LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

RESEARCH ARTICLE

Economic Simulation of Cryptocurrencies and Their Control Mechanisms

Michael R. Mainelli,[*†] Matthew Leitch,[‡] Dionysios Demetis [§]

Abstract. A cryptocurrency needs a relatively stable value if it is to fulfill the traditional functions of money and be useful as a currency. To achieve this, controls are needed within the ecosystem of the cryptocurrency. Although a simulation cannot predict future currency rates or other variables exactly, it is argued that a model that simulates a range of challenging behavior can be a useful testbed for control schemes. To illustrate and explore this idea, an agent-based economic model was used to simulate the early period of a hypothetical cryptocurrency and test two control mechanisms. The results suggest that this approach may be fruitful and that it may be important to include more than just coin minting within the control scheme. An economic simulation model is likely to be a valuable tool in developing and regulating effective cryptocurrency systems.

1. Introduction

Over the last several years cryptocurrencies have generated a lot of excitement and consumed a lot of attention, effort, and electricity. Hundreds have been launched and hundreds are still operating.[1, 2] The best known continue to expand, with (mostly) rising values and continually rising numbers of coins in issue. However, there are many problems to solve and work continues to tackle problems of security, performance, efficiency, legality, and funding.

The focus of this paper is a problem that has had relatively little attention so far: the performance of cryptocurrencies as currencies. Specifically, we discuss the need for economic control, consider how cryptocurrencies can be controlled so that they function well as currencies, and illustrate the idea of using an agent-based economic simulation to test control mechanisms. Such simulations cannot predict the future of cryptocurrencies exactly but might still be useful for simulating challenging conditions and testing economic control mechanisms. The simulation we present is an early exploration of this idea and more recently we have completed the specification of a much more sophisticated simulator. Perhaps future cryptocurrencies will perform better than established fiat currencies do today in part because their economy is controlled by agreed rules embedded in software rather than always by human intervention.

We begin with a brief review of some well-known ideas on the functions of money and comments on the performance of existing cryptocurrencies, then explain the rationale for

[*] MetroGnomo MetroID: ca552760-3e2f-474d-a2da-77262f0947e4
[†] M. Mainelli (michael_mainelli@zyen.com) is the cofounder and executive chairman of Z/Yen Group, UK, and Emeritus Professor of Commerce at Gresham College.
[‡] M. Leitch (matthew_leitch@zyen.com) is an independent consultant and researcher at Z/Yen Group, UK (corresponding author).
[§] D. Demetis (D.Demetis@hull.ac.uk) is a lecturer in Management Systems at the Hull University Business School.

using simulation before specifying the illustrative simulation and its control mechanisms, and finally reporting the results of tests of those control mechanisms.

2. The Functions of Money

The functions of money have been conceptualized in slightly different ways by different theorists, dating back as far as W. S. Jevons's 1875 text *Money and the Mechanism of Exchange* (and beyond).[3] However, it is not controversial to say that money should be an effective medium of exchange, meaning that it should allow two people to make a deal even if they do not have goods or services of exactly equal value to exchange in a barter. In this situation money makes up the difference in value or is the entire value that one side offers in the deal. Money should also be a reliable store of value in the sense that, if two people make two exchanges with each other, but separated by time, neither feels cheated by the fact that the money has changed its purchasing power significantly over time. It should also have a large community of users who accept the currency and are familiar with what it can buy, using the currency for mentally valuing goods and services when making decisions, even when the currency is not being used in a purchase. These are all familiar attributes of established fiat currencies such as Sterling and the US Dollar.

There are arguments in favour of some variations in exchange rates between currencies, but, for the most part, money works best when its purchasing power does not change. This makes it a reliable store of value and allows people to learn the usual prices of goods and services they buy often, allowing them to shop efficiently. Even low rates of inflation are a problem because they erode the purchasing power of savings put aside for old age.

3. Cryptocurrency Performance as Money

Good performance as a currency is not automatic. Currently popular cryptocurrencies have not shown constant purchasing power over time, or anything like it. Even the cryptocurrency that is best known and most widely accepted as a payment method, Bitcoin, has experienced massive rises and falls in value over periods of months and even over periods of a few days. Other cryptocurrencies have experienced large changes in value. Sometimes they peak soon after launch then settle down to low value and low transaction volume later (*e.g.* Vertcoin, Quarkcoin). Sometimes they experience long periods of low value but then rise in value for a few weeks or months before subsiding again (*e.g.* Monero, Vertcoin, Quarkcoin).

Studies of the drivers of Bitcoin prices have often concluded that speculation is one of, or even the main, driver,[4-9] and this is partly driven by publicity, both negative and positive, which is partly spread by social media.[10]

With the value of Bitcoin changing so much it is not surprising that it is hard to find goods and services with an advertised price in BTC, even if BTC is an acceptable form of payment. More often the merchant advertises prices in an established currency that shoppers are familiar with but then offers a BTC price at the point of payment, using the latest exchange rate in some way. This process is described by Luther and White in their 2014 working paper "Can Bitcoin Become a Major Currency?".[11]

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

If cryptocurrencies are to function as currencies then their value needs to be much more stable and that will require economic control mechanisms different from those used by leading cryptocurrencies today.

4. Rationale for Using Simulation

We argue that using simulation to test control mechanisms for stabilizing exchange rates can be a useful strategy even though there is little prospect of producing simulations that make reliable predictions of the exact trajectory of cryptocurrency exchange rates. A familiar analogy is that, while it is impossible in practice to simulate and exactly predict the movement of molecules in a body of air, a bicycle pump will predictably push those molecules into a bicycle tyre. Some effects are more predictable than others.

Market prices, including exchange rates, are unpredictable for a number of reasons. It may be that they are chaotic, in the sense of being extremely sensitive to initial conditions. This was the conclusion of De Grauwe and Vansenten in 1990 after studying the performance of a deterministic model of an economy.[12] Markets are certainly driven by unpredictable factors outside those markets. In the case of Bitcoin, these factors have included hacking incidents and arrests. Attempts to simulate the decision-making of participants will struggle because we do not know how people make their decisions, there are almost certainly large differences between people in how they make their decisions, and people may change their approach over time as they learn from experience and learn more about the significance of events. Even if the decision-making of users of one cryptocurrency was understood there is no guarantee that users of other cryptocurrencies would behave the same way, particularly if the cryptocurrency operates differently.

