and h represents the number of tested lags (we use ten lags). Under the null hypothesis, the test statistic Q follows a $\chi_{(h)}^2$.

The critical rejection of the null hypothesis of randomness is:

$$Q > \chi_{1-\alpha, h}^2 \quad (12)$$

where $\chi_{1-\alpha, h}^2$ is the $1-\alpha$ quantile of the chi-squared distribution with h degrees of freedom. The Ljung-Box test results suggest that the null hypothesis of no autocorrelation in one-minute Bitcoin returns traded by HFTs has not been rejected, indicating randomness and therefore market efficiency. However, both ZI trading groups at the two frequencies and HFTs processing five-minute returns reject the null hypothesis, suggesting a lack of randomly distributed Bitcoin returns.

Also in Table 2, we report the Wald and Wolfowitz (1940) runs test to examine the randomness of a distribution, by taking the returns in a given sequence.¹³ Under the null hypothesis of the test, the number of runs in a sequence of N components is a random variable with a conditional distribution specified by the observation of N_+ positive values and N_- negative values ($N = N_+ + N_-$) is normal, with mean $\mu = \frac{2N_+N_-}{N} + 1$ and standard deviation $\sigma^2 = \frac{2N_+N_-(2N_+N_- - N)}{N^2(N-1)} = \frac{(\mu-1)(\mu-2)}{N-1}$.

The last column of Table 2 reports the p-values and shows that for HFTs at the one-minute frequency, we fail to reject the null of the statistical independence of the returns indicating that this market is efficient. However, we can reject the null hypothesis of randomness for the three other specifications, indicating the inefficiency of ZI traders at the one-minute and five-minute frequency, and the ZI traders at the one-minute frequency.

Instead of comparing the magnitude of each Bitcoin return with its previous samples, Bartels' rank test (Bartels, 1982) ranks all samples from the smallest to the largest. The actual rank is the associated sequential number of $X_i : R(X_i)$. The test has randomness under the null hypothesis and any rank arrangement from all $n!$ possibilities is equally probable. We compute the probability for the test statistic as:

$$NM = \sum_{i=1}^{n-1} [R(X_i) - R(X_{i+1})]^2 \quad (13)$$

The large-sample approximation for $NM / (n(n^2 - 1)/12)$ represents a normal random parameter with mean of 2 and standard deviation of $\sigma_4^2 = 4(n-2)(5n^2 - 2n - 9)/5n(n+1)(n-1)^2$. The critical rejection is:

¹³ Similar to Batten et al. (2016).

$$|NM - 2| \geq Q(\alpha / 2) \sigma_4 \quad (14)$$

Table 2 indicates that the null hypothesis of randomness has not been rejected for one-minute Bitcoin returns traded by HFTs, while randomness can be rejected for the two ZI trading groups at both one- and five-minute timeframes and for the HFTs processing five-minute Bitcoin returns, which is consistent with all of our previous testing procedures.

Since the seminal work of Lo and MacKinlay (1988), the variance ratio test is one of the most popular tests for market efficiency (Hoque et al., 2007) and it examines whether Bitcoin returns represent a collection of IID observations or whether it follows a martingale difference sequence. If Bitcoin returns follow a martingale process, the variance of the k -period holding returns should be equal to k multiplied by the variance of the one-period returns. In other words, the following ratio should be equal to one:

$$V_k = \frac{1}{k} \frac{\sigma^2(r_k)}{\sigma^2(r)} \quad (15)$$

where $\sigma^2(r_k)$ and $\sigma^2(r)$ measure the variance of the k -period and one-period holding returns. Since our two data samples are quite large, we adopted the variance ratio procedure of Lo and MacKinlay (1988), where the sample size is large and k is fixed. The third column of Table 3 (Panels A and B) report the test statistics and the p -values in parentheses. The test statistics (marked with asterisks) for one- and five-minute ZI traders and five-minute HFTs indicate that the corresponding variance ratios are statistically different from one at the five per cent level of significance. However, the variance ratio test for one-minute HFTs' Bitcoin returns is equal to one, which implies that Bitcoin returns in this particular market follow a martingale difference process and the return series represent a collection of IID observations.

For accurate variance ratio test statistics, a choice of holding periods (k values) should be made. However, in some occasions these choices are rather arbitrary and made with insufficient statistical justifications. Choi (1999) developed the automatic variance ratio test in response to this issue. In Choi's (1999) automatic variance ratio test the optimal values of k are determined by using an entirely data-oriented testing procedure:

$$VR(k) = 1 + 2 \sum_{i=1}^{T-1} h(i/k) p(i) \quad (16)$$

where $p(i)$ is the autocorrelation function, and $h(x)$ is the QS window defined as:

$$h(x) = \frac{25}{12\pi^2 x^2} \left[\frac{\sin(6\pi x / 5)}{6\pi x / 5} - \cos(6\pi x / 5) \right] \quad (17)$$

While the null hypothesis of no serial correlation of the test has not been rejected in the market populated with HFTs trading one-minute Bitcoin returns, the null was rejected in all other markets at the two return frequencies (Column Four of Table 3, Panels A and B). These findings are in line with the other empirical results of market efficiency in the experiment so far.

Kim (2009) argues that Choi (1999) outlined small sample properties of the automatic variance ratio test when returns are IID, while its statistical properties under conditional heteroscedasticity are unknown. Multiple variance ratio tests, where the variance ratios of different holding periods are examined with controlled size, seem a reasonable alternative to the automated variance ratio tests. We implement the Chow and Denning (1993) wild-bootstrap test, which allows us to test for the multiple comparison of the set of variance ratio estimates with unity:

$$MV = \max_{1 \leq i \leq I} |M(Y; k_i)| \quad (18)$$

The joint null hypothesis of random Bitcoin returns is $VR(k_i) = 1$ for $i = 1, \dots, l$ against the alternative hypothesis that $VR(k_i) \neq 1$ for holding period k_i . We reject the null hypothesis at α level of significance if MV is greater than the $\left[1 - (\alpha^* / 2)\right]th$ percentile, where $\alpha^* = 1 - (1 - \sigma)^{1/l}$. In other words, the null hypothesis is rejected if any one of the computed variance ratios is significantly different from one.

