

Essays on Bitcoin

By Alex Kroeger

With Advisor Professor Tim Fuerst

Abstract:

The following paper analyzes two distinct topics related to the virtual currency bitcoin. The first is an empirical test of purchasing power parity using volume weighted price data from bitcoin exchanges that facilitate transactions in U.S. dollars, euros, and British pounds. Evidence shows that relative purchasing power parity does indeed appear to hold, but that there is a persistent deviation from absolute purchasing power parity. The second topic is an analysis of bitcoin mining from an economic perspective. A simple model demonstrates that competition in bitcoin mining leads to a great deal of waste compared to the outcome that would be preferred by a central planner seeking to maximize welfare.

Chapter One: Introduction

Topics of Discussion and Structure of the Paper

Since its inception, bitcoin, a virtual currency, has grown in both its popularity and its use. Despite this, there still exists a relative dearth of economic analysis in academia about this new economic phenomenon. Various topics have been researched with regard to bitcoin, including its economic status as a currency (Yermack 2013), the incentives of bitcoin miners (Kroll et al 2013), the economics of bitcoin exchange prices (Ciaian et al 2014), among others. Macroeconomist Paul Krugman weighed in strongly on the normative side of the economic debate with his article “Bitcoin is Evil” published in late 2013.

The object of this paper is to investigate a yet unexplored topic in bitcoin, purchasing power parity, and utilize a different approach to the topic explored by Kroll et al, the economics of bitcoin mining. Therefore, the structure of the paper will be as follows. The remainder of chapter one will provide a brief history of bitcoin and an explanation of how bitcoin operates. Chapter two, entitled “Purchasing Power Parity in the Bitcoin Exchange Market” will analyze bitcoin from the perspective of purchasing power parity across three different currency exchange markets in which bitcoin trades. Chapter three, entitled “An Economic Analysis of Bitcoin Mining,” analyzes bitcoin from the perspective of the users that verify bitcoin transactions. Chapter four concludes. Note that there are separate appendices at the end of chapters two and three. All references are listed at the end.

The History of Bitcoin

Bitcoin (sometimes known by its generally accepted ticker BTC) is an online payment system launched as an open source software in 2009. Its creator (or creators), whose identity to

this day remains unknown, goes by the name Satoshi Nakamoto. Nakamoto published a paper describing his or her creation entitled “Bitcoin: A Peer-to-Peer Electronic Cash System” in 2008.

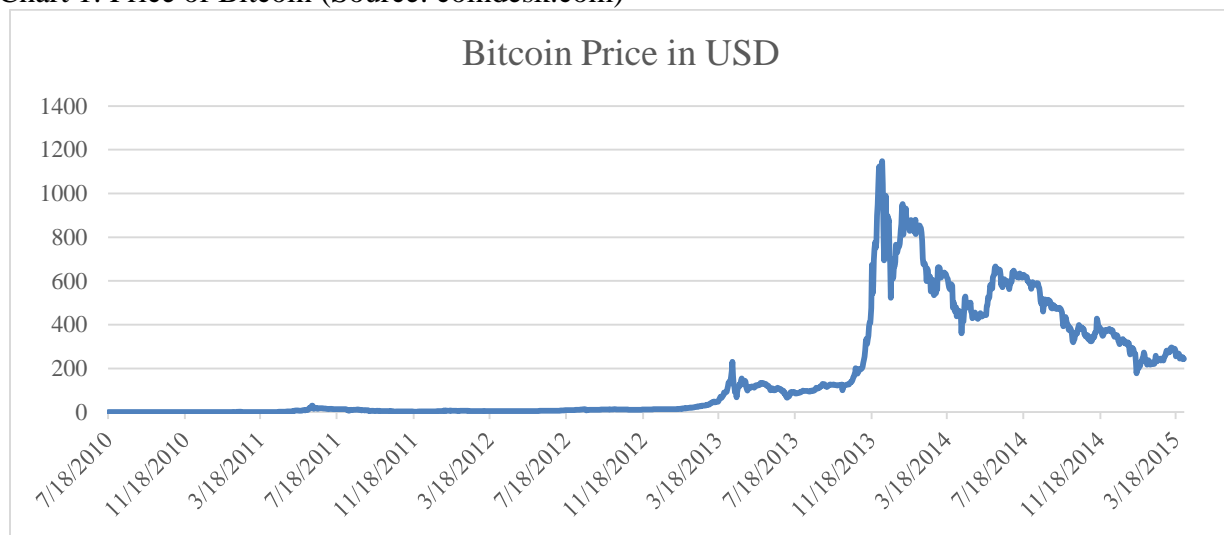
In many ways it functions as a currency, whereby one party can send a unit of currency (in this case a string of code) to another party in exchange for a good or service. As such, bitcoin is often referred to as a “virtual currency” or “cryptocurrency.” All transfers of bitcoin are verified and then recorded on a public ledger known as the *block chain* (Velde 2013).

Potential users can purchase bitcoin by using an online exchange. These exchanges act as either brokers or dealers in allowing users to convert a major currency such as the U.S. dollar into bitcoins. The first of these exchanges, Bitcoin Market, open in February 2010. Another exchange, MtGox, first launched in July of that same year.¹

Based on the dollar values at these exchanges, the market capitalization of bitcoin (the number of bitcoin in circulation multiplied by the market price in dollars) exceeded \$1 million by October 2010. By March 2013, the market cap surpassed \$1 billion.¹ During that time, bitcoin had come to be accepted as payment by a variety of businesses and organizations, from Baidu in China to coffee shops in Palo Alto and antique shops in New Orleans (Fung 2013, Hill 2013). One writer for Forbes, Kashmir Hill, was actually able to live for a week in San Francisco in May 2013 using nothing but bitcoin to make purchases (Hill 2013-2). In October and November 2013, interest in bitcoin in China surged, making BTC China, a Shanghai based bitcoin exchange, the largest in the world for a brief time (Hill 2013-3). The price of a bitcoin surged to over \$1000 as many users in China begin to invest in bitcoin, but has since declined to the 200’s (see Chart 1 below).

¹ Historical events drawn from historyofbitcoin.org.

Chart 1: Price of Bitcoin (Source: coindesk.com)

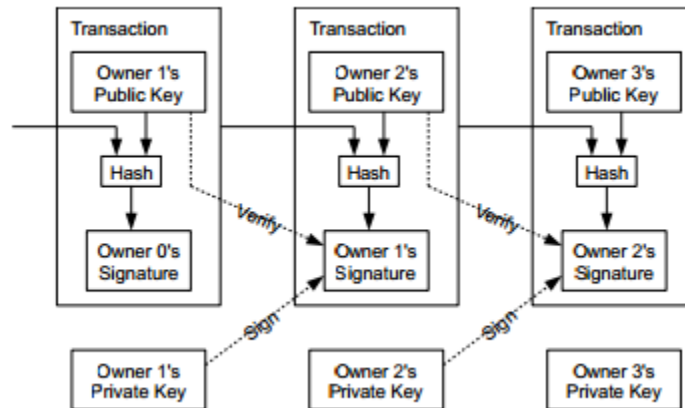


One reason for the growth in bitcoin’s popularity (and notoriety) is the anonymity of its users. Although all transactions are recorded on a public ledger, only one’s public address is associated with the transfer’s one makes. One’s public address contains no identifying information in and of itself, and so as long as the public address is not associated with any identity, transactions remain anonymous. This anonymity has made bitcoin the currency of choice for the so-called “darknet”--websites that sell illegal commodities such as drugs and weapons. One such notable website was “Silk Road,” which was shut down by the United States government in October 2013. Since its closure, a number of new websites have emerged to take its place and have adopted the model of using bitcoin as a medium between buyers and sellers (Power 2014).

How Bitcoin Works

To transact in bitcoin, one broadcasts to the bitcoin network the public key of the payee and the amount of bitcoin one intends to transfer. Every bitcoin address has an associated *private key* that acts as a password to ensure that all transfers are authorized. The private key is

meant to remain secured, and along with one's own public key, it allows one to digitally sign bitcoin transactions. A graphical representation of a series of bitcoin transactions from Satoshi Nakamoto's original paper is reproduced below.



The primary concern of the payee is that the amount of bitcoin being transferred has already been spent, and therefore does not belong to the payer. Another concern is the rate of creation, since a high degree of inflation could reduce the value of one's holdings. What allows bitcoin to be functional is that it overcomes these two major obstacles facing any digital currency: avoiding double spending and controlling creation (Velde 2013). Both of these problems are solved in the process of *mining*.

As transactions are broadcast over the bitcoin network, "miners" work to collect transactions into a group, known as a "block," to be added to the block chain. Every block must be accompanied by a hash (a string of characters of a fixed length generated by a set function) that depends on the list of transactions, the hash of the previous block, and a value called a nonce, which is imputed by the miner. Miners work to find a nonce such that the hash for the block meets the requirements set out by the system. The hash serves as a proof-of-work, since it is difficult to compute (the only usable method is simply to input values until a working nonce is found), but easy to verify using the hash function. Once an acceptable hash is found, the

successful miner broadcasts the block to the network, which is accepted as long as all the transactions are authorized and the hash is valid. As a reward for the computer power expended to verify transactions, bitcoin miners receive newly created bitcoins. The difficulty of mining adjusts every two weeks so that blocks are mined at a rate of one every ten minutes, thus controlling the rate at which new bitcoins are created. Mining will be discussed in greater detail in chapter three.

Chapter 2: Purchasing Power Parity in the Bitcoin Exchange Market

I. Introduction

Given that trading of bitcoin occurs in numerous currencies, it is possible to utilize bitcoin markets to test the theory of purchasing power parity, which states that items sold in several countries should trade for the same price after adjusting for nominal exchange rates. The object of this paper is to analyze whether purchasing power parity holds in bitcoin markets for dollars, euros and British pounds.

The remainder of the paper is organized as follows. Section 2 provides background information on bitcoin exchanges and the study of purchasing power parity. Section 3 summarizes the sources for the data used in the analysis. Section 4 provides a discussion the methodology employed to analyze real exchange rates and the primary findings. Section 5 concludes.

II. Background

Bitcoin Exchanges

An average user can buy and sell bitcoin for major sovereign currencies using online exchanges. Exchanges act as either brokers, serving as a platform over which buyers and sellers come together, or as dealers, taking on an inventory of bitcoin to sell to users and profiting from bid-ask spreads. BTC-E is an example of the former, whereas Coinbase is an example of the latter. Purchased bitcoin can then be stored in an electronic “wallet,” which is a feature often offered by the exchange itself, for later use.

Bitcoin users span the globe, and exchanges do as well. According to bitcoincharts.com, bitcoin currently (as of April 2015) trades in 26 sovereign currencies via exchanges with the

most popular currencies by volume being the Chinese yuan, the U.S. dollar, the euro and the British pound.

Bitcoin exchanges are not without drawbacks, however, since using them exposes users to the risk that the exchange can fail while in possession of one's bitcoins or other currency. Most notable was the collapse of MtGox, an exchange that was based in Tokyo and previously was the world's largest. In February 2013, MtGox went offline, eventually disclosing the reason: hackers had stolen 850,000 bitcoins, the equivalent of \$460 million at the time, leaving many investors unable to recuperate their assets stored with the site (McMillan 2014). In January 2015, hackers stole over \$5 million worth of bitcoin from Bitstamp, a major exchange based in Slovenia, forcing it to shut down briefly (Frey 2015). Another complete shutdown occurred in March of 2015 when a Canadian exchange called Cavirtex experienced a serious security breach (Reader 2015).

While security concerns remain, bitcoin exchanges are generally quite transparent with regard to their operations. Virtually all major exchanges make publicly available real time data about prices and volume for use by investors and web applications. Historical data is collected and stored by many organizations, notably bitcoincharts.com and bitcoinaverage.com. This availability of data makes bitcoin a viable topic of empirical economic investigation.

