An Econometric Analysis of Large Flexible Cryptocurrency-mining Consumers in Electricity Markets

Subir Majumder
Department of Electrical and Computer Engineering,
Texas A&M University,
College Station, TX 77840
subir-em@ieee.org

Ignacio Aravena
Lawrence Livermore National Laboratory,
Livermore, CA
aravenasolis1@llnl.gov

Le Xie
Department of Electrical and Computer Engineering,
Texas A&M Energy Institute,
Texas A&M University,
College Station, TX 77840
le.xie@tamu.edu

Abstract

In recent years, power grids have seen a surge in large data centers and digital currency mining facilities, with individual consumption levels reaching 700MW. This study examines the behavior of these entities in Texas, focusing on how their consumption is influenced by currency conversion, electricity prices, local temperatures, system demand, and other factors. We transform the skewed consumption data, perform correlation analysis, and apply a seasonal autoregressive moving average model for analysis. Our findings reveal that, surprisingly, short-term mining consumption is not correlated with digital currency prices. Instead, the primary influencers are the temperature, energy prices, and demand response strategies deployed by the Electric Reliability Council of Texas (ERCOT). The inferred model can be used to generate public, synthetic datasets to understand the overall impact on power grid operation, while explanations of large-scale miner behaviors can lead to better pricing mechanisms to effectively use the flexibility of these resources towards improving power grid reliability.

Keywords: Demand Response, Econometric Model, Large Flexible Cryptomining Loads, Electricity Markets.

1. Introduction

The Texas energy grid is facing a data-center-driven load growth and management challenge. The Electricity Reliability Commission of Texas (ERCOT)—the market operator in charge of the largest part of the Texas electricity grid—allows both generators and loads to truthfully disclose their price sensitivity and be dispatched through their economic dispatch process. However, we observe that certain data center loads, with an individual capacity greater than or equal to 75.0 MW [1], show price inflexibility in their offer curve for certain settlement intervals, while exhibiting price flexibility across multiple other settlement intervals. As shown in Table 1, these consumers, specifically, cryptocurrency-mining datacenter consumers, significantly reduce their demand during summer months, when high system-wide load stresses the power grid. However, the responsiveness of these consumers is not uniform across multiple facilities when exposed to the same circumstances [2]. Due to difficulties in modeling these behaviors, ERCOT has on occassion resorted to utilize frequency regulation to manage the unexpected variations introduced by data centers. Additionally, these large loads are one of the fastest growing, as illustrated in Figure 1 where one of the load zones in ERCOT saw an effective load growth of 30% within a period of nine months. The challenges in ensuring grid reliability under these emerging circumstances can be concerning for any power grid operator facing a similar stream of data-center connections.

Table 1. Day-time Correlation between Loads and Average Temperatures across Texas

	Large Flexible Load Response	ERCOT-wide Load Response
Non-summer	-0.17	0.78
Summer	-0.40	0.89

Data-center mining facilities generate revenue by selling cryptocurrencies, and earning power curtailment credits through participating in ERCOT's ancillary services markets [3]. Cryptocurrency mining loads typically operate as part of a mining pool[1], where the Bitcoin reward is a function of the hashing power contributed. Hashing power is directly related to the energy consumption of mining loads. The operating costs of these facilities include energy

[1] Refer to: https://www.investopedia.com/terms/m/mining-pool.asp

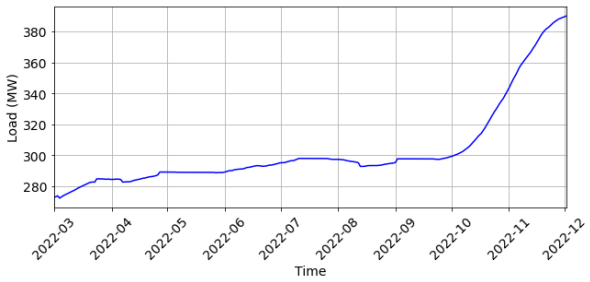

Figure 1. Trend of large flexible loads in a typical Texas load zone.

