The Effects of Bitcoin Mining On The US Electricity Market

Mathias Tveit Bohlin

Master's Thesis

The thesis is delivered to complete the degree

Master's degree in Economics

University of Bergen, Department of Economics

September 2022

UNIVERSITETET I BERGEN

Acknowledgments

Firstly, I would like to thank my supervisor, Anantha Divakaruni. I would not have been able to conduct this thesis without your help and expert knowledge. I would also like to thank my friends at the Institute of Economics for all the great moments throughout my years at UiB. A special thanks to my friends who has helped my thesis with coding and proofreading. Lastly, I would like to thank my mum and dad, your support throughout this time is greatly appreciated.

Abstract

Bitcoin mining is a process which requires vast amounts of electricity. This would impact the electricity market in the respective locations where the mining is taking place. This thesis examines the effects Bitcoin mining has had on the US electricity market in terms of electricity consumption and prices. By using panel data on electricity consumption and prices in 39 states in the US, and four sectors, the effect of Bitcoin mining is estimated in terms of percentage increases. The model estimates an increase in electricity consumption for the aggregated sector to be 0.0083 when mining is increased by one megawatt. With the same specifications, the model estimates an increase for the industrial-specific consumption to be 0.0078 percent. There are lacking sufficient significance of the estimates in the residential-specific consumption. For the commercial-specific consumption, the model estimates an increase of 0.0107 percent. In terms of prices, the model estimates that an increase in one megawatt of Bitcoin, would increase the electricity price in the aggregated sector by 0.0058 percent. The model estimates an increase of 0.0145 percent in the industry-specific price and an increase of 0.0069 percent in the commercial-specific price. For the residential-specific price, the case is the same as for the consumption.

Keywords – Bitcoin, Bitcoin mining, Master's thesis

Contents

List of Figures

List of Tables

1 Introduction

The popularity of cryptocurrencies has been through rapid growth over the last decade. This has led to various debates discussing different subjects of the usage. The oldest and most traded cryptocurrency is Bitcoin. One of the most discussed topics of Bitcoin is the electricity consumption needed to mine new Bitcoins. There have been several papers estimating the electricity consumption and the carbon footprint of Bitcoin, most of them using the same approach, although the results greatly differ.

This thesis aims to estimate the effect Bitcoin mining has had on the US electricity market. For this, I apply the estimates of one paper to analyze the percentage effect this has had on electricity consumption and electricity prices in the US. Panel-data has been used in the analysis, where the data has been gathered from various sources and covers two years, 2020 and 2021. Ideally, the period of the analysis would have been longer, although there is lacking sufficient data on Bitcoin mining previous to 2020. Almost all papers written about this subject, estimate that the electricity consumption of mining is severe. Therefore, it seems reasonable to believe that mining has led to an increase in both consumption and prices of electricity.

Further, the structure of the thesis is as follows: Firstly, I give a summary of papers that have estimated the electricity consumption and the methodology used. Secondly, I will introduce Bitcoin, and the Bitcoin mining procedure, as well as a brief overview of the US electricity market. Then an overview of the data is provided. Here I describe the variables used in the analysis, as well as how they were processed. Further, the summary statistics of the variables are provided. Moreover, a description of the methodology is laid out, followed by the results of the analysis.

2 Literature

There has been a variety of papers estimating the electricity demand and consumption, as well as carbon emissions of Bitcoin mining in recent years. The data and methodology tend to be the same, although the assumptions usually differ. Most papers use a method developed by O'Dwyer and Malone (2014). The method involves using the efficiency of the hardware used to mine Bitcoin and the hash needed to mine a block. This was the first paper that used some form of a profitability function to estimate the electricity consumption, which is known as the bottom-up approach. This methodology has been further developed by several papers. By doing this they estimated that the power demand of Bitcoin was between 0.1 and 10 GW in 2014. The efficiency of mining equipment has been through a rapid development, which makes the power consumption in terms of hashes per second, way less. Küfeoglu and Özkuran (2019) estimate, by using June 2018 as the baseline, that the annual electricity consumption of Bitcoin mining was between 15.47 and 50.24 terawatt hours. They collected data from four different sources, the blockchain itself, the efficiency of mining hardware, historical Bitcoin prices, and power cost data to conduct their calculations. They calculated the upper and lower bound of the consumption, which relied on the hardware used. The lower bound was calculated based on the most efficient hardware at the time. The upper bound was the break-even point of mining revenue and electricity costs. This estimate seems to be higher than what most other papers estimate, so it may be that this is overestimated.

By expanding upon the methodology of O'Dwyer and Malone (2014), Vranken (2017) included more costs in the profitability function. By doing so, they estimated that the electricity consumption of Bitcoin mining was between 100-500 MW in June 2017. This number varies highly, since there is no way to exactly pinpoint the hardware used to mine. Bevand (2017) used the same approach, although he included more levels of calculation, which gave more profitability thresholds. The method of Bevand is the method that the CBECI[1] based their calculation on, which is another site where many papers base their estimates upon (University of Cambridge, 2022). He then took the weighted average of the equipment used, as well as some assumptions regarding the utilized hardware. By doing so, he estimated that the electricity consumption of mining was between 470 and

[1]Cambridge Bitcoin Electricity Consumption Index

540 MW in February 2017, and it increased to between 816 and 944 MW in July 2017, and further to 2100 MW in January 2018. The estimates vary as the weighted average of the most-used equipment varies. Krause and Tolaymat (2018) estimated that the average electricity consumption of Bitcoin mining in 2017 was 948 MW, increasing to an average of 3441 MW in the first half of 2018. They used a bottom-up approach for these estimates. De Vries (2018) used another method to calculate the estimated electricity consumption of Bitcoin mining. The methodology de Vries used was based of the model of marginal product of mining, introduced by Hayes (2017). In this model, the electricity consumption is measured from a more economical standpoint. Where it is assumed that there is an equilibrium for miners. That is, the marginal costs of mining are equal to the marginal product of mining. The marginal product of mining is calculated as the average Bitcoin price multiplied by the mining reward of one day (at the time of de Vries paper, this was equal to 8 351 US dollars times 1837 coins). The marginal costs in Hayes paper, are assumed to be strictly the electricity costs, as he argues that hardware costs and maintenance costs can be ignored. De Vries expands upon this, as he includes hardware costs in his analysis. He then calculates the estimated share of electricity costs from the total costs, by assuming the lowest price recorded of the hardware as well as a life expectancy of two years. He estimates the share to be between 60 and 70 percent. Lastly, he then estimates the electricity costs to be 5 cents per kWh, given an estimate of 7.67 GW, or 7670 MW. This way of estimating is known as the top-down approach, which has been criticized for the tendency to overestimate the electricity consumption.

Another paper by Gallersdörfer et al. (2019) estimated the power consumption, regional power consumption and the carbon emissions of Bitcoin mining. For power consumption, they calculated a lower and an upper bound in the same manner as Küfeoglu and Özkuran (2019). They estimated the power consumption at the end of 2016, the end of 2017, and in November 2018, which they based on ASIC[2] hardware sales. They used the estimates for November of 2018 to create an annual estimate for 2018, by multiplying the number of megawatts by 8760, which is the number of hours in a year. By doing so, they estimated the annual power draw of Bitcoin mining to be 45.8 terawatt hours. One issue with using this methodology is the possibility of inaccurate estimates as the mining difficulty adjusts every two weeks, as described in section 3.2. As a consequence, the extrapolating of

[2]Application-specific integrated circuit, used specifically for Bitcoin mining

the assumed consumption for one month may result in highly over- or underestimated results, depending on the activity in the month used. To increase the accuracy of the carbon emission estimates, they also localized miners by three methods. The first method included accessing mining pool servers IP, in which they found the distribution of the network computing power on a continent basis. Asia stood for 68 percent of the network hash rate, 17 percent contributed from Europe and 15 percent from North America. The second method they used was using the IoT search engine Shodan[3]. Here, they found a more granular distribution at the national level. In the last method they used, they found the IP-addresses from peer-to-peer nodes, as they communicated via a peer-to-peer network. This last method seemed to overestimate the concentration of US miners. They then used the estimates from the power consumption and the regional consumption to calculate the carbon emissions of Bitcoin mining. This was done by multiplying the average and marginal emission factors of power generation in the respective region with the estimated power draw. By doing so, they found that the global carbon footprint of Bitcoin mining could be estimated to be between 22 and 22.9 $MtCO_2$.

This paper was revisited by de Vries et al. (2022). Here, they used updated data regarding the location of miners as well as mining equipment. The authors then used the same approach as the previous paper by matching the updated locations with the carbon intensity of the electricity generation in the area. They estimated that the electricity demand for mining was 13.39 GW in August 2021. They did not extrapolate this estimate to an annual estimation, although following their methodology from the previous paper, this would have been equal to an annual electricity consumption of 117.3 terawatt hours in 2021. They did however estimate the carbon emissions from Bitcoin mining to be 65.4 $MtCO_2$.

Mora et al. (2018) is another paper that focuses on the emissions of Bitcoin. They also used the mining hardware to calculate the electricity consumption of mining. By using the efficency of the hardware and the hashes needed to mine one block, they found an estimate for the electricity needed to mine said block. Accordingly, they extrapolated this estimate to find an estimate for the total consumption, which they estimated to be 13 010 MW in 2017. In addition to this, they estimate the annual carbon emissions of Bitcoin to

[3]An IoT search engine is a search tool that allows identification devices connected to the internet (IoT devices)(Fagroud et al., 2020)

be 69 $MtCO_2$ in 2017. This is more than three times the estimated carbon emission in Gallersdörfer et al. (2019) paper, which indicates that the estimate of Mora et al. (2018) is highly overestimated. The methodology of the paper seems to be highly inaccurate as they extrapolate an already inaccurate estimate. The methodology seems sensitive to changes in both difficulty adjustments and the development of mining hardware. Another critical assumption they made, was that transactions draw power consumption, which is not backed by any other research.

Shan and Sun (2019) did a case study of Bitcoin mining and the CAISO[4]. Here, they argue that by relying on a high level of renewable resources to generate electricity, the grid makes curtailments for reliability reasons. If so, this would reduce both the economic and environmental benefits of the renewable resources. They argue that the curtailment could be mitigated by installing Bitcoin mining facilities at the power plants. By running simulations they estimate that the revenue of installing such a facility could increase the revenue of the power plants by approximately 5.6 – 48.1 million US dollars, while also decrease the curtailment by 50.8-79.9 percent, depending on the Bitcoin price and the mining difficulty. However, this analysis lacks several important aspects, such as cooling costs and additional facilities needed to mine Bitcoin. Bastian-Pinto et al. (2021) creates a case study with wind power and cryptocurrency farms in Brazil. This paper investigates the proposal of wind farm investors to invest in cryptocurrency mining facilities to hedge against electricity price risks. Since the electricity price and Bitcoin price is uncorrelated, the wind farms incentives to keep producing power increase, even though the electricity price is low. They argue that this could significantly increase the revenue of the electricity generator, while also reduce the risk of anticipating the construction. Niaz et al. (2022) provides another study regarding using Bitcoin mining as a source of turning excess energy into profits. They study ERCOT[5], and follow the argument of Shan and Sun (2019) that renewable energy leads to power curtailments, due to the lack of sufficient technology to store the energy supply. They estimate that 93 percent of the curtailed energy could be used to mine Bitcoin at the minimal cost, while generating a revenue of 239 million US dollars to the power plants.