Perhaps most fundamentally, markets are unpredictable because they are driven by perceptions of the perceptions of others.[13] We act on what we think other people will think. Sometimes there is an intrinsic value that can be calculated, but often there is no such intrinsic value. Prices move according to sentiment. In one laboratory study, subjects played the role of investors in technology. In game after game, different people were given the same information about the technology investments, at the same time, yet greatly different values and rankings for the investments emerged in each game.

The simulation study described below was designed to explore and illustrate the potential value of using simulation to test control mechanisms for future cryptocurrencies, most likely designed rather differently from today's. In this study the hypothetical cryptocurrency is simple but has two controls that are not typical for existing cryptocurrencies but might be in future. Although validation against real-world data is not possible with a hypothetical system the simulation aimed to (1) generate behavior with features that are familiar from actual cryptocurrency exchange rate time series, and (2) make that behavior responsive to control mechanisms whose effect is relatively uncontroversial. The model specified below produces large, rapid rises and falls in exchange rate, bursts of activity, and (if parameters are appropriately set) oscillation as participants over-react to recent events. The control mechanisms rely on nothing more than the idea that people will buy coins in a cryptocurrency from the lowest priced source, other things being equal, and that people will buy less cryptocurrency if it has a higher price.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

An agent-based approach was selected because of its ability to produce a rich variety of unpredictable behavior, with detail, and because agent-based simulations of currencies have been produced before.[14-18]

While consistently successful control within a simulation is no guarantee of successful control in reality, failures of control within a simulation are a strong indication that a control scheme is not safe for real use.

5. Specification of the Model

The representation of detail inherent in an agent-based model means that the specification of the model is quite lengthy. Source code written in R is available as a supplemental file.[19]

The agents in the model are a set of merchants, who offer goods for sale, a set of customers, who buy those goods, a mint that produces new cryptocoins, and a market maker that operates a currency exchange between the cryptocurrency (CC) and an established fiat currency (FC). After initialization, the behavior of the agents unfolds over a sequence of discrete days. This model will be explained first without mentioning the control mechanisms, and then the control mechanisms will be explained.

Initialising the model sets up the initial properties of the merchants, their goods, the customers, and the exchange. Each merchant has an initial pot of cash, entirely in FC, is not a participant in the cryptocurrency, and offers a list of goods for sale at advertised FC prices. Each customer has an initial pot of cash, entirely in FC, is not a participant in the cryptocurrency, and has a personal level of daily shopping represented by an average number of purchases per day. The market maker has an initial pool of CC and an initial exchange rate for CC against FC.

In the interests of realism, the merchants are varied, with some having large cash pots and offering many goods for sale, while other have less of everything. Similarly, customers vary from rich to poor, having different sizes of cash pot and different average numbers of purchases per day.

The number of goods offered by each merchant is distributed in the shape of a Zipf law with s being a parameter in the model, as is the highest price of a good offered by each merchant. That is, the price of each good offered by each merchant is calculated from the highest priced good by dividing by the good's number raised to the power s. Each merchant's cash pot is directly proportional to the product of its number of goods offered and its highest price of a good offered, with the constant of proportionality being a parameter of the model.

The number of purchases made by each customer on average is also distributed using a Zipf law. The customer's cash pot is directly proportional to the product of this average number of purchases per day and the average price of all goods on offer from merchants, with the constant of proportionality being another parameter.

Running a trial involves simulating a sequence of days of activity. On each day the following sequence of events occurs.

The merchants decide whether or not to be participants in the CC. If they participate then that means they will advertise CC prices for their goods alongside the FC prices, and they will accept CC in payment. It also means they will hold a stock of CC, revising their stock holding at the end of each day.

ledgerjournal.org
ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

In the absence of an empirically-based model of merchant decisions in this area, a system of functions was chosen to give merchant decision-making some realistic characteristics. In particular, the probability of a merchant opting in or out of the cryptocurrency on any particular day should respond to the merchant's view of the cryptocurrency at that time, which in turn should be influenced by the merchant's personal experience with the cryptocurrency and by general buzz, which is publicity such as news stories and social media activity.

With this in mind, the probability of a merchant opting in is 10^{s-2}, where s is the merchant's sentiment towards the cryptocurrency. Within the simulation various factors affect the sentiment, which starts neutrally at zero. The probability of a merchant opting out is 10^{-s-2}. To illustrate the effect of these formulae, when sentiment is neutral at zero, the probability of opting in is equal to the probability of opting out at 1 in 100. With a sentiment of -1 the probability of opting in reduces to 1 in 1,000 but the probability of opting out rises to 1 in 10. A sentiment of +1 gives a probability of opting in of 1 in 10 and a probability of opting out of 1 in 1,000.

A merchant's sentiment ranges from -2 to +2 and is updated daily according to the formula:

$$s' = (1-\lambda_{ms})s + \lambda_{ms}\left(\left(\gamma_{ms}\max\left[-2,\min\left[2,\frac{S}{S_{ms}}\right]\right]\right) + \left((1-\gamma_{ms})\max\left[-2,\min\left[2,\frac{B}{B_{ms}}\right]\right]\right)\right)$$

where s and s' are the sentiment before and after updating respectively, λ_{ms} and γ_{ms} are weighting factors, S is the merchant's sales that day paid for in CC, S_{ms} is the level of such sales that would yield a sentiment of 1, B is the day's buzz level, and B_{ms} is the level of buzz that would yield a sentiment of 1. In short, the new sentiment is a weighted average of the previous sentiment and a weighted combination of recent sales using CC and current buzz, all clipped to stay between -2 and +2.

When a merchant starts participating then all the merchant's goods are given a CC price alongside the FC price, set to bring in the FC amount if the good was sold at the CC price and then the CC received was immediately exchanged for FC. The exchange rate used for this is the exchange's rate for buying CC. When a merchant stops participating then all CC prices on the merchant's goods are removed.

Merchants decide to adjust their prices daily, weekly, or every 30 days. This schedule is decided on opting in and is selected at random: 10% daily, 18% weekly, 72% monthly. Daily revision is the equivalent to using the current exchange rates, which is a common approach with Bitcoin today, but the model focuses more on infrequent price changes because it is trying to model a cryptocurrency with a relatively stable value used as a practical means of payment

The customers decide whether or not to be participants in the CC. If they participate then that means they will look at advertised CC prices for the goods they want to buy, as well as the FC prices, and will consider paying in CC. It also means they will hold a stock of CC, revising their stock holding at the end of each day.