The Real Exchange Rate and Purchasing Power Parity

Purchasing power parity (PPP) is the theory that the ratio of price levels between two countries should equal the nominal exchange rate between the two countries. In other words, an amount of currency should be able to purchase the same basket of goods in any country once it is exchanged at the nominal rate.

The primary motivation for believing that PPP should hold true is the possibility of international goods arbitrage (Taylor & Taylor 2004). If the price of a good expressed in a common currency differs across countries, an arbitrage opportunity would exist whereby one could purchase the good where it is relatively cheap and sell it where it is relatively expensive, realizing a riskless profit. Similarly, a producer of the good would be able to concentrate shipments to countries in which the good is relatively more expensive. The ability to take advantage of such opportunities should create the pricing pressures to ensure that PPP holds.

There are many reasons to believe that PPP would not hold, however, if there are limitations to international goods arbitrage. Such factors relate to the tradability of the good, such as transaction costs, tariffs, and taxes. Certain services, such as haircuts, are hardly tradable at all. In cases where obstacles to trade exist, one would not expect PPP to hold perfectly (Rogoff 1996).

Another issue is whether or not the basket of goods is exactly comparable in each nation. For this issue it is important to consider two different senses in which PPP could hold. One sense is absolute PPP, whereby the purchasing power of a unit of currency of one nation exactly equals the purchasing power of the amount of currency of another nation after accounting for the nominal exchange rate. The absolute sense would not hold if the basket of goods differs between the two countries, however, since the underlying goods would not be exactly comparable. The other sense that avoids this issue is relative PPP, whereby the proportional change in the exchange rate should offset the difference in inflation rates for the two countries (Rogoff 1996).

Many empirical tests of PPP have been performed. A primary method of testing PPP is to perform unit root tests on the real exchange rate—a method that will be used in this paper (Taylor & Taylor 2004). The real exchange rate is defined as the relative price level adjusted by

the nominal exchange rate. An example calculation of the real exchange rate, q , is shown below for two hypothetically countries, country A and country B. P represents the price of a basket of goods, and C is the respective currency for each country.

$$q = \frac{P_A}{P_B} \times \frac{C_B}{C_A} = \frac{\frac{C_A}{Basket_A}}{\frac{C_B}{Basket_B}} \times \frac{C_B}{C_A} = Basket_B / Basket_A$$

Theoretically, the real exchange rate should be one if the price indexes used reflected the exact same basket of goods (reflecting absolute PPP). In practice, price indices across countries will not reflect the same exact underlying basket of goods, and so the real exchange rate will likely differ from one. If relative PPP holds, however, one would still expect the real exchange rate to remain stationary—any changes in the real exchange rate would represent deviations from relative PPP. Therefore, any changes in the real exchange rate should eventually revert to a common mean. In other words, the process should not have a unit root, whereby shocks become permanent deviations.

The unit root test is estimated by using an autoregression. An example of an autoregression with one lagged term is shown below with q signifying the real exchange rate.

$$q_t = \beta_1 q_{t-1} + \varepsilon_t$$

If the real exchange rate follows a unit root process, then β_1 would be equal to one. This would mean the change in real exchange rates would on average be zero (since no drift term is specified), but the level would not be predictable in the long run. A generalized version could consist of many lagged terms (shown below), and in that case if the series followed a unit root process, the sum of the terms $\beta_1 \dots \beta_n$ would be one.

$$q_t = \beta_1 q_{t-1} + \beta_2 q_{t-2} + \dots + \beta_n q_{t-n} + \varepsilon_t$$

Early empirical tests of a unit root the real exchange rate in the late 1980's suggested that the null hypothesis that the real exchange rate series follows a unit root process cannot be rejected. This was most likely due to the low power of the tests and the limited number of years utilized, however. Tests using longer time spans resulted in successful rejections of the null that real exchange rate series had unit roots (Taylor & Taylor 2004).

III. Data

Data used in the following analysis is drawn from several sources. Bitcoin price data is generated using bitcoinaverage.com's price index. The price index is calculated by collecting data from a multitude of major exchanges for a particular currency. A weighted average price is then estimated using the prices from each exchange and weights according to the volume each exchange experiences. The following analysis utilizes the weighted price index in US dollars, British pounds, and euros. The historical data is made available for convenient download by Quandl, which is the source used for the following analysis. The time period considered is September 2013 to October 2014.

Nominal exchange rate data is also used in order to calculate real exchange rates. The daily U.S. dollar to euro and the daily U.S. dollar to British pound nominal exchange rates were obtained from the Federal Reserve Economic Data (FRED) site maintained by the St. Louis branch of the Federal Reserve. The daily euro to British pound exchange rate is made available by the European Central Bank.

Price index data are drawn from governmental sources for the US and the UK. These data were obtained from the US Bureau of Labor Statistics and the UK Office for National

Statistics, respectively. Price index data in each case is available on a monthly basis. The time period considered is January 1996 to October 2014.

IV. Methodology and Results

Testing for a Unit Root: the Augmented Dickey-Fuller Test

As discussed in the background on purchasing power parity, a common method of testing PPP empirically involves utilizing a unit root test on the time series. One such test is the augmented Dickey-Fuller test. The equation involved in calculating the test statistic deviates somewhat from a standard autoregressive model specified in the background section. For a real exchange rate q and n lagged terms, the model is specified as follows (Cheung & Lai 1995).

$$\Delta q_t = c + \gamma q_{t-1} + \delta_1 \Delta q_{t-1} + \delta_2 \Delta q_{t-2} + \dots + \delta_n \Delta q_{t-n} + \varepsilon_t$$

If γ in the model equals zero, the process follows a unit root. In other words, $\gamma = 0$ would entail the previous term in the series provides no information about the current change. If the process is mean reverting, one would expect that the lagged term would have an effect on the change in the current period. If the process is stationary and mean reverting, then γ should be negative. The test statistic in the augmented Dickey-Fuller test, then, is the estimate of the coefficient $\hat{\gamma}$ divided by its estimated standard error from an ordinary least squares regression. Shown another way, the test statistic, labeled *ADS*, is calculated as follows.

$$ADS = \frac{\hat{\gamma}}{SE(\hat{\gamma})}$$

The closer the test statistic is to zero, the less likely one can reject the null that the series follows a unit root process. The more negative the test statistic, the more likely one can reject

the null. The test statistic and appropriate critical values for subsequent analysis were calculated using Eviews statistical software.

Country-Level Real Exchange Rate

Using data on nominal exchange rates and price indexes, one can construct a real exchange rate and test relative PPP. Using consumer price indexes and nominal exchange rates for the United States and the United Kingdom, a real exchange rate was constructed for these two countries according to the specifications shown in the appendix as Table 1. The resulting value is the price of British goods in terms of U.S. goods. The timeline considered is January 1996 through October 2014. The data series and its corresponding mean (.88) are displayed in Chart 2 below.

Chart 2: US-UK CPI Real Exchange Rate

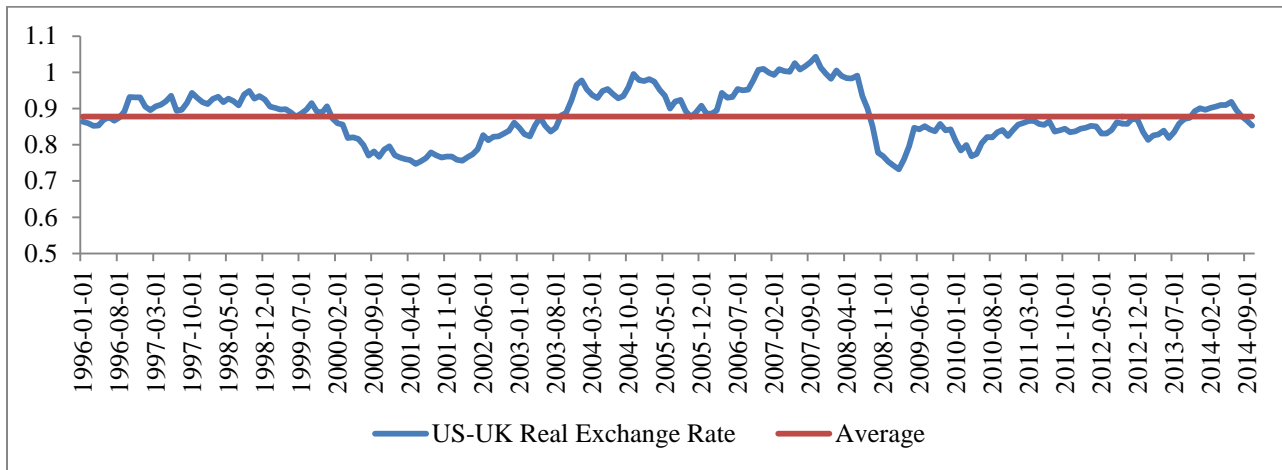


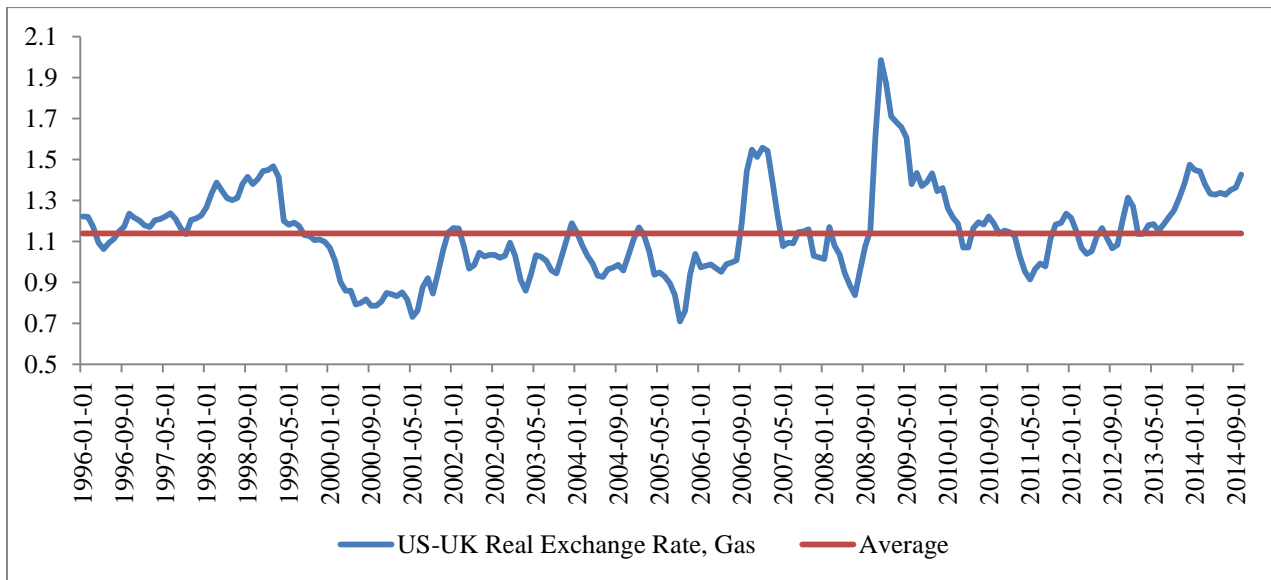
Chart 2 above suggests that the US-UK real exchange rate exhibits long run mean-reverting behavior. Using the augmented Dickey-Fuller test, however, the null hypothesis that

the US-UK exchange rate has a unit root cannot be rejected at a 95% confidence level. The test statistic is roughly -2.48, corresponding to a p-value of .12.

One reason one may not be able to reject the null of unit root is that many of the goods included in consumer price indices are not highly tradable. Consider another example of a real exchange rate involving a more tradable good: gasoline.

Chart 3 below shows the time series for the real exchange rate between the US and UK limited to gas along with its mean (1.14). Again, the timeline considered is January 1996 to October 2014. While the series exhibits significant volatility, it appears to revert to the mean more quickly than the previous total CPI series. A likely reason for this difference is the high degree of trade involved in gasoline markets.