[4]California Independent System Operator

[5]Electric Reliability Council of Texas

3 Bitcoin

3.1 Introduction

Bitcoin is a peer-to-peer electronic cash system, envisioned by Nakamoto (2008). A peer-to-peer system indicates that the participants exchange value directly with each other, without the need for a trusted third-party. For this system to work, it needs a secure system to validate the transactions to remove the problem of double spending, as well as to hinder the tampering of confirmed transactions. This is done via mining, which uses a proof-of-work system that involves computers to solve complex cryptographical puzzles. The confirmed transactions are then stored in a blockchain, which is a public distributed ledger. The ledger is accessible to all nodes in the network.

The blockchain consists of blocks that store information regarding previous transactions. It contains information regarding time and date, the bitcoin address of the seller and buyer, the total value of the transaction, and a unique signature that involves the current and previous blocks(Ashford and Powell, 2022). New blocks are added to the blockchain when miners are completing the puzzles, and new transactions are validated and announced on the network. When this is done, the race for the next block begins. This process is programmed to take around ten minutes for each block. This is also the way that new coins are added to the network, by miners getting a reward for completing the blocks.

3.2 Bitcoin Mining

As the blockchain relies on miners to complete the puzzle, incentives for them to do so are important. The Bitcoin reward they receive upon a puzzle completion functions as such an incentive. The reward for completing a block is 6.25 Bitcoins as of 2022. This reward is halved every 210 000 blocks, which translates to every four years. In addition to this, the miners also get aggregated transaction fees for the transactions in the block they complete. At every moment of time, there are multiple transactions laying in the mempool(Imtiaz et al., 2019). These are transactions waiting to be verified by miners. The higher the transaction fee is for the completed transaction, the shorter is the waiting time for that transaction to be validated, as argued by Easley et al. (2019). The security of the network increases with every new block mined. This is due to the increased costs and difficulty of tampering with previous transactions. Therefore, the mining process needs to consist of sufficient costs to maintain the security of the network.

Each block in the blockchain is minted on the previous block, which is done via a signature. This signature consists of a nonce (number used once) value that satisfies the hash function, SHA-256. This nonce value starts with 0 and increases until the miner finds the solution to the algorithm, which is the case when the hash of the block is less or equal to the target value. The target value changes depending on the difficulty, which adjusts every 2016 blocks, or around every two weeks. The block is then added to the blockchain and broadcasted to all the other nodes on the network, and the nodes will then start working on the next block. Upon completion, all of the transactions within that block are validated and forever stored on the blockchain (Küfeoğlu and Özkuran, 2019). This process is illustrated in figure 3.1.

Figure 3.1: Bitcoin Mining Process

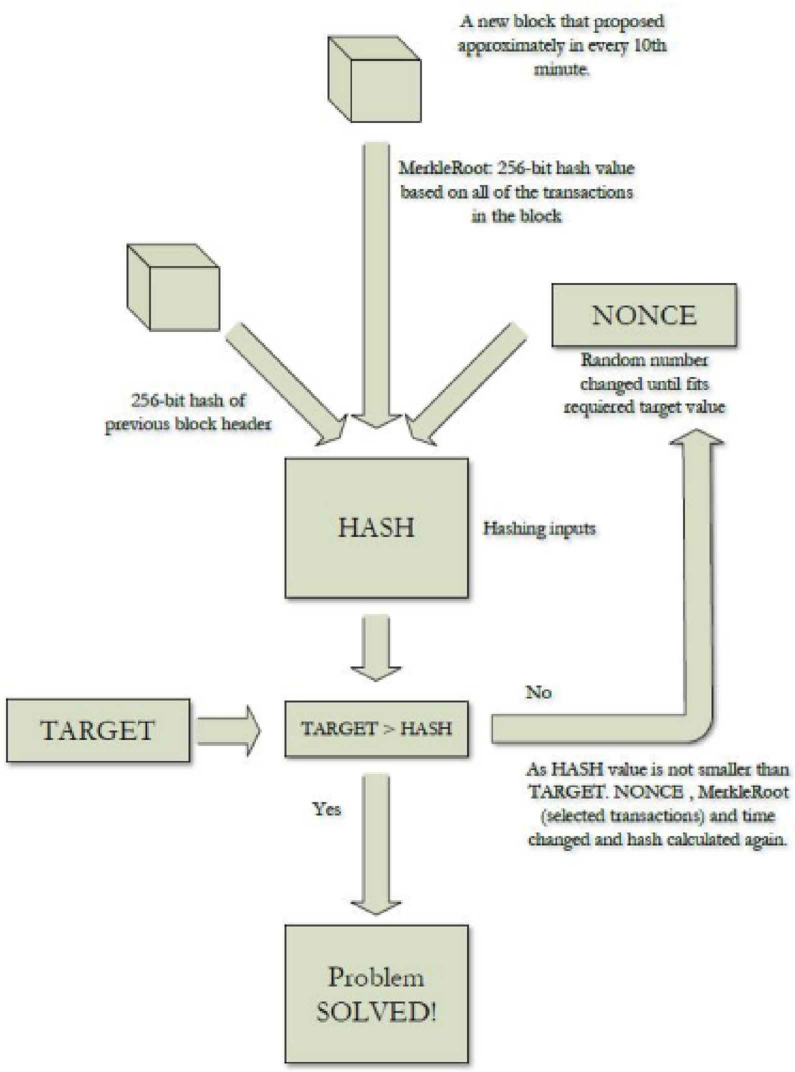

Source: (Küfeoglu and Özkuran, 2019)

Over time, more and more specialized hardware for Bitcoin mining has been developed. In the early stages of Bitcoin mining, miners used central processing units (CPU). The miners then switched to use graphic processing units (GPU), as they found out that this yielded more hash per second. In 2013 Canaan Creative developed application-specific integrated circuits (ASICs) (De Vries and Stoll, 2021). This hardware is specialized to complete one task, which is to solve the bitcoin algorithm. The development of this type of hardware has continued, and as of 2021, one could get several different kinds of this, each providing different types of efficiency in terms of hashing power and electricity consumption. The more efficient hardware, the more hashing power a miner can get for less use of electricity, which reduces costs for miners.

The increase in computing power needed to create a new block in the network has led to miners cooperating and combining their computing power while sharing the rewards. When the miners combine their hash rate in a mining pool, their chances of succeeding to mine new blocks drastically increases. In addition, this also increases the revenues and a more reliable income to the miners. The time it takes for the whole network to generate one block is ten minutes, however for a single mining unit this is not the case. The competition between the miners to complete the blocks is tough. For example, one unit of one of the more popular and efficient mining equipment, Antminer S19 Pro, has a maximum hash rate of 110 TH/s (terahashes per second) (Bitmain, 2022). In comparison, the biggest mining pool Foundry USA has, in June 2022, a hash rate of 49 681 PH/s (petahashes per second) (BTC.com, 2022). One petahash is equal to one thousand terahashes. This translates to the maximum hash rate of one unit of Antminer S19 Pro to be 0.0000024 percent of the Foundry USA mining pool. Therefore, most of the networks hash consists of different mining pools, as shown in figure 3.2.

Figure 3.2: Concentration of mining pools

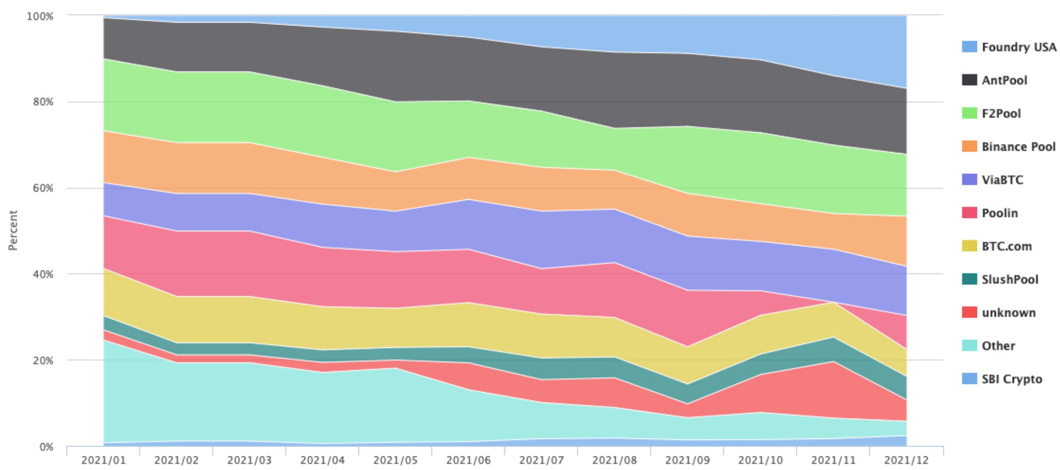

Source: https://btc.com/stats/pool?pool_mode=year

The amount of computing power needed to mine one block is determined by the difficulty and the efficiency of the hardware in use (Gallersdörfer et al., 2019). For instance, as China banned mining in June 2021, much of the networks hash rate went offline. This led to a decrease in the difficulty by 27.9 percent, which is the largest drop of difficulty ever recorded on the network (Blockchain.com, 2022). However, the difficulty normalized shortly after.

3.3 Bitcoin Mining in the US

Historically, most of the mining has taken place in China. Cambridge Center for Alternative Finance has recorded the historical hash rate for a selection of countries since September 2019. Their data by country includes Canada, Iran, Kazakhstan, China, Malaysia, Russia and United States. There are also recorded hash rates in Germany and Ireland, although they argue they likely are Chinese miners using VPNs to hide their locations. Other hash rates not recorded in these countries, are sorted in the category "other". In total, without the inclusion of other, this covers around 90-95 percent of the network hash rate. As seen in figure 3.3 and 3.4, China had the majority of the network hash rate until July, where it suddenly dropped to zero as the government banned mining in June of the same year. In revisiting bitcoin's carbon footprint, de Vries et al. (2022) argues that two events that increases the credibility of the CCAF data. China's ban on mining, along with an internet outage in Kazakhstan, gave some empirical insights to validate the CCAF data. Before the China ban, the data suggested that China represented 44 percent of the total mining activity. After the ban, the hash rate of the entire network decreased by 45 percent. Before the internet outage at the start of January 2022, Kazakhstan represented 18 percent of the total Bitcoin mining activity according to the CCAF data. Immediately after this event, the total hash rate of the network decreased by 15 percent. These two events suggest that the CCAF data is a good proxy for mining locations.