Customer decisions about opting in and out are modelled in a similar way to merchants, and for the same reasons. As with merchants, the probability of a customer opting in is 10^{s-2},

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

where s is the customer's sentiment towards the cryptocurrency, while the probability of opting out is 10^{-s-2}.

A customer's sentiment ranges from -2 to +2 and is updated daily according to the formula:

$$s' = (1-\lambda_{cs})s + \lambda_{cs}\left(\left(\gamma_{cs}\max\left[-2,\min\left[2,\frac{S}{S_{cs}}\right]\right]\right) + \left((1-\gamma_{cs})\max\left[-2,\min\left[2,\frac{B}{B_{cs}}\right]\right]\right)\right)$$

where s and s' are the sentiment before and after updating respectively, λ_{cs} and γ_{cs} are weighting factors, S is the customer's savings that day by paying in CC instead of FC, S_{cs} is the level of such savings that would yield a sentiment of 1, B is the day's buzz level, and B_{cs} is the level of buzz that would yield a sentiment of 1. In summary, the new sentiment is a weighted average of the previous sentiment and a weighted combination of recent personal savings using CC and current buzz, all clipped to stay between -2 and +2. In the experiments reported below buzz was held constant throughout.

The customers then do their shopping. Each day each customer buys a number of goods that is binomially distributed with n being three times the average daily purchases for the customer. Goods are chosen randomly with equal probability from the total list of goods offered by all merchants. There is no attempt to model selection of goods or merchants. The intention was to create varied purchase behavior that would reflect the wealth of each customer and create a stream of purchases to be paid for.

When deciding how to pay, customers who are participating in CC compare the FC price (if there is one) with the CC price, converted using the appropriate current exchange rate or the current price of freshly minted CC, whichever is more favourable, and chooses to pay by the cheapest route. This again reflects the model's focus on a cryptocurrency with a stable value that is used as a practical means of payment. The existence of a body of goods advertised at stable CC prices is a mechanism that could bring some stability to the cryptocurrency.

At the end of each day, merchants and customers decide how much CC they wish to own. The approach is based on a model inspired by Izumi's three phases:

1. Perception: Gathering data about a number of variables the agent believes are relevant to the value of the CC.
2. Prediction: Predicting the distribution of the price of the CC.
3. Decision: Deciding how much to buy/sell.[20]

As with decisions on opting in and out, the absence of an empirically accurate model of behavior in this context required a model with realistic characteristics, even if it is not complete or properly calibrated. The goal is only to replicate the features of realistic behavior.

The variables participants use are shown in Table 1. They cover the areas of market information (variables 1, 2, 3, 4), personal experience (variables 5a and 5b), and other general publicity including word-of-mouth, social media postings, and so on (variable 6). In a real case we imagine the market information being provided by a market quality dashboard on a public website.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

Table 1: Factors considered by participants.

	Definition
1.	Sum of the absolute values of exchange transactions, valued in CC, from the previous day, multiplied by their relative entropy.
2.	Activity of buying real goods on the current day multiplied by relative entropy, in CC.
3.	Holdings of CC by merchants and customers at the end of the previous day multiplied by their relative entropy.
4.	Exchange rate of CC at the start of the current day.
5a.	For merchants, the value of sales made in CC during the current day, valued in FC.
5b.	For customers, the value of savings made by purchasing in CC in the current day, valued in FC.
6.	The Buzz level of the current day.

Relative entropy here means the actual entropy of the distribution (H, in bits) divided by the entropy if the distribution was uniform over the same number of units.

Participants then judge the probability of CC rising or falling by considering a linear function of the exponentially weighted moving average of daily changes in the six variables. The weights used in the linear function and the recency factors used in the exponentially weighted moving averages are individually set at random for each participant, but the recency factors can be constrained to give the participants overall a greater or lesser tendency to react to recent changes.

This determines the parameters of a Normal distribution from which a probability is derived. The final choice of target CC holding level is based on the idea of dividing the participant's cash pot between CC and FC according to the Kelly Betting strategy.[21] This requires the holdings of each to be directly proportional to the probability that each will grow in value relative to the other. So, the more likely it seems that the CC rate will rise, the more of the cash pot is held in CC.

Table 2: Variables used to determine target CC holding.

Variable	Description
T_d	The amount to hold in CC, for day d, expressed in CC.
p_d	The probability of CC rising relative to FC, for day d.
C	The total cash amount held by the person to be split between FC and CC. Their 'cash pot', expressed in FC.
ccr_d	The exchange rate for CC, for day d.
μ_d	The mean inferred from the linear combination of factors.
σ_d	The standard deviation inferred from the linear combination of factors.
I	Set of indexes of the factors in use (1..6).
$e_{i,d}$	Exponentially weighted moving average of ith factor on day d.
w_i	Weight of ith factor.
λ_i	Recency parameter for ith factor.
$f_{i,d}$	Value of ith factor on day d.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

Mathematically, the variables used are listed in Table 2. These are all from the point of view of a single participant.

In summary, the target holding of CC for the participant on day d is:

$$T_d = \frac{p_d \times C}{ccr_d}$$

Where p_d is the probability of CC rising relative to FC, inferred from the distribution of changes of factors:

$$p_d = 1 - \text{Normal}(0, \mu_d, \sigma_d)$$

Where Normal is the cumulative normal distribution, μ_d is the mean and σ_d the standard deviation inferred from the linear combination of factors.

$$\mu_d = \frac{\sum_{i \in I} e_{i,d} \times w_i}{\sum_{i \in I} w_i},$$

$$\sigma_d = \sqrt{\frac{\sum_{i \in I} \left(e_{i,d} - \mu_d\right)^2 \times w_i}{\sum_{i \in I} w_i}},$$

where I is the set of indexes of the factors in use, w_i is the weight for each factor, and e_i is the exponentially weighted moving average of changes in the factors considered.