Chart 3: The Real Exchange Rate in Gasoline, US-UK



As one would expect, an application of the augmented Dickey-Fuller test to the gasoline real exchange rate series shows that one is able to reject the null that the series has a unit root at a 95% confidence level. The test statistic is -3.42, corresponding to a p-value of roughly .01.

Bitcoin Real Exchange Rate

Even more so than gasoline, bitcoins are highly tradable. As such, one would expect that relative PPP should hold in bitcoin markets. Real bitcoin exchange rates between U.S. bitcoins, UK bitcoins, and Euro Zone bitcoins are calculated using the methodology specified in Table 2 of the appendix.

One of the issues considered first, however, is whether beginning of the week or end of the week effects would affect the analysis. Bitcoin data is available for every day including weekends, since the electronic exchanges run at all hours. Nominal exchange rate data, however, is only available for weekdays when active trading is occurring. For each real exchange rate, an autoregression was estimated by treating the series as continuous (despite weekend gaps) and including four lagged terms and dummies for Monday and Friday. In each case, the coefficients on Monday and Friday were miniscule: less than .001 in absolute value relative to a series centered around roughly one (see means in table 1 below). Furthermore, only one coefficient was statistically significant at a 95% confidence level. Given the weak effect both in terms of magnitude and statistical significance, beginning and end of week effects are ignored and the time series is treated as continuous across weekends.

Given that the underlying goods (bitcoins) in the calculation are identical, one would expect the real exchange rate to be centered on one. A real exchange rate of one would reflect

absolute purchasing power parity, signifying that a bitcoin bought in either U.S. dollars, British pounds, or euros would cost the same after currency conversions.

The first interesting result is that the real exchange rate differs persistently from one for each currency pair. One would expect that PPP would not hold consistently (given exogenous shocks), but that the mean should not deviate systematically. A mean less than one implies that the bitcoins bought in the currency stated in the denominator are relatively cheaper than bitcoins bought in the currency specified in the numerator of the real exchange rate. The mean exchange rates are shown in the following table and graphs.

Table 1: Means of Bitcoin Real Exchange Rates

Real Exchange Rate	Mean Time Period: 9/3/13-10/6/14	Implication
BTC_{GB}/BTC_E	.972	Euro Zone bitcoins are relatively cheap in terms of British bitcoins.
BTC_{GB}/BTC_{US}	.958	American bitcoins are relatively cheap in terms of British bitcoins.
BTC_E/BTC_{US}	.986	American bitcoins are relatively cheap in terms of Euro Zone bitcoins.

Chart 4: Time Series of BTC_{GB}/BTC_E

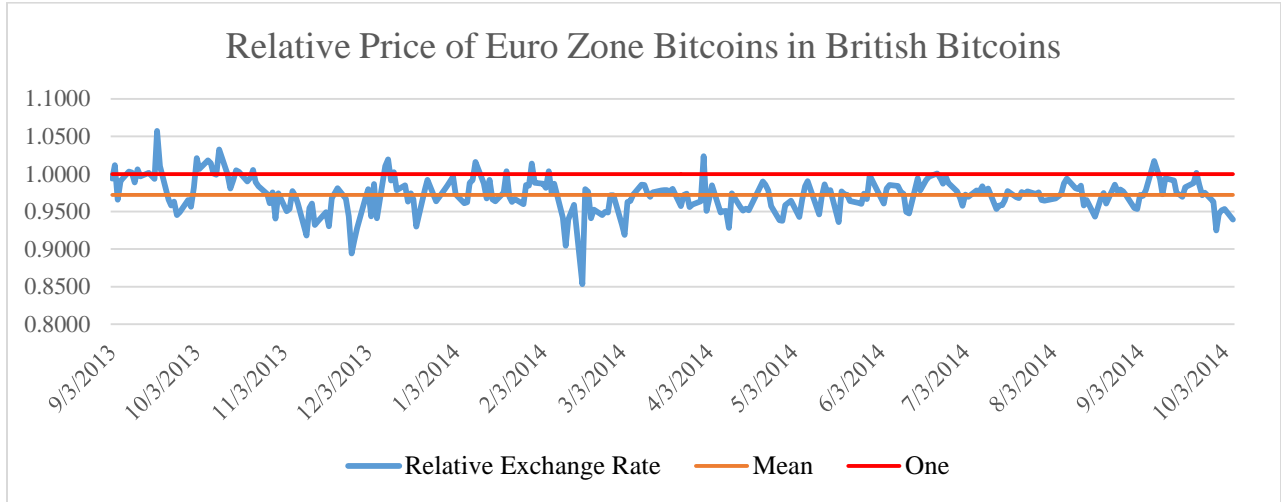


Chart 5: Time Series of BTC_{GB}/BTC_{US}

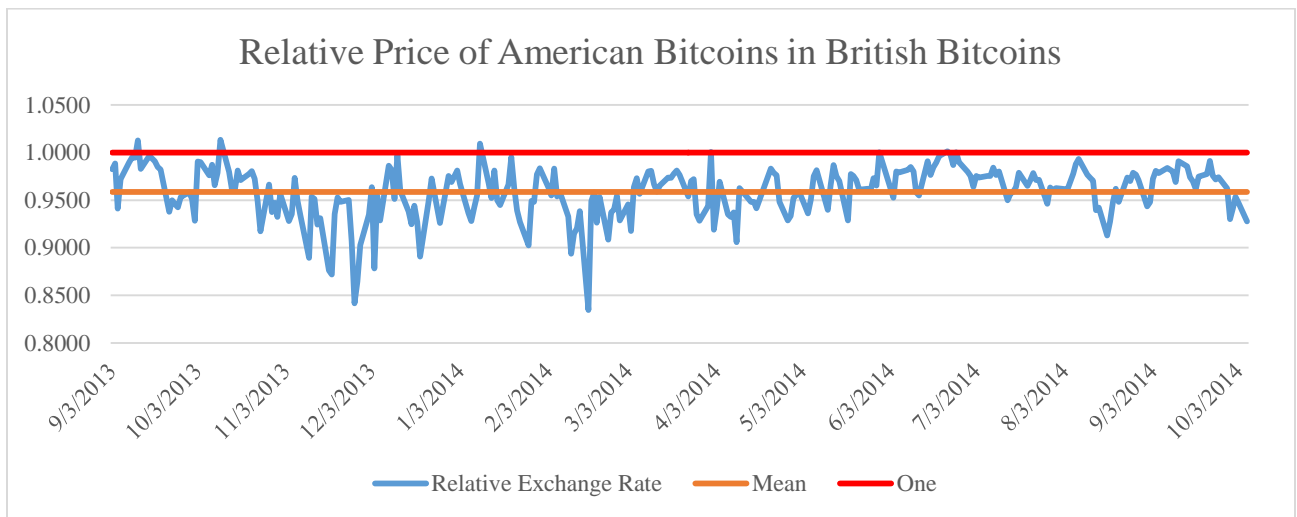
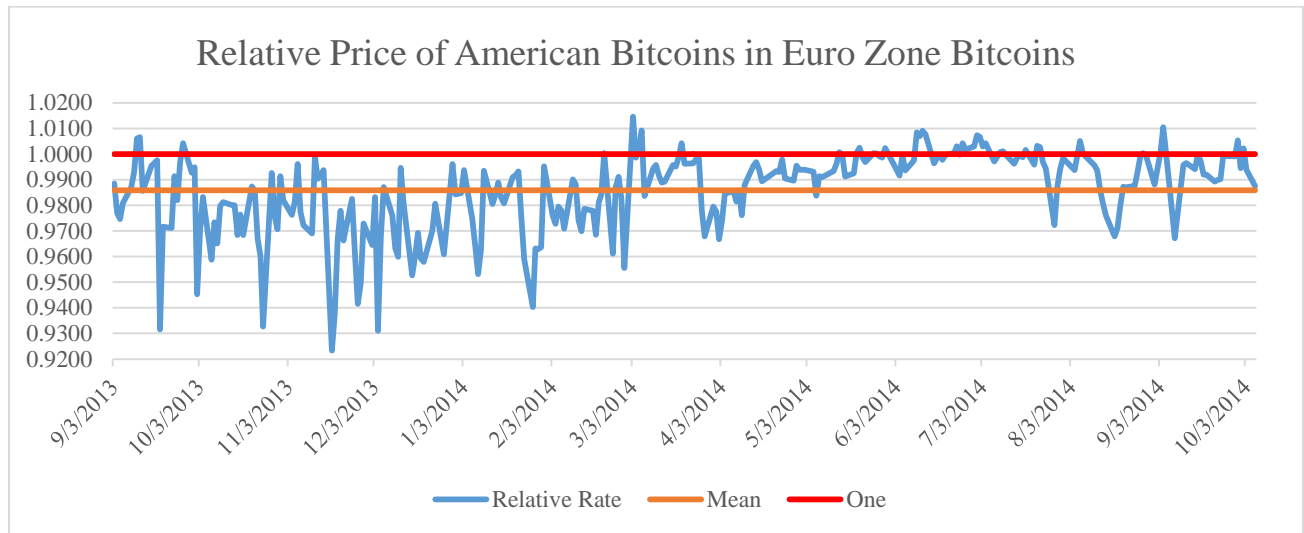


Chart 6: Time Series of BTC_E/BTC_{US}



As explained in Table 1, the mean real exchange rates imply that of the three currencies considered, bitcoins are cheapest when bought in U.S. dollars. They are most expensive when bought in British pounds. This result is unexpected given that there are no differences between bitcoins bought in any particular currency, so no differences in the underlying item being bought can account for the deviation from one. As mentioned earlier, a primary justification for PPP is international goods arbitrage. Bitcoins are actively traded on many exchanges across many currencies, and so in theory an arbitrage opportunity exists. On average across the exchanges, one should in theory be able to buy bitcoins in markets where they are cheap (in the bitcoin-U.S. dollar market in this case) and sell them where they are expensive (such as in the bitcoin-British pound market). Theoretically one could earn riskless profit utilizing such a strategy.

Given that these deviations exist, there must be limits to the arbitrage opportunity described above on a persistent basis. Many frictions could exist. For example, there could be a great deal of counterparty risk in selling on an exchange that specializes in a currency in which it

bitcoin is expensive. Analyzing these frictions and determining the source of the persistent deviations from absolute PPP warrants further investigation.

Although absolute PPP does not seem to hold, this does not preclude the possibility of consistency with relative PPP. As discussed, the method to assess relative PPP is with a statistical test for a unit root. For each bitcoin real exchange rate series, the null that the series has a unit root can be rejected at a 99% confidence level based on a four lagged term Augmented Dickey-Fuller test, shown as model (1) below. For the regression results, refer to Tables 3, 4 & 5 in the appendix. As expected, the bitcoin real exchange rate behaves as a stationary series. In other words, the series appears to be mean reverting as one might expect, even if the mean is not one.

$$\Delta q_t = c + \gamma q_{t-1} + \delta_1 \Delta q_{t-1} + \delta_2 \Delta q_{t-2} + \delta_3 \Delta q_{t-3} + \delta_4 \Delta q_{t-4} + \varepsilon_t \quad (1)$$

Also consistent with bitcoin's high degree of tradability, shocks dissipate very quickly. The following graphs show impulse response functions for each series. Each graph is based on a linear autoregression with four lagged terms.

Each graph shows the projected behavior of the corresponding series to a positive one standard deviation shock. In each case the magnitude of the shock is roughly .02, which is approximately a 2% deviation from mean. The vertical axis shows the value of the real exchange rate minus the mean. The horizontal axis is time in days. One can see that shocks to the series in each case are reduced by 50% within one day, and are nearly eliminated entirely by day 10.