procurement through various mechanisms, including long-term power purchase agreements with generating facilities, and transactions at ERCOT's cleared energy markets. ERCOT also has a fixed-cost recovery mechanism, where it identifies the four highest 15-minute electricity usage intervals each month from June to September—during peak demand times—and allocates the fixed transmission and distribution network costs among all load participants proportionally to their consumption in these 4 intervals for the next year. These usage intervals are calculated on an *ex-post* basis. For a 500 MW crypto-mining farm, consuming 500MW on these intervals, the annual recovery cost is calculated as $500\text{MW} \times \$4.96/4\text{CP kW} \times 1000 \times 12 = \$29.76\text{M}$ (the rates for 4CP charges are taken from [4]), which constitutes a significant portion of the facility's operation cost.

Thus, the profit of a mining facility can be defined as:

$$P^M = \sum_{\forall t} \left(\pi_t^B k^B E_t^H - \pi_t^R E_t^R - \pi_t^D E_t^D + \gamma(E_t^M)\right) \tag{1}$$

where,

$$E_t^M = E_t^P + E_t^D + E_t^R = E_t^H + \psi(E_t^H, T_t) \tag{2}$$

Here, π_t^B represents the exchange rate of the cryptocurrency, while π_t^D and π_t^R are the day-ahead and real-time electricity market prices, respectively, for interval t. Parameter k^B is the efficiency of cryptocurrency miners' power supply. The function $\gamma(\cdot)$ represents the opportunity cost by avoiding 4CP charges, in terms of the miner's energy consumption E_t^M. Ignoring miners' revenue from ancillary services participation, (1) identifies the total profit for the miners. Eq. (2) indicates that the total energy procured by a mining facility (E_t^M), which corresponds to the sum of power purchase agreements (E_t^P), day-ahead (E_t^D), and real-time (E_t^R) energy procurement. The miners use a portion of their procured energy for hashing (E_t^H) and another part for cooling, which is a function $\psi(\cdot)$ of hashing power and ambient temperature (T_t). Using this model to predict or generate synthetic data for miners, however, poses multiple challenges. First, aside from power purchase agreements, all energy prices are only known on an *ex-post* basis. Second, as one crypto miner reported in its U.S. Securities and Exchange Commission (SEC) annual report [3], miners do not sell their entire cryptocurrency inventory, implying that the expected value of holding cryptocurrency must be higher than the current exchange rate. Third, the opportunity cost for cryptocurrency miners can be extremely complex to compute. These uncertainties in revenues and costs make it extremely difficult to determine the energy consumption of miners that maximizes their profit.

The response of industrial facilities to electricity prices has already been thoroughly discussed in the literature, and researchers have already extensively worked in this regard [5]. For example, in [6], researchers have discussed 'arbitrage price,' which determines farms' profitability from energy usage. However, contrary to other industries, the exchange rate of cryptocurrencies is highly volatile. Additionally, cryptocurrencies can be stored in infinite quantities and for indefinite periods, which implies that cryptocurrency miners may not be subjected to same market forces as in other industries. Furthermore, as prices are known on an *ex-post* basis, and, based on market rules, cryptocurrency miners must bid at least 24 hours in advance for the day-ahead market and at least 1 hour in advance for the real-time market, it is important to investigate how cryptocurrency miners have behaved in electricity markets and responded to policies in the past, in order to understand how increasing numbers of these consumers will affect market performance for all participants in the long term.

High-level behavioral analyses of cryptocurrency mining facilities are available in the literature [7]. Researchers have also studied the impacts of various demand response programs for cryptocurrency mining loads in Texas [8] and mechanisms for miners to participate in the ERCOT market for profit maximization [9]. However, there is a lack of large-scale, data-driven analyses that provide predictive insights into why and to what extent mining facilities respond to electricity prices beyond simple heuristics. In this article, we use predictive modeling to gain insights into the energy consumption of mining facilities. As highlighted in (2), mining consumption can be regressed against ambient temperature, cryptocurrency prices, and day-ahead and real-time energy prices. During summer months, mining facilities may use predictors other than electricity prices to hedge against consumption during 4CP hours, which will also impact energy consumption. Additionally, there could be other endogenous factors based on historical operating experience that may not be explained by the exogenous factors discussed before.