We report the test statistics and the corresponding p -values in the fifth column of Table 3, Panels A and B. We observe that the Chow and Denning (1993) wild-bootstrap test failed to reject the null hypothesis of randomness for one-minute Bitcoin returns generated by the HFTs at the five per cent level of significance. All other markets for both ZI traders and the HFTs operating at the five-minute frequency rejected the null hypothesis which implies market inefficiency.

Finally, we adopt the spectral shape test of Durlauf (1991), which is completely different to the variance ratio tests that focus only on zero frequency deviations of the spectral density function from the null hypothesis. In contrast, the spectral shape test is appropriate for detecting deviations from the null hypothesis at all frequencies in the range of $[0, \pi]$. Therefore, the spectral shape test is appropriate for examining the randomness of Bitcoin returns against a broader set of alternatives. We compute the test in the following way:

$$U_T(t) = \int_0^{\pi t} \left[\frac{I_T(\lambda)}{\sigma^2} - \frac{1}{2\pi} \right] d\lambda \quad t \in (0, 1) \quad (19)$$

where $U_T(t)$ measures the cumulative deviations of the normalised sample from the null hypothesis.

We find no evidence against the null hypothesis of random Bitcoin returns for the one-minute returns of HFTs. However, the null hypothesis of the spectral shape test is rejected for both ZI traders at one- and five-minute returns and for the HFTs trading at the five-minute frequency.

Overall, our experimental results of twelve different econometric tests demonstrate that the higher the frequency of trading the more efficient the market. This might indicate that more information is

instantaneously reflected in the prices in the Bitcoin market at higher frequencies when HFTs are present. It seems that HFTs operating at higher frequencies react to Bitcoin price changes in a timely manner, making the entire market more efficient. This could potentially be explained by the greater liquidity of the one-minute market, but also by the fact that the structure of the decision-making mechanism of HFTs is probably better adapted to the structure of the information in the one-minute Bitcoin market. However, we find that this is not the case for the ZI traders, who document strong, significant evidence against market efficiency for all tests at the one-minute and the five-minute frequency.

4.2. The implications of traders' cognitive abilities and Bitcoin microstructure on market efficiency

In this sub-section, we examine in detail the role of Bitcoin microstructure and traders' cognitive abilities in market efficiency. The empirical findings regarding the extent to which traders' cognitive abilities or market microstructure influence market efficiency has been controversial.

Some studies suggest that market structure is the main driving force behind market efficiency. Becker (1962) and Gode and Sunder (1993) study simulated markets with ZI traders and report that traders who behave completely randomly (they do not observe, remember or learn) manage to achieve close to 100 per cent efficient allocation of resources. ZI traders had the same impact as intelligent traders on market efficiency and therefore the authors concluded that traders' cognitive abilities play a secondary role in market efficiency and that market structure is the main driving force.

In contrast, the supporters of the individual rationality doctrine criticised the previous work on the topic on the basis of inappropriately tested market structure. Cliff and Bruten (1997) argue that the study of Gode and Sunder (1993) was biased on a very specific assumption of symmetry in both supply and demand. ZI traders no longer efficiently allocate resources once this particular condition is breached. Brewer et al. (2000) suggest that ZI traders achieve high allocative efficiency because previous studies on the topic follow a Marshallian path. The author observes much lower allocative efficiency of ZI traders when the Marshallian path is abolished.

Our experimental results within laboratory Bitcoin market settings suggest that ZI traders operating in one- and five-minute frequency Bitcoin markets are unable to achieve market efficiency in any of the

investigated markets. In other words, traders who behave completely randomly cannot strive for market efficiency and, therefore, traders' cognitive abilities should play a more significant role in Bitcoin market efficiency. However, the five-minute frequency market populated with HFTs with specific trading rules that maximise profits, remember, observe and learn also failed to achieve market efficiency. Hence, we can conclude that Bitcoin market efficiency dynamics are influenced by both traders' cognitive abilities and the microstructure of the market.

4.3. Examining the Hayek hypothesis within smart electronic Bitcoin market settings

Hayek (1945, 1968) argues that markets can efficiently operate even when market participants have limited knowledge of the surrounding environment or of the other participants. The validity of the Hayek hypothesis has been examined in a number of studies producing vexed results. Smith (1982) provides strong support for the hypothesis in stationary double-auction markets with constant repetitive conditions of supply and demand. Holt et al. (1986) test the validity of the hypothesis in a double-auction market and observe substantial price deviation from the competitive equilibrium, which resulted, however, in only a relatively small loss of market efficiency and thus giving support for the hypothesis. Davis and Williams (1991) broadly confirm the findings of Holt et al. (1986). Othman and Sandholm (2010) and Shachat and Zhang (2012) find some empirical support for the Hayek hypothesis in the market environments they investigated.

Our empirical findings indicate that the Hayek hypothesis does not hold in Bitcoin markets populated with ZI traders. We demonstrate that market participants equipped with limited information about the environment are unable to strive for market efficiency. Zero intelligence traders in our experiment behave entirely randomly and therefore they are unable to experience learning, observation and adaptation processes over time. The ZI market is characterised with conditions of closed economy and there is no breeding (creating new market participants to replace the underperforming ones) or any broker fees.

4.3. Robustness checks

In this sub-section we examine whether the efficient market populated with HFTs trading at a one-minute frequency is characterised by long-memory for the Bitcoin return series. A random process long-memory in cases when the autocorrelation function decays asymptotically as power-law in the structure of $\tau^{-\alpha}$ with $\alpha < 1$ (Lillo and Farmer, 2004). In economic words, past Bitcoin return values could have important implications on the present, suggesting abnormal diffusion under stochastic conditions and highlighting the existence of long-memory. If $\alpha < 1$, the process is characterised by long-memory and the smaller the exponent value α , the longer the memory of the Bitcoin return series. We follow Lillo and Farmer (2004) to define the long-memory procedure:

$$\gamma(k) \sim k^{-\alpha} L(k) \text{ if the limit } k \rightarrow \infty \quad (20)$$

where $0 < \alpha < 1$ and $L(x)$ measure a slowly varying function at infinity when $\lim_{x \rightarrow \infty} L(tx)/L(x) = 1$.