Chart 8: Impulse Response Function for BTC_{GB}/BTC_E

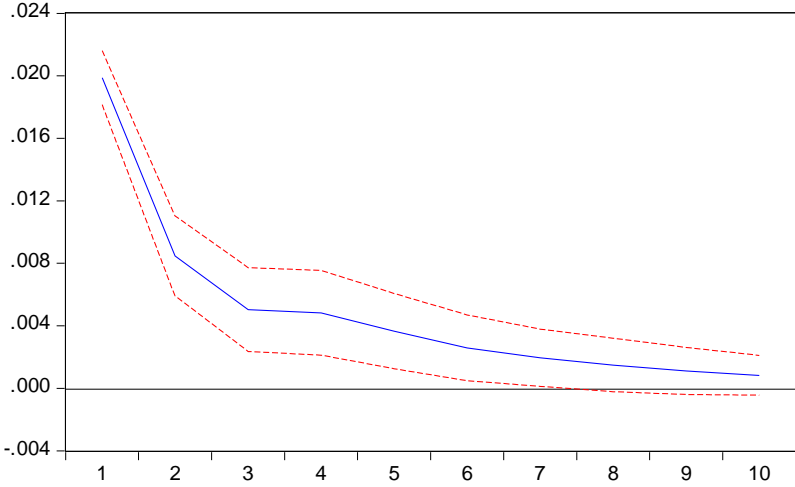


Chart 9: Impulse Response Function for BTC_E/BTC_{US}

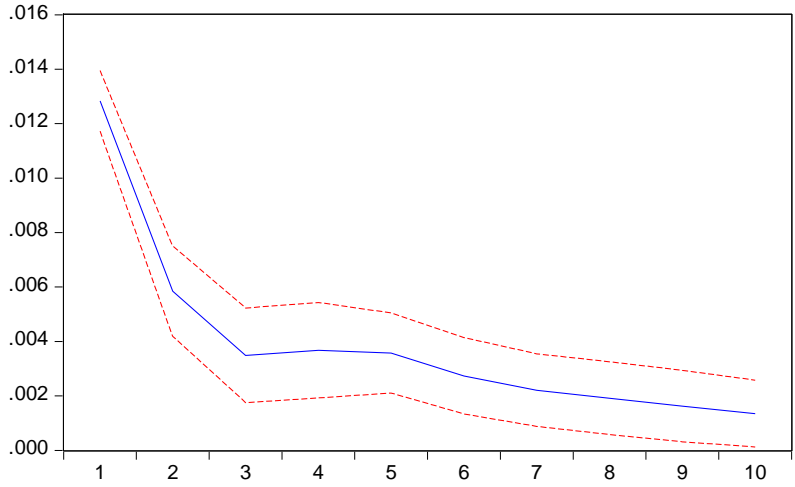
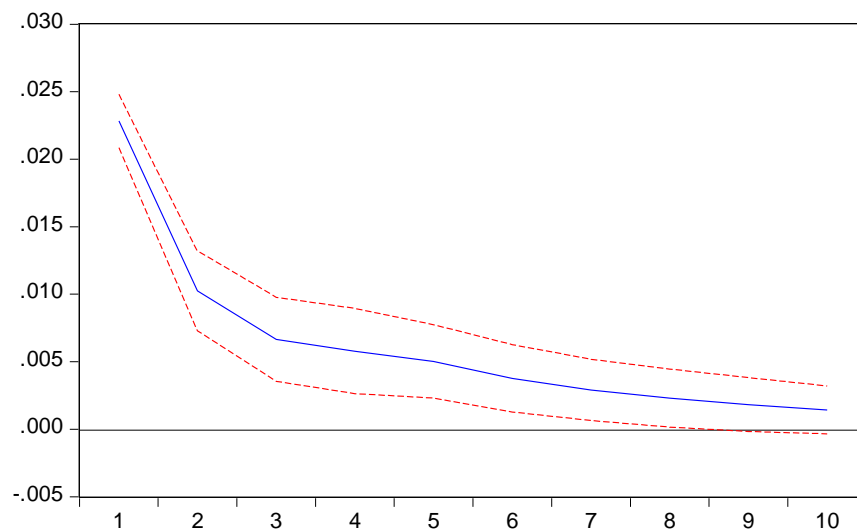


Chart 10: Impulse Response Function for BTC_{GB}/BTC_{US}



Attempts to Explain the Bitcoin Real Exchange Rate

In addition to testing whether PPP held generally, the other question is to discover what factors explain deviations in relative PPP—what pushes the series away from the mean. Among the explanatory factors examined are volatility in the respective markets and relative volume of the markets.

One would expect volatility in a currency market to reduce its value relative to other currencies. For example, if the price of bitcoin in terms of bitcoin in British pounds began fluctuating more widely, one would expect bitcoins bought in British pounds to lose value relative to the bitcoins bought in another market—bitcoins bought in US dollars, for example.

To test the effect of volatility, rolling 20-day standard deviations of the percentage changes in price of bitcoin of each of the two reference currency were added to an autoregression of each real exchange rate with four lagged terms. The results of the regressions suggested the opposite of the expected effect. For example, a 10 percentage point increase in the 20 day rolling standard deviation percentage price change in the bitcoin-dollar market corresponds to an

increase of .007 in the relative price of U.S. bitcoins in terms of Euro Zone bitcoins. While the coefficients are of the opposite sign than expected, only one of the coefficients (the 20 day rolling standard deviation of the price changes in the bitcoin-pound market) was found to be statistically significant at a 95% confidence level.

Relative transaction volume could also presumably affect the real exchange rate. The expected effect is that greater relative market liquidity, measured here as relative transaction volume, would increase the relative value of the currency. Relative transaction volume is defined as the total volume in one reference currency divided by the total transaction volume of the other.

Again, the regression results suggest the opposite in two of three cases. For example, the regression results showed that a ten basis point increase in the ratio of bitcoin volume in British pounds to bitcoin volume in euros, the price of Euro Zone bitcoins in terms of British bitcoins increases by roughly .001. The effect is miniscule and not statistically significant. The coefficient is the expected sign for the Euro-US real bitcoin exchange rate, but again the effect is not statistically significant.

V. Conclusion

Consistent with bitcoin's high degree of tradability, bitcoin prices exhibit relative PPP. The more surprising result is the deviation from absolute purchasing power parity: the mean of the real exchange rates persistently differ from one. This is a topic that warrants further investigation to determine what inter-country factors present obstacles to arbitrage that could explain such a phenomenon.

Furthermore, more work could be done to analyze what factors explain interday deviations from relative PPP. Rolling average volatility and relative volumes did little to explain this variation.

VI. Appendix, Chapter Two

Table 1: Real Exchange Rate Calculations

Real Exchange Rate	Description	Calculation
$Goods_{US}/Goods_{UK}$	The U.S.-UK real exchange rate. It is the price of British goods in terms of U.S. goods.	The real exchange rate is calculated as follows: $\frac{CPI_{UK}}{CPI_{US}} \times \frac{\$}{\pounds} = Goods_{US}/Goods_{UK}$
Gas_{US}/Gas_{UK}	The US-UK real exchange rate for gasoline. It is the price of British gas in terms of U.S. gas.	GPI signifies <i>gas price index</i> . The real exchange rate is calculated as follows: $\frac{GPI_{UK}}{GPI_{US}} \times \frac{\$}{\pounds} = Gas_{US}/Gas_{UK}$

Table 2: Bitcoin Real Exchange Rate Calculations

Real Exchange Rate	Description	Calculation
BTC_E/BTC_{US}	The relative price of an American bitcoin in terms of Euro Zone bitcoins. It should equal one given no limits to arbitrage. If it is less than one, the American bitcoin is relatively inexpensive.	For each day observation, the real exchange rate is calculated by dividing the price of a bitcoin in USD by the price of a bitcoin in Euros, and then dividing that by the actual nominal exchange rate. The formula appears as follows: $\frac{\frac{\$}{BTC_{US}}}{\frac{\text{€}}{BTC_E}} \times \frac{\text{€}}{\$} = BTC_E/BTC_{US}$
BTC_{GB}/BTC_{US}	The relative price of an American bitcoin in terms of British bitcoins. It should equal one given no limits to arbitrage. If it is less than one, the American bitcoin is relatively inexpensive.	For each day observation, the real exchange rate is calculated by dividing the price of a bitcoin in USD by the price of a bitcoin in Pounds, and then dividing that by the actual nominal exchange rate. The formula is as follows: $\frac{\frac{\$}{BTC_{US}}}{\frac{\text{£}}{BTC_E}} \times \frac{\text{£}}{\$} = BTC_{GB}/BTC_{US}$
BTC_{GB}/BTC_E	The relative price of a Euro Zone bitcoin in British bitcoins. It should equal one given no limits to arbitrage. If it is less than one, a Euro Zone bitcoin is relatively inexpensive.	For each day observation, the real exchange rate is calculated by dividing the price of a bitcoin in Euros by the price of a bitcoin in Pounds, and then dividing that by the actual nominal exchange rate. The formula is as follows: $\frac{\frac{\text{€}}{BTC_E}}{\frac{\text{£}}{BTC_{GB}}} \times \frac{\text{£}}{\text{€}} = BTC_{GB}/BTC_E$

Table 3: Model (1) for BTC_E/BTC_{US}

Variable	Coefficient	Std. Error	t-Statistic	Prob.
q(-1)	-0.253311	0.064738	-3.912852	0.0001
D(q(-1))	-0.297224	0.076817	-3.869243	0.0001
D(q(-2))	-0.243011	0.075485	-3.219335	0.0014
D(q(-3))	-0.113859	0.070455	-1.616043	0.1073
D(q(-4))	-0.082885	0.061920	-1.338586	0.1819
Constant	0.249846	0.063826	3.914492	0.0001
			t-Statistic	Prob.*
Augmented Dickey-Fuller test statistic			-3.912852	0.0022
Test critical values:				
1% level			-3.454812	
5% level			-2.872203	
10% level			-2.572525	

*MacKinnon (1996) one-sided p-values.

Table 4: Model (1) for BTC_{GB}/BTC_{US}

Variable	Coefficient	Std. Error	t-Statistic	Prob.
q(-1)	-0.295348	0.069205	-4.267722	0.0000
D(q(-1))	-0.264758	0.077594	-3.412081	0.0007
D(q(-2))	-0.178159	0.075234	-2.368049	0.0186
D(q(-3))	-0.113439	0.069821	-1.624718	0.1054
D(q(-4))	-0.129340	0.061112	-2.116436	0.0353
Constant	0.282873	0.066357	4.262868	0.0000
			t-Statistic	Prob.*
Augmented Dickey-Fuller test statistic			-4.267722	0.0006
Test critical values:				
1% level			-3.454812	
5% level			-2.872203	
10% level			-2.572525	

*MacKinnon (1996) one-sided p-values.

Table 5: Model (1) for BTC_{GB}/BTC_E

Variable	Coefficient	Std. Error	t-Statistic	Prob.
q(-1)	-0.387619	0.075090	-5.162082	0.0000
D(q(-1))	-0.189001	0.080940	-2.335091	0.0203
D(q(-2))	-0.110418	0.078050	-1.414714	0.1583
D(q(-3))	-0.012501	0.071543	-0.174738	0.8614
D(q(-4))	0.007804	0.061847	0.126174	0.8997
Constant	0.376578	0.073029	5.156575	0.0000

	t-Statistic	Prob.*
Augmented Dickey-Fuller test statistic	-5.162082	0.0000
Test critical values:		
1% level	-3.454812	
5% level	-2.872203	
10% level	-2.572525	

*MacKinnon (1996) one-sided p-values.

Chapter Three: An Economic Analysis of Bitcoin Mining

I. Introduction

As mentioned in chapter one, the process of verifying transactions and creating new bitcoins is called “mining.” In the market for bitcoin mining, agents face a decision of whether or not to enter based on costs and potential rewards. The following essay analyzes this decision on the part of potential miners and uses that analysis to draw conclusions about the desirability of the current mining scheme in terms of social welfare.

The rest of the chapter is organized as follows. The remainder of the introduction will provide a description of bitcoin mining and background information necessary for the subsequent sections. Section II introduces a static equilibrium model of the decision faced by potential bitcoin miners. Section III uses actual data to analyze bitcoin markets and test the model outlined in section II. Section IV concludes.

