Hotelling Meets Crypto-Currency: Do Bitcoin Rents Follow Hotelling's Rule?

Craig E. Landry, Dylan Turner, & Jeffrey H. Dorfman*

August 16, 2022

Abstract

Hotelling's Rule is a central premise of natural resource economics, but there are difficulties in empirical assessment. We test Hotelling's Rule using bitcoin rents. Similar to natural resources, new bitcoins enter the market solely through the efforts of agents expending computer resources. Bitcoin has a fixed resource stock, exhibits a fairly homogeneous technology, and aspects of investment behavior are observable over time. These factors make bitcoin an ideal test bed for Hotelling's Rule. We find evidence of co-integration of bitcoin rents with several market indices, providing support for Hotelling's Rule.

Key Words: asset pricing, bitcoin, natural resources, Hotelling's Rule

JEL Classification: G1, Q3

*Landry: Professor, Department of Agricultural and Applied Economics (AAEC), University of Georgia. `clandry@uga.edu`. Turner: Ph.D Student, AAEC, University of Georgia. `dylan.turner25@uga.edu`. Dorfman: Professor, AAEC, University of Georgia. `jdorfman@uga.edu`.

1 Introduction

In laying out a theory of non-renewable resource utilization that could allay concerns over myopic depletion of society's scarce natural resources, Harold Hotelling (1931) created a central and enduring paradigm of modern resource economics. Since miners of a finite, known stock of non-renewable resources will treat the stock as natural capital, they should extract such that real returns from a marginal increment of the resource are equalized across time. As such, the optimal extraction path will lead to economic rents rising at the rate of market interest. This extraction plan is akin to that of the benevolent social planner - a special application of the First Fundamental Theorem of Welfare Economics.

While theoretically sound, strong empirical evidence of Hotelling's Rule has proven rather elusive. Hotelling's initial discussion of non-renewable resource utilization focused on a handful of stylized cases. Early empirical explorations found limited evidence for exhaustible resource prices trending upwards (Barnett and Morse 1963, Smith 1979, Slade 1982, Berck and Roberts 1996), but such price dynamics apply only to the very specialized case of negligible extraction costs (among other assumptions). Subsequent studies have sought to explore more realistic formulations of non-renewable resource extraction, producing many variants of the Hotelling rule. For example, Pindyck (1978) introduced exploration and discovery to augment the natural resource base; repeated discovery of new resource stocks can induce a saw-tooth-shaped price path as the basis for asset pricing is updated (Dasgupta and Heal 1979; Arrow and Chang 1981; Lasserre 1984). Other innovations in Hotelling's theory include aspects of market structure (e.g. longevity and malleability of capital (Campbell 1980; Farzin 1984) and imperfect competition (Stiglitz, 1976)), technological change (Slade 1982), taxation (Heaps 1985), and most recently, capital investments and stock effects (Anderson et al., 2018). Given this variability, empirical tests of Hotelling are typically joint tests of the theoretical formulation and empirical validity. To complicate matters, additional empirical issues include assessment of stationarity in time-series data, potential endogeneity of key parameters, and measurement of relevant shadow values (Slade and Thille 2009).

Interest in exhaustible resources, however, has waned somewhat during the last few decades. Stavins (2011) observes that while exhaustible resources have not become more scarce (in an economic sense), many renewable resources have been plagued by open-access and common pool problems, in some cases resulting in extinction. Thus, while academic interest in, for example, fisheries has generated novel and insightful research papers that address fundamental issues in natural resource dynamics (Kirkley, Paul, and Squires 2002; Huang and Smith 2014), exhaustible resources have primarily played secondary roles in analysis of other issues, such as real options analysis of large capital investments (Kellogg 2014), local conflict (Berman et al. 2017) or market inefficiencies (Asker, Collard-Wexler, and De Loecker 2019). In the only recent paper of which we are aware that directly addressed the Hotelling Rule, Anderson, Kellogg, and Salant (2018) produce a modification that predicts little price response for crude oil extraction (due to well capacity constraints), but intertemporal optimization in well drilling (affecting rig rental prices). Our primary motivation in this paper is to contribute to the literature on Hotelling's Rule by focusing on a novel commodity that permits a much cleaner test.

The administrative structure surrounding the crytpo-currency bitcoin presents a unique set of properties that abide many of the theoretical conditions laid out by Hotelling's theory of exhaustible resources. At the most basic level, crypto-currency miners expend computing power and electricity to obtain bitcoins from an exhaustible common pool of known dimensions. The mining process produces a time profile of economic rents that depend upon market demand, the stock of bitcoins, and aggregate mining efforts. In this way, we view bitcoin mining as a *de facto* exhaustible resource allocation problem as defined by Hotelling. Interestingly, the original bitcoin white paper includes an analogy to natural resources, stating: "The steady addition of a constant amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended." (Nakamoto 2008).

In this paper, we utilize a variant of the Hotelling model developed by Levhari and Pindyck (1981) to account for resource durability. Like gold and other non-fuel minerals, bitcoins can be traded many times and may be stockpiled. For bitcoin and durable minerals, persistence of the mined resource introduces a stock effect that can put downward pressure on price, as the demand equation reflects the amount of bitcoin in circulation. In addition, the bitcoin protocol is designed to mint new bitcoins at a roughly constant rate, rendering the extraction path exogenous to the optimization problem.

Aside from these complications, bitcoin mining offers many desirable properties for testing Hotelling's Rule: the *in situ* stock of bitcoins is known with certainty; bitcoins are a common pool resource that can be easily accessed by anyone with a computer; the technology for mining bitcoin is based on computing power, which is fairly simple to model; mining "difficulty" follows a pre-defined schedule that produces predictable changes in the extraction rate; and bitcoins represent a decentralized currency that largely escapes regulatory scrutiny and is not directly subject to subsidies or taxes. Moreover, important covariates necessary to analyze Hotelling's rule–the cost of electricity, energy consumption per unit of mining effort, market price of bitcoin, and difficulty level of the bitcoin "mining" algorithm–are, in principle, observable and may be tracked over time. As such, objective decisions for bitcoin mining can be explicitly modeled (Hayes 2018), permitting a direct test of Hotelling based on demand, cost, and the market equilibrium implied by dynamic profit maximization (Slade and Thille 2009). Using this structure, we compare bitcoin mining rents to two common market indices that reflect varying levels of risk, the S&P 500 (GSPC) and a Junk Bond Index (HYG). We find a significant co-integrating relationship between mining rents and each market index suggesting that, although the bitcoin spot price can be volatile at times, mining rents are being governed by the overall market rate of return. This result is generally robust to sensitivity analysis and constitutes what we believe is some of the strongest and most compelling evidence in support of Hotelling's Rule of exhaustible resources.

The remainder of this paper is organized as follows. Section 2 details the technical aspects of bitcoin mining and provides an overview of market structure. Section 3 sets up a theoretical variant of Hotelling's rule specific for bitcoin mining. Section 4 discusses data for our empirical analysis, and section 5 details our empirical methodology. Section 6 reports results. Section 7 subjects results to a sensitivity analysis. Section 8 discusses our findings, and section 9 concludes.

2 Bitcoin Mining

Crypto-currencies like bitcoin are predicated on a set of computer protocols (easily verified, but computationally difficult) that issue virtual "coins", securitize their ownership, and verify transactions made within the network. The entire histories of coinage distribution and transactions are contained within the "blockchain", which is recorded and stored on every computer in the network, providing for security that increases with the size of the network. A full and complete description of the intricacies of the mining process are beyond the scope of this paper, but we summarize the relevant aspects of mining below and identify why it creates an ideal testing environment for Hotelling's Rule. For more detailed discussions on mining see Nakamoto (2008), Kroll, Davey,and Felton (2013), Sapirshtein, Sompolinsky, and

Zohar (2016), and Hayes (2016, 2018).

2.1 Purpose of Mining

Bitcoin mining accomplishes two critical functions for crypto-currency. First, mining verifies transactions that users initiate, and second, mining allows new coins to enter the market in a non-arbitrary manner. When a bitcoin transaction is initiated (for example sending $10 to a friend), it gets bundled with other transactions into what is known as a "block." At this point all of the transactions in the block are pending, much like a pending transaction on a credit card. Once the block is full, it gets sent to active miners for verification. Before it is sent out, however, a unique alpha-numeric string known as a "header" is applied to the block.

The process of mining includes procedures to verify that the bitcoin transactions that have been submitted to the network are, indeed, legitimate, meaning that participants on the network are trading bitcoin that they actually own. Each miner has a full history of the public ledger of transactions, known as the "block chain." The miner's task is to verify that pending transactions are consistent with the history of previous transactions in the block chain. As long as the majority of miners agree that a transaction is valid, then verification proceeds.

Verifying all individual transactions in a block is trivial, but verification of the block is not complete until a miner produces a SHA-256 hash[1] of the block's header that is lower than a target value. Doing so is computationally expensive and generates the artificial scarcity that limits the rate at which bitcoins enter the market. A full description of this process

[1] A hash is the output of a hash function. A hash function is a one-way cryptographic function that turns an input into a fixed length alpha-numeric string. It is "one-way" in the sense that the input cannot be recovered from viewing the output. The only way to verify the input used to create the output is through trial-and-error until an input is passed through the function that matches the output in question.

is outside the scope of this paper, but what is important for our purposes is to understand that verifying a block of bitcoin transactions entails guessing random numbers until one is found that allows verification of the block to be completed.