Considering the Hurst exponent, the long-memory process can be rearranged as:

$$\alpha = 2 - 2H \quad (21)$$

And the short-memory procedure can be expressed as:

$$H = 1/2 \quad (22)$$

We consider the Hurst exponent because there is a strong relationship between the diffusion components of the un-segregated process. As the Hurst exponent of this particular market is 0.588, we calculate the associated (α) exponent by using Equation 21. We estimate that $\alpha = 0.824$, which indicates a very insignificant presence of long-memory processes in one-minute Bitcoin return series processed by HFTs. This finding is consistent with our previous empirical results. Therefore, we can conclude that the one-minute Bitcoin return series traded by HFTs are efficient – they possess very weak long-memory with $H=0.588$, and $\alpha = 0.824$.

5. Conclusions

Since they were first proposed by Nakamoto in 2008, cryptocurrencies have attracted more and more media and investor attention. Recently, the academic literature is catching up to examine the potential benefits, uses and risks of cryptocurrencies. Although the literature regarding cryptocurrencies is growing, it is still in its infancy and few comprehensive studies have been carried out.

Therefore, this paper adds to the literature by studying the market efficiency of Bitcoin in a high-frequency context. Although some papers – for example, Urquhart (2016), Nadarajah and Chu (2017), Bariviera (2017) and Jiang et al. (2018) – have examined the efficiency of Bitcoin, they have all employed methods that represent rationale agents while also examining the market efficiency at the daily level.

In this paper, we therefore implement the STGP-based learning algorithm to design four smart electronic Bitcoin markets and obtain historical one- and five-minute data in order to simulate the interactions of multiple heterogeneous traders. We populate the smart electronic markets with high-frequency traders as well as zero-intelligence traders in order to examine the impact these traders have on the Bitcoin market. After we extract the returns from the one-minute and five-minute markets populated by zero-intelligence traders, we find strong evidence of market inefficiency. This is because we find that returns are predictable and do not follow a random walk, in contrast to the Hayek hypothesis which postulates that markets can operate in an efficient manner even when market participants possess a limited knowledge of the surrounding environment.

Once we examine the market of high-frequency traders with one-minute data, we find that all of our testing procedures indicate strong evidence of market efficiency and that returns do follow a random walk. This suggests that investors at this high frequency are reflecting all current information into the prices of Bitcoin and therefore Bitcoin prices are unpredictable. However, the five-minute market populated with high-frequency traders indicates significant evidence of inefficiency, suggesting that investors are unable to process information in a timely and rational manner and that some noise may be left in the prices of Bitcoin.

References

- Ahmed, W. M. A. (2020). Is there a risk-return trade-off in cryptocurrency markets? The case of Bitcoin. *Journal of Economics and Business*, 108, 1-21.
- Anandalingam, G., Day, R. W., Raghavan, S. (2005). The landscape of electronic market design, *Management Science*, 51(3):316-327.
- Arifovic, J. (1996). The behaviour of the exchange rate in genetic algorithm and experimental economies. *Journal of Political Economy*, 104, 510-541.
- Arthur, W. B., Holland, J., LeBaron, B., Palmer, R., Tayler, P. (1997). Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W. B., Durlauf, S., Lane, D. (Eds.), *The Economy as an Evolving Complex System II*. Addison-Wesley, Reading, MA, pp.15-44.
- Atsalakis, G., Atsalaki, I. G., Pasiouras, F., Zopounidis, C. (2019). Bitcoin price forecasting with neuro-fuzzy techniques. *European Journal of Operational Research*, 276: 770-780.
- Aslan, A., Sensoy, A. (2020). Intraday efficiency-frequency nexus in the cryptocurrency markets. *Finance Research Letters*, 35:1-6.
- Barnett, W. A., Gallant, A. R., Hinich, M. J., Jungeilges, J. A. (1998). A single-blind controlled competition among tests for nonlinearity and chaos. *Journal of Econometrics*, 82, 157-192.
- Bartels, R. (1982). The rank version of von Neumann's ratio test for randomness. *Journal of the American Statistical Association*, 77, 40-46.
- Banzhaf, F. W., Nordin, P., Keller, R. E., Francone, F. D. (1998). *Genetic Programming: An introduction*. New York Publishers, USA.
- Bariviera, A. F. (2017). The inefficiency of Bitcoin revisited: A dynamic approach. *Economics Letters*, 161, 1-4.
- Batten, J. A., Lucey, B. M., Peat, M. (2016). Gold and silver manipulation: What can be empirically verified? *Economic Modelling*, 56, 168-176.
- Bouri, E., Molnár, P., Azzi, G., Roubaud, D., Hagfors, L. I. (2017). On the hedge and safe haven properties of Bitcoin: Is it really more than a diversifier? *Finance Research Letters*, 20, 192-198.
- Becker, G. (1962). Irrational behaviour and economic theory. *Journal of Political Economy* 70(1), 1-13.
- Beltratti, A., Margarita, S. (1992). Evolution of trading strategies among heterogeneous artificial economic agents. In: Meyer, J. A., Roitblat, H. L., Wilson, S. W. (Eds.), *From Animals to Animats 2*. MIT Press, Cambridge, MA.
- Bichler, M., Gupta, A., Ketter, W. (2010). Research commentary – designing smart markets. *Information Systems Research* 21(4), 688-699.
- Borri, N. (2019). Conditional tail-risk in cryptocurrency markets. *Journal of Empirical Finance*, 50, 1-19.
- Brewer, P., Huang, M., Nelson, B., Plott, C. 2002. On the behavioural foundations of the law of supply and demand: human convergence and robot randomness. *Experimental Economics* 5(3), 179-208.
- Brock, W. A., Dechert, W. D., Schieinkman, J. A., LeBaron, B. (1996). A test for independence based on the correlation dimension. *Econometric Reviews*, 15, 197-235.
- Brock, W. A., Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8-9), 1235-1274.
- Brock, W. A., LeBaron, B. (1996). A dynamic structural model for stock return volatility and trading volume. *Review of Economics and Statistics*, 78, 94-110.
- Catania, L., Grassi, S., Ravazzolo, F. (2019). Forecasting cryptocurrencies under model and parameter instability. *International Journal of Forecasting*, 35, 485-501.
- Chaim, P., Laurini, M. P. (2018). Volatility and return jumps in Bitcoin. *Economics Letters*, 173, 158-163.
- Cheah, E. T. Fry, J. M. (2015). Speculative Bubbles in Bitcoin Markets? An Empirical Investigation into the Fundamental Value of Bitcoin. *Economics Letters*, 130, 32-36.
- Chen, S. H., Lux, T., Marchesi, M. (2000). Testing for non-linear structure in an artificial financial market. *Journal of Economic Behaviour and Economic Organization*, 46(3), 327-342.