The Process of Bitcoin Mining

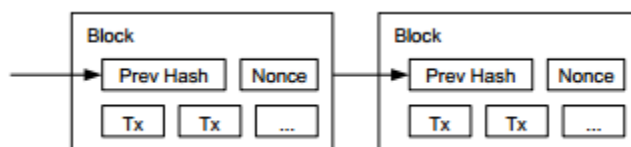
As stated in the introduction, bitcoin mining serves two purposes: avoiding double spending and creating new bitcoins in a controlled manner. Transactions are constantly being broadcasted over the bitcoin network. Miners gather the data for these transactions and attempt to append a group of them, known as a “block,” to the block chain.

A key concept in bitcoin mining is a hash function. A hash function is a function that maps a string of text and/or numbers of arbitrary length into a string of characters of a fixed length called a *hash*. The key to a hash function is that it is nearly impossible to infer the contents of the original string given a hash, but it is relatively easy to determine whether or not a hash is correct given the original string by applying the hash function (Mironov 2005).

Each block of bitcoin transactions must be accompanied by a hash. The hash is the result of all the information of the block transformed by a hash function, which in the case of bitcoin is the function SHA256. The network, however, specifies that the hash must be less than a particular number. The job of the miner is to find a value called a *nonce* that will be included with the block that yields a hash that fits the requirements of the network. Finding this nonce is very difficult, and the only effective way to do so is through brute force computing, i.e. trial and error using a sequence of different values (Velde 2013).

Once a miner successfully finds a hash for the block, the miner broadcasts the block chain that includes his or her addition to the transaction history. Other miners verify the authenticity of the transactions and check the hash, and once a majority of miners have done so, it is officially added to the block chain, and the race to mine the next block begins.

The process of verifying and adding blocks to the public ledger helps maintain the integrity of the system. The convention is for miners to accept the longest chain of transactions as valid. Therefore, since each block has an associated hash that depends on the hash of the previous block (see graphic below), in order to falsify a transaction history, a miner would need to create the proof-of-works needed for the target block and all subsequent blocks and then overtake the current longest record of transactions. Such an action would likely require the control of over 50% of all of the hashrate of the network (an enormous amount of computing power), and would likely destroy the value of any bitcoin gained in the scheme due a loss of confidence in the system (Nakamoto 2008). As such, it is unlikely there is any economic incentive to alter the transaction history.



The miner who successfully mined the newly added block is rewarded in the form of brand new bitcoins (currently 25 bitcoins per block) as well as any transaction fees. Transaction fees are amounts of bitcoin that users can voluntarily add to their transactions in order to incentivize the miner to include the given transaction in the miner's block, thus potentially reducing the amount of time required to verify and finalize the transaction. The number of new bitcoins awarded began at 50 and halves every 210,000 blocks (roughly every four years). This ensures a steady decay in the rate of creation, so that the number of bitcoins created will approach but not exceed 21 million (Velde 2013).

Since miners are rewarded in bitcoin, the incentive to mine is tied to its value. As the price of bitcoin increases, there is a greater opportunity to capture revenue through the mining of bitcoin. Price is the main driver of this incentive, since the number of bitcoin rewarded is constant with the exception of predictable changes every four years. One variable component in the amount of bitcoin rewarded is transaction fees, but they currently make up a relatively small portion of miners' revenue, as will be discussed in Section 3.

One item to note is the winner-take-all nature of mining rewards. Although many miners are all at once attempting to add a block to the blockchain, there is only one miner who ultimately takes credit for the block, which entails a great deal of uncertainty for the individual miner. This has led to the creation of "mining pools"—collections of individual miners that agree to distribute any rewards received among all the miners in the pool in proportion to some measure of each miner's contributed computing power.²

In order to maintain the expected rate of creation of new bitcoin, the difficulty of mining adjusts every 2016 blocks (roughly every two weeks) to ensure that roughly six blocks are mined every hour. The difficulty adjusts due to changes in the level of computing power available on

² See <https://www.btcguild.com/> and <https://ghash.io/> for two examples of bitcoin mining pools.

the network. As more miners enter or as existing miners devote more computing power to mining, blocks are mined at a faster rate. As computing power increases, the hash requirement gets more stringent—the target value under which the hash must lie is pushed lower in order to reduce the chance that any given nonce satisfies the requirement (Velde 2013). See Charts 1 & 2 in the appendix for a representation of how the proportional difficulty of mining has changed over time.

Participating in bitcoin mining entails significant costs. One of the primary expenditures in a Bitcoin mining operation is the amount of electricity consumed by machines searching for the correct nonce to mine a block, a process consisting of trying as many values as is possible. Mining also entails significant fixed costs, since investments must be made to purchase and maintain the equipment used (Tiller 2014).

II. A Simple Model of the Decision to Mine Bitcoin

A potential bitcoin miner would seek to maximize potential economic profit. Each miner chooses his or her individual hashrate h to satisfy this objective. The rewards of mining are a function of a known quantity q awarded for a successfully mined block multiplied by a price p that is determined by market forces by bitcoin users and over which miners have no control. As such, p is considered a stochastic variable. Given the winner-take-all nature of mining rewards, a miner receives those awards with a probability σ , which depends on the hashrate of the individual miner h versus the hashrate of the overall network H . As such, mining revenue r can be specified as:

$$r = \sigma pq,$$

$$\text{where } \sigma = \frac{h}{H}$$

There are certain costs associated with supplying a hashrate h . The associated variable costs can be specified as a function of h multiplied by some cost factor c . There will also be some fixed costs fc associated with entering the mining market. Total costs tc of entrance can thus be modeled as:

$$tc = ch + fc$$

Expected profit Π is then a function of revenue less total costs. Expected profit is:

$$\Pi = r - tc$$

$$\Pi = \frac{h}{H} pq - ch - fc$$

Each miner is unaware of the behavior of other miners. Therefore, the overall hashrate supplied can be thought of as a function of the individual contribution h and the amount of effort k contributed by each other miner. The market will consist of n active miners, and so H can be specified as follows.

$$H = h + (n-1)k$$

$$\Pi = \frac{hpq}{h + (n-1)k} - ch - fc$$

Three conditions will hold in equilibrium. First, each individual miner will be profit maximizing, and therefore the first order condition *foC* of profit with respect to h will equal 0. This sets marginal revenue, which in this case is the increase in probability of being the successful miner from increasing h , equal to marginal cost, which in this case is the cost of supplying the increase in h . Second, in equilibrium it will be assumed that each miner has access

to the same technology, and so each miner will make produce put forth the same amount of effort, so h will equal k . Free entry will also be assumed, so potential miners will produce as long as there is a profit incentive to do so. Therefore, in equilibrium, economic profit will be zero.

$$foc = \frac{pq}{h+(n-1)k} - \frac{hpq}{(h+(n-1)k)^2} - c = 0$$

$$k = h$$

$$\Pi = 0$$

Solving these for n and h simultaneously given the above conditions results in the following equilibrium values for n and h , denoted as h^* and n^* , respectively.

$$h^* = \frac{\sqrt{fcpq} - fc}{c}$$

$$n^* = \frac{\sqrt{fcpq}}{fc}$$

Implications of the model

That economic profit in the bitcoin market is zero in equilibrium is a well-justified assumption given low barriers to entry in the market for mining. The reasoning for believing that bitcoin mining has low barriers to entry is the non-proprietary nature of mining software³ and mining equipment and the inability of miners to influence the price of bitcoin, which is determined independently by users in exchange markets. If this assumption holds true, the model has significant welfare implications.

Since miners do not affect the prices faced by users of bitcoin, it is reasonable to assume that an economic planner focused on maximizing social welfare would seek to carry out bitcoin

³ For example, see <http://www.bitcoinx.com/bitcoin-mining-software/> for a variety of open source mining applications.

mining in such a way that minimizes the economic resources expended in the process. For illustrative purposes, let us assume that p and q equal 10 (normalizing total industry revenue to 100), fixed costs fc equal 5 per firm, and the cost factor c equals 1. In this case, in equilibrium, h^* would equal ~ 17.36 and n^* would equal ~ 4.72 , implying total industry costs of 100, satisfying the zero profit condition. The overall network hashrate would ~ 77.64 .

Let us assume the planner with the power to dictate who mined and with what hashrate wanted to hit a target hashrate of 77.64 (that achieved in a competitive market), but at the lowest cost. Given the specification of constant marginal cost, a planner would view bitcoin mining as a natural monopoly. In that case, in order to save on fixed costs, the planner would specify that one miner would provide the entire hashrate for a cost of ~ 82.64 , which is a 17.36% cost reduction in this simple example.

However, recall that the difficulty of bitcoin mining adjusts based on the amount of computing power present on the network. Also note that changes in difficulty do not affect the usefulness of the currency for users, but rather serve merely to keep the creation of bitcoin on a predictable path. Therefore, the planner would not simply seek to provide the same hashrate at a lower cost, but the planner would seek to keep the difficulty of mining as low as possible, reducing the hashrate and the amount of resources expended by a great deal more.

Wasteful Competition—A Thought Experiment

The model suggests that bitcoin that if there is positive economic profit to be made, then over the long run miners will enter and put pressure on the level of difficulty until there is no profit incentive for further entrance. This competition allows the verification process of bitcoin to be decentralized and allows bitcoin to follow a predictable path in terms of increasing the

monetary base. However, since mining costs are driven by competitive forces, there is a great deal of waste in terms of electricity expended and computing equipment utilized.

For the sake of illustration, let us assume that bitcoin will soon become the de facto currency of trade for many countries, having an overall real value to rival that of the U.S. dollar. That would imply that (in real terms) the amount of bitcoin outstanding would hold a value something similar to that of the total currency in circulation for the U.S. dollar. In this case, using the U.S. dollar money supply would be inappropriate, since there is no banking system utilizing bitcoin that could serve as a comparison. According to the Federal Reserve, there was roughly 1.36 trillion U.S. dollars in circulation on March 11, 2015, and according to Blockchain (a bitcoin wallet website), the number of bitcoin outstanding at the time was roughly 13.9 million.

US Dollars Outstanding [A]	1.36 trillion
Bitcoin Outstanding [B]	13.9 million
Value per Bitcoin [A/B]	~ \$ 98,000

And so if bitcoin were to replace the dollar, this would imply a real value of roughly \$98,000 per bitcoin in today's dollars. Due to the competitive forces outlined above, the cost of mining each bitcoin would approach \$98,000 in terms of electricity, computing equipment and opportunity costs. At the current rate of 25 bitcoins mined per block, and assuming that a block is mined roughly every 10 minutes, this would imply an overall equilibrium economics cost of mining of roughly \$351 million per day or \$128 billion per year.

Some of these economic costs would be the opportunity costs associated with using the computer equipment for its next most profitable activity. Let us assume that approximately 40% of this cost consists of electricity consumed by mining machines. At a U.S. commercial energy price of \$.1013/kilowatt-hour,⁴ this would imply that bitcoin miners would consume about 500,000 gigawatt hours of electricity per year. To put this number in context, this would comprise roughly 1/8 of the net electricity production in the United States in 2014—an enormous amount of energy.⁵

III. An Empirical Analysis

A great deal of information about mining is publicly available. As stated in the introduction, one of the primary variable expenditures involved in mining is the considerable electricity cost associated with running mining machines. The following is an empirical analysis of how to estimate these costs.

Calculating Mining Energy Costs

In order to estimate the costs related to bitcoins, I estimate the power consumption required to generate the level of computing power seen on the bitcoin network. For definitions of relevant terms, see the appendix.

The estimate of the daily network hashrate (a measure of computing power in terms of calculations performed) was obtained from blockchain.info.⁶ The data was truncated to include only data post-2011 when activity on the network became sizable.

⁴ The quoted U.S. energy price was taken from the U.S. Energy Information Administration website. It is an estimate of the nationwide average cost of industrial electricity in March 2015.