The first miner to achieve verification of the header is awarded the transaction fees that were paid by all the users who had transactions in the block, plus an additional cache of bitcoin known as the "block reward". The block reward is the only method by which new bitcoins enter circulation. Transaction fees are paid from currently circulating coins, whereas the block reward consists of new bitcoins, that have never been in circulation.

The underlying mechanisms of the hash function render verification of the header akin to guessing random numbers from a uniform distribution. Consequently, the probability of a given miner solving a block is proportional to their share of the total computational power on the network. If a miner controls 1% of the aggregate computational power being directed at mining, they have a 1% chance of solving each block. In reality, most miners control only small fractions of a percent of the aggregate computational power, meaning they solve blocks very rarely. To help smooth revenues, it is common for miners to pool resources. When one miner in the pool solves a block, everyone in that pool splits the reward proportionally to how much computational power they contribute to the pool.

All block chain information is publicly available and can be viewed through a number of methods. Here we use the block explorer available at www.blockchain.com/explorer, to demonstrate; consider Figure 1. The block number (orange rectangle) is 564,595 indicating that 564,594 blocks have been verified previous to this one. The red rectangle shows that there are 2816 individual transactions in this block that were verified, and the blue rectangle shows that the sum of the transaction fees on those 2816 transactions is .132 bitcoins. The green rectangle shows that the block reward is 12.5 bitcoins. This particular block was solved by “AntPool” (purple rectangle), a well known mining pool. In total, the miners in this pool earned 12.632 bitcoins: 12.5 bitcoins from the block reward plus an additional .132 bitcoin in transaction fees, all of which was distributed among the miners actively participating in the “Antpool” mining pool when the block was solved.

2.2 Extraction Dynamics

Because bitcoin exists in a digital medium, the mining protocol can regulate, at virtually no cost, the market dynamics in a way that is not possible with physical natural resources. This regulation comes in the form of a dynamic block reward and a dynamic difficulty parameter, which alter returns to mining based on well known and pre-specified criteria. The block reward was designed to encourage miners to verify transactions during the emergence of bitcoin when transaction volume was low, and transaction fees alone were insufficient to incentivize verification (Nakamoto, 2008). Additionally, given that there is no central authority, the block reward serves as the mechanism to distribute the entire bitcoin stock. A common goal of crypto-currencies is creation of a decentralized currency. In the case of bitcoin (and others), users pay “miners” to verify transactions; the block reward decreases over time to ween miners off initial distribution, and the protocol transitions towards transaction fees being the sole source of mining revenues. Initially, the block reward was set to 50 bitcoins, and it decreases by a factor of two every 210,000 blocks. In practice, this implies that the block reward is cut in half approximately every four years.

Since solving a block is essentially a matter of guessing a correct random number, more computational power both from improvements in technology and increased popularity of mining can vastly alter the rate that blocks are solved and new coins enter circulation. To counteract this, the bitcoin protocol has a built in difficulty parameter that adjusts based on the aggregate computational power on the network at any particular time. The difficulty adjusts by altering the probability that a randomly guessed number will result in verification of a block which in turn increases or decreases the expected amount of time it takes to solve each block.

The difficulty parameter is defined to target a solution rate of 1 block, on average, every 10 minutes. Difficulty adjusts either upwards or downwards every 2016 blocks based on the total time it took to solve the previous 2016 blocks. At 1 block every 10 minutes, difficulty adjusts approximately every two weeks. In theory this means that regardless of bitcoin demand or computational power, the aggregate rate that blocks are solved is constant. In practice, it is possible that short term changes in the aggregate mining effort cause blocks to be solved at a faster or slower rate then the target 1 block per 10 minute target. Figure 2 shows the theoretical extraction path to date (based on 1 block being solved exactly every 10 minutes) in comparison to the actual observed extraction path.

The deterministic nature of the block reward and difficulty adjustment make it possible to define the rate of extraction for any given period. At the current block reward of 12.5 bitcoins per block, approximately 1800 bitcoins are mined per day. Given that the block reward decreases by a factor of two every 210,000 blocks, the total number of bitcoins extracted can be represented by an infinite geometric series that converges in the limit to 21,000,000[2]. Since the "*in situ*" stock of bitcoin is a function of a convergent geometric

[2]Why bitcoin was set up with a total stock of 21 million as opposed to any other amount was never explicitly discussed in Nakamoto (2008). A correspondence by Satoshi Nakamoto on an internet forum, however, suggests that the stock level was chosen so that if bitcoin was widely adopted, the bitcoin prices of goods could reach a certain level of parity with the US dollar. In particular, if the unit price of bitcoin ever reaches $1 million, then the smallest divisible bitcoin unit (1 bitcoin can be split into 100 million discrete units, each of which is colloquially known as a "satoshi") would have parity with the U.S penny.

series, the total supply is fixed and precisely known[3]. This is in contrast to other exhaustible resources which, by definition, have fixed stocks but the precise size of that stock is usually unknown with asset pricing markets working off a constantly evolving series of estimates. Additionally, technological innovation can unexpectedly increase the economically viable stock of traditional natural resources.

2.3 Common Pool Resource

Bitcoin mining can be conducted from anywhere in the world, and there is no subdivision of the stock of bitcoins that are paid out to miners independent of the block reward, making it a common pool resource with uncertain individual returns. Although, Hotelling did not specifically address extraction of common pool exhaustible resources, the topic has since been given attention in the natural resource literature. Sinn (1984) offers a theoretical model of oligopolistic firms that have access to oil reserves which exhibit interconnectivity, potentially leading to seepage from one reserve to another. Seepage introduces elements of a common pool resource. If above-ground storage is available, firms have incentives to over-extract oil to prevent seepage into rivals' reserves. Sinn shows that under reasonable conjectures about rival firms' extraction behavior, dynamic equilibria can be characterized by immediate extraction of all reserves, which are then stored above ground and sold off on the market at the firm's discretion, or gradual movement toward above ground storage, with equilibrium extraction patterns depending upon marginal storage costs (and ignoring

[3]The total stock being a consequence of the decreasing block reward raises the question of whether the stock of bitcoins will be exhausted, or if the extracted stock will just get arbitrarily close to the 21 million-unit cap. It is technically possible to extract the last bitcoin since bitcoin isn't infinitely divisible. Figure 3 shows the theoretical extraction path, again based on solving exactly 1 block every 10 minutes. Regardless, its clear that the vast majority of all bitcoin will be extracted rather quickly, followed by a slow asymptotic approach towards 21 million. The theoretical harvest path suggetss that approximately 99% of all bitcoin will be extracted by 2035.

extraction costs). In this framework, resource price increases faster than the market interest rate due to the added costs of storage. Only when above ground store costs are zero do we recover the standard Hotelling result (despite the common pool nature of the exhaustible resource).

Gaudet, Moreaux, and Salant (2002) consider a similar situation, but focus on dynamics of extraction cost (ignoring storage costs). For an exhaustible common pool resource that may be stored after extraction and is subject to open access, the price path will eventually rise at the rate of interest, but may have an initial phase during which price is constant (similar to competitive equilibrium). The competitive outcome phase only appears if initial *in situ* reserves are sufficiently large to saturate market demand for a period of time before storage begins, at which time the price begins rising at the rate of interest. This is particularly relevant to bitcoin, since storage costs are virtually zero and there is no quality degradation of the resource stock while in storage. Given the set adjustments in the difficulty parameter, however, the rate of extraction for bitcoin is effectively exogenous.

2.4 Supply Constraints

Since the bitcoin mining protocol adjusts the difficulty parameter to render the extraction rate roughly constant, miners are not able to select an extraction path. Rather, they can make investments in mining hardware while utilizing electricity to augment computational power applied to the bitcoin network. The supply constraints implied by the cryptocurrency protocol are similar to well capacity constraints analyzed by Anderson, Kellogg, and Salant (2018). In their application, geological constraints render oil extraction from individual wells exogenous, but they find evidence of intertemporal optimization, consistent with Hotelling's Rule, in well drilling and rig rental prices, which are capital investments. Similarly, bitcoin miners have no control over the rate of coinage, but make investments in mining hardware to maximize returns.

2.5 Returns from Mining

Individuals engaging in bitcoin mining choose the amount of computing hardware to commit to the endeavor. At any point in time, the lifetime of the hardware (T) and the efficiency (θ) are assumed given, so the choice variable is individual computing power (in hashes per second or "hash rate"), x_i.[4] Additional components of cost include the price of electricity (γ_{st}), which varies across space and time, and the fixed cost of capital used for computing which is expressed as a rental rate per unit of computational power, k_{it}. Revenues from bitcoin depend upon the block reward at time t, b_t, and the spot price of bitcoin, $P(B_t)$, which in turn depends upon the stock of bitcoin in circulation, $B_t = \sum_{h=1}^{t} b_h$, and market demand for bitcoin as a currency, as a portfolio asset, or potentially for speculative purposes.

Under this formulation, the individual miner's optimization problem is:

$$\max_{x_i} \pi_i = \sum_{t=1}^{T} \delta^{t-1} \left[P(B_t) b_t \left(\frac{x_i}{x_i + X_{-it}} \right) - x_i \left(\theta_i \gamma_t + k_{it} \right) \right] \tag{1}$$

where t denotes block chain verification periods, and X_{-it} is the total computing power applied by other miners on the bitcoin network at time t. Equation 1 represents expected returns from mining, with the probability of winning the block reward increasing in own efforts and decreasing in the total effort of other miners. Without loss of generality, assume individual hash-rate is small relative to the total computing power at any given time, so we can replace $x_i + X_{-it} = X_t$. Lastly, we will assume risk neutrality, so miners simply want to

[4]One hash is equivalent to 1 guess at the correct answer needed to solve a block. Mining hardware with a hashrate of 1 terahash per second indicates the hardware is capable of making 1 trillion attempts at solving the current block every second.

maximize equation 1.