- Chen, S-H., Yeh, C-H. (2001). Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market. *Journal of Economic Dynamics and Control*, 25, 363-393.
- Chen, S-H., Yeh, C-H. (2002). On the emergent properties of artificial stock markets: the efficient market hypothesis and the rational expectations hypothesis. *Journal of Economic Behaviour and Organization*, 49, 217-239.
- Choi, I. (1999). Testing the random walk hypothesis for real exchange rates. *Journal of Applied Econometrics*, 14, 293-308.
- Chow, K.V., Denning, K.C.(1993). A simple multiple variance ratio test. *Journal of Applied Econometrics*, 58, 385-401.
- Chu,J.,Zhang,Y.,Chan,S.(2019). The adaptive market hypothesis in the high frequency cryptocurrency market. *International Review of Financial Analysis*, 64: 221-231.
- Cincotti, S., Focardi,S. M., Marchesi, M., Raberto, M. (2003). Who wins? Study of long-run trader survival in an artificial stock market. *Physica A*, 324, 227-233.
- Cliff D, Bruton J.(1997). Zero is not enough: on the lower limit of agent intelligence for continuous double action markets. HP Technical Report, HPL-97-141.
- Cohen, K. J., Maier, S. F., Schwartz, R. A., Whitcomb, D, K. (1983).A simulation model of stock exchange trading, *Simulation*, 41,181-191.
- Corbet, S., Lucey, B. M., Urquhart, A., Yarovaya, L. (2019). Cryptocurrencies as a Financial Asset: A Systematic Analysis. *International Review of Financial Analysis*, 62, 182-199.
- Corbet, S., Lucey, B., Yarovaya, L. (2018). Datestamping the Bitcoin and Ethereum bubbles. *Finance Research Letters*, 26, 81-88.
- D'Amico,G.,Petroni,F.(2018).Copula based multivariate semi-Markov models with applications in high-frequency finance. *European Journal of Operational Research*, 267, 765-777.
- Davis,D. D., Williams, A. W. (1991). The Hayek hypothesis in experimental auctions: institutional effects and market power. *Economic Inquiry*, 29 (2), 261-278.
- Daurlauf, S., N. (1991). Spectral based testing of the martingale hypothesis. *Journal of Econometrics*, 50, 355-376.
- Dempster, M., A., H., Jones, C. M. (2001). A real-time adaptive trading system using Genetic Programming, *Quantitative Finance*, 1, 397-413.
- Dunis, C.L., Laws, J., Karathanasopoulos, A.(2013). GP algorithm versus hybrid and mixed neural networks. *European Journal of Finance*, 19 (3): 180-205.
- Dyrhberg, A. H. (2016). Hedging capabilities of bitcoin. Is it the virtual gold? *Finance Research Letters*, 16, 139-144.
- Ehrentreich, N. (2008). Lecture notes in Economics and Mathematical Systems. Springer-Verlag Berlin Heidelberg.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25, 383-417.
- Frankel, J. A., Froot, K. A. (1988).Explaining the demand for dollars: International rates of returns and the expectations of chartists and fundamentalists. In: Chambers, R., Paarlberg, P. (Eds.), *Agriculture, Macroeconomics, and the Exchange Rate*, Westview Press, Boulder, CO.
- Gode, D. K., Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *The Journal of Political Economics*, 101, 119-137.
- Grossman, S., Stiglitz, J. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70, 393-408.
- Guo, Z., Koehler, G. J., Whinston, A. B. (2012). A computational analysis of bundle trading market design for distributed resource allocation. *Information Systems Research* 23(3), 823-843.
- Haynes, T., R. Wainwright, S. Sen, D. Schoenefeld (1995). Strongly Typed Genetic Programming in evolving cooperation strategies. Proceedings of the Sixth International Conference on Genetic Algorithms.https://www.researchgate.net/profile/Sandip_Sen4/publication/220885710_Strongly_Typed_Genetic_Programming_in_Evolving_Cooperation_Strategies/links/54dbfb1b0cf23fe133afd3a5.pdf.