⁵ The estimate for mining electricity consumption was 506,000 gigawatt hours per year. The U.S. net production total for 2014 was 4,092,935 gigawatt hours for 2014, according to the U.S. Energy Information Administration. <http://www.eia.gov/totalenergy/data/monthly/pdf/sec7_5.pdf>.

⁶ Blockchain.com was founded in 2011 and is the world's leading Bitcoin wallet provider. Its site blockchain.info maintains up to date data on activity on the bitcoin network.

To get power consumption of bitcoin miners relative to hashrate, I used averages across several machines produced as of January 2015 specifically for the purpose of bitcoin mining.⁷ The end result is an estimate of the amount of power needed (in watts) to maintain a hashrate of 1 GH/s: about .89 watts. This can be applied to the overall network by multiplying the average network hashrate each day by the estimate of power necessary to generate that computing power. The watt estimate was then converted to kilowatts and multiplied the power by the amount of time in a day to end up with an estimate of daily energy consumption in terms of kilowatt hours.

Estimates of energy prices were drawn from the EIA (the U.S. energy information association) at a monthly frequency. The “commercial” data series is used as a proxy cost. The estimate of power consumption for each day was multiplied by the estimated power cost for the corresponding month. To get an estimate of the cost per bitcoin, the calculated daily cost is divided by the number of bitcoins awarded during that day.

Mining Revenue Estimates

Mining revenue comes in the form of newly created bitcoins as well as any bitcoin obtained from transactions fees. Values of both (these are accessible data from the public ledger and are therefore not estimates) were obtained from blockchain.info as well. It should be noted that transactions fees have on average made up less than .5% of miners’ revenue since 2011, and so the vast majority of mining revenue comes from new bitcoins.

⁷ Bitcoin Wiki maintains a list of specialized mining machines and many of their relevant characteristics. Information was drawn from the page on January 29, 2015. <https://en.bitcoin.it/wiki/Mining_hardware_comparison>.

Profit Over Time

Throughout much of its history, the electricity cost of mining a bitcoin is estimated to be negligible relative to the price at which a bitcoin can be sold over an exchange (see Charts 3, 4, & 5 in the appendix). Note that for Charts 3 & 4, price is specified on a separate axis than the electricity cost. One should also note that the calculation of the amount of electricity required to maintain a particular hashrate was obtained using modern machines, and so it is likely that less efficient machines were used early on in bitcoin's history, and so the electricity cost estimates are likely low for earlier points in time.

Recently (since around the beginning of 2014), however, the electricity cost of mining per bitcoin can be seen converging with market price. Given the model specified before, this is expected, since miners should enter the market until there is no profit incentive to do so. This convergence also appears with the relatively recent stagnation of the overall network hashrate (see chart 6), which suggests a significant slowdown in the rate of entry in the mining market.

Regression Analysis

It is reasonable to believe, as was specified in the model shown previously, that the market price of bitcoin is unaffected by mining activity. However, mining activity is dependent on the market price of bitcoin, since increases in market price increase the incentive to mine.

A natural way to measure the effect of price on mining costs is with an elasticity—to what degree does a percentage change in price yield a percentage change in costs. The data available are market price for each day and the total variable costs TVC normalized by the quantity of bitcoin awarded q . For shorthand, let this be represented as $VCpB$, the total variable

cost per bitcoin. The elasticity could be measured using a simple regression model such as model (1).

$$VCpB_t = \frac{TVC_t}{q_t}$$

$$\ln(VCpB_t) = \beta_0 + \beta_1 \ln(p_t) + \varepsilon_t \quad (1)$$

β_1 , as specified, would measure the elasticity between price and variable costs. As a measure of elasticity, the coefficient β_1 can be written as the first derivative of total variable costs per bitcoin with respect to price multiplied by the levels of each variable.

$$\beta_1 = \frac{d(VCpB)}{dp} \left(\frac{p}{VCpB} \right)$$

Regression model (1) is a somewhat naïve specification, however. It is reasonable to believe that there are factors, such as growing public knowledge for instance, not contained in the model that are contributing to a time trend in each variable. The “true” model could be thought of as including a linear time trend t to parse out the unexplained factors affecting both variables. The “true” regression model would thus be specified as follows.

$$\ln(VCpB_t) = \beta_0 + \beta_1 \ln(p_t) + \beta_2 t_t + \varepsilon_t \quad (2)$$

The estimated coefficient $\hat{\beta}_1$ from model (1) can be found in the appendix in regression table 1. Model (1) yields an estimate of 1.96. One can consider the bias of this estimate using the following specification for the expected value of $\hat{\beta}_1$.

$$E[\hat{\beta}_1] = \beta_1 + \beta_2 \hat{\delta}_1$$

$$t_t = \delta_0 + \delta_1 \ln(p_t) + \nu_t \quad (3)$$

The above two equations show that the bias on $\hat{\beta}_1$ is related to the coefficient β_2 from model (2) as well as the OLS estimate of the linear effect of log price on a linear time trend. First, one would expect a positive coefficient on β_2 since there is a general upward trend in total variable costs over time. Second, one would expect $\hat{\delta}_1$ from model (3) to be positive as well given the upward movement in price over time (with the exception of the post 2014 downward trend) due to unobserved demand factors. These two facts combined would suggest that the OLS estimate $\hat{\beta}_1$ would be too high of an estimate.

Indeed, the estimate $\hat{\beta}_1$ from model (2), found in regression table 2 in the appendix, is significantly lower at 0.652. The exact magnitude of the difference (1.308) is equivalent to the product of the OLS estimate $\hat{\beta}_2$ (0.007) and the OLS estimate $\hat{\delta}_1$ (182.9).

Another method to estimate β_1 that takes into account the bias generated by unobserved time effects is to difference the data. Consider for a moment two separate data points in time: t and $t-1$. The relationship specified by model (1), which is theorized to persist over time, can be written thusly for these two time periods.

$$\ln(\text{VCp}B_t) = \beta_0 + \beta_1 \ln(p_t) + \varepsilon_t$$

$$\ln(\text{VCp}B_{t-1}) = \beta_0 + \beta_1 \ln(p_{t-1}) + \varepsilon_{t-1}$$

Let us call the unobserved factors that bias the coefficient $\hat{\beta}_1$ x . These unobserved factors are part of a “true” model and affect both variable costs as well as price. Taking x into account leads to the following specifications for the two points in time written above.

$$\ln(\text{VCp}B_t) = \beta_0 + \beta_1 \ln(p_t) + \gamma x_t + \varepsilon_t$$

$$\ln(\text{VCp}B_{t-1}) = \beta_0 + \beta_1 \ln(p_{t-1}) + \gamma x_{t-1} + \varepsilon_{t-1}$$

Taking the difference (the model at t minus the model at $t-1$) yields the following. β_1 can now be estimated by regressing the difference in the log level of $VCpB$ on the difference in log level of p . The key assumption is that while x biases the coefficient β_1 , the differenced x does not. This is certainly the case if one can capture the true model using a linear time trend t , since the difference would simply become a time invariant constant.

$$\ln(VCpB_t) - \ln(VCpB_{t-1}) = \nu_0 + \beta_1[\ln(p_t) - \ln(p_{t-1})] + \gamma(x_t - x_{t-1}) + \mu_t$$

Model (4) is shown below. The differenced unobserved effects are assumed to be absorbed into a new constant τ_0 . To take into account the relatively low short term variability in $VCpB$, a long differencing time interval of 180 days (roughly 6 months) is used to remeasure β_1 .

$$\ln(VCpB_t) - \ln(VCpB_{t-180}) = \tau_0 + \beta_1[\ln(p_t) - \ln(p_{t-180})] + \xi_t \quad (4)$$

The OLS estimates from model (4) are reported in regression table 4 in the appendix. The regression coefficient is for β_1 is 0.714, which is contained in a 95% confidence interval of the previous estimate of 0.652 from model (2).

Application of the Empirical Analysis to the Model

What was obtained through the empirical analysis was an estimate of the elasticity of changes in total variable costs (normalized by the quantity of bitcoin created) to changes in the price of bitcoin. In the economic model described previously, variable costs are the additional costs associated with supplying a greater individual hashrate h . These variable costs per miner vc have been specified as:

$$vc = ch$$

To calculate total variable costs for the entire market, one simply multiplies individual variable costs by the number of miners n . The result is total variable costs TVC . The regression analysis utilizes total variable costs normalized by the amount of bitcoin created q .

$$TVC = nch$$

$$VCpB = TVC / q = nch / q$$

Assuming that the equilibrium conditions hold means that one can substitute the equilibrium values of n and h , denoted previously as n^* and h^* , respectively. Replacing the values of n and h yields the following value for $VCpB$, which will be denoted $VCpB^*$, since it is an equilibrium value. Note importantly that the variable cost factor c drops out of the equation.

$$VCpB^* = \frac{\sqrt{fc pq} \left(\frac{\sqrt{fc pq}}{fc} - 1 \right)}{q}$$

Recall that since the estimated coefficient is an elasticity, it is equivalent to the derivative of $VCpB$ with respect to p multiplied the ratio of the levels of the two variables. Therefore, β_1 can be written as the following.

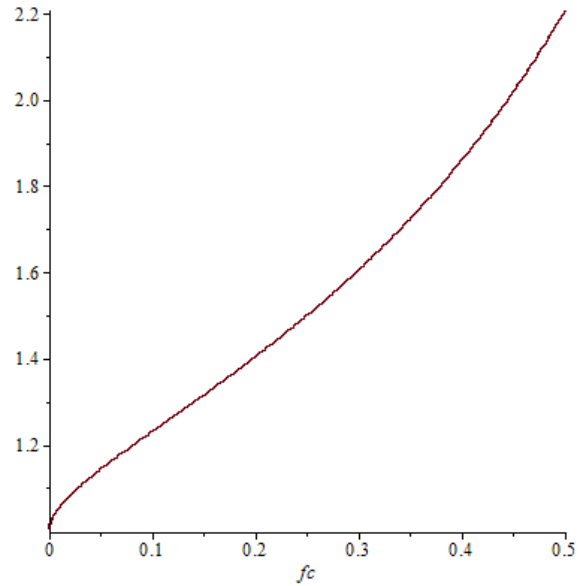
$$\beta_1 = \frac{d(VCpB)}{dp} \left(\frac{p}{VCpB} \right)$$

These terms in can be calculated as a function of fc , p and q using the equilibrium value of $VCpB$ shown above as $VCpB^*$. Calculating the product of the derivative of $VCpB^*$ with respect to p and the ratio of the levels of p and $VCpB^*$ yields the following value for β_1 .

$$\beta_1 = \frac{\sqrt{fc pq} - 0.5 fc}{\sqrt{fc pq} - fc}$$

Note that the value of β_1 should be greater than one for all positive fc , p and q . The economic intuition for this fact is relatively simple. Consider that the market for mining is at a steady state at time t (although the model is static, time is used for illustrative purposes) for a given p and q . At time $t+1$, the price has increased by 5%. Since q is fixed, total revenues summed across all miners will increase by 5%. Furthermore, since free entry is allowed, there will be no economic profit in equilibrium. Therefore, total costs summed across all miners, including incumbents and new entrants, should increase by 5%. Fixed costs for the market will increase by the number of new entrants multiplied by the required fixed costs. Variable costs will increase by the amount of variable costs contributed by new entrants as well as the increase contributed by incumbents increasing their own hashrate h . The result is that total variable costs for the market will move proportionally greater than fixed costs for the market, resulting in a coefficient of β_1 that is greater than one. The similar coefficient on a measure of total costs would be exactly one, whereas a coefficient on fixed costs would be less than one.

The following graph shows values of β_1 on the vertical axis plotted against values for fixed costs fc ranging from 0 to 0.5 on the horizontal axis. Both p and q are set equal to one. Fixed costs can then be interpreted as a percentage of total revenue (since total revenue is the product of p and q and is therefore normalized to one). For example, a value of 0.3 for fixed costs would imply that fixed costs are 30% of potential industry revenue.