First order conditions for an optimum investment in computing power produce the following variant of the Hotelling Rule:

$$\left[\frac{P(B_t)b_t}{X_t} - \theta_i\gamma_t - k_{it}\right] = \delta\left[\frac{P(B_{t+1})b_{t+1}}{X_{t+1}} - \theta_i\gamma_{t+1} - k_{it+1}\right] \tag{2}$$

for all t = 1 to T-1. Equation 2 indicates that expected returns from mining should rise at the rate of market interest. If this were not the case, miners would be better off making other investments in the economy.

3 Theoretical Model

While it looks familiar, equation 2 is quite different from the standard Hotelling first-order conditions. Equation 2 describes optimal capital investment, which provides the intertemporal link across block verification periods. This is distinct from a typical natural resource mining problem where extraction quantity gives rise to the condition that preempts intertemporal arbitrage. Given the exogenous rate of crypto-currency coinage created by the difficulty parameter and the pre-specified block reward schedule, the primary determinant of market dynamics that remain under the control of miners is the aggregate mining effort. Summing equation 2 over aggregate computing power (X) produces the market conditions for Hotelling's Rule applied to crypto-currency:

$$\left[P(B_t)b_t - X_t\left(\bar{\theta}_t\bar{\gamma}_t + \bar{k}_t\right)\right] = \delta\left[P(B_{t+1})b_{t+1} - X_{t+1}\left(\bar{\theta}_{t+1}\bar{\gamma}_{t+1} + \bar{k}_{t+1}\right)\right] \tag{3}$$

for all t = 1 to T-1, where the $\bar{\theta}_t$, $\bar{\gamma}_t$, and $\bar{k}_t$ parameters represent industry averages. Substituting for the discrete discount factor $\delta = \frac{1}{(1+r)}$ gives rise to the predicted rent path for bitcoin:

$$r = \frac{\left[\dot{P}(B_t)b_t - \dot{X}_t\left(\bar{\theta}_t\bar{\gamma}_t + \bar{k}_t\right)\right]}{\left[P(B_t)b_t - X_t\left(\bar{\theta}_t\bar{\gamma}_t + \bar{k}_t\right)\right]} \tag{4}$$

Equation 4 indicates that rents from bitcoin mining should rise at the rate of market interest. The primary mechanism by which this happens is aggregate effort, which influences the total costs of computing power. In this case, marginal miners make critical decisions about entry/exit and hardware upgrades. Also important is the market price of bitcoin, which we assume is updated by rational expectations, but may exhibit short-term speculative trends.

Since time is measured in block chain verification periods, we can express equation 4 as a daily expectation by aggregating over approximately 144 verifications that occur everyday (one block mined every ten minutes for 24 hours). This produces our estimate of daily rents that will form the basis for our empirical investigation:

$$\text{Daily Asset Rent } = P(B_t)\sum_{h=1}^{144} b_h - X_t\left(\bar{\theta}_t\bar{\gamma}_t + \bar{k}_t\right) \tag{5}$$

where the summation represents the total addition of bitcoins on a given day. Equation 5 provides an empirical basis for assessing the dynamic behavior depicted in equation 4. The supply restrictions created by the crypto-currency protocol give rise to an exogenous resource flow that should influence individual capital investment patterns (Anderson, Kellogg, and Salant 2018). We note that analyzing bitcoin rents that have accrued since mining began includes periods where the block reward exhibits a discrete downward jump, which should induce further adjustments in effort. We turn next to the data that we assemble to analyze equation 5.

4 Data

Bitcoin data for our analysis come from data.bitcoinity.org, a website that aggregates data from the bitcoin blockchain and compiles it into a daily time series. Relevant to our analysis is the spot price of bitcoin ($P(B_t)$ in equations 2 - 5), the contemporaneous block reward (b_t in equations 2 - 5), and the total computational power being directed towards mining (X_t in equations 2 - 5). Although bitcoin originated in January of 2009, a consistent market price was not available until bitcoin exchanges became commonplace. Consequently, sufficient price data are available only since July 18, 2010. Our data report daily price for every major bitcoin exchange which we average to get a single daily price.

The block reward is directly observable and relatively static, changing only every 210,000 blocks (approximately every 4 years). To date there have been three block reward eras including the initial era with a block reward of 50 bitcoins which lasted until November, 28, 2012. After this the block reward dropped to 25 bitcoins until July 10, 2016, when it dropped again to its current level of 12.5 bitcoins. The block reward is currently anticipated to fall to 6.25 bitcoins sometime around May, 2020, ushering in the fourth block reward era.

The aggregate computational power directed at bitcoin mining (hash rate) is also directly observable and can be plugged into equation 5. Unlike aggregate market data, the price of electricity that miners face and the energy efficiency of the mining hardware being employed are not directly observed and must be approximated. To get an estimate of hardware efficiency, we follow a method used by Hayes (2018). We utilize archival web page data[5] to access previous versions of wiki pages that are consistently updated by the bitcoin mining community and show all contemporaneously available mining hardware for a selected date.

[5]Using `www.waybackmachine.com` allows us to view historical records of the following wiki page: `https://en.bitcoin.it/wiki/Mining_hardware_comparison`

We web scrape all available entries (388 total, spread out from May 31, 2011 to October 31, 2019) of the wiki page to obtain a comprehensive list of what mining hardware was available throughout our time series. Along with availability, this data contains precise information on retail price for each unit, computational power, and electricity usage.

To infer which hardware a miner is likely to select, we use data on computational power and electricity consumption of mining units to construct the profits that each mining unit is expected to generate from one day of mining. In addition, we calculate the time each mining unit would need to run to cover it's purchase price through generated profits. Doing so yields a list for each day in our time series of available mining units and the estimated "repayment period" for each unit. We assume bitcoin miners always select hardware that has the shortest repayment period,[6] which allows us to define the optimal hardware choice for each day in our time series. Using the mining hardware characteristics identified through this selection criteria, we construct estimated average variable costs and capital costs incurred from bitcoin mining.

The energy efficiency of the selected hardware is used to infer electricity consumption per unit of computational power ($\bar{\theta}$ in equation 5). To construct a rental rate on capital that miners face ($\bar{k_t}$ in equation 5) we use the retail price of the selected mining hardware on each day to calculate the hardware's fixed costs per unit of computational power that the unit possess. Figure 4 shows the results from this calculation where capital costs per megahash per second (i.e. capital costs per unit of computational power) are on the y-axis. We divide these capital costs by 730 (a typical 2 year life span for a mining unit expressed in

[6] We choose this selection criterion since it takes into account both fixed and variable costs as opposed to picking the most efficient unit which only optimizes with respect to variable costs.

days)[7] which yields an estimate for capital costs incurred on a daily basis. Given that we are defining capital costs as a rental rate, we apply an interest rate on daily capital costs by using the 12 month LIBOR rate plus two percentage points (i.e. "LIBOR plus 2"), which is a reasonable basis for global interest.

At this point, our data construction is consistent with a narrative describing a bitcoin miner who rents mining hardware for the day, returns it at the end of the day and pays the fixed costs associated with their use of the equipment plus interest. They can then rent the same hardware for the next day, or if a more profitable offering is released that day, they will rent that hardware.

Since bitcoin miners are unlikely to be swapping out mining hardware on a daily basis, however, we apply a 180-day-moving-average to our time series of hardware characteristics (computational power, electricity consumption, and retail price), to account for capital inflexibility, before calculating the aggregated variable and capital costs. This formulation is consistent with a 6 month window for phasing out old hardware.[8] All subsequent analysis is based on this construction of mining costs.

Given that electricity pricing varies based on location, and bitcoin mining takes place all over the world, defining the electricity price that miners face is non-trivial. Hileman and Rauchs (2017) make an attempt to geographically locate most major mining facilities along with estimates of how much energy they use. By their account, they are able to locate

[7]To assess the typical useful lifespan of mining hardware we calculate the expected daily profits for each hardware unit from its release date to the end of our time series. We apply a 3-month-moving-average to these daily profits and define the end of the unit's lifespan as the point where the value of the moving-average drops below zero. Doing so suggests that the average useful lifespan for all hardware meeting our "shortest repayment" criteria is 794 days, or about 2 years and 2 months.

[8]The release of more efficient and profitable mining hardware is likely to trigger a rapid transition to the new hardware as the miners who obtain it first can realize extra-ordinary profits until the rest of the mining industry also obtains the new hardware. Thus the order of the moving average applied here should be relatively short compared to the actual lifespan of a mining unit. Regardless, we subject this assumption to sensitivity analysis later on.

the origin of the electricity for 63% of all bitcoin mining as of 2017. We assume that this geographic distribution of mining facilities is constant throughout our time series. Based on this, we collected country specific data on residential electricity prices across time. We then construct a weighted (by electricity usage) average world electricity price that bitcoin miners face. Figure 5 shows the bitcoin price over time along with the estimated electricity costs and capital costs per bitcoin mined.

4.1 Assumptions and Caveats

4.1.1 Mining Hardware Characteristics

Constructing the rental rate on mining hardware requires assumptions on the hardware selection criteria, the lifespan of a mining unit, and the interest rate a miner would face in capital markets. We conduct sensitivity analysis on these assumptions and report the results in a later section. Overall, our empirical results hold for a reasonable range of values for these critical parameters.