Figure 3.3: Hash Rate by Country in Percent

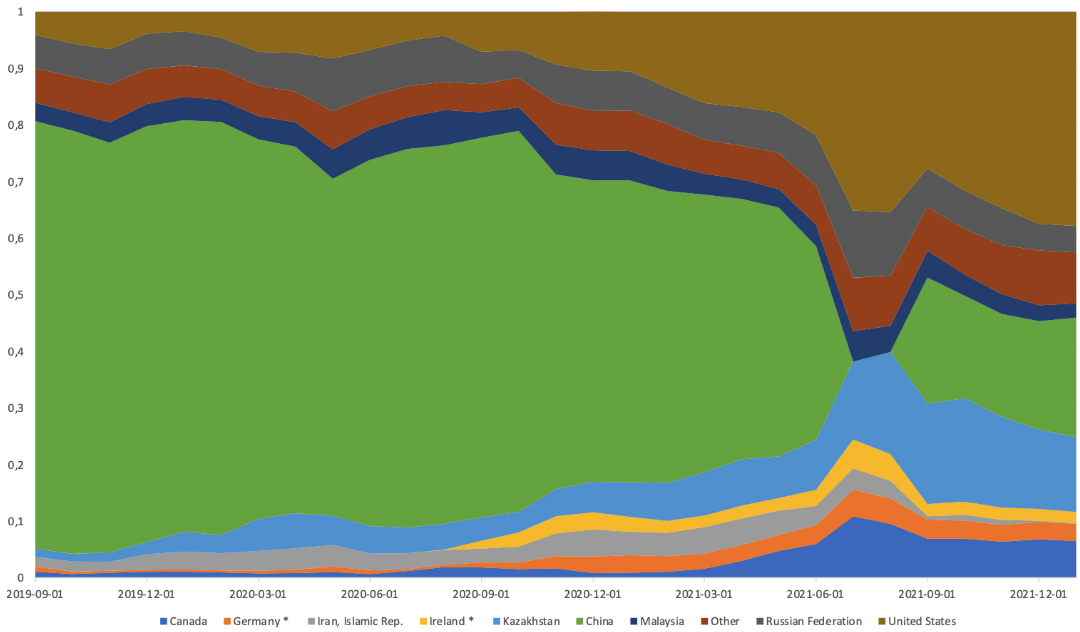

Figure 3.4: Hash Rate by Country in TH/s

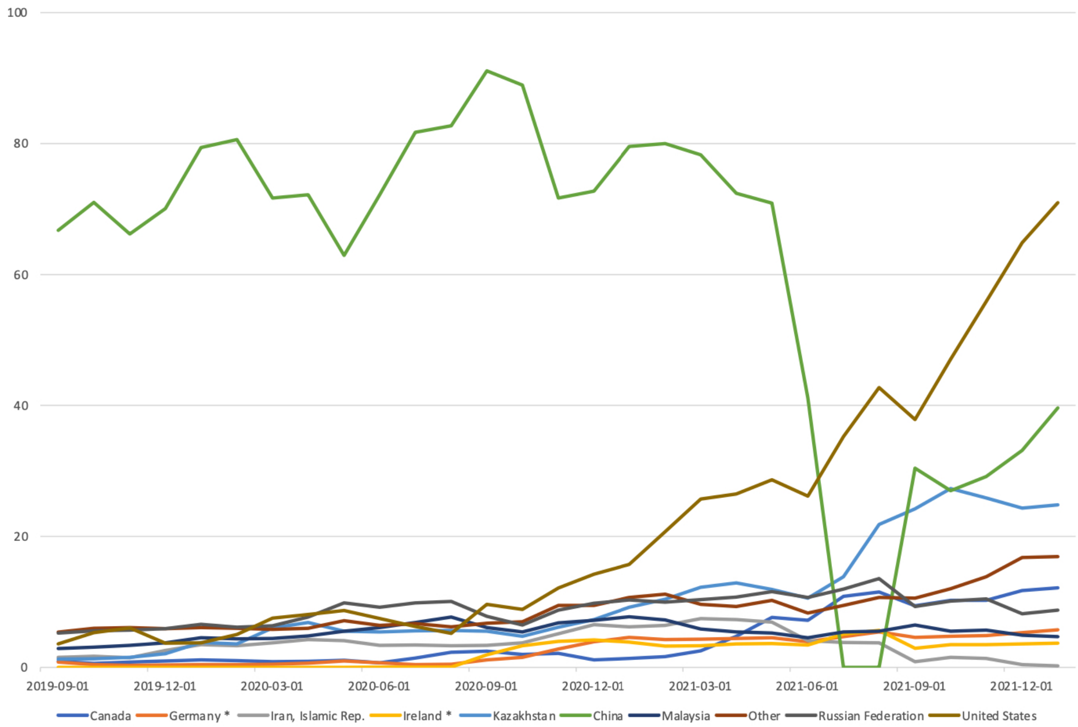

As seen in figures 3.3 and 3.4, the share of mining in the US started to rise towards the end of 2020 and continued to rise until the end of the data set, which stretched to January 2022. This increase provides a better foundation for the analysis in this thesis as it shows the impact of mining on the US electricity market.

4 US Electricity Market

This section provides a brief overview of the electricity market. This market is quite complex, and I will not go in depth on the operation of this.

4.1 Generation and transmission

The electricity market consists of the trade of electricity. Electricity is generated by using various fuels, such as fossil fuels[6], nuclear energy, or renewable energy[7]. These units of input produce electricity by utilizing different types of turbines. Steam turbines are the most common turbine in terms of electricity production, in which fossil fuels, nuclear energy, biomass, geothermal and solar thermal energy are used to produce steam that powers the turbine which again produces energy (U.S. Energy Information Administration, 2022d). The historical electricity mix is shown in figure 4.1, where natural gas was the largest source of electricity in 2021.

Figure 4.1: Generation mix of inputs

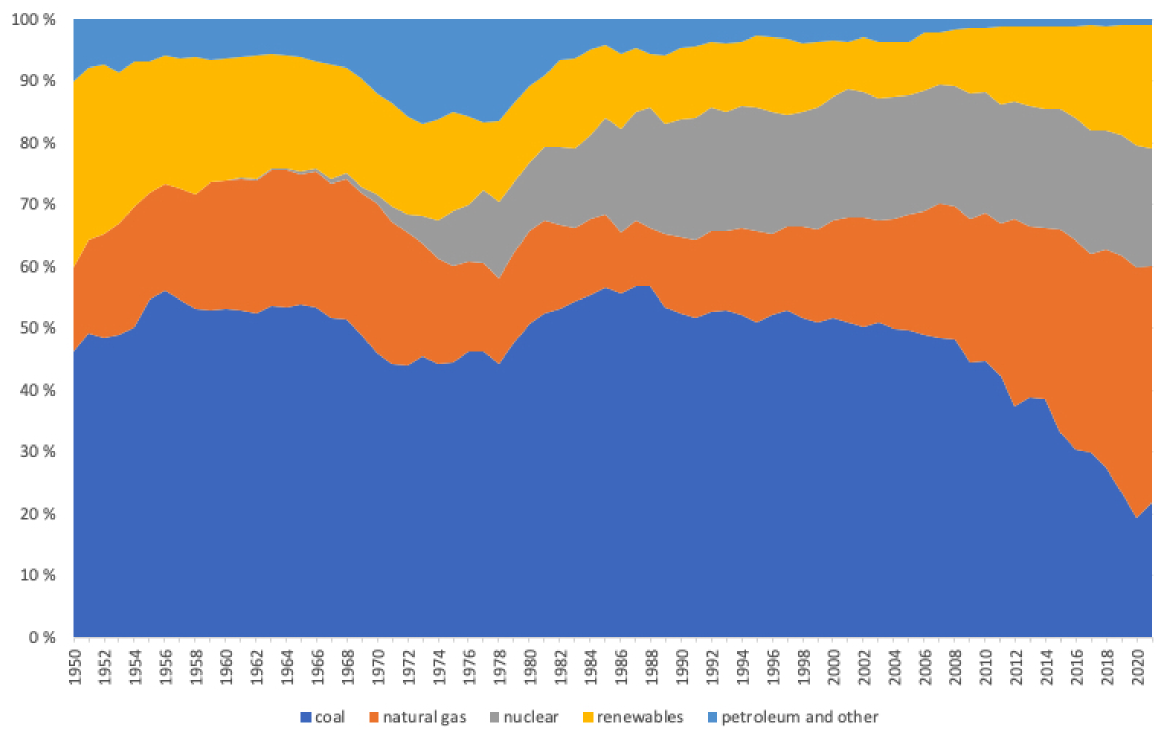

[6]Fossil fuels are in general coal, natural gas, and petroleum.
[7]Renewable energy are generally wind, solar and hydropower

When electricity is generated, it needs to be transferred to the customer, as electricity is generally not storable unless energy storage technology is used. Generally, electricity is not storable unless energy storage technology is used. Therefore, it needs to be transferred efficiently. For supply and demand to be met, electricity transmission is operated by grid operators. This is useful for power plants, as they do not need to produce at their full capacity at all hours of the day (U.S. Energy Information Administration, 2022c). These use different algorithms to calculate the most efficient route for the electricity to travel from the generators to the customers. This is done via high voltage transmission lines from the generators, and, depending on the type of customer, the high voltage line is switched into lower voltage transmission lines before the electricity are transferred into the load area[8] (Hogan, 2022; U.S. Energy Information Administration, 2022f).

4.2 Wholesale and retail sales markets

The U.S. electricity market consists of both a wholesale and a retail sale market. The wholesale market is a market where electric utility companies and electricity traders sell and buy electricity before it is sold to end-customers (United States Environmental Protection Agency, 2022b). The end-customers are the consumers of electricity, and they buy the electricity in the retail market. Over the last thirty years, many states have been deregulating the wholesale market, by introducing competitive markets. These markets are called independent system operators (ISO), and they are a marketplace where trade of electricity in the wholesale market occurs (Federal Energy Regulatory Commission, 2022). Here, utility companies are usually only responsible of selling and distributing the electricity to consumers, while independent power producers generate the electricity that is sold. States that have not incorporated this system, are using traditional markets. In these markets, there are usually vertically integrated utilities that are responsible for the generation, transmission, and distribution systems of the electricity (United States Environmental Protection Agency, 2022b).

The retail market is also, as the wholesale market, split into two types. The traditional retail market is a market where customers do not possess a choice in what utility provider they use. This is given by the location of the generators and customers. In these markets

[8] Load area are the end destination of the electricity.

there is a high concentration of market power, and the utility companies are usually vertical integrated, as they provide generation of electricity and transmission and distribution of the electricity (United States Environmental Protection Agency, 2022b).

The other type of market is a competitive market, where the power plants and utility companies are independent of each other. Here, the role of the utility companies is to transmit and distribute the electricity generated by independent power plants. Most of these markets include consumer choice, although not all (United States Environmental Protection Agency, 2022b).

4.3 Sectors of consumers

Then consumers of electricity are divided into four main sectors, the industrial, residential, commercial, and transportation sectors.

The industrial sector consists of the agriculture, forestry, fishing, hunting, mining, and construction industries (U.S. Energy Information Administration, 2022h). In this sector most of the electricity consumption is used for processing, producing, and assembling goods, while some electricity is consumed for other matters such as heating, cooling, lightning and so forth. What separates these electricity consumers from the other sectors, is that they can receive electricity at a higher voltage. Therefore, the distribution of electricity to these sectors is more efficient, and less electricity is lost in transmission (United States Environmental Protection Agency, 2022a). This makes electricity prices for the industrial sector lower than for the other sectors. In addition to this, electricity consumption from the industrial sector is usually more stable than the other sectors, as the demand is usually around the same at all hours of the day. This sector consumes less electricity than the commercial and residential sector, although they consume the most energy of all the sectors.