Based on these insights, this paper proposes an autoregressive model with exogenous variables (AR-X)

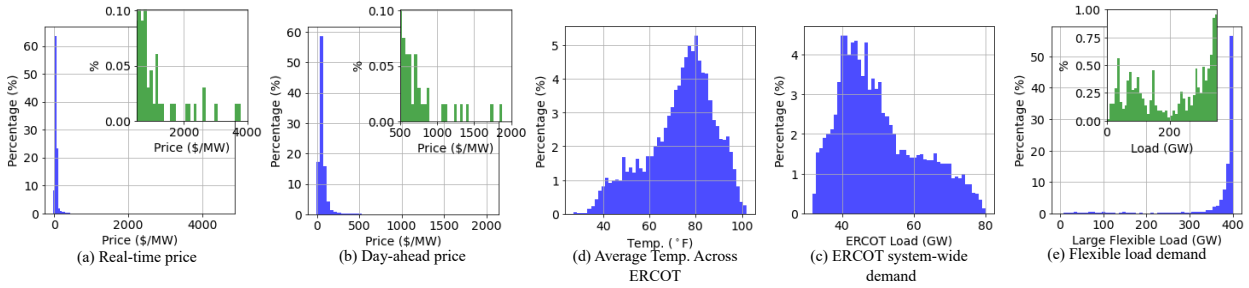

Figure 2. Histogram of the various hourly datasets for Apr.-Oct. 2022.

Table 2. Summary Statistics for All Variables Contributing to Miners' Energy Consumptions

Metric	Real Time Price	Day Ahead Price	ERCOT System Demand	Texas Average Temperature	LFL Demand
Mean	65.31	68.48	50218.1	73.4	370.59
Std. Dev.	136.56	79.5	11168.27	14.58	70.67
Minimum	-26.2	-6.48	31796.35	27.25	0.1
Maximum	4645	2045	80037.84	102.02	400.34
Skewness	**18.58**	**10.52**	0.67	-0.6	**-3.57**
Kurtosis	**457.82**	**178.53**	-0.44	-0.24	**12.16**
J-B Test p-Values	**0**	**0**	**0**	**0**	**0**
ADF Statistic p-Values	0	0	0.03	0.04	0
BP Test p-Values	0.47	0.44	**0.01**	0.79	**0**
Durbin Watson Test	**0.26**	**0.12**	**0**	**0**	**0**

for capturing the influence of these factors on miners' energy consumption. We develop two AR-X models, one describing the demand during summer and the other during non-summer months. The use of developed models can generate synthetic data for large-scale power system simulations under various environmental and market scenarios, advancing the state-of-art in modeling these types of consumers in electricity markets.

2. Exploratory Data Analysis

Fig. 2 depicts the histogram of the hourly time-series panel data for cryptocurrency miners' demand and related explanatory parameters from March to October 2022. The electricity price data includes average real-time and day-ahead prices across ERCOT load zones. ERCOT system demand represents the total demand across all ERCOT-managed regions in texas[2]. For average temperature, we collected weather data from several weather stations across ERCOT-managed regions in Texas[3], and calculated the average temperature across these stations. The flexible load demand dataset, representing hourly load data aggregated across an ERCOT load zone, is not publicly available and can be obtained upon request.