4.1.2 Electricity Price

With respect to our estimate on the price of electricity that bitcoin miners face, it's possible that this price is biased upwards since we use residential rates, when in reality large mining operations may potentially be able to negotiate cheaper industrial rates depending on the country. This would amount to an intercept shift in our mining rents time series. Given that our main results are the product of co-integration testing, this should have no effect on our conclusions since co-integrating relationships are invariant to a change in intercept.

Although we capture some limited electricity price dynamics by updating the price of electricity for each year in the data series, there is concern that the average electricity price could be endogenous with respect to the dynamics of bitcoin price. For example, if mining profitability is sufficiently low, it could spur endogenous geographic location of mining facilities into areas with lower electricity prices. We find that for the majority of bitcoin's existence, however, the maximum price of electricity that miners could pay and maintain profitability is much higher that almost all sources of electricity. This means that until recently, the price of electricity has not been the primary factor that determining profitability. Figure 6 depicts the natural log of the "break-even" price for electricity. More specifically, this is the price for electricity (in USD per kilowatt hour) that a bitcoin miner could pay and have their revenues from mining exactly offset variable and fixed costs.

To limit the effect of any endogenous relocation into cheaper electricity markets, we apply a buffer of 150% of our estimated world average price (red line in figure 6), with the assumption that as long as the "break even" electricity price is is consistently above this threshold, there would insufficient pressure to force miners to geographically relocate mining operations solely for cheaper electricity. The break even price dropped below the 150% buffer for approximately 4 months in early 2016, however we do not believe that this is long enough to prompt the relocation of mining operations to cheaper electricity markets. On June 15, 2018, however, the break even price dropped below the 150% buffer and has stayed below that threshold almost continuously since. Additionally, the break even price stayed below the weighted world average electricity price for much of late 2018 and early 2019. We suspect endogenous relocation of mining facilities to be an issue in this portion of the data. Accordingly, we drop data after June 15, 2018 and conduct our analysis on data before this point in time.

Finally, our analysis of mining hardware profitability revealed that no profitable mining hardware existed prior to Feburary 16, 2012. This is primarily due to the price of bitcoin being very low in the crypto-currency's infancy and mining taking place on consumer-grade CPUs. As such, we can safely assume that bitcoin mining before then was regarded as a hobby for those interested in blockchain technology and not an investment venture. Accordingly, we drop data before this point, meaning our final time series for our analysis runs from February 16, 2012 to June 15, 2018.

4.1.3 Growth Rate of Mining Rents

At various points in our time series, mining rents hover near zero due to short-term fluctuations in the price of bitcoin. This means that periodically the one-year growth in mining-rents-time-series has large one-day "spikes" caused by large percentage changes (i.e. mining rents were \$.02 per bitcoin 1 year ago, and are \$25 per bitcoin now, resulting in a percentage change of about %1200). We smooth out these spikes by applying a 30-day-moving-average (which we subject to sensitivity analysis) to the mining rents growth rate time series. We believe this is a fairly innocuous procedure given that miners are unlikely to be altering business decisions based on a single days return, but instead on the basis of monthly or quarterly returns. Figure 7 shows the unadjusted mining rents and figure 8 shows mining rents with the moving average applied.

5 Empirical Methodology

Hotelling's rule provides us with the theoretical path that bitcoin rents should follow. We recognize that noise traders, neophytes, and speculators may hold sufficient sway in the bitcoin markets, thus causing the price of bitcoin to deviate from the theoretical prediction. Such deviations, however, represent a potential profit opportunity for any bitcoin traders who spot the over- or under-pricing of bitcoin and that profit potential gets larger the greater the deviation between the price of bitcoin and the theoretical path specified by Hotelling's Rule. This story precisely fits the idea of two co-integrated time series following an error correction model (Granger 1981; Engle and Granger 1987). If bitcoin rents, as depicted in equation 5, are shown to be co-integrated with the returns of, for example, the S&P 500, the interpretation is that some economic force is pulling bitcoin rents back in line with Hotelling's Rule; whenever the deviation between the two asset returns grows large enough to produce potential profit, arbitrage will restore the equilibrium relationship. The mechanism for restoring bitcoin prices to the Hotelling price path is simply to divest mining hardware when bitcoin rents are below the equilibrium path and to invest in mining hardware when bitcoin rents are above the path, so if Hotelling's Rule is correct it should be straightforward for the marketplace to enforce it.

This presents us with a simple approach to testing Hotelling's Rule: test the returns series for both bitcoin and a chosen market index for co-integration. If the two are co-integrated, then Hotelling's Rule would be validated in a manner analogous to either the weak or semi-strong form of Fama's market efficiency hypothesis; the strong form analogy would have bitcoin prices never straying from the path, a condition anyone can see does not hold simply by looking at a graph of bitcoin rents over time. Therefore, we first test the two return series for unit roots and then test the two series together for co-integration.

The series were tested for unit roots using the Phillips-Perron test (Phillips and Peron, 1988) due to its robustness to heteroscedasticity and its superior power compared to the Dickey-Fuller test. As demonstrated by Lee, List, and Strazicich (2006), however, conclusions about the stationarity of natural resource prices can be conditional on whether structural breaks are accounted for in the unit-root test. Accordingly, in addition to the Phillips-Perron test, we employ the unit root test introduced by Lee and Strazicich (2003) which tests for stationarity in the presence of either one or two endogenously selected structural breaks.

Table 1 reports results for both the Phillips-Perron test and the one and two break Lee-Strazicich tests. Phillips-Perron test statistics indicate no evidence of stationarity for any series used in our analysis.[9] Lee-Strazicich test statistics indicate no evidence of stationarity for any of the series in the specification that allows for one structural break. The two break version of the test rejects the null of non-stationarity for the junk bond ETF (at 5% significance level) and the S&P 500 (at 10% significance level). The test finds no evidence of stationarity in the mining rents series. Overall, the collection of unit root tests conducted suggest all series are non-stationary and of order I(1); thus, the series are eligible to be tested for co-integration.

We utilize the commonly used methodology of Johansen (1988) to test for co-integration, which is predicated on a vector error correction model (VECM), of which the basic form is defined in equation 6.

$$\Delta \mathbf{X}_t = \Pi \mathbf{X}_{t-1} + \sum_{j=1}^{p-1} \Gamma_j \Delta \mathbf{X}_{t-p+1} + u_t \tag{6}$$

$\mathbf{X}$ is a vector containing the time series under consideration, Γ is a parameter matrix and $\Pi = \alpha\beta'$ where β is a matrix that contains the co-integrating vectors and α contains adjustment parameters which describe the speed at which the system adjusts to a new equilibrium. In our case the VECM is specified without a deterministic trend as supported by the likelihood ratio test of Johansen and Juselius (1990). Thus, u_t takes the form of a Gaussian white noise process with zero mean. Testing for co-integration amounts to determining the rank of Π, by sequentially testing, via a likelihood ratio test, the null of $H_0 : rk(\Pi) = \phi$ against the alternative $H_1 : \phi < rk(\Pi) \leq \phi + 1$ for increasing values of ϕ until the test fails to reject. Once the test fails to reject, the current value of ϕ is inferred to be the number of

[9] The lag orders used in all stationarity tests were chosen according to the Newey-West criterion.

co-integrating relationships in the system. Given that our test only consists of two time series at a time, the rank of Π will be at most one.

6 Results

Test statistics (both trace and max eigenvalue) from the Johansen tests are presented in table 2. The tests were specified with a lag order of 8 as determined by the AIC. As previously mentioned, the test is specified without a constant or deterministic trend (although the results are robust to either of these specifications). The test rejects the null that there are no co-integrating factors between mining rents and the S&P 500. Similarly, the test suggests there is a co-integrating relationship between mining rents and the junk bond index. The co-integrating vectors for each market index tested are reported in the following matrices.

$$\beta_{GSPC} = \begin{bmatrix} 1 & 1 \\ -263.94 & 2959.90 \end{bmatrix} \qquad \beta_{HYG} = \begin{bmatrix} 1 & 1 \\ -320.67 & 5006.52 \end{bmatrix}$$

An advantage of the Johansen procedure over the other popular test for co-integration, the Engle-Granger test, is the ability to hypothesis test restrictions on individual parameters in the co-integration vector. This includes the ability to test each series in the system for exogeneity which permits additional assessment of validity. Intuitively, mining rents should be endogenously determined from within the system (i.e. mining rents should respond to changes in the market rate) while the market rate of interest should be exogenously determined. Testing either series for exogeneity entails setting its adjustment parameter, α equal to zero and testing, via a likelihood ratio test, if the restriction holds. Table 3 reports test statistics with significance levels for this procedure. Using either the S&P 500 or the Junk Bond ETF as the comparison market rate results in a rejection (1% level) of the null that bitcoin mining rents are exogenous. With respect to the comparison indices, the test rejects exogeneity for the S&P 500, but only at the 10% level, while the test indicates no evidence of the Junk Bond ETF being endogenous. Overall tests for exogeneity are compatible with the Hotelling narrative applied to bitcoin rents.