- Haynes, T., D. Schoenefeld, R. Wainwright (1996). Type inheritance in Strongly Typed Genetic Programming, in: K. Kinnear and P. Angeline, eds. *Advances in Genetic Programming 2* (Morgan Kaufmann Publisher, Pittsburgh, PA).
- Hayek FA. 1945/1984. The use of knowledge in society. In *The Essence of Hayek*, Nishiyama C, Leube KR (eds). Hoover Institution Press: Stanford, CA; 211–224. (First published in *American Economic Review*, 35, 519–530.)
- Hayek F. A. 1968/1984. Competition as a discovery procedure. In *The Essence of Hayek*, Nishiyama C, Leube KR(eds). Hoover Institution Press: Stanford, CA; 254–265. (First presented as a lecture and published in *Kieler Vorträge* No. 56.)
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, MI.
- Holt, C. A., Langan, L. W., Villamil, A. P., (1986). Market power in oral double auctions. *Economic Inquiry*, 24 (1), 107–124.
- Hoque, H. A. A. B., Kim J. H., Pyun, C. S. (2007). A comparison of variance ratio tests of random walk: A case of Asian emerging stock markets. *International Review of Economics and Finance*, 16(4), 488-502.
- Hu, A. S., Parlour, C. A., and Rajan, U. (2018). Cryptocurrencies: stylized facts on a new investible instrument. Working paper. Available from: <http://webuser.bus.umich.edu/urajan/research/crypto.pdf>
- Hu, Y., Valera, H. G. A., Oxley, L. (2019). Market efficiency of the top market-cap cryptocurrencies: Further evidence from a panel framework. *Finance Research Letters*, 31:138-145.
- Hurst, H. E. (1951). Long-term storage of reservoirs: An experimental study. *Transactions of the American Society of Civil Engineers*, 116, 770–799.
- Ibikunle, G., McGroarty, F., Rzaev, K. (2020). More heat than light: Investor attention and bitcoin price discovery. *International Review of Financial Analysis*, 69, 101459.
- Jiang, Y., He, N., Weihua, R. (2018). Time-varying long-term memory in Bitcoin market. *Finance Research Letters*, 25, 280-284.
- Kajtazi, A., Moro, A. (2019). The role of bitcoin in well diversified portfolios: A comparative global study. *International Review of Financial Analysis*, 61, 143-157.
- Kaplan, D. T. (1994). Exceptional events as evidence for determinism. *Physica D*, 73, 38–48.
- Karageorgos, A., Thompson, S., Mehandjiev, N. (2002). Agent-based system design for B2B electronic commerce. *International Journal of Electronic Commerce*, 7(1): 59-90.
- Kareken, J., Wallace, N. (1981). On the indeterminacy of equilibrium exchange rates. *Quarterly Journal of Economics*, 96, 207-222.
- Katsiampa, P. (2017). Volatility estimation of Bitcoin: A comparison of GARCH models. *Economics Letters*, 158, 3-6.
- Katsiampa, P., Corbet, S., Lucey, B. (2019). High frequency volatility co-movements in cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money*, 62, 35-52.
- Kim, K.-J., Ahn, H. (2011). Collaborative filtering with a user-item matrix reduction technique. *International Journal of Electronic Commerce*, 16(1):107-128.
- Kim, J. W., Lee, B. H., Shaw, M. J., Chang, H.-L., Nelson, M. (2001). Application of decision-tree induction techniques to personalized advertisements on internet storefronts. *International Journal of Electronic Commerce*, 5(3):45-62.
- Kim, J. H. (2009). Automatic variance ratio test under conditional heteroscedasticity. *Finance Research Letters*, 6, 179-185.
- Kim, G., Markowitz, H. (1989). Investment rules, margin, and market volatility. *Journal of Portfolio Management*, 16(1), 45-52.
- Kim, Y. S., Street, W. N., Russell, G. J. and Menczer, F. (2005). Customer Targeting: A Neural Network Approach Guided by Genetic Algorithms. *Management Science*, 51(2):264-276.
- Koza, J. R. (1992). *Genetic Programming. On the programming of computers by means of natural selection*. MIT Press.
- Kristoufek, L. (2018). On Bitcoin markets (in)efficiency and its evolution. *Physica A: Statistical Mechanics and its Applications*, 503, 257-262.
- Kristoufek, L., Vosvrda, M. (2019). Cryptocurrencies market efficiency ranking: Not so straightforward. *Physica A: Statistical Mechanics and its Applications*, 531:1-8.

- LeBaron, B. Arthur, W. B., Palmer, R. (1999). Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23, 1487-1516.
- LeBaron, B. (2000). Agent-based computational finance: suggested readings and early research. *Journal of Economic Dynamics and Control*, 24, 679-702.
- LeBaron, B. (2006). Agent-based computational finance. *Handbook of Computational Economics*, 2, 1187-1233.
- Lettau, M. (1997). Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economic Dynamics and Control*, 21, 1117-1148.
- Lillo, F., Farmer, J. D. (2004). The long memory of the efficient market. *Studies in Nonlinear Dynamics and Econometrics*, 8(3), 1-20.
- Lim, K-P., Hooy, C-W. (2013). Non-linear predictability in G7 stock index futures. *The Manchester School*, 81(4), 620-637.
- Ljung, G. M., Box, G. E. P. (1978). On a measure of the lack of fit in time series models. *Biometrika*, 65 (2), 297-303.
- Lo, A. W., MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *The Review of Financial Studies*, 1(1), 41-66.
- Mai, F., Shan, Z., Bai, Q., Wang, X. S., Chiang, R. H. (2018). How does social media impact bitcoin value? A test of the silent majority hypothesis. *Journal of Management Information Systems*, 35(1):19-52.
- Lucey, B. M., Vigne, S. A., et al. (2018). Future directions in international financial integration research-a crowdsourced perspective. *International Review of Financial Analysis*, 55, 35-49.
- Ma, F., Liang, C., Ma, Y., Wahab, M. I. M. (2020). Cryptocurrency volatility forecasting: A Markov regime-switching MIDAS approach. *Journal of Forecasting*, <https://doi.org/10.1002/for.2691>
- Makarov, I. and Schoar, A. 2018. Trading and arbitrage in cryptocurrency markets. Working paper. Available from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3171204
- Manahov, V., Hudson, B., Gebka, B. (2014). Does high frequency trading affect technical analysis and market efficiency? And if so, how? *Journal of International Financial Markets, Institutions and Money*, 28, 131-157.
- Manahov, V. (2016). Front-running scalping strategies and market manipulation: why does high-frequency trading need stricter regulation? *The Financial Review*, 51, 363-402.
- Manahov, V., Zhang, H. (2019). Forecasting financial markets using high-frequency trading data: Examination with Strongly Typed Genetic Programming. *International Journal of Electronic Commerce*, 23(1);12-32.
- Manahov, V. (2020). Cryptocurrency liquidity during extreme price movements: is there a problem with virtual money? *Quantitative Finance*, <https://doi.org/10.1080/14697688.2020.1788718>
- Midgley, D. F., Marks, R. E. and Cooper, L. G. (1997). Breeding competitive strategies. *Management Science*, 43(3):257-275.
- Montana, D. J. (1994). Strongly Typed Genetic Programming. Technical report 7866, Bolt Beranek and Newman, Inc., Cambridge, MA.
- Montana, D. J. (2002). Strongly Typed Genetic Programming [online]. Available from <http://personal.d.bbn.com/~dmontana/papers/stgp.pdf>. Accessed 03/06/2018
- Nadarajah, S., Chu, J. (2017). On the inefficiency of Bitcoin. *Economics Letters*, 150, 6-9.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system [online]. Available from https://s3.amazonaws.com/academia.edu.documents/54517945/Bitcoin_paper_Original_2.pdf?AWSAccessKeyId=AKIAIWOWYYGZ2Y53UL3A&Expires=1551881943&Signature=Xm36mMXpk%2FaexupkF9XyX0oIs%2BU%3D&response-content-disposition=inline%3B%20filename%3DBitcoin_A_Peer-to-Peer_Electronic_Cash_S.pdf
- Neely, C., Weller, P., Dittmar, R. (1997). Is technical analysis in the foreign exchange market profitable? A genetic programming approach. *Journal of Financial and Quantitative Analysis*, 32(4), 405-426.
- Noda, A. (2020). On the evolution of cryptocurrency market efficiency. *Applied Economics Letters*, <https://doi.org/10.1080/13504851.2020.1758617>
- Phillip, A., Chen, K., Peiris, S. (2018). A new look at cryptocurrencies. *Economics Letters*, 163, 6-9.
- Platanakis, E., Urquhart, A. (2020). Should investors include Bitcoin in their portfolios? A portfolio theory approach. *British Accounting Review*, 52(4): 1-19.