As shown in Fig. 2 with summary statistics detailed in Table 2, cryptocurrency mining load demand and prices exhibit significant skewness. From the p-values of the Breusch-Pagan (BP) test, the dataset, particularly cryptocurrency miners' energy consumption and electricity prices, displays non-constant variance (heteroscedasticity). If not addressed, this skewness and heteroscedasticity can cause inaccuracies in regressive models. According to the Gauss-Markov assumption, for linear regression estimators to remain unbiased, the error terms must have zero conditional means and be homoskedastic [10]. Additionally, ensuring normality in the error distribution is essential for applying the Central Limit Theorem, which aids in inferential statistics, including hypothesis testing and constructing confidence intervals.

2.1. Data Transformation

While performing the time-series analysis, it is essential to remove all diurnal and seasonal patterns in the datasets. Additionally, transformations are applied to achieve an approximately Gaussian distribution [11], especially in those cases where the panel data are heavily skewed. While this transformation is not mandatory, it helps in ensuring that the residuals satisfy the Central Limit Theorem (i.e., namely that a model constructed from sequential addition of random variables will, under mild assumptions, inevitably exhibit Gaussian characteristics). There is no specific sequence for applying these steps. Given the exponential growth in crypto-miners penetration in the ERCOT grid, as highlighted through Fig. 1, we

[2] All these datasets are sourced from `www.ercot.com`.
[3] Available from `www.wunderground.com`.

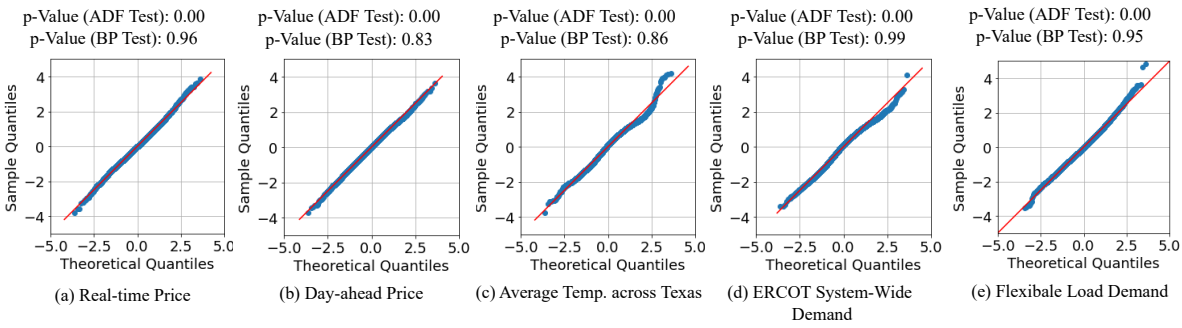

(a) Real-time Price (b) Day-ahead Price (c) Average Temp. across Texas (d) ERCOT System-Wide Demand (e) Flexibale Load Demand

Figure 3. Q-Q plots for transformed datasets.

first extract responsive components from the general trend. Here, we assumed that the daily peak mining load demand remains constant within a rolling window, which also provides us with the trend component. The hourly time-series miners' consumption is obtained by dividing the trend component from the actual time-series data. The transformation and standardization steps are given below:

i. We apply a non-parametric transformation to make the dataset, $y_{s,d}$, approximately follow a Gaussian distribution. The inverse quantile transform, a non-parametric technique, sorts the dataset in monotonic order, estimates cumulative probabilities, and identifies discrete quantiles for transformation. The transformation process is as follows:

$$y^G_{s,d} = Q^{-1}\left(y_{s,d}\right) \tag{3}$$

ii. We remove seasonality and diurnal effects by normalizing the dataset using the sample mean and standard deviation:

$$\tilde{y}_{s,d} = \frac{y^G_{s,d} - \hat{\mu}_{s,d}}{\hat{\sigma}_{s,d}} \tag{4}$$

The quantile plots, Augmented Dickey-Fuller (ADF) statistic p-values, and BP test p-values in Fig. 3 show that all transformed datasets are normally distributed, stationary and homoskedastic.