The exponentially weighted moving average is calculated from daily differences, except initially. Some rely on the situation the day before, while others are more up to date, reflecting the current day. For factors that use the previous day the formulae are:

For $d = 1$,

$$e_{i,1} = 0$$

For $d = 2$,

$$e_{i,2} = \lambda_i \times f_{i,d-1}$$

For $d \geq 3$,

$$e_{i,d} = \lambda_i \times (f_{i,d-1} - f_{i,d-2}) + (1 - \lambda_i) \times e_{i,d-1}$$

Where λ_i is the recency factor for variable i and $f_{i,d-1}$ is the value of variable i for day $d - 1$. For factors that use the current day the formulae are:

For $d = 1$,

$$e_{i,1} = \lambda_i \times f_{i,1}$$

For $d \geq 2$,

$$e_{i,d} = \lambda_i \times (f_{i,d} - f_{i,d-1}) + (1 - \lambda_i) \times e_{i,d-1}$$

The overall effect of this approach is to make participants herd followers rather than contrarians.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

If participants wish to own more CC than they currently have then they choose between buying more on the exchange and buying freshly minted CC, taking the cheapest option. If participants wish to own less CC than they currently have then they sell the surplus on the exchange.

At the end of the day the exchange market maker decides on a new price for the next day, based on the market maker's remaining pool of CC and the demand and supply experienced that day. The rule for doing this is explained under Control mechanisms, below.

6. Control Mechanisms

There are two control mechanisms in this model. One is to offer freshly minted CC at a fixed price. This is not currently a widely implemented feature, though it could be argued that Bitcoin miners are in effect buying freshly minted bitcoin by providing mining services and some 'stable coins' involve creating a unit of cryptocurrency on deposit of a unit of fiat currency. Our simulated cryptocurrency is not necessarily one that requires miners and it is likely that successful cryptocurrencies in future will not use miners in order to achieve competitive efficiency.

Changing the price of freshly minted CC changes the demand for it from participating merchants and their customers. This then affects demand for CC on the exchange, and so affects the exchange rate. Participants consider the minting price when buying goods and when deciding how to obtain more CC.

To control inflation the money supply needs to rise in line with use of, and demand for, the cryptocurrency. This is another reason for considering this control mechanism in preference to the more familiar constant supply of new coins regardless of participant growth and market activity.

This control relies only on the participants preferring to buy their CC by the cheapest means. We imagine that in reality this option would be offered by a website and would be as convenient as using the exchange. It might even be offered by the same website.

The other control mechanism is the rule used by the market maker to adjust the exchange rate for the next day. The market maker's discretion is limited but still potentially important. It makes the difference between a rate that changes too sluggishly and one that over-reacts, introducing unnecessary instability. The market maker could fix the currency rate by making no change from day to day, but the pool might be exhausted or, alternatively, the pool might end up containing all issued CC. To keep the pool in a reasonable range the market maker has to make adjustments to the rate. However, if those adjustments are too great then instability could be increased. In effect, the market maker's stock of CC provides a buffer that can be used to smooth exchange rate movements but not to limit them. The simulation has a rule for rate adjustment, as follows:

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

$$0 > P: \qquad r' = r \times 1.01^{\frac{D-S}{a}} \times \left(1 + \frac{g}{p}\right) \times \left(1 - \frac{g}{w+p}\right)$$

$$0 \le P \le w: \qquad r' = r \times 1.01^{\frac{D-S}{a}} \times \left(1 + \frac{g}{P+p}\right) \times \left(1 - \frac{g}{w-P+p}\right)$$

$$w < P: \qquad r' = r \times 1.01^{\frac{D-S}{a}} \times \left(1 + \frac{g}{w+p}\right) \times \left(1 - \frac{g}{p}\right)$$

where D is demand for CC that day, S is supply of CC that day, P is the size of the market maker's pool of CC, a is an attenuation factor that reduces the size of price adjustments, g and p limit the maximum adjustment in either direction, and w is the width of the range within which the multiplier varies with P. The effect of this is to provide an adjustable function of pool size and net demand for CC that reacts to imbalances in supply and demand but also tries to keep the pool size within a range (Figure 1).

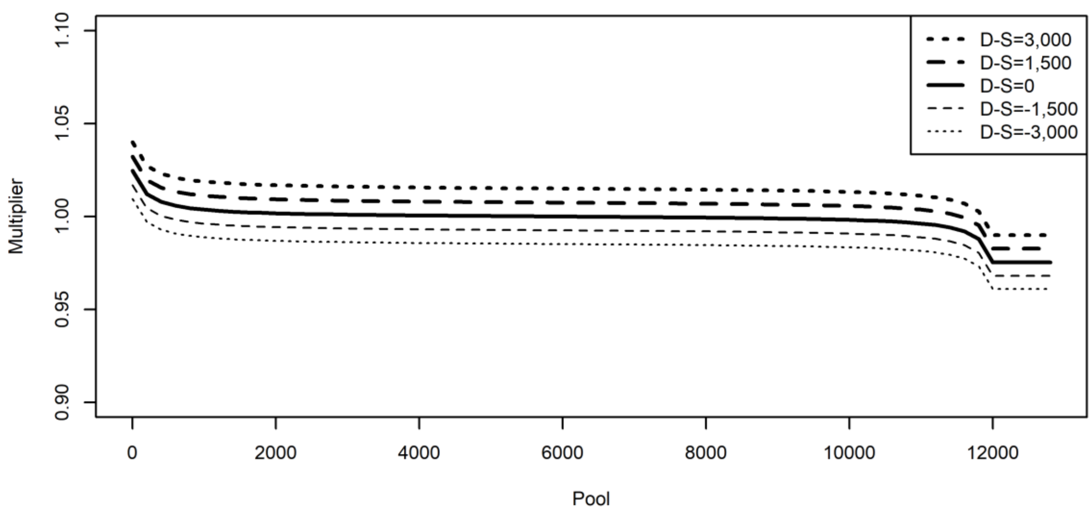

Figure 1: Example of the exchange rate multiplier function used to adjust the rate according to current pool size and Demand – Supply (D-S).

This rule is designed for a system where the exchange rate is revised daily, which would in practice only be feasible for a cryptocurrency that already had a very stable exchange rate and where the customers were ordinary shoppers rather than currency traders. More complex and computationally intensive simulation is required to model the more likely situation of a market maker who revises prices much more frequently, or an order-driven market.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

7. Typical Simulation Behavior

A qualitative understanding of what typically happens in trials will provide useful context for reporting the effect of the control mechanisms. In a typical trial, customers' and merchants' CC use and holdings tend to rise as customers and merchant opt in. This effect reduces as they approach an approximate equilibrium point where joiners and leavers more or less balance. As they join they demand CC and this draws on the market maker's pool, stimulating an increased exchange rate. Activity is typically sporadic and participants' views on how much CC to hold fluctuate wildly at times, while at other times hovering at a fairly constant level. As intended, controlling this behavior is extremely challenging.