- Polasic, M., Piotrowska, A. I., Wisniewski, T. P., Kotkowski, R., Lightfoot, G. (2015). Price fluctuations and the use of bitcoin: An empirical inquiry. *International Journal of Electronic Commerce*, 20(1):9-49.
- Poli, R., Langdon, W. B., McPhee, N. F. (2008). A field guide to Genetic Programming. San Francisco, USA.
- Othman, A., Sandholm, T. (2010). Automated Market-making in the Large: The Gates Hillman Prediction Market. Cambridge, USA.
- Qian, B., K. Rasheed. 2004. "Hurst Exponent and Financial Market Predictability." In *IASTED Conference on Financial Engineering and Applications, FEA 2004*. Athens, GA: University of Georgia, 203–209.
- Raberto, M., Cincotti, S., Focardi, S., Marchesi, M. (2001). Agent-based simulation of a financial market. *Physica A: Statistical Mechanics and its Applications*, 299, 319-327.
- Rieck, C. (1994). Evolutionary simulation of asset trading strategies. In: Hillebrand, E., Stender, J. (Eds.), *Many-Agent Simulation and Artificial Life*. IOS Press.
- Rognone, L., Hyde, S., Zhang, S. S. (2020). News sentiment in the cryptocurrency market: An empirical comparison with Forex. *International Review of Financial Analysis*, 69, 101462.
- Routledge, B. R. (1994). Artificial selection: genetic algorithms and learning in a rational expectations model. Technical Report, GSIA, Carnegie Mellon, Pittsburgh, Penn.
- Schlereth, C., Barrot, C., Skiera, B., Takac, C. (2013). Optimal product-sampling strategies in social networks: How many and whom to target? *International Journal of Electronic Commerce*, 18(1):45-72.
- Sermpinis, G., Stasinakis, C., Theofilatos, K., Karathanasopoulos, A. (2015). Modeling, forecasting and trading the EUR exchange rates with hybrid rolling genetic algorithms—Support vector regression forecast combinations. *European Journal of Operational Research*, 247, 831-846.
- Shachat, J., Zhang, Z. (2012). The Hayek hypothesis and long run competitive equilibrium: an experimental investigation, Available from: http://mpra.ub.uni-muenchen.de/39579/1/MPRA_paper_39579.pdf (accessed 01.04.18)
- Shen, D., Urquhart, A., Wang, P. (2019). Does twitter predict Bitcoin? *Economics Letters*, 174, 118-122.
- Shen, D., Urquhart, A., Wang, P. (2020). Forecasting the volatility of Bitcoin: The importance of jumps and structural breaks. *European Financial Management*, <https://doi.org/10.1111/eufm.12254>.
- Smith, V. L., (1982). Markets as economizers of information: experimental examination of the 'Hayek Hypothesis'. *Economic Inquiry*, 20 (2), 165–187.
- Tay, N.S.P. (2013). Agent-based models of financial markets. In *Encyclopedia of Finance*, eds, C-F. Lee and A.C. Lee, Springer Science+ Business Media, New York.
- Tiwari, A. K., Jana, R. K., Das, D., Roubaud, D. (2018). Informational efficiency of Bitcoin – An extension. *Economics Letters*, 163, 106-109.
- Urquhart, A. (2016). The inefficiency of Bitcoin. *Economics Letters*, 148, 80-82.
- Urquhart, A. (2017). How predictable are precious metal returns? *The European Journal of Finance*, 23(14), 1390-1413.
- Urquhart, A. (2017). Price clustering in Bitcoin. *Economics Letters*, 159, 145-148.
- Urquhart, A., Hudson, R. (2013). Efficient or adaptive markets? Evidence from major stock markets using very long historic data. *International Review of Financial Analysis*, 28, 130-142.
- Urquhart, A., McGroarty, F. (2016). Are stock markets really efficient? Evidence from the adaptive market hypothesis. *International Review of Financial Analysis*, 47, 39-49.
- Urquhart, A., Zhang, H. (2019). Is Bitcoin a hedge or safe haven for currencies? An intraday analysis. *International Review of Financial Analysis*, 63, 49-57.
- Wald, A., Wolfowitz, J., 1940. On a test whether two samples are from the same population. *Annals of Mathematical Statistics*, 11 (2), 147–162.
- Wolf, W. J. and Sorensen, S. E. (2000). Three Scheduling Algorithms Applied to the Earth Observing Systems Domain. *Management Science*, 46(1):148-166.
- Wu, D. J. (2000). Artificial agents for discovering business strategies for network industries. *International Journal of Electronic Commerce*, 5(1):9-36.
- Wu, L., Chen, S. (2020). Long memory and efficiency of Bitcoin under heavy tails. *Applied Economics*, 52, 5298-5309.