2.2. Correlation Analysis

2.2.1. Value of cryptocurrencies First, we only have access to historical daily Bitcoin exchange rate data, making it difficult to compare it against hourly cryptocurrency miners' responses. Secondly, the panel data is for the year 2022, when Bitcoin prices generally exhibited a downward trend, while, as shown in Fig. 1, there is an overall upward trend in cryptocurrency miners' daily peak energy consumption. This could lead to incorrect conclusions about the relationship between cryptocurrency miners' energy consumption and Bitcoin exchange rate. To address these limitations, we calculated the daily net energy consumption for miners using detrended electricity consumption data. Instead of using actual Bitcoin prices, we employed the Relative Strength Index (RSI) [12], a momentum measure describing the speed and magnitude of a security's price

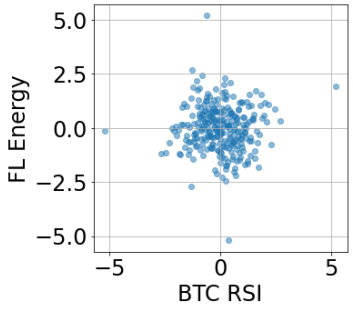

Figure 4. Comparing RSI of Bitcoin and energy consumption of flexible loads.

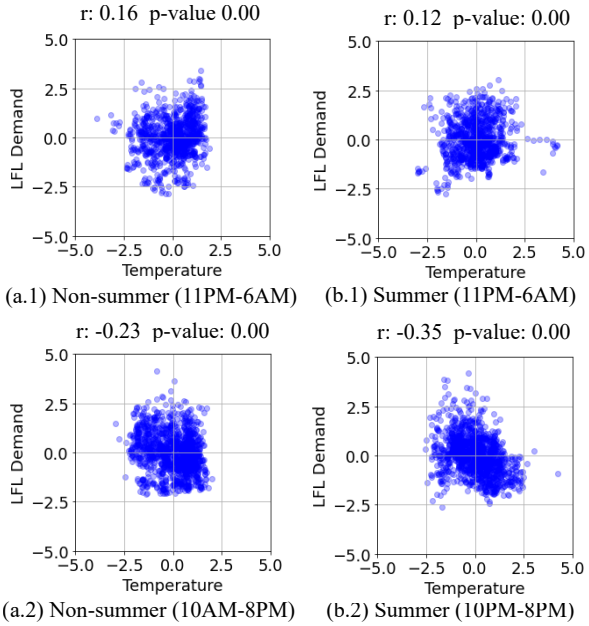

(a.1) Non-summer (11PM-6AM) (b.1) Summer (11PM-6AM)

(a.2) Non-summer (10AM-8PM) (b.2) Summer (10PM-8PM)

Figure 5. Correlation identifying how much of crypto-miners energy consumption is responsible for cooling.

changes. The RSI is a short-term measure of overvalued or undervalued security conditions that can potentially be used by crypto miners to determine the degree of demand response.

The scatter plot of the RSI of the Bitcoin exchange rate and the energy consumption of flexible loads, depicted in Fig. 4, shows a p-value of correlation coefficient 0.97. This suggests that, given the panel data concerned, cryptocurrency miners are agnostic to Bitcoin prices in the short term. Note that the energy consumption panel data we have corresponds to large flexible loads, including data centers, crypto mining facilities, and hydrogen production facilities. Hydrogen production capacity in Texas is currently in its early stages, with an installed capacity of only half a megawatt (MW) [13]. However, as of May 2024, there are plans for 1.6 gigawatts (GW) of capacity, which will be responsive to electricity prices. Electrolysis facilities, unlike Bitcoin mining operations, may respond to short-term hydrogen prices. Additionally, hydrogen production facilities have limited storage capabilities compared to Bitcoin mining.