8. Effect of Controlling the Minted Price

The effect of offering freshly minted CC should be to cap the rate offered by the exchange. If participants can buy the CC they want more cheaply from the mint than from the exchange, then demand at the exchange will disappear that day, leading the exchange to revise its rate downwards.

The effect of this control was tested in the simulation by comparing 100 trials of 730 days each with a minted price of 10 (too high to have any effect) against a minted price of 1.2 (just above the starting exchange rate of 1). The parameters held constant over all conditions reported in this paper are given in Table 3.

Table 4 shows comparative performance for two fixed price levels for newly minted CC, with moderately responsive exchange rate adjustments by the market maker. Figures 2 and 3 show the dramatic difference in distribution of the exchange rate at the end of each trial (after 730 days). The minted price of 10 is too high to have an effect and in a typical trial the exchange rate rises most of the time, with no new CC being issued. With the mint price of 1.2, typically, the exchange rate rises, hits the cap of 1.2, there is a burst of new CC issued, the exchange rate drops, and then tends to oscillate below 1.2, gradually tending to fall a bit further on average.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

Table 3: Parameter values held constant between conditions.

Name	Description	Value
r_seed	Random number seed – or -1 for default	-1
m_max	Number of merchants	10
m_gnave	Average number of goods offered by each merchant	5
m_gineq	Inequality of number of goods offered by each merchant	1
m_pave	Average price of first good for a merchant	3
m_pineq	Inequality factor for price of first good for each merchant	1
c_max	Number of customers	40
c_pave	Mean purchases per day of each customer	3
c_ineq	Inequality of number of purchases per day per customer	1
m_pot_mult	Multiplier used to set merchants' cash pots	10
c_pot_mult	Multiplier used to set customers' cash pots	20
ccr_start	r, Initial FC/CC exchange rate (mid-point)	1
cc_fxl	Fraction of currency lost in a round trip at the exchange	0.03
pool_start	P, Initial CC pool held by the currency exchange	5000
t_days	Number of days to do in one trial	730
s_trials	Number of trials to do in one MC Simulation	100
m_now_prob	Probability of setting 'now' prices, given opting in	0.1
m_wkly_prob	Probability of setting weekly prices, given not setting 'now' prices	0.2
fmask	Mask for including factors in the probability judgements	[1,1,1,1,1,1]
rcncy	Range for lambdas of merchants and customers	[0,0.1]
exc_type	Exchange rate revision method	9
exc_width	w, Intended range of exchange's pool	12000
exc_greatest	g, Controls multiplier at pool=0	5
exc_peak	p, Also controls multiplier at pool=0	200

Table 4: Performance metrics with and without a minting cap. Sample standard deviations in parentheses.

	Not Capped	Capped
Price of newly minted CC	10.0	1.2
Attenuation parameter, a	2,000	2,000
Mean final exchange rate	5.37 (1.05)	0.96 (0.17)
Mean maximum exchange rate	5.37 (1.05)	1.19 (0.01)
Mean minimum exchange rate	1.00 (0.00)	0.92 (0.15)
Mean average exchange rate	2.56 (0.31)	1.07 (0.08)
Mean standard deviation of exchange rate	1.250 (0.296)	0.071 (0.053)
Mean fractal dimension of exchange rate	1.047 (0.022)	1.349 (0.095)
Mean absolute exchange rate movements	0.0345 (0.018)	0.0060 (0.001)
Mean absolute daily exchange rate changes	0.0081 (0.0022)	0.0025 (0.0006)
Mean maximum pool size	4,934 (199)	8,859 (858)
Mean minimum pool size	338 (147)	743 (253)
Mean number of pool outs	0 (0)	0 (0)
Mean total CC in issue at trial end	5,000 (0)	9,938 (824)

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

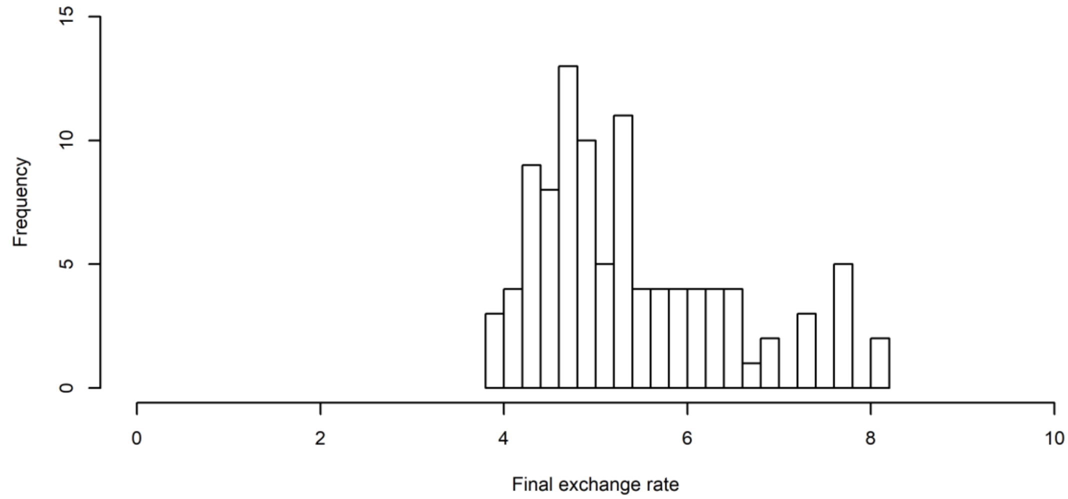

Figure 2: Distribution of final exchange rate with no minting due to high minted price.