- Wu,L., Chen,S.(2020). Long memory and efficiency of Bitcoin under heavy tails. *Applied Economics*, 52(47):5298-5309.
- Zargar,F,N.,Kumar,D(2019). Informational inefficiency of Bitcoin: A study based on high-frequency data. *Research in International Business and Finance*, 47: 344-353.
- Zhang,Y.,Chan,S.,Chu,J.,Nadarajah,S.(2019). Stylised facts for high frequency cryptocurrency data. *Physica A: Statistical Mechanics and its Applications*, 513:598-612.
- Zhang,Y.,Chan,S.,Chu,J.,Suliman,H.(2020).On the Market Efficiency and Liquidity of High-Frequency Cryptocurrencies in a Bull and Bear Market. *Journal of Risk and Financial Management*, 13(8): 1-8.

Appendix A

Table 1. The descriptive statistics for Bitcoin returns generated by HFTs and ZI traders from 1st January 2015 to 3rd August 2020. We extract the processed data from STGP algorithm. “Std.Dev” refers to the standard deviation of returns. While “max” and “min” refer to the maximum and minimum returns. “Skew.” denotes the skewness of returns while “Kurt.” refers to the kurtosis of returns. Finally, the “JB”, “PSC” and BDS” refer the Jarque-Bera, test, the Rissanen’s PSC criterion and the BDS test respectively. ^a indicates a failure to reject the null hypothesis that Bitcoin returns are identically and independently distributed (IID).

	Mean	Std.Dev.	Max	Min	Skew.	Kurt.	JB	PSC	BDS
Panel A: 1-minute frequency									
HFTs	0.5728	6.4378	47.8240	-64.2543	-1.8835	7.28	884.21	(0,0)	0.18 ^a
ZI traders	0.7793	8.2325	81.3663	-93.7566	-0.9994	14.18	2780.36	(1,0)	0.69
Panel B: Five-minute Frequency									
HFTs	0.6996	7.5530	73.1173	-77.3416	-0.8552	13.06	1893.27	(0,1)	0.46
ZI traders	0.8909	9.0107	92.3651	-96.9952	-0.8773	18.22	7423.28	(1,0)	0.92

Table 2. The market efficiency econometric results for Bitcoin returns generated by HFTs and ZI traders from 1st of January 2015 to 3rd August 2020. We extract the processed data from STGP algorithm. Specially, we report the GARCH, Kaplan, Hurst exponent, Ljung-Box test and Wald-Wolfowitz runs test results. For the GARCH results, we report the GARCH ordering, while for the L-B test and the Runs test, we report the p-values. ^a indicates a failure to reject the null hypothesis that Bitcoin returns are identically and independently distributed (IID) at 5 percent level of significance.

	GARCH (p,q)	Kaplan	Hurst	L-B test	Runs test
Panel A: One-minute Frequency					
HFTs	-	0.101 ^a	0.588	(0.52)	(0.48)
ZI traders	(1,1)	0.999	0.394	(0.00)	(0.00)
Panel B: Five-minute Frequency					
HFTs	(0,1)	0.906	0.411	(0.00)	(0.00)
ZI traders	(1,2)	1.561	0.363	(0.00)	(0.00)

Table 3. Market efficiency econometric statistics for Bitcoin returns generated by HFTs and ZI traders from 1st January 2015 to 3rd August 2020. We extract the processed data from STGP algorithm. We report the test statistics and p-values for all tests. ^a indicate a failure to reject the null hypothesis of randomness at the 5 percent level of significance.

^b indicate that the corresponding variance ratios are statistically different from 1 at the 5 percent level of significance. ^c indicate a failure to reject the null hypothesis of no serial correlation at the 5 percent level of significance. ^d indicate a failure to reject the null hypothesis of randomness at the 5 percent level of significance. ^e indicate a failure to reject the null hypothesis of randomness at the 5 percent level of significance.

Test	Bartels test	VR test	AVR test	W-B AVR test	Spectral test
Panel A: One-minute Frequency					
HFTs	0.3874 (0.49) ^a	0.188 (0.57)	0.979 (0.77) ^c	0.036 (0.31) ^d	0.204 (0.12) ^e
ZI traders	3.224 (0.01)	1.751 (7.46) ^b	2.988 (0.01)	2.113 (6.58)	3.219 (2.36)
Panel B: Five-minute Frequency					
HFTs	2.546 (0.02)	1.863 (5.68) ^b	1.347 (0.01)	1.046 (2.82)	2.231 (0.99)
ZI traders	5.202 (0.01)	3.726 (9.11) ^b	5.338 (0.10)	4.790 (7.47)	3.996 (3.02)

Table 4. Smart electronic Bitcoin market parameter settings. This table presents the artificial Bitcoin market parameter settings for HFTs and ZI traders. All parameters are set at optimum level to ensure greater replication of real-life Bitcoin trading.

<i>Smart electronic Bitcoin market parameters</i>	
Total population size of HFTs and ZI traders in each market	100,000 traders
Initial wealth(equal for all traders)	\$100,000
Significant Forecasting range	0% to 10%
Number of decimal places to round quotes on importing	2
Minimum price increment for prices generated by model	0.01
Minimum position unit	20%
Maximum genome size	4096*
Maximum genome depth	20**
Minimum initial genome depth	2
Maximum initial genome depth	5
Breeding cycle frequency (quotes)	1
Minimum breeding age (quotes)	80***
Initial selection type	random
Parent selection (percentage of initial selection that will breed)	5%****
Mutation probability (per offspring)	10%
Seed generation from clock	Yes
Creation of unique genomes	Yes
Offspring will replace the worst performing traders of the initial selection	Yes

* Maximum genome size measure the total number of nodes in an trader's trading rule. A node is a gene in the genome such as a function or a value.

** Maximum genome depth measures the highest number of hierarchical levels that occurs in an trader's genome (trading rule). The depth of a trading rule can be an indicator of its complexity.

*** This is the minimum age required for traders to qualify for potential participation in the initial selection. The age of a trader is represented by the number of quotes that have been processed since the trader was created. This measure also specifies the period over which agent performance will be compared. Our minimum breeding age is set to 80, which means that the trader's performance over the last 80 quotes will be compared.