Finally, examining the adjustment parameters from the VECM provides insight into the speed at which a divergence from the Hotelling rent path is corrected. Table 4 reports adjustment parameters and the corresponding estimated time needed for the market to correct a disequilibrium. A negative and significant adjustment parameter indicates that series will adjust towards the other series in the VECM until equilibrium is restored. Given the results from the exogeneity test, adjustment parameters are expected to be negative for mining rents and positive for the comparison market rates. The adjustment parameter for mining rents in the specification using the S&P 500 (Junk Bond ETF) as the comparison rate is -0.0041 (-0.0039). This suggests that .41% (.39%) of a deviation from equilibrium is corrected after one time period (1 day in this case). These rates suggest that following a deviation, equilibrium is expected to be restored after 246 days (254 days).

7 Sensitivity Analysis

During construction of our time series on mining costs we make four assumptions; we define: i) the criteria by which miners select mining hardware, ii) the amortization length for a hardware purchase, iii) the interest rate on borrowed capital, and iv) the rate at which old hardware is phased out upon the release of more efficient hardware. In addition, we assume that a 30-day moving average is the appropriate level of smoothing on the growth rate of mining rents to approximate the scale on which miners make business decisions.

We subject all of these assumptions to sensitivity analysis, by letting the level of each presumed parameter vary over a reasonable range and noting the percentage of the time the results hold at a 5% significance level. Figure 5 reports results from the sensitivity analysis. Altering the hardware selection criterion from “shortest repayment” to “most efficient” (i.e. the unit the has the lowest variable cost) has no effect on results. We let the hardware amortization length vary from 18 months to 2.5 years in 90 day increments and find that the results hold between 77%-92% of the time depending on which unit-root test is used.[10] We let the interest rate vary from LIBOR plus 1% to LIBOR plus 3% and find that the results hold 86% of the time. The order of the moving average for hardware characteristics (the speed that the industry adopts new hardware) is assessed over a range of 3 months to 2 years in 90-day increments. The results hold between 87.5%-100% depending on the unit-root test. Finally, our results hold 69% of the time when we let the moving-average-order for the growth rate of mining rents vary from 10-90 days, in increments of five. Overall, the vast majority of alternative specifications lead to qualitatively equivalent results.

8 Discussion

Efficiency of competitive market forces is a cornerstone of economic thought. While theoretical proofs and classroom experiments provide a foundation for understanding and demonstrating important economic concepts, empirical support is vital for testing economic ideas, verifying contextual elements, and designing effective economic and social policy. Empirical

[10] We test for stationarity in our sensitivity analysis using the Phillips-Perron test and one break version of the Lee-Strazicich test. The two break version of the Lee-Strazicich test is extremely computationally intensive for long time series and is not feasible to run (even on a distributed cluster) many times as is required for the sensitivity analysis.

evidence also provides an important feedback to theory, permitting further development and refinement of economic ideas.

Economists' understanding of the exploitation of non-renewable resources is heavily influenced by Harold Hotelling's seminal paper published in *The Journal of Political Economy* in 1931. In this paper, Hotelling lays out how firms should respond to scarcity in their decisions to extract and sell exhaustible resources. Notably, he shows that if firms seek to maximize profits, they should choose an extraction path that leads economic rents to rise at the rate of interest, just as a social planner would dictate. When Hotelling's Rule was published, however, it was not widely appreciated (partly due to the complexity of the underlying derivations). It wasn't until the 1970s, when concerns over depletion of natural resources and limits to economic growth were peaking, that Hotelling's work gained widespread attention. Since that time, the importance of Hotelling's Rule as a central premise in the economics of non-renewable resources cannot be overstated.

In energy economics, Hotelling's Rule plays a primary role in transitioning from non-renewable to renewable energy sources (Dasgupta and Heal 1979). Essentially, cheap non-renewables will dominate the energy market in early periods; energy prices will rise due to scarcity rents (following Hotelling's Rule), and the energy sector will transition to renewable resources (i.e. the backstop technology) when non-renewables become too expensive. Placing this paradigm in the context of externalities, climate change, Pigouvian taxes, and tradable permits, one becomes quickly oriented with the recent literature on economics of energy and climate change (e.g., Eichner and Pethig 2011; Jensen et al. 2015). The validity of Hotelling's Rule is particularly valid for debates on climate and energy policy (Hart and Spiro 2011; Kronenberg 2008). Announcements (like subsidies for renewable energy) that create expectations of lower future demand for a resource governed by Hotelling's Rule lower scarcity rents and can cause a decrease in price and increase in current consumption. This has come to be known as the "green paradox" (Sinn 2008) and is an active strand of current research. Hotelling's rule also plays a role in sustainable development and inter-generational equity (Hartwick 1977), technical change in capital-resource economies (Dimaria and Valente 2008), and the natural resource curse of developing economies (Boyce and Emery 2006).

Clearly, the validity of the Hotelling Rule is important for our understanding of how the economy interacts with the natural environment. This relationship has major implications for economic welfare, poverty, environmental quality, efficiency of market institutions, and the role of government policy. A primary difficulty in assessing Hotelling's Rule is the diversity of circumstances that one encounters in empirical analysis. In typical empirical evaluations of Hotelling's Rule, researchers must deal with unknown resource stocks; unanticipated discoveries; exploration costs; diversity of technologies; technological change; various market structures, including vertical and horizontal integration, and complex contracting relationships; market failures; and other difficult, real-world aspects of non-renewable resource exploitation. Joint tests of theory and empirical aspects can produce divergent conclusions about Hotelling's Rule (e.g., Hart and Spiro 2011; Anderson, Kellogg, and Salant 2018). In addition, empirical analysis is often complicated by endogeneity of key parameters, measurement problems related to assessment of shadow values, and time-series data issues (Slade and Thille 2009).

Realizing the complications inherent in applications to natural resources, we take a distinct and novel approach to assessing Hotelling's Rule. Our analysis is built upon a recognition of the similarities between crypto-currency mining and non-renewable resource exploitation. Bitcoin miners expend computer power and electricity to compete for probabilistic returns (composed of transaction fees for verification and block rewards of new virtual coins) on a regular basis (approximately every ten minutes). While the flow of bitcoins is restricted by the mining protocol's difficulty parameter, expected returns are proportional to an individual miner's share of computing power relative to total effort of all miners. As such, individual investments in mining hardware should recognize opportunity costs of capital that give rise to bitcoin rents rising at the rate of interest.

Focusing an analysis of Hotelling's rule on crypto-currency markets is advantageous for several reasons. The *in situ* stock of bitcoins is the sum of a geometric series and is known with certainty. As such, there are no discoveries that would lead to a resetting of the asset rent path (as miners update beliefs about stock). Miners incur no exploration costs that could complicate calculation of rents. The technology for mining bitcoin, based on computing power, is relatively homogeneous, well documented on the internet, and relatively easy to model. Mining returns are determined by a predetermined protocol that produces exogenous and predictable changes in the extraction rate. Bitcoins are a common pool resource that can be easily accessed by anyone with a computer, producing a competitive environment for crypto-currency returns. Lastly, as a decentralized asset, bitcoin is not directly subjected to regulatory or other policy restrictions that could complicate asset pricing and rent paths.

Gathering information on the cost of electricity, energy consumption per unit of mining effort, the market price of bitcoin, and difficulty level of the bitcoin "mining" algorithm, we are able to construct a time series of bitcoin returns since the crypto-currency's inception. We focus on the period of time after which bitcoin was viewed as a viable asset (starting in Feburary of 2012 which marked the first time profitable mining hardware was available) until rents became constrained by locally available electricity prices (starting in June of 2018), likely leading firms to geographically relocate in search of cheaper inputs. Using these data, we compare bitcoin mining rents to two common market indices - the S&P 500 (GSPC) and a Junk Bond Index (HYG). We find necessary evidence of unit-roots for each series and significant co-integrating relationships between mining rents and each market index. While the spot price of bitcoin can be volatile at times, our results indicate that mining rents are governed by the market rate of return. Our results are robust to sensitivity analyses that vary assumptions about capital investment decisions and adoption rates of new technology over reasonable ranges. We contend that the dynamic pattern of bitcoin rents provides some of the strongest and most compelling evidence in support of Hotelling's Rule of exhaustible resources.

9 Conclusion

Hotelling's rule is one of oldest and longest-standing principles in the natural resource literature and is the central premise for most modern theory on the economics of exhaustible resources. Ironically, despite its simple and elegant representation, Hotelling's original theory, has yet to be empirically validated in a satisfying way, largely due to the difficulty in recreating the stylized environment described in the original 1931 article. Empirical studies to date are typically focused on testing variants of the original theory, and although valuable in their own right, assess the validity of Hotelling's rule in an indirect way.

By leveraging the unique properties of the bitcoin mining process, we revisit Hotelling's canonical framework for exhaustible resource extraction and empirically asses its validity in the simplified environment afforded by the crypto-currency mining industry. Notably, our results are the first test of Hotelling's Rule where the assumption of a precisely known resource stock is tenable, technology is relatively homogeneous, and exploration costs are not relevant. We use publicly available data to construct a data set consisting of bitcoin mining rents at the industry level and test for co-integration with several broad market indices. Despite significant volatility in bitcoin spot prices, we find a significant co-integrating relationship between mining rents and the market rate of return, thus lending empirical support to Hotelling's rule.

Although not touched on in this paper, other crypto-currencies exhibit properties different from bitcoin, such as alternative methods of distribution and generating scarcity and some of which have dynamic stocks governed by varying principles and protocols. Utilizing these unique environments for otherwise difficult tests of theory remains a potentially fruitful area for future research.

References

Anderson, S., Kellog, R., & Salant, S. (2018). Hotelling under pressure. *Journal of Political Economy*, *126*(4).