The residential sector is the sector which covers living quarters for households (U.S. Energy Information Administration, 2022i). These customers usually consume electricity for purposes such as heating, cooling, water heating, lightning, and electronic appliances (United States Environmental Protection Agency, 2022a). This is the largest sector in terms of consumption from the retail market, as well as number of customers (U.S. Energy Information Administration, 2022k). The consumption in this sector is highly seasonal, as shown in section 6.1. Compared to the industrial sector, this sector receives electricity at a lower voltage. This decreases the efficiency of transporting the electricity, as well as it increases the losses of electricity due to transmission. This leads to a price mark-up, which increases the prices for these customers.

The commercial sector consists of service-providing facilities and equipment for business (U.S. Energy Information Administration, 2022b). This covers hotels, restaurants, offices, governments, hospitals, and so forth. Here, the consumption is somewhat similar to the residential sector, as most of the electricity consumption stems from heating, lightning, and electronic appliances (United States Environmental Protection Agency, 2022a). The electricity price in this sector is usually somewhere between the price of the industrial and residential sector.

The transportation sector covers vehicles (U.S. Energy Information Administration, 2022m). This sector consumes a lot of energy, by directly burning fossil fuels to power the vehicles. In terms of electricity consumption, this sector accounts for less than 1 percent of the total electricity consumption (United States Environmental Protection Agency, 2022a). Therefore, I have excluded this sector in my analysis.

5 Data

5.1 Raw Data

This section describes the raw data as well as the variables used in my analysis.

5.1.1 Retail Sales of Electricity

Finding the correct measurements and data for electricity consumption was a difficult task. After a while, I found a panel data set regarding retail sales of electricity from the data browser of U.S. Energy Information Administration (2022j). Retail sales data is a measure for sold electricity, which makes it a proxy for consumed electricity. This data is separated into sectors and states. The sectors included in this data set are the sectors described in section 4.3, as well as a sector named other[9]. Following the argument from section 4.3, I have excluded the transportation sector as well as the other sections from my data set, as those sectors consumption is miniscule.

I downloaded data for 2020-2021 from 39 states. I excluded Alaska, Hawaii, and District of Columbia for reasons that they are not comparable to the rest of the U.S. I also excluded states that computed for more than 0.5 percent of the total Bitcoin hash rate in December 2021 on Cambridge Center for Alternative Finance's Bitcoin Electricity Consumption Index and did not have any reported mining in my Bitcoin mining data set (University of Cambridge, 2022). I was then left with data for the aggregated sectors as well as three sectors, for two years and 39 states.

5.1.2 Electricity Prices

Data regarding electricity prices was found in the same data browser as the electricity consumption data (U.S. Energy Information Administration, 2022a). Hence, the structure of the data was similar. It included the same sectors, as well as for all the sectors aggregated. For prices, as well as for electricity consumption, I only study the industrial, commercial, and residential sectors. The time interval and the states are the same as for

[9]the other covers sales for public streets, highways and other sales to public authorities, railroads, railways, and interdepartmental sales (U.S. Energy Information Administration, 2022g).

the electricity consumption.

5.1.3 Bitcoin Mining Data

The Bitcoin mining data was provided to me exclusively by the digital assets provider CoinShares, which they used to estimate the updated carbon impact of Bitcoin mining after the China ban (CoinShares, 2022). The data set consisted of estimates of regional power draw for mining in megawatt, the regional hash rate in terms of TH/s, and the regional hash rate percent. All these measures were given in monthly averages, hence why my analysis is done in this measure. My focus in the analysis is on the power draw.

The structure of the dataset is quite complicated, although they explain the methodology in their paper (CoinShares, 2022). They base their power draw estimations on network efficiency, which is calculated by the sum of all functional ASIC hardware contributes to the total hash rate of the network. The efficiency of the hash rate is calculated based on the efficiency of the hardware utilized. On the contrary to most former estimation practices, as described in section 2, they calculate the amount of ASIC hardware in use rather than scaling up from one unit. This makes the estimations more precise. From the hardware efficiency they calculate the average efficiency factor of the network. This measure is then used to estimate the power draw, by calculating the average number of watts drawn by the entire network per TH/s of hash rate generated. They gather information regarding hash rates from the Bitcoin blockchain itself and scales them up to the terms I use in my analysis.

By using hash rates and hash rate percent as a base, enables them to calculate the regional power draw. They use locational data that is verifiable from public data or private data provided to them by miners themselves, while adding publicly accessible data from CCAF and Foundry USA mining pool. Having access to data in terms of the regional hashing percent, allows them to calculate the power draw for each region by using the network efficiency model. This gives the most up-to date Bitcoin mining data in terms of regional power draw.

5.1.4 Control Variables

A control variable is a regressor included in the analysis to hold constant factors that could lead to bias in the regressor of interest, stemming from omitted variables (Stock et al., 2003).

5.1.4.1 Population

The population of the United States is not counted each year. Therefore, the United States Census Bureau publish population estimates each year (United States Census Bureau, 2021) For my time interval, the Bureau used April 2020 as their foundation for the estimates for July in 2020 and 2021. Therefore, I will be using the estimates for 2020 throughout the whole of 2020 and the same approach for 2021 in my analysis. The reasoning for population as a control variable is that population correlates positively with electricity demand. Therefore, population may have some impact on electricity consumption and prices. Lin and Zhu (2020) argue that most of the electricity consumption in China can be explained by population and per capita GDP, and there is reason to believe that this is also the case in the US.

5.1.4.2 Temperature

The weather affects electricity consumption and prices in different ways. The weather directly affects the ability to generate electricity through renewable energy like solar panels, hydropower, or wind farms. The electricity consumption is also directly affected by the weather through the need for heating or cooling. In the US, on the contrary to Norway, higher temperatures increase the electricity demand, and thereby also increase the prices. So, the seasonal changes are the opposite of Norway. To control for the weather effects of electricity, I therefor include average temperatures. This data was collected from NOAA National Centers for Environmental information (2022), and it is reported in the average monthly temperature in Fahrenheit. As some states are quite large areawise, this measure might not be entirely accurate, but it is as close as I can get to control for weather effects.

5.1.4.3 Per Capita Income

The third control variable included in the analysis is per capita income. This is used to control for the economic state of the different states, as this is shown to impact the demand for electricity. The data was downloaded from the U.S. Bureau of Economic Analysis (2022), who publish quarterly estimates for the per capita income by state. As this was downloaded with quarterly data, I used the data for the quarter within the months included in that respective quarter, making four estimates within each year.

5.2 Data Processing

All the data processing is done in either Rstudio or in Excel.

Structuring a complete data set for my analysis was a time-consuming task, since a lot of converting was necessary. Firstly, data regarding electricity consumption was downloaded in million kilowatt hours. I converted this to match my data on Bitcoin mining, which was in megawatt hours. To do this, I multiplied the data by one thousand. The reasoning for this is that I first had to convert it from a million kilowatt hours into kilowatt-hours, hence multiplying by one million. I then converted the variable from kilowatt-hours into megawatt-hours, which I did by dividing the numbers by one thousand.

As the Bitcoin mining data was given in monthly averages, I had to convert the data for consumption into monthly average. This was done by using Rstudio and computing a list of the total number of days for each month. I then computed the monthly average by dividing the total number of electricity consumption and the sectoral consumption by the number of days in each respective month.

Secondly, units of electricity prices were downloaded in cents per kilowatt-hour. I converted this to dollars per megawatt-hour, to better match the overall data set. This was done by multiplying by ten. The reasoning for this, is that I first had to convert the data into dollars per kilowatt hours, which is done by dividing the number by one hundred. I then had to convert it from kilowatt hours into megawatt hours, which is done by multiplying by one thousand. This data was already given in monthly averages, so I did not need to process the data any further.

Thirdly, the Bitcoin data was only filtered out for the states in the US. The unit was already the template I was tweaking the other data to, so I did not need to convert this any further. For average temperatures, the data were already given in monthly averages. However, I had to download a file for each month for every state. This was then assembled together. In addition, it was not necessary to convert the population data. Here, I used the estimates for each year and held them constant. Lastly, the data for per capita income was added to their respective months within the quarter.

5.3 Finalized Data

After the processing was done, I was left with a panel dataset which consisted of 39 states in the time frame of 2020-2021. It consisted of data on electricity consumption for the average megawatt hours for the total of all sectors, and the industrial, commercial, and residential sectors. The case was the same for electricity prices, which consisted of the average price for the total of all sectors, and for industrial, commercial, and residential sectors. It also consisted of the amount of power draw used in Bitcoin mining, given in monthly average megawatts, and control variables, average temperature, population, and per capita income.

6 Summary Statistics

This section provides insight to the raw data used in the analysis.

Table 6.1: Summary statistics for the main analysis sample

Statistic	N	Mean	St. Dev.	Min	Median	Max
Average Electricity Consumption	936	182,445.1	193,883.0	12,774.2	143,886.6	1,422,097.0
Average Industrial Electricity Consumption	936	50,790.8	57,007.4	1,533.3	38,406.3	378,733.3
Average Residential Electricity Consumption	936	70,702.6	76,003.9	4,871.0	53,838.7	608,612.9
Average Commercial Electricity Consumption	936	60,619.5	68,151.4	4,200.0	44,903.2	474,233.3
Average Eletricity Price	936	109.3	31.6	69.1	98.1	205.4
Average Industrial Eletricity Price	936	75.5	27.3	44.2	66.0	169.4
Average Residential Eletricity Price	936	135.5	35.4	90.7	123.9	251.1
Average Commercial Eletricity Price	936	108.7	25.9	69.5	102.8	184.9
Power Draw	342	148.3	196.6	6.3	68.8	961.1
Average Temperature	936	53.7	16.8	6.6	54.5	85.2
Population	936	5,488,056.0	5,487,826.0	577,267	4,243,850.0	29,527,941
Per Capita Income	936	58,028.7	9,001.3	40,120	56,256	83,982

In table 6.1 the summary statistics of the data is presented. It shows the number of observations, mean, standard deviation, minimum, median, and maximum of the variables used in the analysis. The first variables are the dependent variables, then the regressor, and the last three variables are the control variables. Power draw only has 342 observations in the data set. The reason for this is that there is recorded mining in fifteen of the thirty-nine states included in the analysis. Wyoming has no recorded mining in January and August 2020, and Nevada has no recorded mining from January 2020 until Mars 2021, as well as May 2021. This accounts for the missing observations in the data set. For the 594 missing observations, I have used a logarithmic technique consisting of adding a minuscule number (in my instance, 0.001). This is a commonly used practice to have the observations observed on the logarithmic scale (Burbidge et al., 1988).