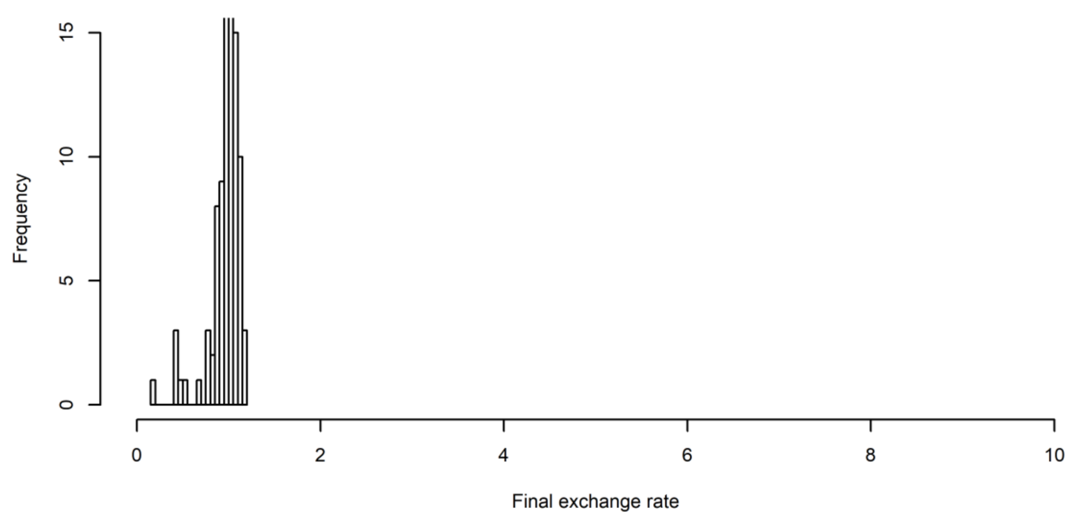

Figure 3: Distribution of final exchange rate with price of newly minted coins set at 1.2.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

9. Effect of Different Exchange Rate Revision Rules

The effects of changing the responsiveness of the exchange rate revision rule were again tested using 100 trials for each condition, with simulations over 730 days. Performance data are given in Table 5 for five levels of responsiveness.

Responsiveness	**High**				**Low**
Price of newly minted CC	1.2	1.2	1.2	1.2	1.2
Attenuation parameter, a	500	1,000	2,000	4,000	8,000
Mean final exchange rate	1.08	1.05	0.96	0.77	0.49
Mean maximum exchange rate	1.21	1.20	1.19	1.19	1.20
Mean minimum exchange rate	1.00	0.98	0.93	0.76	0.49
Mean average exchange rate	1.12	1.11	1.07	0.96	0.79
Mean standard deviation of exchange rate	0.042	0.047	0.071	0.135	0.233
Mean fractal dimension of exchange rate	1.522	1.465	1.349	1.170	1.035
Mean absolute exchange rate movements	0.0178	0.010	0.006	0.0073	0.0225
Mean absolute daily exchange rate changes	0.0089	0.005	0.0025	0.0016	0.0014
Mean maximum pool size	8,107	8,354	8,859	9,818	11,103
Mean minimum pool size	1,065	850	743	711	743
Mean number of pool outs	0	0	0	0	0
Mean total CC in issue at trial end	8,108	9,447	9,9396	10,904	12,148

Table 5: Performance metrics with different levels of responsiveness, set by the attenuation parameter, a.

Two interesting effects can be observed. First, as the responsiveness increases, the exchange rate tends to vary by more each day (Figure 4). This is revealed by the increasing average daily absolute difference in exchange rate, and by the fractal dimension, which was calculated using a refinement of Higuchi's (1988) method.[22] If our goal is a stable exchange rate that does not suffer from moment-to-moment changes then this increasing responsiveness is a bad thing.

ledgerjournal.org
ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

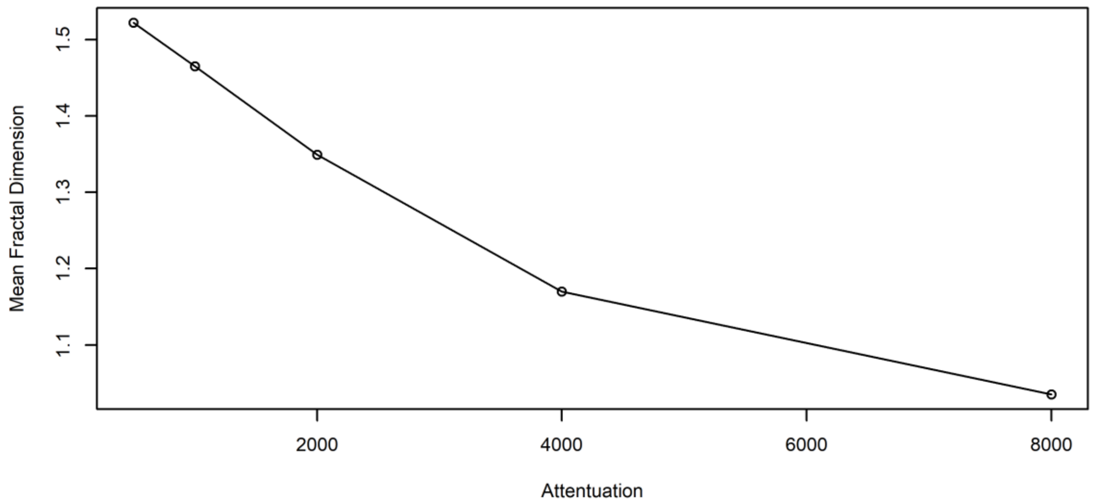

Figure 4: Reducing responsiveness (increasing attenuation, a) decreases daily changes as shown by the fractal dimension.

However, the second effect is a problem that arises when the responsiveness is low. With a minting cap in operation and very low responsiveness the exchange rate often rises initially, hits the cap, then slowly declines to zero. The reason for this steady decline is probably that our simulated participants are herd-following investors and, if a trend is established for long enough, it tends to become permanent unless something powerful reverses it.