10 Tables

Table 1: Test Statistics for Stationarity Tests

Data Series	Phillip-Perron	Lee-Strazicich 1	Lee-Strazicich 2
Mining Rents	-13.34	-4.12	-5.31
S&P 500	-16.37	-4.64	-5.52*
Junk Bond ETF	-8.83	-2.87	-6.24**

Note: For All Tests H_0: Unit Root is Present

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

Table 2: Johansen Test Results

Rank	Eigenvalue	Trace Stat.	Max Eigenvalue Stat.
Panel A: S&P 500 as Comparison Index			
rk(Π) $\leq$1	0.01	4.88	4.88
rk(Π) = 0	0.00	28.33***	23.46***
Panel B: HYG as Comparison Index			
rk(Π) $\leq$1	0.01	4.56	4.56
rk(Π) = 0	0.00	24.25***	19.69***

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Test for Exogeneity

Data Series	Test Stat
Panel A: S&P 500 as Comparison Index	
Mining Rents	13.95***
S&P 500	3.67*
Panel B: HYG as Comparison Index	
Mining Rents	15.06***
Junk Bond ETF	0.01

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Rate of Adjustment to Equilibrium

Data Series	Adjustment Rate (α)	Correction Time
Panel A: S&P 500 as Comparison Index		
Mining Rents	-0.0041***	246 Days
S&P 500	0.0000**	Exogenous
Panel B: HYG as Comparison Index		
Mining Rents	-0.0039***	254 Days
Junk Bond ETF	0.0000	Exogenous

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Sensativity Analysis (5% Significance Level)

				Percentage of Time Results Hold		
Assumption	Assumed Value	Assesed Range	Step Size	Non-Stationary	Cointegrated	Non-Stationary & Cointegrated
Panel A: Phillips-Perron Test for Stationarity and S&P 500 as Comparison Index						
Hardware Selection Criteria	Shortest Repayment	Most Efficient		100%	100%	100%
Amoritization Length	2 Years	1.5 - 2.5 Years	90 Days	76.92%	100%	76.92%
Interest Rate	LIBOR + 2%	LIBOR + 1% - 3%	.5%	100%	85.71%	85.71%
Hardware Characteristics MA Order	180 Days	3 Months - 2 Years	90 days	100%	100%	100%
Mining Rents MA Order	30 Days	10 - 90 Days	5 Days	87.5%	81.25%	68.75%
Panel B: Lee-Strazicich Test (1 break) for Stationarity and S&P 500 as Comparison Index						
Hardware Selection Criteria	Shortest Repayment	Most Efficient		100%	100%	100%
Amoritization Length	2 Years	1.5 - 2.5 Years	90 Days	92.31%	100%	92.31%
Interest Rate	LIBOR + 2%	LIBOR + 1% - 3%	.5%	100%	85.71%	85.71%
Hardware Characteristics MA Order	180 Days	3 Months - 2 Years	90 days	100%	100%	100%
Mining Rents MA Order	30 Days	10 - 90 Days	5 Days	87.5%	81.25%	68.75%
Panel C: Phillips-Perron Test for Stationarity and Junk Bond ETF as Comparison Index						
Hardware Selection Criteria	Shortest Repayment	Most Efficient		100%	100%	100%
Amoritization Length	2 Years	1.5 - 2.5 Years	90 Days	76.92%	100%	76.92%
Interest Rate	LIBOR + 2%	LIBOR + 1% - 3%	.5%	100%	85.71%	85.71%
Hardware Characteristics MA Order	180 Days	3 Months - 2 Years	90 days	100%	87.5%	87.5%
Mining Rents MA Order	30 Days	10 - 90 Days	5 Days	87.5%	81.25%	68.75%
Panel D: Lee-Strazicich Test (1 break) for Stationarity and Junk Bond ETF as Comparison Index						
Hardware Selection Criteria	Shortest Repayment	Most Efficient		100%	100%	100%
Amoritization Length	2 Years	1.5 - 2.5 Years	90 Days	92.31%	100%	92.31%
Interest Rate	LIBOR + 2%	LIBOR + 1% - 3%	.5%	100%	85.71%	85.71%
Hardware Characteristics MA Order	180 Days	3 Months - 2 Years	90 days	100%	87.5%	87.5%
Mining Rents MA Order	30 Days	10 - 90 Days	5 Days	87.5%	81.25%	68.75%

11 Figures

Figure 1: Example of a Block

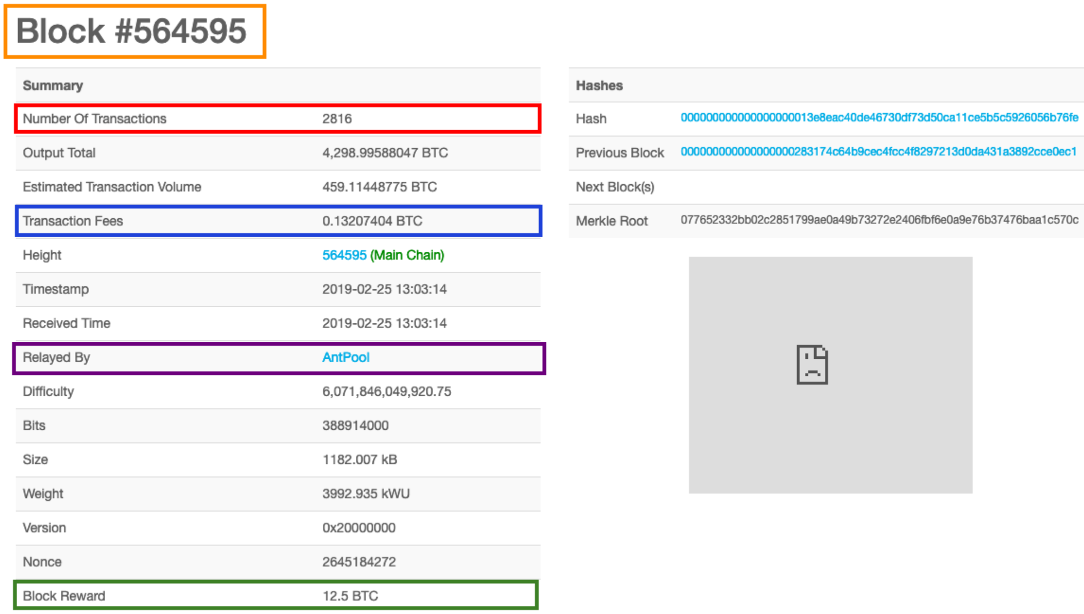

Block #564595

Summary	
Number Of Transactions	2816
Output Total	4,298.99588047 BTC
Estimated Transaction Volume	459.11448775 BTC
Transaction Fees	0.13207404 BTC
Height	564595 (Main Chain)
Timestamp	2019-02-25 13:03:14
Received Time	2019-02-25 13:03:14
Relayed By	AntPool
Difficulty	6,071,846,049,920.75
Bits	388914000
Size	1182.007 kB
Weight	3992.935 kWU
Version	0x20000000
Nonce	2645184272
Block Reward	12.5 BTC

Hashes	
Hash	000000000000000000013e8eac40de46730df73d50ca11ce5b5c5926056b76fe
Previous Block	000000000000000000283174c64b9cec4fcc4f8297213d0da431a3892cce0ec1
Next Block(s)	
Merkle Root	077652332bb02c2851799ae0a49b73272e2406fbf6e0a9e76b37476baa1c570c