The mean value of electricity consumption in table 8.1 shows the mean of all the monthly average electricity consumption in megawatt-hours across all the states included in the sample. The number of 182 445.1 shows that the mean is quite high, which indicates that the general electricity consumption in the states is high. The mean is also higher than the median, which may be why the mean is "pushed" up by states with much higher electricity consumption than others, as the maximum observation is 1 442 097. For the high maximum value and the low minimum value, seasonal variation also needs to be considered. The standard deviation of this variable is also high, which shows that there is a high spread in the sample. To account for this in the analysis, I have used logarithms. For the sectors included for electricity consumption, the sector with the highest mean is the residential sector, which shows that most of the electricity consumption is consumed by residents. The industrial sector consumes the least electricity in this data set. The reason for this, is that the industrial sector consumes most energy, rather than electricity. This indicates that they mostly use primary energy sources in production, while electricity is a secondary energy source (U.S. Energy Information Administration, 2022n). As for the other measures, the tendency in sector-specific consumptions is also highly spread, which is accounted for by logarithms.

For the electricity price, the mean in table 8.1 shows the mean for the monthly average electricity price in US dollars per megawatt-hours of all the observations in the sample. The mean is 109.3 dollars, compared to the median value of 98.1, it shows that some states with high electricity prices may push the mean up. The standard deviation of 31.6 shows that the spread is lower than for consumption, but that it still needs to be accounted for. The difference between the observed min and max values is quite large, although this may be due to seasonal variation in prices as well as differences between states. For the sectors, the industrial sector has considerably lower prices than the other sectors, as the residential sector has considerably higher prices than the other sectors. As mentioned earlier, the main reason for the industrial sector having such low prices, is that they can receive electricity at a much higher voltage than the other sectors, which increases the efficiency of transmitting the electricity greatly.

For power draw, the mean of 148,3 and the median of 68.8 shows that there also are some states here that drive the mean up. The standard deviation also shows that there is a high spread in this variable. The case is the same for the control variables included in the analysis.

6.1 Seasonality of Electricity Consumption and Prices

Figure 6.1: Consumption Seasonality

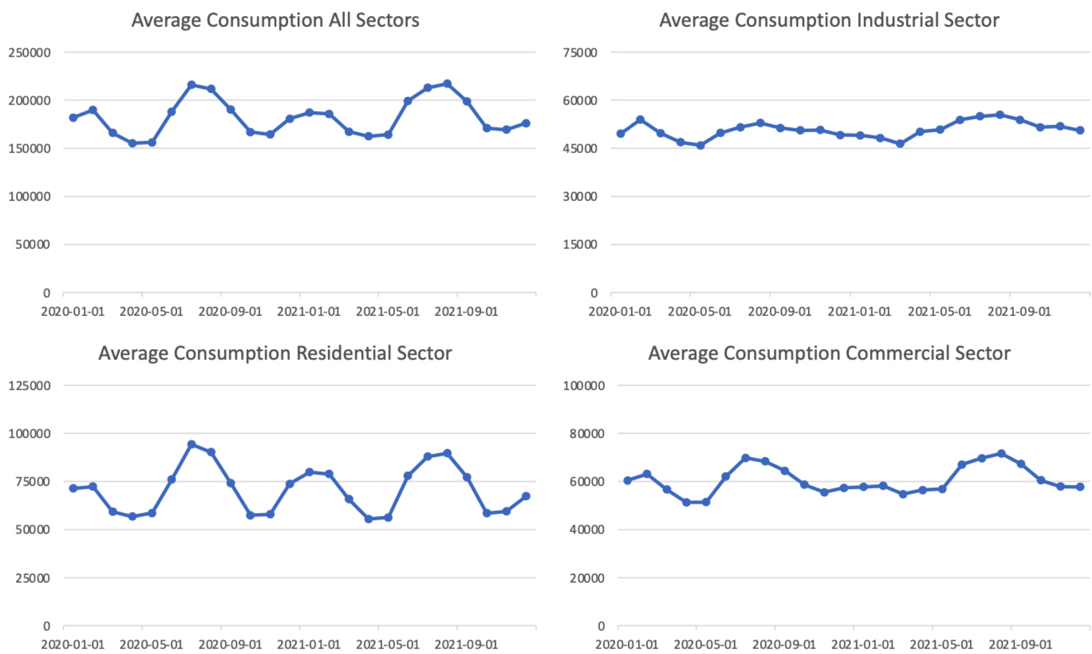

Figure 6.1 portrays the trends in electricity consumption. It shows that electricity is highly seasonal, where consumption in the all-sector is at its most in the summer months, and least in the late autumn with a small bump in the winter. For the industrial sector this trend is flatter, as the need for electricity is not dependent on the season and is therefore stable throughout the year. However, as the residential sector is highly seasonal, it may account for most of the overall seasonality in the all-sector consumption. As for the all-sector, the consumption in the residential is at the most in the summer and winter months. Most likely, the need for cooling and heating is the reason for these spikes. For the commercial sector, consumption is flatter than the consumption sector, but with more spikes than the industrial sector. This may be because there is less need for cooling and heating, which leads to more stable consumption throughout the year.

Figure 6.2: Prices Seasonality

The price and consumption trends are similar, which can be seen in figure 6.2. For the all-sector prices, the prices tend to be higher in the summer months compared to the rest of the year. However, prices are higher in 2021 than in 2020. Some of this is explained by the rise in fuel prices, especially natural gas, for which the costs more than double in 2021 compared to 2020 (U.S. Energy Information Administration, 2022l). Furthermore, the winter of 2021 was severely harsh in some areas of the US. In Texas, major winter storms led to significant energy disruptions, which contributed to an increase in the average price. During the extreme winter, wind turbines froze, and the flow of natural gas was restricted. This caused constraints on the electricity supply, which resulted in higher prices (U.S. Energy Information Administration, 2022l). The constraint is mainly reflected in the industrial and commercial sector. The residential sector has not been affected as much in general of these events. The reason may be the fact that most of the highly affected areas are Texas and the central US, where the concentration of industries is high relative to population.

7 Methodology

This section provides an overview of the method used in the analysis.

7.1 Models

The objective of this thesis is to estimate the effects Bitcoin mining has had on electricity consumption and prices. This is done using a common method for panel data analysis, the fixed effect model. This model consists of the following equations:

$$\ln(C_{s,i,t}) = \beta_0 + \beta_1 * ln(PD_{i,t}) + \delta_x * ln(x_{i,t})' + \lambda_i + \sigma_t + \epsilon_{i,t} \tag{7.1}$$

$$\ln(P_{s,i,t}) = \beta_0 + \beta_1 * ln(PD_{i,t}) + \delta_x * ln(x_{i,t})' + \lambda_i + \sigma_t + \epsilon_{i,t} \tag{7.2}$$

Equation (7.1) represents electricity consumption, while equation (7.2) represents electricity prices. $C_{s,i,t}$ measures the electricity consumption. Electricity consumption is measured in monthly average megawatt-hours and t is denoted with s, i and t, which are sector, state, and time. I investigate three different sectors, the industrial, commercial, and residential sectors, as well as the overall effect on all sectors of the market. The i represent each individual state, which includes 39 states in this analysis. The t represents time, which is every month during 2020-2021. $P_{s,i,t}$ represents electricity prices. These are also measured in monthly average US dollars per megawatt-hours. It includes the same denotation as electricity prices. The regressor is $PD_{i,t}$, which represents the power draw from Bitcoin mining. This is measured in monthly average megawatts. This regressor is used to estimate β_1 in the analysis for both consumption and prices. The consumption, price and power draw are all log-transformed to reduce the skewness in the variables.

There are also some control variables included in the analysis, represented by $x'_{i,t}$, which are a vector of the control variables. These are population, as a proxy for the size of the state, temperature as a proxy for the weather, and lastly, per capita income, as a proxy for the economic factors in the states. They are measured in respectively residents, Fahrenheit, and dollars.

The lambda (λ_i) is the state-fixed effects. This includes all time-invariant characteristics within each state that affects the consumption and prices. The sigma (σ_t) is a factor that catches events that is out of the ordinary, this only varies over time, and not over states. This covers events such as extreme weather that causes electricity outages. The last term included in the analysis is the error term epsilon ($\epsilon_{i,t}$). This captures variation in consumption and prices, which is not explained by the other variables in the analysis.

7.2 Fixed effects Model

There are a few key assumptions to this model that need to be met for the analysis to be valid (Stock et al., 2003). (1) $\epsilon_{i,t}$) does not correlate with the regressor:

$$\mathrm{cov}(PD_{i,t}, \epsilon_{i,t}) = 0 \tag{7.3}$$

This assumption implies that there are no omitted variables in the analysis. If the error term does correlate with power draw in any form, the assumption would be violated. (2) The entities in the sample are identically and independently distributed (i.i.d), which means that the sample used in the analysis are drawn randomly. If this does not hold, there could be selection bias in the analysis. (3) Large outliers in the analysis are unlikely. (4) There is no perfect multicollinearity. In other words, if the dependent variable and the regressor were exactly the same, it would be impossible to compute the estimates.

The electricity consumption model (equation 7.1) will be estimated first. The first step in the fixed effects model is to compute the demeaned average of the included variables.

$$\ln(\bar{C}_{s,i,t}) = \bar{\alpha}_i + \beta_1 * ln(P\bar{D}_{i,t}) + \delta_x * ln(\bar{x}_{i,t})' + \bar{\lambda}_i + \bar{\sigma}_t + \bar{\epsilon}_{i,t} \tag{7.4}$$

This is shown in equation (7.4), where the average over the whole sample period is calculated, which are represented by the bar. Here the term α_i are the unobserved heterogeneity (Wooldridge, 2015). Further, this is then deducted from each individual variable in the next step.

$$\begin{aligned}(ln(C_{s,i,t}) - ln(\bar{C}_{s,i,t})) = (\alpha_i - \bar{\alpha}_i) + \beta_1 * (ln(PD_{i,t}) - ln(P\bar{D}_{i,t})) \\ +\delta_x * (ln(x_{i,t})' - ln(\bar{x}_{i,t})') + (\lambda_i - \bar{\lambda}_i) + (\sigma_t - \bar{\sigma}_t) + (\epsilon_{i,t}\bar{\epsilon}_{i,t})\end{aligned} \tag{7.5}$$

Equation (7.5) is called the within transformation. This indicates that all the variation between states is filtered out, and the only variation we are left with are the variation within states over time (Stock et al., 2003). This means that the constant term and the state fixed effects zero out since they do not vary over time and are therefore eliminated from the equation (Verbeek, 2017). Here, the constant variation between states be included in the analysis, which reduces the omitted variables in the analysis. This can be simplified to:

$$ln(\tilde{C}_{s,i,t}) = \beta_1 * (ln(\tilde{PD}_{i,t})) + \delta_x * ln(\tilde{x}_{i,t})' + \tilde{\sigma}_t + \tilde{\epsilon}_{i,t} \tag{7.6}$$

Here β_1 is estimated by OLS regression on the demeaned variables $ln(\tilde{C}_{s,i,t})$ and $ln(\tilde{PD}_{i,t})$.

7.3 Fixed effects versus Random Effects

One alternative method I could have used, is the random effects model. If there was clear evidence that constant characteristics in the error term were uncorrelated with the regressor, then the random effects model would have been a better fit. Although, since we are dealing with a large geographical entity, we cannot treat our sample as a random sample from a large population (Wooldridge, 2015)). The random effects model includes both between and within variation in the data. This is done by including both variation in a stochastic error term and adds a new assumption that there are no constant characteristics, only random (Angrist and Pischke, 2009). This would then lead the random effect estimator to be inconclusive, and therefore the fixed effect is efficient.