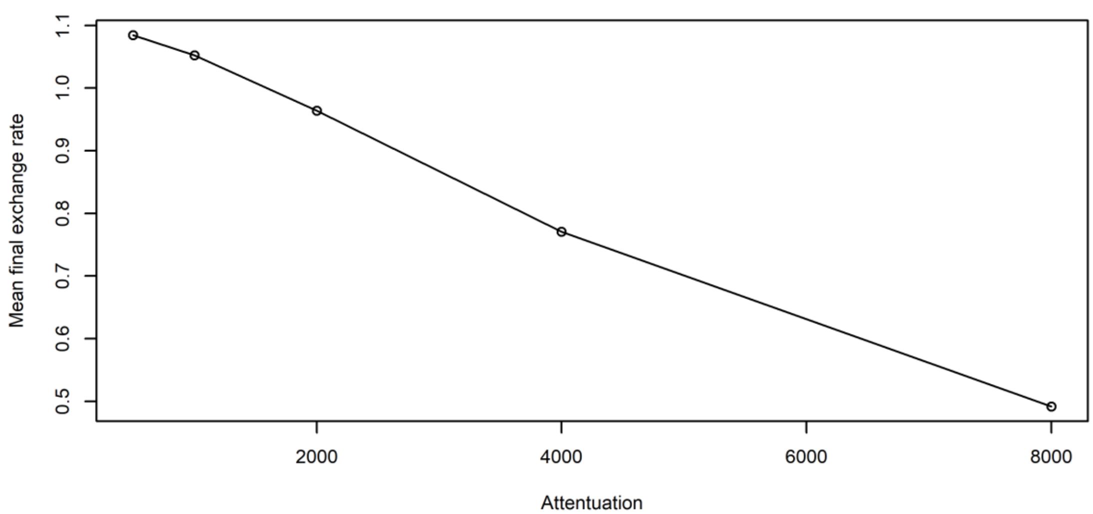

Figure 5: Reducing responsiveness (increasing attenuation, a) decreases mean final exchange rate as more and more simulations end in terminal decline.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

10. Discussion

The simulation trials described above show how a simulator might be used to test control mechanisms for future cryptocurrencies. This approach might complement more mathematical analysis of mechanisms that perhaps tries to establish facts about highest and lowest exchange rates, or about asymptotic behavior.

The results of any simulation should be interpreted with caution, and that is especially true of a simplified model of a hypothetical system. In this case the system being simulated is inherently one where exact predictions are not feasible; the precise mechanisms driving outcomes are not fully understood, there are many factors outside the scope of the simulated agents, and users of different cryptocurrencies might think differently.

Observations that agree with expectations will naturally be more credible than others. With the minting price effect, it is a logical prediction based on an uncontroversial assumption that the minting price will cap the price on the exchange.

The capping effect of selling freshly minted cryptocurrency at a controlled price could be useful with real cryptocurrencies as a means of preventing large price increases. It involves releasing new coins in response to demand. This contrasts with the more familiar approach of minting new coins according to a fixed schedule that is not responsive to the level of demand for the coins.

This raises a general question: should we expect to achieve good economic control of a cryptocurrency using fixed schedules for any control element, or is some kind of responsiveness to events required in all or most cases? Observations on real cryptocurrencies, with their characteristically large, rapid changes in value, suggest that responsive controls are needed.

The effect of changing the rule by which the cryptocurrency exchange revises the exchange rate raises another general question: which elements of a cryptocurrency ecosystem should be included in the overall design and brought to bear on economic control? The familiar approach from most existing cryptocurrencies is to control only the minting of coins, leaving others to provide exchanges independently. If the exchange's rules are important for the stability of the currency then perhaps, as illustrated by the simulation, the exchange should be part of the control system too, and designed as such from the start.

Another potential control mechanism is to offer some goods or services at a controlled price expressed in the cryptocurrency. For example, one coin might be worth a burger, or a carwash. In the 1960s and 1970s Green Shield Stamps were offered in the UK. Shoppers enthusiastically collected the stamps and stuck them in the collecting books provided in order to trade them for goods advertised in a large catalogue. The Green Shield Stamp prices in the catalogue gave the stamps a real value that was understood by users of the stamps and stable while each year's catalogue was in use.

The goods or services involved need not be those offered by a company that has issued the cryptocoins and is intent on controlling them for its own ends. Cryptocurrency communities might agree that, as part of the design of a cryptocurrency, the coins will always be worth some standardized unit of labour or other resource.

Of all the control mechanisms discussed so far, this is the one with the best chance of preventing a cryptocurrency's exchange rate from declining to zero and staying there. This is a

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

problem that appeared in the simulations and has also occurred many times with real cryptocurrencies.

Control might also be exercised by providing, or not providing, information to participants via a public market quality dashboard. The simulation described above gives participants knowledge of the distribution of CC holdings and transaction values because this provides insight into the true popularity of the cryptocurrency. At present this distribution information is not usually available for cryptocurrencies.

All these controls might be tested using a suitable simulator capable of representing the design of a cryptocurrency and its ecosystem and, since completing the trials described above, we have developed a detailed design for such a simulator.

11. Areas for Further Research

Further research might improve on this initial work in a number of ways.

Alternative objective functions. In the trials reported above, the overall objective was to stabilize the exchange rate of the cryptocurrency against an established fiat currency. This involved limiting the overall range of exchange rates and the pace of variation. The single established fiat currency was a convenient alternative to modelling a basket of such currencies and trying to control the CC's rate against them. However, in a more advanced model additional currencies might be introduced with the objective being to maintain a stable exchange rate against a basket of those currencies. More sophisticated objectives might also pay attention to other factors crucial to the sustained success of a cryptocurrency. In particular, this might include the electricity consumption of the whole network, and resulting transaction costs. With Bitcoin, the huge number of full nodes now participating means that the work of recording a single transaction is duplicated over thousands of computers, making the system unsustainably inefficient.

More realistic simulation of challenging behavior. Although it is not possible to exactly predict the evolution of a cryptocurrency over a long period of time, it should be possible to develop simulations that offer a wider range of more realistic challenging behaviours. This might include rapid rises and falls, accelerating rises followed by steep falls (a bubble bursting), herd following and contrarian reactions, over- and under-reaction to events, more heterogeneous agents, unpredictable external influences (*e.g.* promotional activity, successful hacking, arrests, illegal activity, regulatory threats or changes), and even deliberate attempts to push exchange rates in a particular direction. More sophisticated simulations should also allow exchange rates to be adjusted more frequently and could also simulate an order-driven exchange.

More sophisticated control schemes. In the above simulations controls had fixed parameters, but much more is possible. The controls could be improved by testing a greater variety of mechanisms, using different rules for controlling responses, using the mechanisms in combinations, and even giving the control system the ability to learn over time.

Larger volume of experimentation. To test more control schemes and discover approaches that work effectively it might be possible to open up the simulation to make it available to more people. Another direction might be to attack the problem with machine learning

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

techniques, perhaps using the Markov Decision Problem paradigm with Reinforcement Learning methods.