2.2.2. The cooling energy requirements A significant portion of the energy consumed by cryptocurrency miners is dedicated to cooling (see eq. (2)), making it a major factor in their overall energy consumption. The cooling requirements are influenced by factors such as ambient temperature, the efficiency of the cryptocurrency miners, and hashing energy consumption. In this context, we focus on the impact of ambient temperature on cooling needs. During the daytime, the strong correlation between temperature and system-wide energy prices can obscure the cooling energy consumption and, even in non-summer months, temperatures can remain high into the late evening. As illustrated in Fig. 5, from 10 PM to 6 AM, both in non-summer and summer periods, we observe weak positive correlations with p-values close to 0. This confirms the physical principle that higher ambient temperatures necessitate more energy for cooling.

2.2.3. Price responses If we ignore a few price peaks, historically in the ERCOT market—as shown in Fig. 2(a,b)—day-ahead prices are statistically higher than real-time prices and have a comparatively narrower standard deviation. This implies that day-ahead prices remain elevated for longer periods. Therefore, cryptocurrency miners' response to day-ahead prices will be stronger than their response to real-time prices, especially during the summer. Prices tend to be statistically lower at night, suggesting that cryptocurrency miners may not be incentivized to respond to either day-ahead or real-time prices during both summer and non-summer months during late-night hours. During the summer, prices remain higher than during non-summer months, as shown in Fig. 7. We observe that cryptocurrency miners respond more vigorously to both day-ahead and real-time prices during the summer months. These price-responsive behaviors are depicted in Fig. 6. While not shown for brevity, cryptocurrency miners respond further vigorously during peak demand hours (3 PM-7 PM). The correlation coefficient for day-ahead prices during non-summer times increases to -0.29 (p-value 0.00), and during summer times to -0.42 (p-value 0.00).

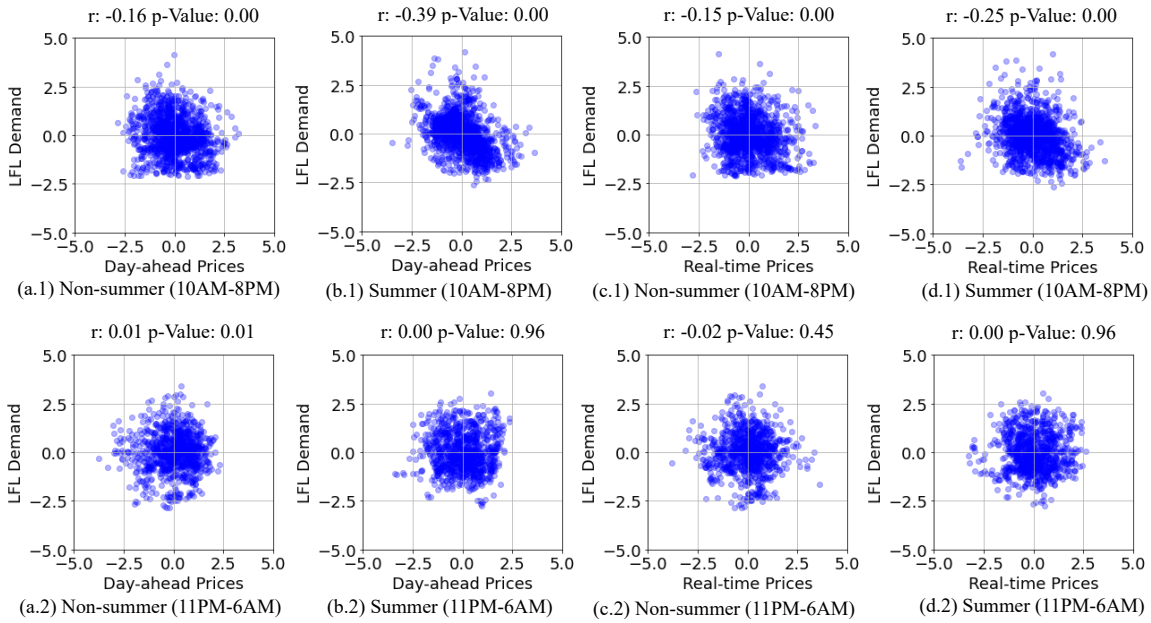

Figure 6. LFL responses to prices.