7.4 Standard Errors

When dealing with panel data, clustered standard errors are common to use. Clustered standard errors are a form of heteroskedasticity- and autocorrelation-robust standard errors. This type of standard errors allows for autocorrelation within states, but not across (Stock et al., 2003). This makes clustered standard errors the preferred standard errors, as regular heteroskedasticity-robust standard errors only allow heteroskedasticity and not autocorrelation.

8 Results

This section provides the results from the regression analysis that has been done in order to estimate the effect Bitcoin mining has on electricity consumption and prices.

8.1 Effect on Consumption

Table 8.1 and 8.2 displays the regression run on the four types of consumption in this thesis, with power draw from Bitcoin mining as the regressor. Table 8.1 is run without control variables, while table 8.2 includes average temperature, population, and per capita income as control variables. All models included in this section are run with state and time fixed effects. Clustered robust standard errors on state have also been used in all the models to account for potential heteroskedasticity.

Model (1) in table 8.1 is run on the all-sector consumption. Here, the effect power draw has on electricity consumption is estimated to be 0.0109 percent. This is in line with my hypothesis that the relationship is positive, although the actual percent is lower than expected. Given the annual electricity consumption estimates presented in section 2, this effect seems to be underestimated. Possible reasons for this will be discussed in section 9. When including control variables, as presented in table 8.2, this estimate drops minimally to 0.0096. The small drop in the estimates indicates a high stability in the regression, and the estimates seem to be robust. The estimates in both tables are statistically significant at the one percent level, which indicates that the estimates are not a result of random variation in the data. The explanatory power in both models is sufficiently high, and when including control variables it slightly increases, which indicates that the control variables do have a effect on consumption.

Models (2), (3), and (4) in table 8.1 and 8.2 are run on sector-specific consumption, namely the industrial, residential, and commercial sectors. The estimated effect on the industrial sector is 0.0083 percent in the fully specified model. Also in this case, the estimates are robust, as the difference between the models with and without control variables is small. In table 8.1, the estimate lies within the one percent significant level, while the estimates in table 8.2 lies within the five percent significance level, which is acceptable. The estimate for the industrial sector is even smaller than for the all-sector consumption, which is the opposite of what I expected. As more mining has become more industrialized, as described in section 3.2, I would have expected this effect to be illustrated in the consumption. One explanation, may be that the retail sales of electricity are underestimated for the industrial sector. Since some of the mining facilities may be located at power plants, this would not be represented in the retail sales of electricity, as they would have consumed the electricity directly at the power plants.

Model (3) in table 8.1 represents the residential sector. The estimated effect on consumption in the residential sector is reported to be 0.0123 percent. Comparing this estimate to the fully specified model in table 8.2, this estimate drops to 0.0078 percent. In the fully specified model, this estimate is not statistically significant within the five percent significance level, which makes it hard to predict an accurate effect. As for the residential sector, this sector has the highest sector-specific consumption of all the sectors, which may indicate that other omitted factors influence this consumption more. These factors may not be picked up by the state and time fixed effects, such as the individuals preferred temperature at home, the average time of showers, or other factors which influence the residential consumption that are hard to measure.

Model (4) in table 8.1 reports the estimated effect on the commercial sector, which is estimated to be 0.0102 percent. This estimate increases slightly when including control variables in table 8.2, to 0.0107 percent. In addition to this, the explanatory power of the model is sufficiently high. Therefore, it is reasonable to assume high stability in this estimate. However, as the commercial sector covers mostly service-providing facilities, I would argue that this effect may be a type two error[10]. The reason for this, is that the impact of the power draw from Bitcoin mining should mainly affect the industrial and residential sector, as the most reported mining activity are registered in either specific firms dedicated to mining, or within residential areas through mining pools[11]. In addition to this, the fixed effects model is more prone to type two errors than regular OLS, since the coefficients only use within-state changes over time (Allison, 2009). I therefore suspect the mining activity in the commercial sector to be low and almost non-existent.

[10]A type two error is when the null hypothesis is not rejected, when it is in fact false (Becker and Greene, 2001)

[11]Both industrial mining and "single-mining" engage in mining pools.

Table 8.1: Electricity Consumption (1)

Dependent Variables:	logC	logC_i	logC_r	logC_c
Model:	(1)	(2)	(3)	(4)
Variables				
logPD	0.0109***	0.0072***	0.0123***	0.0102***
	(0.0011)	(0.0026)	(0.0019)	(0.0015)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.66214	0.23385	0.60705	0.66605

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 8.2: Electricity Consumption (2)

Dependent Variables:	logC	logC_i	logC_r	logC_c
Model:	(1)	(2)	(3)	(4)
Variables				
logPD	0.0096***	0.0083**	0.0078*	0.0107***
	(0.0011)	(0.0032)	(0.0040)	(0.0020)
logPop	0.9141***	1.357*	-0.5440	2.546***
	(0.3307)	(0.7638)	(0.5915)	(0.5511)
logAvg_Temp	-0.1298***	0.0357	-0.2763***	-0.1413***
	(0.0352)	(0.0354)	(0.0741)	(0.0322)
logPer_Cap_Inc	0.0066	-0.1777**	0.3266***	-0.2581***
	(0.0428)	(0.0701)	(0.0601)	(0.0475)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.68227	0.24540	0.64576	0.70501

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

8.2 Effect on Prices

Tables 8.3 and 8.4 show the estimated effect of power draw on electricity prices. It follows the same structure as the two previous tables, where model (1) uses the all-sector price as the dependent variable, and (2), (3) and (4) covers the industrial, residential and commercial sectors. All models are run with state and time fixed effects, and with cluster robust standard errors, clustered around states. Table 8.3 is run without control variables, while table 8.4 includes the same control variables as table 8.2.

The effect power draw has on electricity prices in all sectors is estimated to be 0.0101 percent when power draw increases with one unit in table 8.3. This follows my hypothesis that there is a positive relationship between increases in prices when the level of mining increases. When including control variables, the relationship stays positive, although the estimated effect is reduced by almost half. This shows that results may be sensitive to changes when including more specifications. However, the estimate is still statistically significant within the five percent level. This may indicate that there are more factors that influence electricity prices than consumption from Bitcoin mining. Factors like fuel costs are likely to affect electricity prices heavily. Ideally, this should have been included as a control variable for prices, but I have not been able to find a sufficient data set regarding this. This follows for all sectors.

For the industrial sector, the estimated effect of an increase of one unit of mining, would be 0.0145 percent on the electricity price. Without control variables, this effect is estimated to be 0.0266 percent. This shows that the estimates do not have the desired stability. However, the within r-squared value increases as the control variables are included, which indicates that the estimates could be closer to the true value. This is the highest estimated effect among all sectors for prices, which is the opposite of the estimated effects in electricity consumption. As shown in section 6.1, there is less seasonality in the industrial specific electricity price than for the all-sector price and the residential price. This may be shown in the estimates, as the estimates for the industrial sector are considerably higher than for the other sectors included in the analysis.

On the other side, the commercial sector follows a similar seasonality as the industrial sector, even though the estimated effects in this sector are considerably lower than the industrial sector. The estimated effect is 0.0069 percent in the fully specified model, although, as previously argued, I believe this estimate may be a victim of a type two error.

For the residential sector, the estimates do not satisfy the required five percent significance level. Therefore, the effect of mining on the residential price cannot be explained by my model, if there is any effect at all. There may be reason to believe that the increase in electricity prices for this sector are mostly described by the increase in fuel prices. This might be caused by the fact that most of the electricity in the residential sector is used for heating or cooling of houses, which is mostly done via natural gas, and not included in the retail sales of electricity.

Table 8.3: Electricity Prices (1)

Dependent Variables:	logP	logP_i	logP_r	logP_c
Model:	(1)	(2)	(3)	(4)
Variables				
logPD	0.0101***	0.0226***	0.0039	0.0108***
	(0.0033)	(0.0069)	(0.0025)	(0.0035)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.21161	0.22246	0.23701	0.11844

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

8.3 Standardized coefficients

Another way to estimate this effect is by using standardized coefficients. These coefficients are standardized, which indicates that they have been transformed, so that the mean of the variable is zero and the standard deviation is one (Verbeek, 2017). By doing so, one could estimate the effect on the dependent variable when the standard deviation of the regressor increases by one unit. This results in the coefficients being measured in the same unit, which makes it easier to interpret across different measurements. Here, this is done since power draw is measured in megawatts, while consumption is measured in megawatt-hours. I have calculated them to get the coefficients T by using equation (8.1) for each beta value and sector consumption and prices in R.

$$\hat{b} = \beta_1 * \frac{sd(ln(PD))}{sd(ln(C_s))} \tag{8.1}$$

Table 8.4: Electricity Prices (2)

Dependent Variables:	logP	logP_i	logP_r	logP_c
Model:	(1)	(2)	(3)	(4)
Variables				
logPD	0.0058**	0.0145***	0.0018	0.0069***
	(0.0022)	(0.0053)	(0.0018)	(0.0021)
logPop	-0.1049	1.427	-0.0334	-0.5273
	(0.6773)	(1.462)	(0.5812)	(0.8150)
logAvg_Temp	0.0139	0.0050	0.0471**	0.0700***
	(0.0234)	(0.0357)	(0.0198)	(0.0245)
logPer_Cap_Inc	0.3356***	0.5123***	0.1733***	0.3488***
	(0.0605)	(0.0967)	(0.0497)	(0.0656)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.27489	0.29527	0.26441	0.18935

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The standardized coefficients for electricity consumption are illustrated in table 8.5, where B is the unstandardized coefficient and β is the standardized. By comparing the coefficients, it seems that Bitcoin mining has the greatest impact on the commercial sector. For one increase in the standard deviation for power draw, the increase in the standard deviation for electricity consumption in the commercial sector increases by 0.0281. For the all-sector consumption, this increase results in a 0.0253 increase in standard deviation for the all-sectors electricity consumption. From these coefficients, the industrial sector is the sector which is the least affected sector by mining, slightly less than the residential sector.

The standardized coefficients for prices are illustrated in table 8.6. The overall notion here is that prices seem to be more affected by mining than consumption. The standard deviation of the all-sector electricity price increases with 0.0560 when the standard deviation of power draw increases by one unit. Compared to the other coefficients, this is less than for both the industrial sector and for the commercial sector. The industrial sector standard deviation increases by 0.1178, when the standard deviation of power draw increases by one unit. The commercial sector standard deviation increases by 0.0771 when the standard deviation for power draw increases by one unit. As the coefficient for the residential sector does not satisfy the significance requirement, this estimate cannot be interpreted.

Table 8.5: Standardized Coefficients for Electricity Consumption

Sectors	B	β
All sectors	0.0096	0.0253
Industrial sector	0.0083	0.0187
Residential sector	0.0078	0.0200
Commercial sector	0.0107	0.0281

Table 8.6: Standardized Coefficients for Electricity Prices

Sectors	B	β
All sectors	0.0058	0.0560
Industrial sector	0.0145	0.1178
Residential sector	0.0018*	0.0191*
Commercial sector	0.0069	0.0771

* represents regression coefficents that is not stastically significant.