12. Conclusion

We have argued that cryptocurrencies that aim to be successful as currencies need to be controlled in such a way that their value does not fluctuate greatly over time. Without this control they are suitable only for speculation and have a role similar to online poker – an entertaining test of skill and nerve with no practical use. That control need not be provided by ad hoc human intervention, though this may still need to be the ultimate control. Instead, it could be provided by open and pre-agreed software mechanisms and rules. However, such control mechanisms need to be well designed and should be tested in advance through some kind of simulation. The simulation trials described in this paper illustrate in a very simple way the potential of this approach. If a control scheme does not work in a simulation then it is probably not safe to release it into the real world.

Future control mechanisms are likely to be responsive to events – especially demand – rather than relying on fixed plans for minting coins. They may also involve bringing more of the elements of a cryptocurrency ecosystem into the control scheme.

Acknowledgement

Cardano Foundation, along with Z/Yen Group, supported this research financially.

Notes and References

[1] Hileman, D. G., Rauchs, M. "2017 Global Cryptocurrency Benchmarking Study." Cambridge Centre for Alternative Finance, University of Cambridge Judge Business School (2017) `https://ssrn.com/abstract=2965436`.

[2] ElBahrawy, A., Alessandretti, L., Kandler, A., Pastor-Satorras, R., Baronchelli, A. "Evolutionary Dynamics of the Cryptocurrency Market." *Royal Society Open Science* **4.11** 170623 (2017) `https://doi.org/10.1098/rsos.170623`.

[3] Jevons, W. S. *Money and the Mechanism of Exchange*. New York: D. Appleton and Co. (1875). Text is available online at: `https://www.econlib.org/library/YPDBooks/Jevons/jvnMME.html`.

[4] Baek, C., Elbeck, M. "Bitcoins as an Investment or Speculative Vehicle? A First Look." *Applied Economics Letters* **22.1** 30-34 (2014) `https://doi.org/10.1080/13504851.2014.916379`.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

[5] Bolt, W., van Oordt, M. R. "On the Value of Virtual Currencies." *Bank of Canada*. Staff Working Paper (2016) https://www.bankofcanada.ca/wp-content/uploads/2016/08/swp2016-42.pdf.

[6] Bouoiyour, J., Selmi, R., Tiwari, A. "Is Bitcoin Business Income or Speculative Bubble? Unconditional vs. Conditional Frequency Domain Analysis." *MPRA* (2014) https://mpra.ub.uni-muenchen.de/id/eprint/59595.

[7] Cheah, E., Fry, J. "Speculative Bubbles in Bitcoin Markets? An Empirical Investigation into the Fundamental Value of Bitcoin." *Economics Letters* **130** 32-36 (2015) http://eprints.whiterose.ac.uk/84108/

[8] Cheung, A., Roca, E., Su, J. "Crypto-Currency Bubbles: An Application of the Phillips–Shi–Yu (2013) Methodology on Mt. Gox Bitcoin Prices." *Applied Economics* **47.23** 2348-2358 (2015) https://doi.org/10.1080/00036846.2015.1005827.

[9] Ciaian, P., Rajcaniova, M., Kancs, d'A. "The Economics of BitCoin Price Formation," *Applied Economics* **48.19** 1799-1815 (2016) https://doi.org/10.1080/00036846.2015.1109038.

[10] Matta, M., Lunesu, I., Marchesi, M. "Bitcoin Spread Prediction Using Social and Web Search Media," *UMAP Workshops* (2015) https://ceur-ws.org/Vol-1388/DeCat2015-paper3.pdf.

[11] Luther, W.J., White, L.H. (2014) "Can Bitcoin Become a Major Currency?" *GMU Working Paper in Economics* 14-17 (5 June 2014). https://dx.doi.org/10.2139/ssrn.2446604.

[12] De Grauwe, P., Vansenten, K. "Deterministic Chaos in the Foreign Exchange Market," *Center for Economic Policy Research* Discussion Paper 370 (1990) https://ideas.repec.org/p/cpr/ceprdp/370.html.

[13] Harris, I., Mainelli, M. R. *The Price of Fish*. London: Nicholas Brealey Publishing 156-182 (2011).

[14] Chatagny, R., Chopard, B. "A Parallel Model for the Foreign Exchange Market," *Parallel Computing* **26.5** 587-600 (2000).

[15] Delage, V., Brandlhuber, C., Tuyls, K., Weiss, G. "Multi-Agent Based Simulation of FOREX Exchange Market," in P. De Causmaecker, J. Maervoet, T. Messelis, K. Verbeeck, T. Vermeulen (Eds.), *Proceedings of the 23rd Benelux Conference on Artificial Intelligence (BNAIC)* (2011) http://lirias.kuleuven.be/1675903.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130

[16] Setzu, A. "A Framework for Financial Markets Modeling and Simulation," *Exchange Organizational Behavior Teaching Journal* (2007) `https://www.diee.unica.it/DRIEI/tesi/19_setzu.pdf`.

[17] Usami, A., Tsuya, R., Iba, T., Takayasu, H. "Building a Simulation Model of Foreign Exchange Market: Reproduction of Yen Dollar Market." 9th Joint International Conference on Information Sciences (JCIS-06) (2006) `https://doi.org/10.2991/jcis.2006.319`.

[18] Cocco, L., Concas, G., Marchesi, M. "Using an Artificial Financial Market for Studying a Cryptocurrency Market." *Journal of Economic Interaction and Coordination* **12.2** 345-365 (2017) `https://doi.org/10.1007/s11403-015-0168-2`.

[19] `http://ledger.pitt.edu/ojs/index.php/ledger/rt/suppFiles/130/0`.

[20] Izumi, K. "An Artificial Market Model of a Foreign Exchange Market," *Citeseer^X* (accessed 11 May 2019) `https://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.143.5360`.

[21] Kelly, J. L., "A New Interpretation of Information Rate." *Bell System Technical Journal* **35.4** 917-926 (1956).

[22] Higuchi, T. "Approach to an Irregular Time Series on the Basis of the Fractal Theory." *Physica D: Nonlinear Phenomena* **31.2** 277-283 (1988).

Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.

ISSN 2379-5980 (online)
DOI 10.5915/LEDGER.2019.130