Given that the estimate for β_1 from regression model (2) 0.652, one should conclude that either the model in some way does not reflect the realities of the bitcoin mining market, or the regression model and the data utilized do not give an accurate estimate of β_1 . Either option is possible. The model is limited in its applicability in several respects. First, the model is static. Therefore, applying it to data must come with the assumption that each separate data point represents an equilibrium. In reality, reaching equilibrium is a dynamic process. This is evidence in chart 5 in the appendix. Despite the secular decline in bitcoin price over the time period from roughly the beginning of 2014 to the end of the data in January 2015, the total electricity costs expended to mine appeared to continue to increase. One explanation for this phenomenon is that entry occurs over long periods.

Furthermore, the data utilized in the regression is flawed as well. For instance, it does not capture changes in mining technology. To make estimation feasible, a point estimate had to suffice for calculating energy costs, although developments in the equipment available for mining has certainly changed over time, which would be a driver of variable costs.

IV. Conclusion

Discussion and Implications

Despite the limitations of the model discussed, it is reasonable nonetheless to draw some basic conclusions from it. Most notable is the equilibrium condition that economic costs will converge with revenue, which is reasonable if one believes in the competitive nature of the market. While the empirical analysis did not yield a regression coefficient consistent with a static model, it was clear that variable costs do indeed react positively to changes in price.

This result has significant welfare implications, since the scalability of bitcoin depends on the extent to which society finds the outlay of resources for mining to be acceptable. The thought experiment proposed suggested that if bitcoin were to be valued overall near the U.S. dollar, the electricity consumed in mining could be an eighth of all the electricity produced in the United States. From a sustainability standpoint, this makes bitcoin, and any cryptocurrency of similar design, an unattractive alternative to major fiat currencies in the long run. The only way to solve this problem in bitcoin have a central authority that can control the mining process, which is at odds with bitcoin's very nature of being an open-source, decentralized payment system.

Opportunities for Further Research

Many topics of possible research are not included in the analysis. One topic for further investigation is the possibility of a "51% attack" in which one large miner (or a group of colluding miners) seeks to falsify the block chain by commanding over half of the computing power of the bitcoin network. This problem is briefly mentioned by Kroll, Davey and Felten

(2013), but remains a topic for additional inquiry as to whether there is an economic incentive to take control of the network.

Another topic not explored by this paper is the effect of risk aversion the decisions of potential miners. Given the probabilistic nature of mining rewards but the certainty of costs, bitcoin mining lends itself to a utility framework that discounts uncertain outcomes. Risk aversion indeed explains the creation of bitcoin mining pools that split mining rewards, and so this trend in industry structure could potentially be supported by such a model.

More realism could be applied to the model that incorporates some dynamic features. These would include a discounting of future expected profit as well as a time dependent move toward equilibrium. Such adjustments would help better ground the model discussed herein.

VII. Appendix, Chapter Three

Definitions of key terms:

Power	The first derivative of energy with respect to time, measured in energy/time (a watt, for example, is joules per second)
Kilowatt-Hour	A common unit of energy. It is the amount of energy used in a 1 kilowatt machine running for one hour.
Hashrate	The number of calculations (hashes) a machine can perform per unit of time. The relevant measure is Gigahashes per second (GH/s).

Chart 1

Source: blockchain.info

Note: Difficulty is expressed as relative to a difficulty of 1 at inception in 2009.

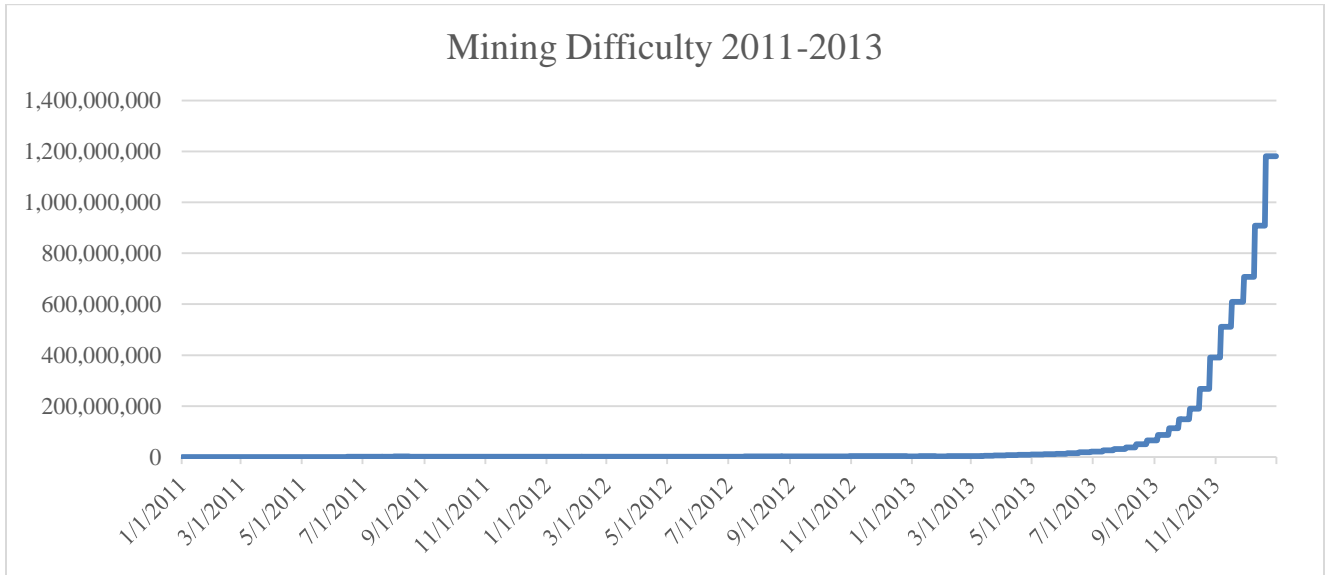


Chart 2

Source: blockchain.info

Note: Difficulty is expressed as relative to a difficulty of 1 at inception in 2009.