9 Discussion

This section provides some of the limitations of the empirical methodology and the availability of data.

The strength of the fixed effects model is that the model controls for unobservable characteristics that do not change over time. In my model, this could include individual characteristics, such as individual preferences for cooling/heating, shower duration, lightning etc. All these characteristics are likely to not vary over time, at least in the time frame of this sample, and are therefore filtered out of the model. Although, there could arise problems regarding unobservable characteristics that do vary over time, as this is could be a source of unobserved heterogeneity (Hill et al., 2020). For instance, if these preferences did change within the duration of the sample, these variables would be omitted. This could be a source of unobserved heterogeneity, which again would lead to some bias in the estimates. For these variables to be filtered out, it is essential that they do not vary with time.

There can be several omitted variables in this analysis, with some previously mentioned in section 8. For electricity consumption in the residential sector, the presence of preferences that do change over time could be present. Noeurn (2021) found that in Cambodia, most of the electricity consumption in the residential sector stems from the high-income group. I have not included variables to control for this on the individual level, only for the state level by per capita income. There could also be other factors that are difficult to control for, such as home characteristics like insulation, drafts etc. This should be filtered out by the model, although if there are big changes in this during the sample period, this would be a time-varying characteristic which would be omitted. Merlin and Chen (2021) found that the level of unemployment had a significant negative effect on electricity consumption in DR Congo. This is also not controlled for in my analysis. For the industrial sector, there are other factors that could be omitted. Reitler et al. (1987) found that the leading effect to increased electricity demand in the industrial sector stems from the increase in energy-intensive industries. I have not found sufficient data regarding this to include this in my analysis, which makes it a possible omitted variable. However, one could assume that the time it takes to increase this sector, especially regarding infrastructure, would make this sector time-invariant and therefore filtered out by the model. Another source of changes in electricity consumption in the industrial sector could stem from changes in the demand and thereby supply of individual firms included in the sector. This is hard to measure, and therefore not included in the analysis, and could be a source of omitted variables.

For prices, possible sources for omitted variables include can be fuel prices, extreme weather, generation costs, regulation and transmission and distribution systems(U.S. Energy Information Administration, 2022e). Extreme weather, regulation and transmission and distribution systems would likely be filtered out of the model. Extreme weather would be filtered out by time dummies, while regulation and transmission and distribution systems would be filtered out by state fixed effects. However, problems could arise by fuel prices and generation costs, which would not be picked up by either the time dummies or the state fixed effects. This could lead to a severe bias in the estimates for prices in all sectors.

Another concern may be measurement errors. Mainly, this could stem from the data I have used regarding Bitcoin mining. As mentioned in section 5.1.3, this data is estimated by using regional hash rates and transforming it into megawatts by calculation. This calculation relies on heavy assumptions, which may cause some problems regarding the accuracy of this measure. In appendix A1, the analysis is done by using hash rates measured in TH/s [12] instead of power draw as the regressor. The results do not vary much, although these estimates are in general lower than those with power draw. Therefore, there could be arguments being made that the power draw estimations are overestimated. The assumptions made for the calculation of power draw, are mainly drawn around mining hardware and network efficiency. If network efficiency is not calculated correctly, this would directly lead to inaccurate estimates of power draw, as this is the main factor that transforms hash rates into megawatts. They base the network efficiency of the hardware in use at any given time. There is no way to calculate this precisely, so they assume that the hardware in use is a mixture of existing hardware based on efficiency, production, and breakdown rates of the hardware (CoinShares, 2022). There are most likely differences between miners in what sort of mining equipment they use, and therefore it could be regional differences in efficiency. This would again translate into differences in the power

[12]EH/s is terahashes per second. One terahash is equal to one trillion hashes per second.

draw of regions. This calculation could be biased in both directions, and most likely it affects both ways. Where some regions are being overestimated, and others are being underestimated. However, despite these factors, this data set is currently the most detailed available regarding this topic.

The data set regarding Bitcoin does not include the electricity used to cool down the computers. There are multiple reasons for miners to avoid the equipment to overheat. If they do so, this would lead to a decrease in hash rates, lower power efficiency and higher maintenance costs, and a lower lifespan of the hardware. Mainly there are two ways Bitcoin miners cool down the facilities, either air cooling[13] or immersion cooling[14]. The technology for immersion cooling is rapidly growing, which is a less electricity dependent technique than air cooling. Although, as of now, air cooling is wider used than immersion cooling. By not having any estimates of the usage of electricity towards this, the true power draw of Bitcoin mining could be severely underestimated.

Another concern regarding the data set is the time frame of the sample. Ideally, the data would be more granular, preferably at as little as five-minute intervals. Although, there is no current available for this sort of data. If this were the case, one could transform the mining data from megawatts into megawatt-hours, and directly see changes in demand and prices of electricity when there are changes in Bitcoin mining. This would also correct for the potential missing effect of power draw on the electricity market, as using monthly averages is not ideal. There could be events where the monthly mining are either shut down during the month, or heavily increased within the month. By using monthly averages, these effects are averaged out throughout the month, which could be hiding the actual effect of mining.

Ideally, there would be more specific data on mining location as well as time intervals, with location being as precise as ZIP-codes of mining facilities or information regarding which county within the state mining is present. This would also open up for other potential methodologies, that could estimate the effect more directly. For example, one method that could have been used is spatial regression discontinuities. If this were available, one could compare the changes in electricity specific factors in counties where mining is present,

[13]Air cooling is a cooling method that consists of cooling down the air around the mining rig.

[14]Immersion cooling is a method that consists of submerging the mining rig in thermally conductive liquid with greater insulation properties than air (Gerasymovych, 2022))

with neighboring counties where mining is not present. This would lead to more accurate estimations of the effect Bitcoin mining has on the US electricity market.

Another worry regarding the data used, would be that some mining has occurred outside of the US electricity market. This is mining where the miners have used electricity directly from the generators, rather than using the markets to buy electricity. Mining companies like Greenidge, Digihost and Marathon have all bought previously decommissioned power plants and restarted them to operate mining facilities (Corey, 2021; Milman, 2021). In addition to this, they also contribute to the electricity market as they sell their excess electricity. Making them net contributors to the electricity market. To control for such factors is not possible with the available data, and therefore this could also hinder the true estimates of the analysis.

10 Concluding Remarks

This thesis presents potential effects of Bitcoin mining on electricity consumption and prices in the US. By utilizing a fixed effects model to account for potential omitted variables, this analysis provided results according to the hypothesis that there is a positive relationship between mining and electricity consumption and prices. For consumption, the sector with the highest increase in percentage was the commercial sector. However, this is suspected to be a victim of a type two error, as this sector would most likely not have mining of significance. The residential sector is not significant within the five percent level, which is another potential error. Assumably, this sector should have significant mining, as there are most likely many miners mining from their residence. For prices, the most affected sector was the industrial sector. This follows the argument that large mining facilities receive electricity from the industrial sector at a high voltage. As this sector could be argued to be energy-intensive, it is reasonable to believe that this will impact the prices on electricity in this sector, as demand increases. The estimation for the residential sector is not significant within the five percent level (or within the ten percent). This could also stem from a type one error, as there is reason to believe aggregated mining from households takes up a significant portion of total mining. Although, this could be spread over the whole sample, which would not be picked up from the estimated power draw.

However, there are many concerns regarding this analysis, with the biggest being the data used in the analysis. The data regarding bitcoin mining stems from estimated power draw in megawatts, which is likely to be biased as there is no way of accurately measuring the exact power draw as of now. In addition, there are also likely unaccounted factors in mining that draw electricity, such as cooling. Another data concern is the sample duration. As there is no available data dating back longer than 2020, this makes a big concern. Small time-samples increase the possibility of biased estimators. There are also some concerns regarding potential omitted variables, such as fuel prices, unemployment rates, and other unobservable characteristics.

Further research on this topic is encouraged. Depending on the granularity of future data, there could be various ways to estimate the impact of miners on the US electricity market.

This would require more transparency from the largest mining pools, mainly about the locations of facilities and the source of electricity generators. As the share of total miners locating in the US continued to rise throughout 2021, the impact of mining in the future could be larger than what is estimated today.

References

Allison, P. D. (2009). *Fixed effects regression models*. SAGE publications.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Ashford, K. and Powell, F. (2022). What is bitcoin and how does it work? *Forbes*. https://www.forbes.com/advisor/investing/cryptocurrency/what-is-bitcoin/ Retrieved: 2022-08-10.

Bastian-Pinto, C. L., Araujo, F. V. d. S., Brandão, L. E., and Gomes, L. L. (2021). Hedging renewable energy investments with bitcoin mining. *Renewable and Sustainable Energy Reviews*, 138:110520.

Becker, W. E. and Greene, W. H. (2001). Teaching statistics and econometrics to undergraduates. *Journal of Economic Perspectives*, 15(4):169–182.

Bevand, M. (2017). Electricity consumption of bitcoin: a market-based and technical analysis. *Retrieved February 4, 2019, from mrb's blog: http://blog. zorinaq. com/bitcoin-electricity-consumption*.

Bitmain (2022). Bitmain antminer s19 pro (110th). https://www.asicminervalue.com/miners/bitmain/antminer-s19-pro-110th Retrieved: 2022-08-10.

Blockchain.com (2022). Network Difficulty [Statistics]. https://www.blockchain.com/charts/difficulty Retrieved: 2022-08-10.

BTC.com (2022). Foundry USA [Statistics]. https://btc.com/stats/pool/Foundry\%20USA Retrieved: 2022-08-10.

Burbidge, J. B., Magee, L., and Robb, A. L. (1988). Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association*, 83(401):123–127.

CoinShares (2022). The bitcoin mining network energy and carbon impact. *CoinShares Research*.

Corey, K. (2021). A Bitcoin Boom Fueled by Cheap Power, Empty Plants and Few Rules. *The New York Times*. https://www.nytimes.com/2021/12/05/nyregion/bitcoin-mining-upstate-new-york.html Retrieved: 2022-08-13.

De Vries, A. (2018). Bitcoin's growing energy problem. *Joule*, 2(5):801–805.

de Vries, A., Gallersdörfer, U., Klaaßen, L., and Stoll, C. (2022). Revisiting bitcoin's carbon footprint. *Joule*, 6(3):498–502.

De Vries, A. and Stoll, C. (2021). Bitcoin's growing e-waste problem. *Resources, Conservation and Recycling*, 175:105901.

Easley, D., O'Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109.

Fagroud, F. Z., Ajallouda, L., Toumi, H., Achtaich, K., El Filali, S., et al. (2020). Iot search engines: exploratory data analysis. *Procedia Computer Science*, 175:572–577.

Federal Energy Regulatory Commission (2022). Rtos and isos. https://www.ferc.gov/power-sales-and-markets/rtos-and-isos Retrieved: 2022-8-15.

Gallersdörfer, U., Klaaßen, L., and Stoll, C. (2019). The carbon footprint of bitcoin. *Joule*, 3(7):1647–1661.