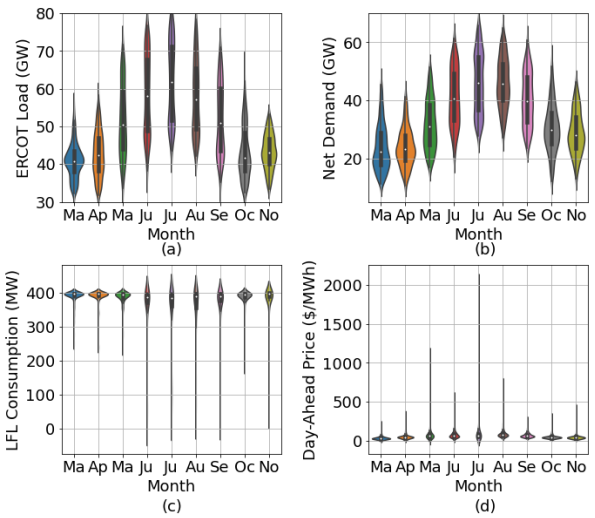

Figure 7. Comparing Demand across ERCOT, net-load, LFL consumption, and day-ahead price.

Selecting a narrower window for real-time prices did not significantly increase the correlation coefficients.

2.2.4. The predictors contributing to 4CP responses There are three main issues with using simple correlation analysis to understand the direct impact of energy prices on demand. Firstly, prices are not known *a priori*. Market entry closes for the day-ahead market about 24 hours in advance, and for the real-time market about 1 hour before actual consumption. Consequently, cryptocurrency miners must decide whether and how much to shut down their facilities latest with the real-time market. Secondly, as observed in Fig. 7(c-d), the day-ahead prices in June, August, and September were not significantly higher, yet cryptocurrency miners responded as vigorously as they did in July. Thirdly, as shown in Fig. 4, when energy prices are low, it is trivial for miners to operate at full capacity. This implies that cryptocurrency miners are likely using factors other than electricity prices to hedge against 4CP charges.

4CP peaks are calculated based on ERCOT-wide demand and are price-agnostic. For example, in August 2023, the peak demand occurred on the 10th, while the price peaked at approximately $4000/MWh on the 11th. Except for a few instances, 4CP peaks in ERCOT generally arise between 4 PM and 6 PM. However, it is also implied that a higher load leads to higher electricity prices. To capture the impact of 4CP hedging, we need to focus on months when energy prices were low, such as June and September, within hours 4 PM-6 PM. We use ERCOT-wide system demand as the predictor. As depicted in Fig. 8, the correlation appears strongest when considering months with lower energy prices alone.

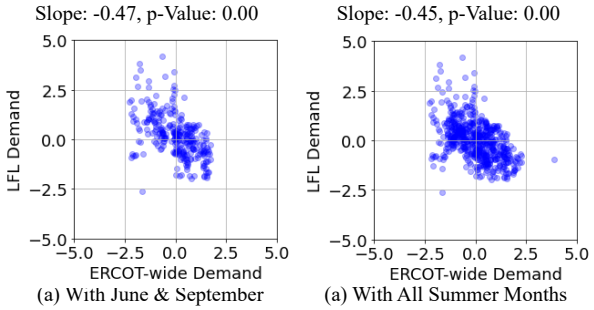

Figure 8. Factors responsible for 4CP demand response in addition to prices.

2.2.5. Auto-regressive Model The Durbin-Watson tests in Table 2 indicate a significant presence of autocorrelation within the cryptocurrency miners' energy consumption dataset. Autocorrelation occurs when variables are correlated with their own past values, suggesting that the energy consumption of cryptocurrency mining facilities is influenced by their historical operational patterns.