Figure 2: Theoretical vs Actual Harvest Path

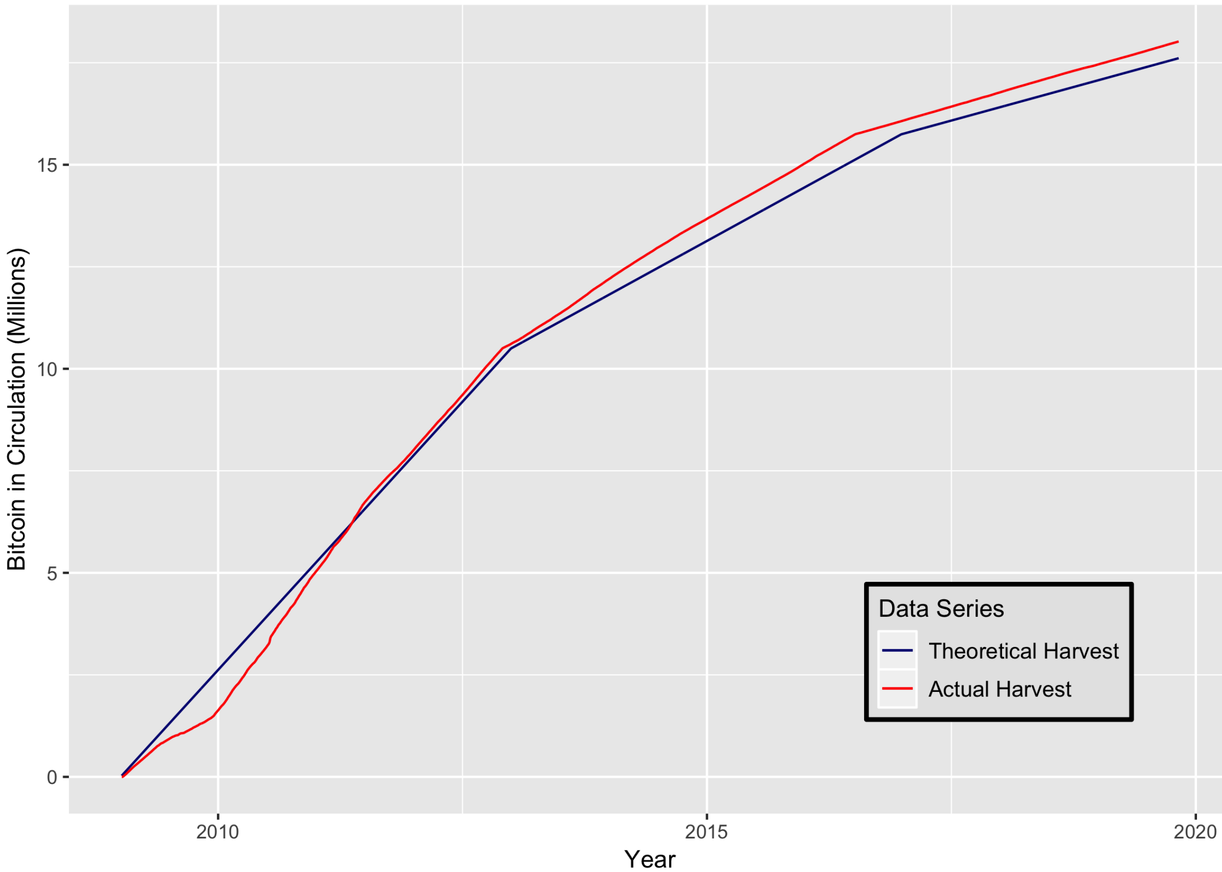

Figure 3: Theoretical Harvest Path

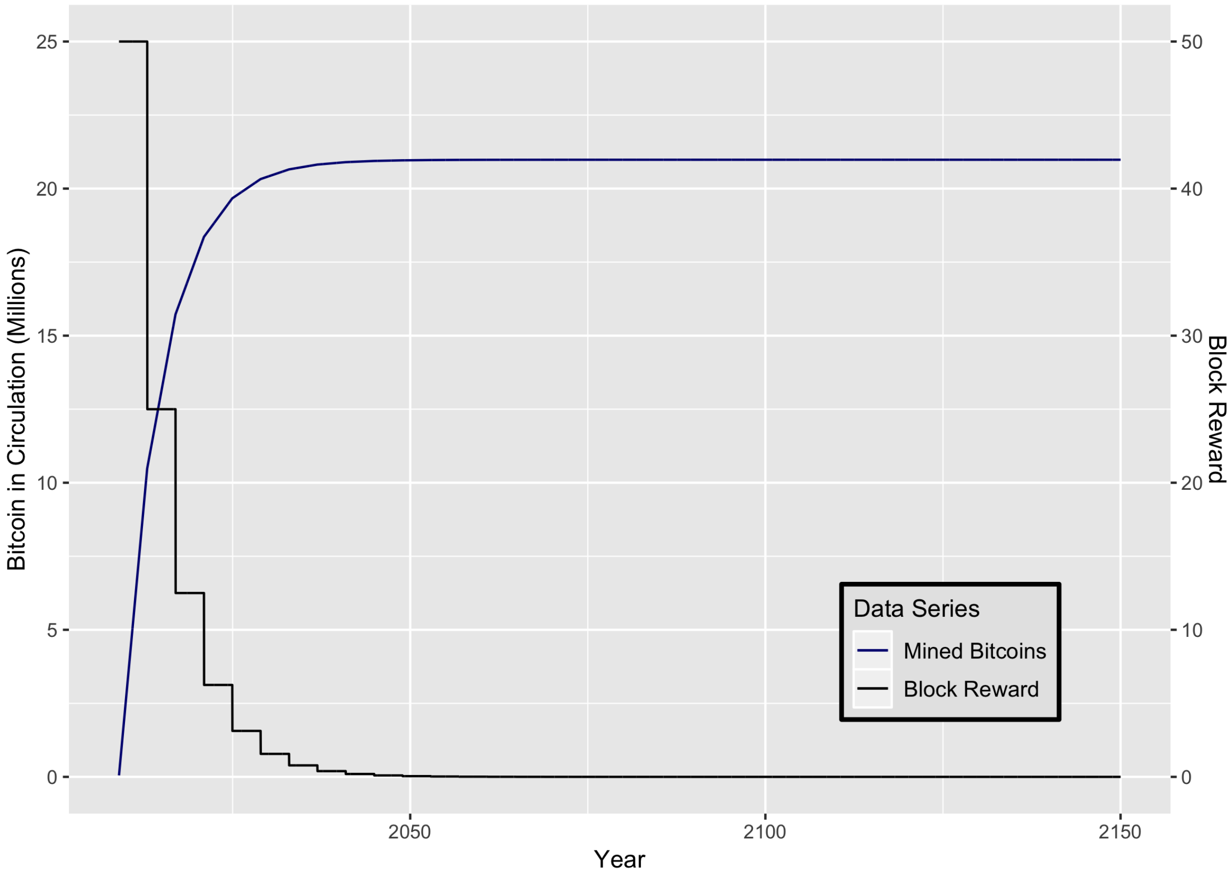

Figure 4: Capital Costs per Unit of Computing Power

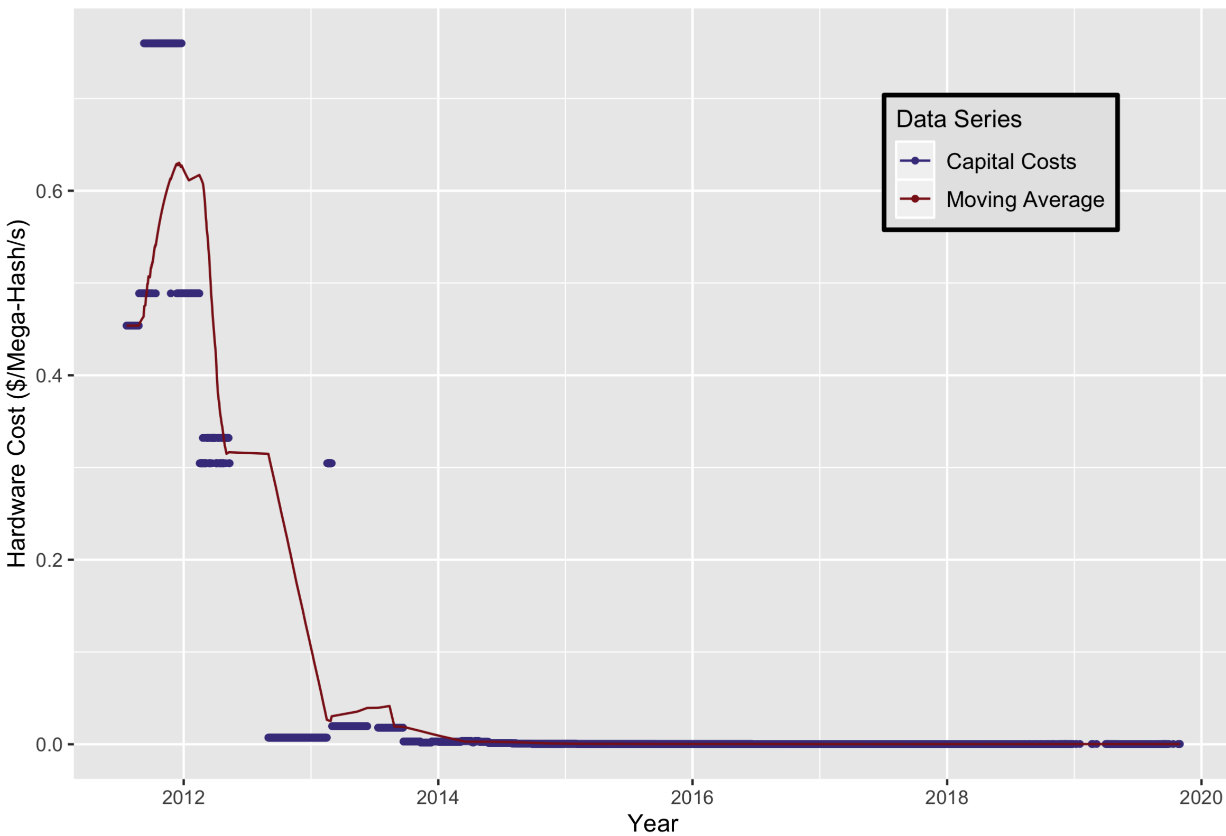