Gerasymovych, S. (2022). AIR COOLING VS IMMERSION COOLING CONTAINER: WHAT'S BETTER FOR BITCOIN MINING? *ezblockchain*. https://ezblockchain.net/article/immersion-cooling-vs-air-cooling-technologies-in-crypto-mining/ Retrieved: 2022-08-12.

Hayes, A. S. (2017). Cryptocurrency value formation: An empirical study leading to a cost of production model for valuing bitcoin. *Telematics and informatics*, 34(7):1308–1321.

Hill, T. D., Davis, A. P., Roos, J. M., and French, M. T. (2020). Limitations of fixed-effects models for panel data. *Sociological Perspectives*, 63(3):357–369.

Hogan, W. W. (2022). Electricity market design and zero-marginal cost generation. *Current Sustainable/Renewable Energy Reports*, pages 1–12.

Imtiaz, M. A., Starobinski, D., Trachtenberg, A., and Younis, N. (2019). Churn in the bitcoin network: Characterization and impact. In *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 431–439. IEEE.

Krause, M. J. and Tolaymat, T. (2018). Quantification of energy and carbon costs for mining cryptocurrencies. *Nature Sustainability*, 1(11):711–718.

Küfeoğlu, S. and Özkuran, M. (2019). Bitcoin mining: A global review of energy and power demand. *Energy Research & Social Science*, 58:101273.

Küfeoglu, S. and Özkuran, M. (2019). Energy consumption of bitcoin mining.

Lin, B. and Zhu, J. (2020). Chinese electricity demand and electricity consumption efficiency: Do the structural changes matter? *Applied Energy*, 262:114505.

Merlin, M. L. and Chen, Y. (2021). Analysis of the factors affecting electricity consumption in dr congo using fully modified ordinary least square (fmols), dynamic ordinary least square (dols) and canonical cointegrating regression (ccr) estimation approach. *Energy*, 232:121025.

Milman, O. (2021). Bitcoin miners revived a dying coal plant – then CO2 emissions soared. *The Guardian*. https://www.theguardian.com/technology/2022/feb/18/bitcoin-miners-revive-fossil-fuel-plant-co2-emissions-soared Retrieved: 2022-08-13.

Mora, C., Rollins, R. L., Taladay, K., Kantar, M. B., Chock, M. K., Shimada, M., and Franklin, E. C. (2018). Bitcoin emissions alone could push global warming above 2 c. *Nature Climate Change*, 8(11):931–933.

Nakamoto, S. (2008). Bitcoin whitepaper. *URL: https://bitcoin. org/bitcoin. pdf-(: 17.07. 2019)*.

Niaz, H., Liu, J. J., and You, F. (2022). Can texas mitigate wind and solar curtailments by leveraging bitcoin mining? *Journal of Cleaner Production*, 364:132700.

NOAA National Centers for Environmental information (2022). Climate at a Glance:

Statewide Time Series, Average Temperature [Data set]. https://www.ncdc.noaa.gov/cag/ Retrieved: 2022-03-28.

Noeurn, V. (2021). Factors affecting electricity consumption of residential consumers in cambodia. In *IOP Conference Series: Earth and Environmental Science*, volume 746, page 012034. IOP Publishing.

O'Dwyer, K. J. and Malone, D. (2014). Bitcoin mining and its energy footprint.

Reitler, W., Rudolph, M., and Schaefer, H. (1987). Analysis of the factors influencing energy consumption in industry: a revised method. *Energy Economics*, 9(3):145–148.

Shan, R. and Sun, Y. (2019). Bitcoin mining to reduce the renewable curtailment: a case study of caiso. *Available at SSRN 3436872*.

Stock, J. H., Watson, M. W., et al. (2003). *Introduction to econometrics*, volume 104. Addison Wesley Boston.

United States Census Bureau (2021). Annual Estimates of the Resident Population for the United States, Regions, States, District of Columbia and Puerto Rico: April 1, 2020 to July 1, 2021 [Data set]. https://www.census.gov/data/tables/time-series/demo/popest/2020s-state-total.html#par_textimage_1574439295 Retrieved: 2022-03-28.

United States Environmental Protection Agency (2022a). Electricity Customers. https://www.epa.gov/energy/electricity-customersRetrieved: 2022-8-15.

United States Environmental Protection Agency (2022b). U.S. Electricity Grid Markets. https://www.epa.gov/green-power-markets/us-electricity-grid-markets Retrieved: 2022-8-15.

University of Cambridge (2022). Bitcoin mining map. *https://ccaf.io/cbeci/mining$_m$ap*.

University of Cambridge (2022). Cambridge bitcoin electricity index. *https://www.cbeci.org/*.

U.S. Bureau of Economic Analysis (2022). GDP and Personal Income: SQINC1 Personal Income Summary: Personal Income, Population, Per Capita Personal Income [Data set]. https://apps.bea.gov/itable/iTable.cfm?ReqID=70&step=1 Retrieved: 2022-03-28.

U.S. Energy Information Administration (2022a). Average retail price of electricity [Data set]. https://www.eia.gov/electricity/data/browser/#/topic/5?agg=2,0,1&geo=g&freq=M&start=200101&end=202205&ctype=linechart&ltype=pin&rtype=s&maptype=0&rse=0&pin= Retrieved: 2022-03-25.

U.S. Energy Information Administration (2022b). Commercial sector. *In U.S. Energy Information Administration*. https://www.eia.gov/tools/glossary/index.php?id=C.

U.S. Energy Information Administration (2022c). Electricity Explained: Electricity generation, capacity, and sales in the United States. https://www.eia.gov/energyexplained/electricity/prices-and-factors-affecting-prices.php Retrieved: 2022-5-25.

U.S. Energy Information Administration (2022d). Electricity Explained: Electricity in the United States. https://www.eia.gov/energyexplained/electricity/electricity-in-the-us.php Retrieved: 2022-5-25.

U.S. Energy Information Administration (2022e). Electricity explained: Factors affecting electricity prices. https://www.eia.gov/energyexplained/electricity/prices-and-factors-affecting-prices.php Retrieved: 2022-08-12.

U.S. Energy Information Administration (2022f). Electricity explained: How electricity is delivered to consumers. https://www.eia.gov/energyexplained/electricity/delivery-to-consumers.php Retrieved: 2022-05-25.

U.S. Energy Information Administration (2022g). Electricity sales. *In U.S. Energy Information Administration*. https://www.eia.gov/tools/glossary/index.php?id=E.

U.S. Energy Information Administration (2022h). Industrial sector. *In U.S. Energy Information Administration*. https://www.eia.gov/tools/glossary/index.php?id=I.

U.S. Energy Information Administration (2022i). Residential sector. *In U.S. Energy Information Administration*. https://www.eia.gov/tools/glossary/index.php?id=R.

U.S. Energy Information Administration (2022j). Retail sales of electricity [Data set]. https://www.eia.gov/electricity/data/browser/#/topic/5?agg=2,0,1&geo=g&freq=M&start=200101&end=202205&ctype=linechart&ltype=pin&rtype=s&maptype=0&rse=0&pin= Retrieved: 2022-03-25.

U.S. Energy Information Administration (2022k). Table 2.1. Number of Ultimate Customers Served by Sector, by Provider, [Statistics]. https://www.eia.gov/electricity/annual/html/epa_02_01.html Retrieved: 2022-08-20.

U.S. Energy Information Administration (2022l). Today in Energy: During 2021, U.S. retail electricity prices rose at fastest rate since 2008. https://www.eia.gov/todayinenergy/detail.php?id=51438 Retrieved: 2022-08-23.

U.S. Energy Information Administration (2022m). Transportation sector. *In U.S. Energy Information Administration*. https://www.eia.gov/tools/glossary/index.php?id=T.

U.S. Energy Information Administration (2022n). Use of energy explained: Energy use in industry. https://www.eia.gov/energyexplained/use-of-energy/industry.php Retrieved: 2022-08-25.

Verbeek, M. (2017). *A guide to modern econometrics*. John Wiley & Sons.

Vranken, H. (2017). Sustainability of bitcoin and blockchains. *Current opinion in environmental sustainability*, 28:1–9.

Wooldridge, J. M. (2015). *Introductory econometrics: A modern approach*. Cengage learning.

Appendix

A1 Analysis With Hash Rate

In this appendix section, the same analysis and model presented in section 7 is used, only with hash rate measured in TH/s (terahashes per second) as the regressor instead of power draw. This is done in order to show that the results does not vary much, as the calculation behind power draw might be imprecise, as described in section 9.

The main takeaway from these results, is that the estimated effect is in general lower than for power draw. The way power draw is calculated could be the reason for this difference. However, the prediction power of the estimates do not change over the course of these tables, the within r-squared are around the same as for the models in section 8.

Table A1.1: Electricity Consumption with Hash Rate (1)

Dependent Variables:	logC	logC_i	logC_r	logC_c
Model:	(1)	(2)	(3)	(4)
Variables				
logHR	0.0053***	0.0029***	0.0065***	0.0049***
	(0.0006)	(0.0009)	(0.0008)	(0.0008)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.66124	0.23175	0.60713	0.66500

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A1.2: Electricity Consumption with Hash Rate (2)

Dependent Variables:	logC	logC_i	logC_r	logC_c
Model:	(1)	(2)	(3)	(4)
Variables				
logHR	0.0046***	0.0031***	0.0048***	0.0048***
	(0.0006)	(0.0011)	(0.0016)	(0.0009)
logPop	0.9407***	1.421*	-0.5669	2.592***
	(0.3286)	(0.7803)	(0.5836)	(0.5779)
logAvg_Temp	-0.1292***	0.0363	-0.2759***	-0.1406***
	(0.0353)	(0.0354)	(0.0742)	(0.0324)
logPer_Cap_Inc	0.0226	-0.1589**	0.3344***	-0.2382***
	(0.0409)	(0.0689)	(0.0556)	(0.0472)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.68172	0.24245	0.64624	0.70317

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A1.3: Electricity Prices with Hash Rate (1)

Dependent Variables:	logP	logP_i	logP_r	logP_c
Model:	(1)	(2)	(3)	(4)
Variables				
logHR	0.0042***	0.0097***	0.0012	0.0045***
	(0.0012)	(0.0028)	(0.0009)	(0.0013)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.20215	0.20660	0.23416	0.10872

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A1.4: Electricity Prices with Hash Rate (2)

Dependent Variables:	logP	logP_i	logP_r	logP_c
Model:	(1)	(2)	(3)	(4)
Variables				
logHR	0.0025***	0.0064**	0.0004	0.0030***
	(0.0009)	(0.0024)	(0.0006)	(0.0008)
logPop	-0.0746	1.496	-0.0059	-0.4924
	(0.6950)	(1.524)	(0.5847)	(0.8345)
logAvg_Temp	0.0143	0.0059	0.0472**	0.0704***
	(0.0235)	(0.0358)	(0.0198)	(0.0246)
logPer_Cap_Inc	0.3470***	0.5401***	0.1790***	0.3623***
	(0.0613)	(0.0989)	(0.0490)	(0.0674)
Fixed-effects				
Time	Yes	Yes	Yes	Yes
State	Yes	Yes	Yes	Yes
Fit statistics				
Observations	936	936	936	936
Within R^2	0.27261	0.29044	0.26370	0.18653

Clustered (State) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*