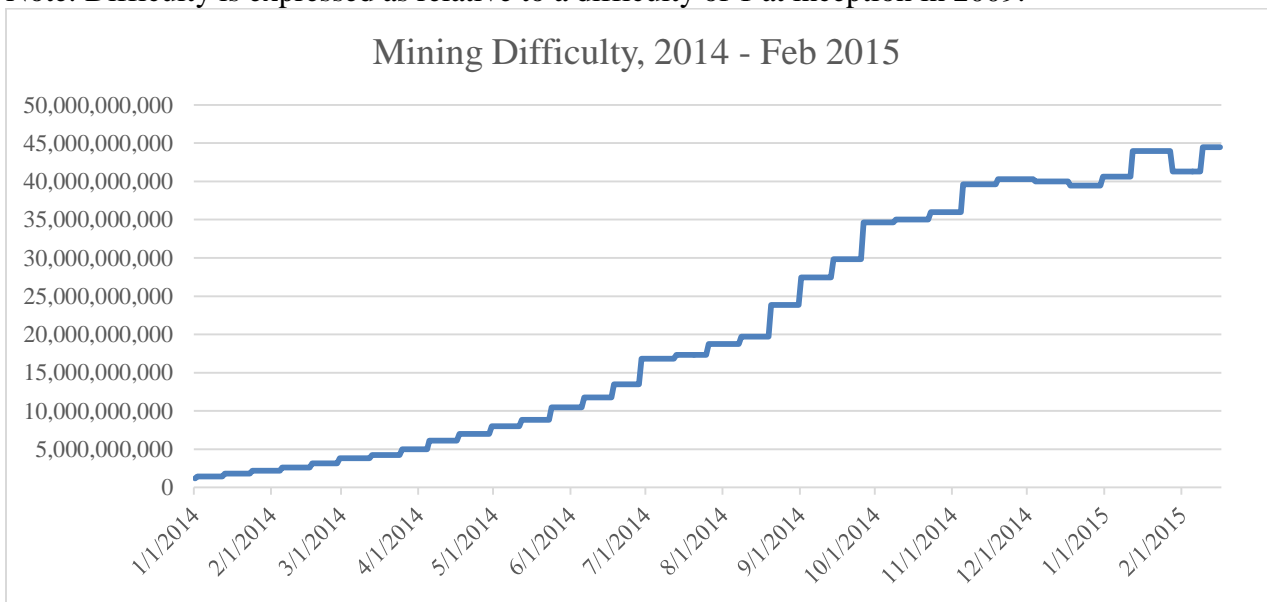


Chart 3

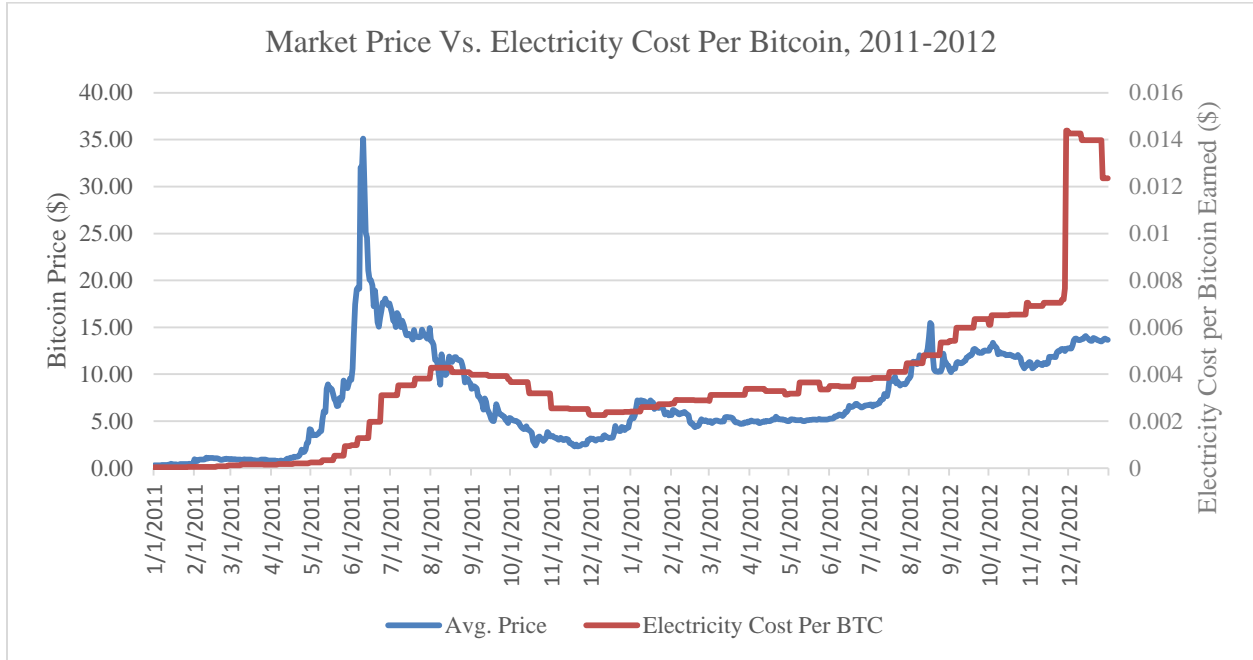


Chart 4

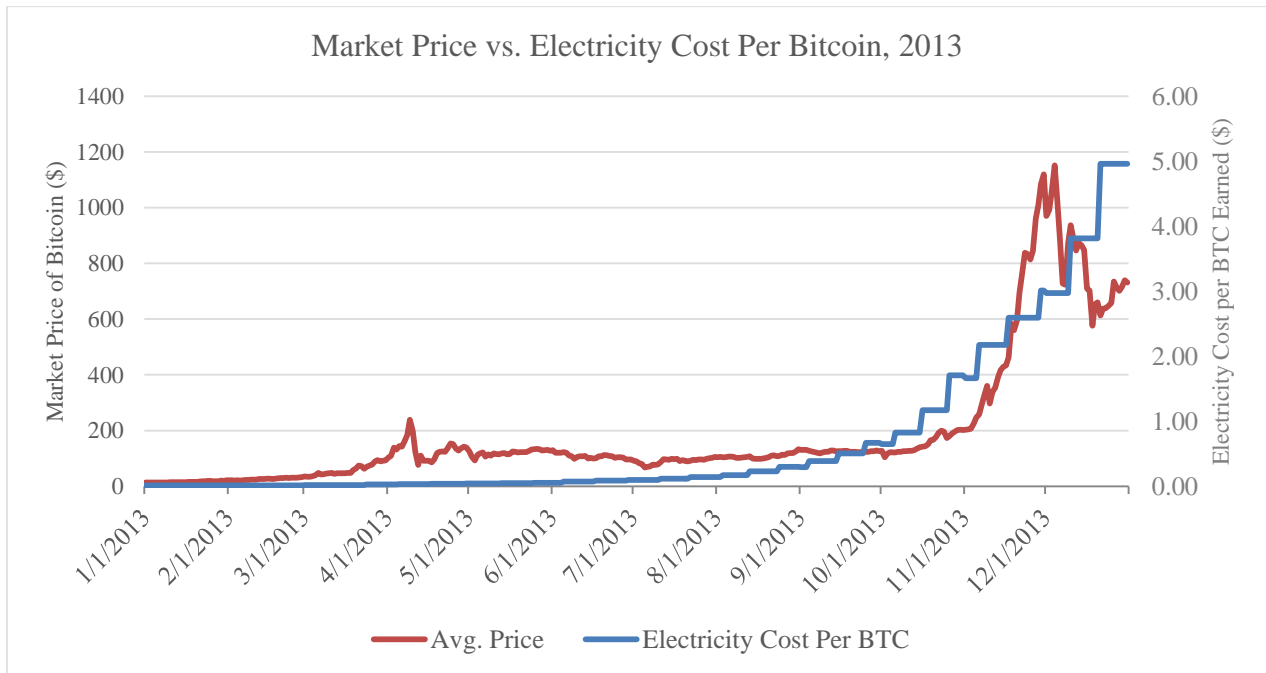


Chart 5

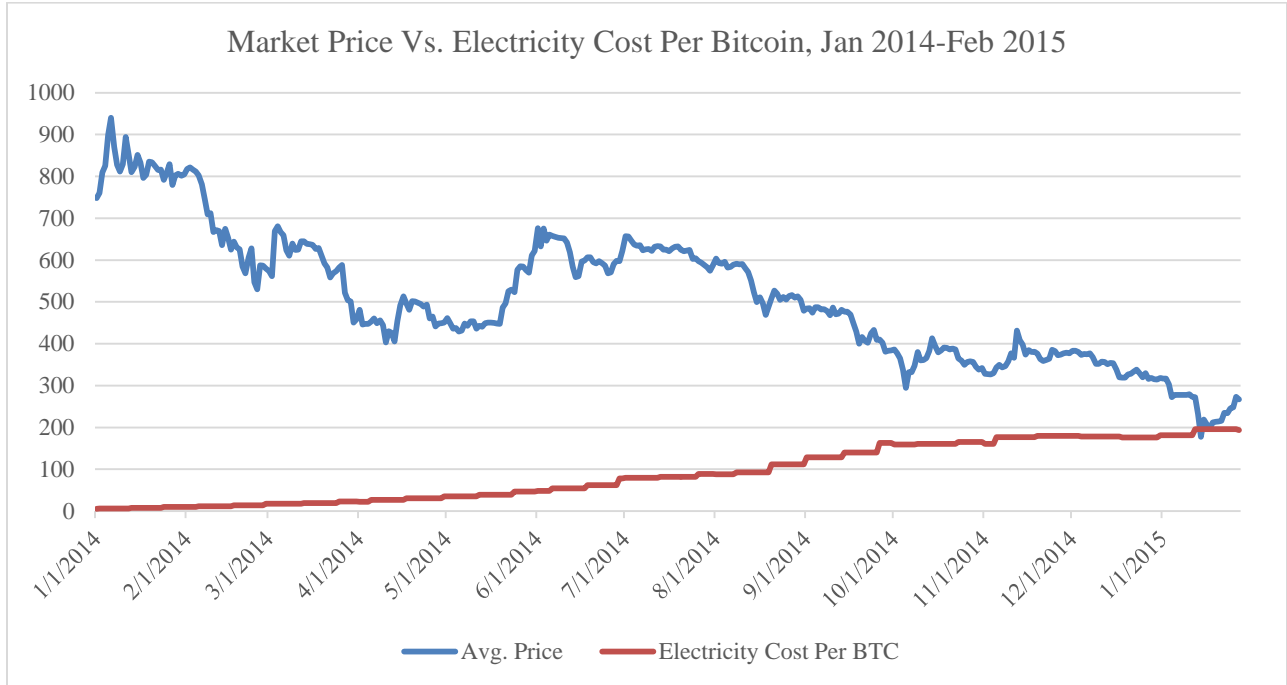
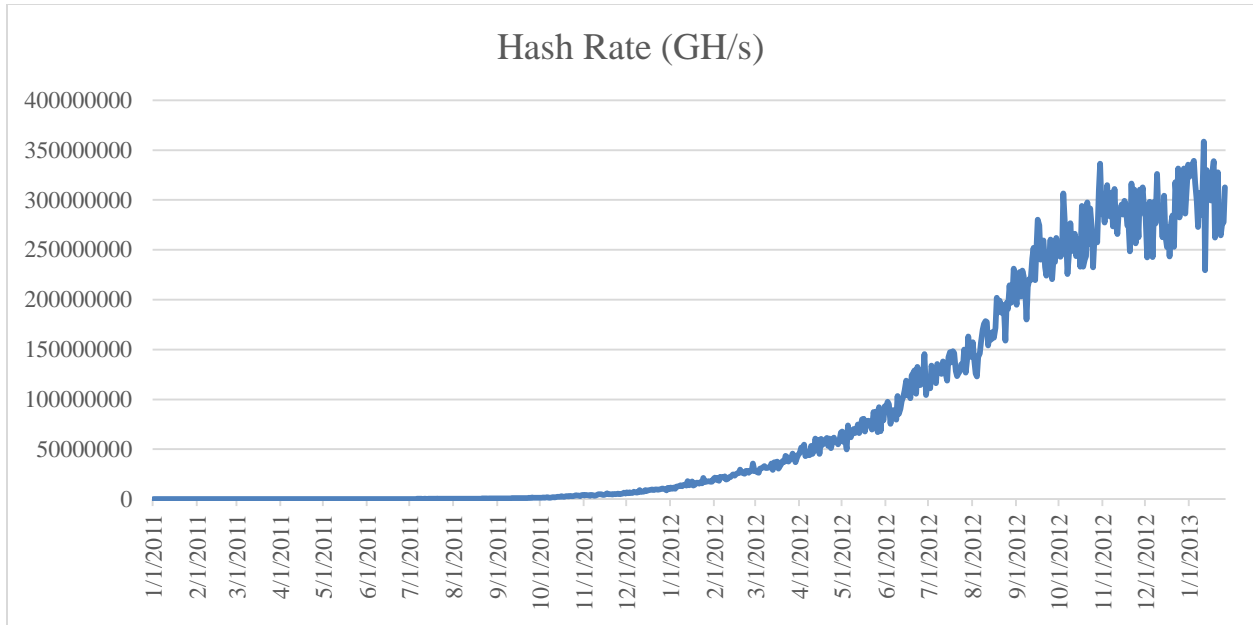


Chart 6

Source: blockchain.info



Regression Table 1: Model (1)

VARIABLES	ln(Electricity Cost per Bitcoin)
ln(price)	1.960** (0.0195)
Constant	-9.497*** (0.0817)
Observations	1,488
R-squared	0.872

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Regression Table 2: Model (2)

VARIABLES	ln(Electricity Cost per Bitcoin)
ln(price)	0.652*** (0.0361)
Linear Time Trend	0.00715*** (0.000182)
Constant	-10.14*** (0.0596)
Observations	1,488
R-squared	0.937

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Regression Table 3: Model (3)

VARIABLES	Linear Time Trend
ln(price)	182.9*** (1.947)
Constant	89.83*** (8.149)
Observations	1,488
R-squared	0.856

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Regression Table 4: Model (4)

VARIABLES	Difference in ln(Electricity Cost per Bitcoin)
Difference in ln(price)	0.714*** (0.0270)
Constant	1.354*** (0.0367)
Observations	1,308
R-squared	0.348

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Chapter Four: Concluding Remarks

This paper analyzed bitcoin from two perspectives: the user market and the mining market. The interesting result from the user market is a persistent deviation from absolute purchasing power parity despite a bitcoin not being qualitatively different regardless of where and how it is purchased. Several factors were used to attempt to explain movements in the real exchange rate—relative transaction volume and price volatility—but were shown to be ineffective. This phenomenon warrants further investigation as to the cause of this discrepancy, whether it is an idiosyncrasy related to the exchanges themselves or some other fundamental economic driver related to legal or macroeconomic factors.

The second topic called into question the viability of bitcoin as a scalable currency. Bitcoin mining is a competitive market, and so the resources expended line up with the opportunities to earn revenues. Competition in this market is almost purely destructive, however. Although mining taken as a whole serves the purpose of controlling inflation and verifying transactions, this process could theoretically be done with much fewer resources by an honest central authority. These processes can be accomplished in a decentralized way by mutually disinterested agents, but at the cost of consuming many more resources.

References

- Cheung, Yin-Wong, and Kon S. Lai. "Lag order and critical values of the augmented Dickey–Fuller test." *Journal of Business & Economic Statistics* 13, no. 3 (1995): 277-280.
- Miroslava, Rajcaniova, Kancs d'Artis, and Ciaian Pavel. "The Economics of BitCoin Price Formation." EERI Research Paper Series (2014).
- Frey, George. "Hackers Steal \$5 Million from Major Bitcoin exchange." *Fortune*, January 5, 2015. Web.
- Fung, Brian. "Shut up and Take My Bitcoins! A Map of Bitcoin-friendly Businesses." *The Washington Post*, Nov. 20, 2013. Web.
- Hill, Kashmir. "Bitcoin's Incredible Year." *Forbes*, December 31, 2013. Web.
- Hill, Kashmir (2). "21 Things I Learned About Bitcoin From Living On It For A Week." *Forbes*, May 9, 2013. Web.
- Hill, Kashmir (3). "From Walmart To Bitcoin: The CEO Behind The Chinese Exchange Sending BTC To New Highs." *Forbes*, November 8, 2013. Web.
- Kroll, J. A., Davey, I. C., & Felten, E. W. "The economics of Bitcoin mining, or Bitcoin in the presence of adversaries." *In Proceedings of WEIS* (2013).
- McMillan, Robert. "The Inside Story of Mt. Gox, Bitcoin's \$460 Million Disaster." *Wired*, March 3, 2014. Web.
- Mironov, Ilya. "Hash Functions: Theory, Attacks, and Applications." *Microsoft Research, Silicon Valley Campus* (2005).
- Nakamoto, Satoshi. "Bitcoin: A Peer-to-Peer Electronic Cash System." (2008). Accessed via bitcoin.org.

- Power, Mike. "Life after Silk Road: How the Darknet Drugs Market Is Booming." *The Guardian*, May 30, 2014. Web.
- Reader, Ruth. "Bitcoin Exchange Coinsetter Acquires Canadian Exchange Cavirtex." *VentureBeat*, April 8, 2015. Web.
- Rogoff, Kenneth. "The purchasing power parity puzzle." *Journal of Economic literature* (1996): 647-668.
- Taylor, Alan M. & Taylor, Mark P. "The Purchasing Power Parity Debate." *Journal of Economic Perspectives* 18, no. 4 (2004): 135-15.
- Tiller, Martin. "Bitcoin Basics: The Economics of Mining." *Nasdaq*, October 28, 2014. Web.
- Velde, François. "Bitcoin: A primer." *Chicago Fed Letter* (2013).
- Yermack, David. "Is Bitcoin a real currency? An economic appraisal." No. w19747. National Bureau of Economic Research (2013).