**** 5% of the best performing traders of the initial selection that will act as parents in crossover operations for creating new traders.

****This is the minimum age required for traders to qualify for potential participation in the initial selection. The age of a trader is represented by the number of quotes that have been processed since the trader was created. This measure also specifies the period over which agent performance will be compared. Our minimum breeding age is set to 80, which means that the trader's performance over the last 80 quotes will be compared.

****5% of the best performing traders of the initial selection that will act as parents in crossover operations for creating new traders.

Appendix B

Description of Genetic Programming (GP) and Strongly Typed Genetic Programming (STGP)

Although, Schlereth et al. (2013) designed agent-based models that provide a promising link to individual behaviour, most existing techniques to agent-based system design fail to deal with the complexity of design (Karageorgos et al., 2002). Genetic Programming (GP) represents a machine-learning technique to facilitate the enhancement of computer programs on a basis of natural evolution (Banzhaf et al., 1998). Evolutionary genetic algorithms (GAs) deploy computer programs that better match the designated tasks. Agents in the population compete and breed with each other in order to generate better agents, where GP parse trees, also known as tree genomes, are utilised as the basis of the algorithms. Opposite to neural networks, decision-tree structures represent specific rules that can be expressed in English (Kim et al. 2001). Under GAs' approach the economy is seen as evolving complex system in which artificial traders perform the activities of the real-life economy (Wu, 2000). The procedure enhances search output by performing different solutions with genetic operators (Kim and Ahn, 2011).

The GP design procedure states that one would need to specify all of the programs and different variables that can be used as nodes in a parse tree. However, Strongly Typed Genetic Programming (STGP) is able to reduce the size of the searching space to a greater degree (Montana, 1994). The STGP search space consists of the set of all parse trees, which implies that all functions have the correct number and type of parameters. The STGP parse tree is purposefully limited to a certain maximum depth of 20 in our experiment (Table 4). We fix the maximum depth to 20 in order to maintain a finite and manageable search space, and to prevent the trees from growing to a very large size.

Haynes et al. (1995) highlight that the generic functions and generic data classification for these functions are the two critically important factors in the design of STGP algorithms. The authors also point out that data classification needs to be specified at the early stages of programming. Consequently, the process of initialisation and the different genetic operators are allowed to create parse trees. The

same parse trees play a positive role in the programming process because the search space can be substantially reduced (Haynes et al. 1996).

The process of creating genetic programs involves the following five main steps:

1. Initially, create a randomly generated population of trees (trading rules) with the only requirement that they are well defined and aim to generate results that are appropriate to the issue under investigation. These trading rules involve the application of the basic principles of biological evolution in order to develop a new and enhanced population of trading rules. The development of this new population is based on a domain-independent system conducted by the Darwinian theory of natural selection and the *survival of the fittest* principle.
2. Estimate the fitness of trading rules in the initial population according to an appropriately selected set of criteria.
3. Develop a new population by using the following programming techniques:
 - (i) Perform a crossover process (copy existing traders into the new population).
 - (ii) Perform a mutation process. Mutation represents a random selection of a pair of existing trading rules and their recombination to generate a new trading rule. While crossovers mix population sub-trees, mutations substitute sub-trees with new sub-trees. The crossover (Figure 1) and mutation (Figure 2) techniques are performed with the probability of selection for the operations, and are skewed towards selecting cryptocurrency market participants with higher levels of fitness.
4. Estimate the fitness of each cryptocurrency market participant in the new population.
5. Repeat the above steps and keep a record of the overall fitness of the cryptocurrency market participants.




Figure 1. The crossover process for developing new trading rules in GP and STGP (Dunnis et al. 2013).

Figure 1 illustrates the process of crossover, where randomly selected subtrees are exchanged. The entire sub-trees (highlighted areas) are swapped from Point 1* to Point 2* and from Point 3* to Point 4*, creating two trading rules (offspring models). Trading rules are chosen on the basis of their fitness, with the crossover process selecting future areas of the search space where trading rules consists of parts from the superior trading rules. The initial selection of the best trading rules is based on the Breeding Fitness Return to act as parents in the process of crossover. Each pair of parents is able to produce two offspring trading rule and, therefore, the number of parents and the number of offspring are equal. Consequently, the newly created trading rules replace those that perform poorly.

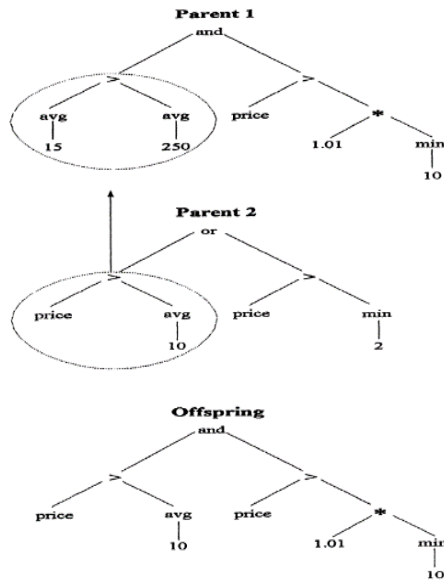


Figure 2. Mutation process for creating new trading rules in GP and STGP (Neely et al., 1997).

Figure 2 illustrates the process of mutation. A pair of trading rules is randomly chosen with the probability of selection weighted towards higher fitness trading rules. The sub-trees of the two parent rules are also randomly chosen, with one of the selected sub-trees abandoned afterwards and substituted with a different sub-tree to generate the offspring rule. The STGP explores parts of the search space by maturing a population of different trading rules, with the trading rules in each consecutive generation aiming to solve the issue under investigation. As a full technical explanation of crossover and mutation is beyond the scope of this paper, interested readers may refer to Koza (1992) for more details.

Figure 3 indicates that the trading rule of HFTs generates a buy (sell) signal if the average Bitcoin price over the past one-minute (five-minutes) is greater (less) than the current price and the current volume is less than 500 where the volume function protects HFTs from sweep risk exposure as previously pointed out.




Figure 3. Example of the HFTs' trading rule.