Figure 5: Price and Estimated Average Mining Cost

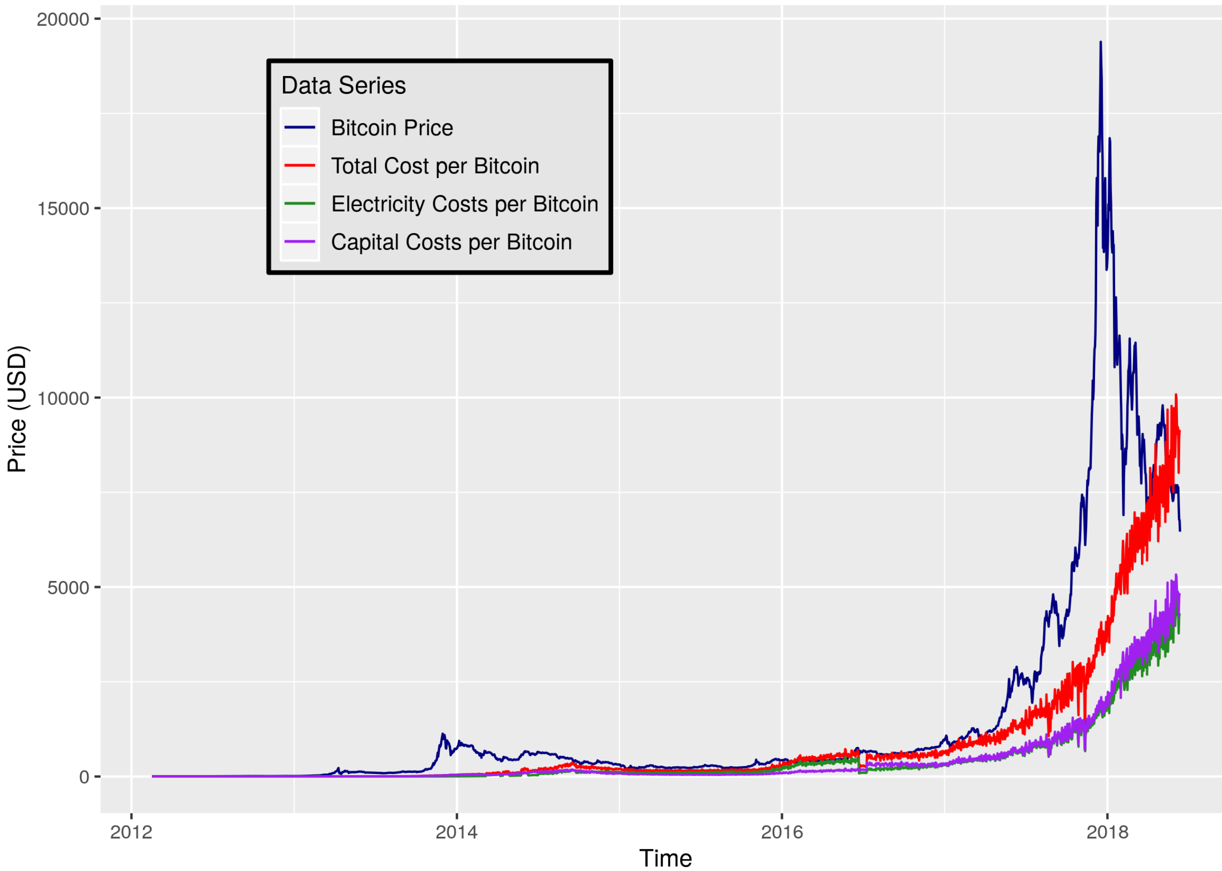

Figure 6: Log break-even Price

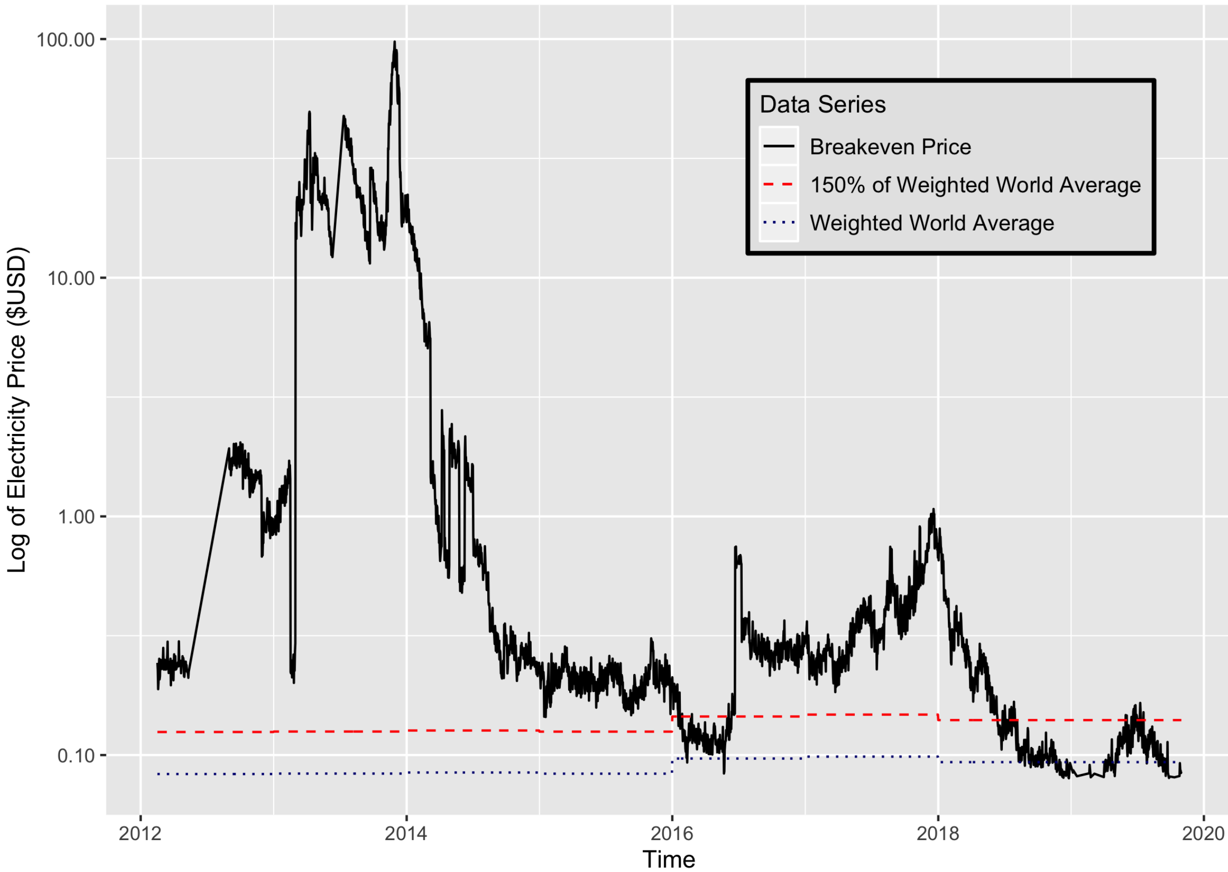

Figure 7: Percent Change From 1 Year Ago

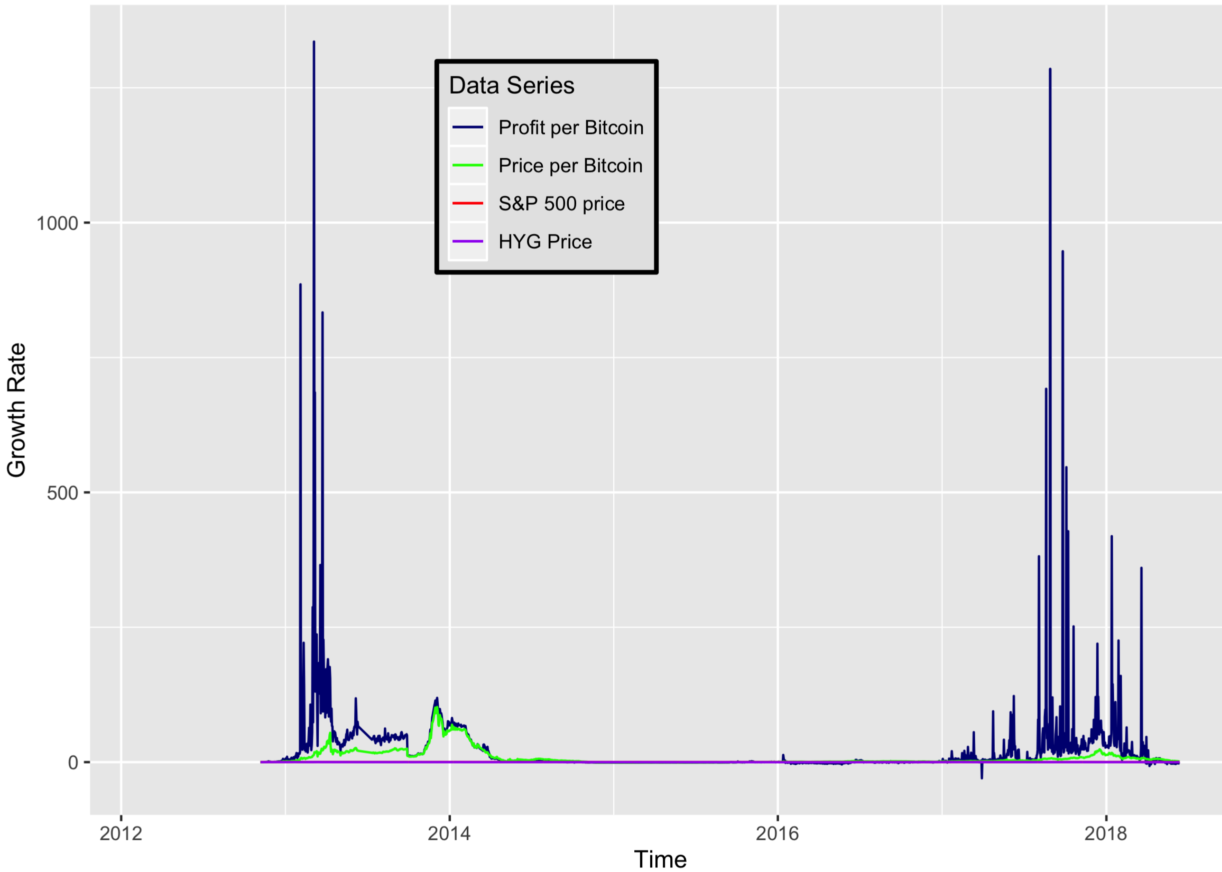

Figure 8: Percent Change From 1 Year Ago