Auto-Regressive Integrated Moving Average (ARIMA) processes are a class of stochastic processes used to analyze time series data. The ARIMA process, attributed to Box and Jenkins [14], hypothesizes that the residual term is randomly drawn from a normal distribution with zero mean and constant variance, known as a white noise process. However, ARIMA models can be robust to the non-normality of residuals. As with other time-series analyses, the residuals need to be homoskedastic, and the time series itself must be stationary. A general ARIMA model is formally defined as follows:

$$\Phi(B^S)\phi(B)\nabla^d\nabla_s^D y_t = \Theta(B^S)\theta(B)\epsilon_t \qquad (5)$$

where y_t is the modeled cryptocurrency miners energy consumption data. Here, B is the backshift operator,

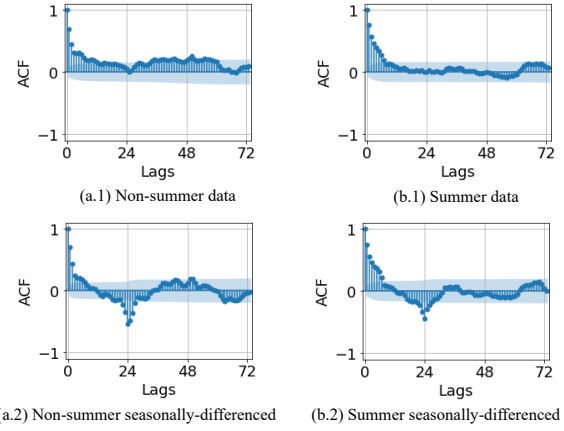

Figure 9. ACF plots without and with seasonal differencing considering both non-summer and summer months.

where $B^l r_t := r_{t-l}$, and S is the seasonality of the time series. The functions representing auto-regressive, moving average, and differences and their seasonal forms are defined as $\phi(B) = 1 - \sum_{i=1}^{p} \phi_i B$, $\theta(B) = 1 + \sum_{i=1}^{q} \theta_i B$, $\nabla^d y_t = (1-B)^d y_t$, $\Phi(B^S) = 1 - \sum_{i=1}^{P} \Phi_i B^S$, $\Theta(B^S) = 1 + \sum_{i=1}^{Q} \Theta_i B^S$, and $\nabla_S^D y_t = (1-B^S)^D y_t$. The parameters p, d, q, P, D, Q and S identify the specific ARIMA process.

Autocorrelation factors (ACF) for both non-summer and summer months are plotted in Fig. 9. Spikes around a lag of 24 in the ACF plots, which become prominent with seasonal differencing of 24 periods, suggest that the data exhibits seasonality, which is expected since the time-series dataset is hourly. The diminishing ACF plots indicate the presence of moving average components.

3. Empirical Observation of Cryptocurrency Miners' Behavior

The correlation analysis indicates that factors such as energy market prices, average temperatures across Texas, and ERCOT-wide energy demand influence the energy consumption of cryptocurrency mining facilities in a complex manner. We observe that these factors can affect each other, necessitating a focus on specific time slots to capture the underlying physics-based relationships. The objective of this section is to perform multivariable linear regression to develop mathematical models describing the energy consumption of aggregated cryptocurrency mining facilities. We hypothesize the models to be as follows:

$$\begin{aligned} E_t^{M,\text{ns}} = N^{-1}\Bigg(&\psi^{\text{ns}} T_t \\ &+ \mathbb{I}^d(t)\left(\sum_{\forall n\geq 0} \delta_n^{D,\text{ns}} \pi_{t-n}^{D} + \sum_{\forall n\geq 1} \rho_n^{D,\text{ns}} \pi_{t-n}^{R}\right) \\ &+ \mathbb{I}^p(t)\left(\sum_{\forall n\geq 0} \delta_n^{P,\text{ns}} \pi_{t-n}^{D} + \sum_{\forall n\geq 1} \rho_n^{P,\text{ns}} \pi_{t-n}^{R}\right) \\ &+ \text{ARMA}^{\text{ns}}(p,d,q)(P,D,Q,[24])\Bigg) \end{aligned} \tag{6}$$

$$\begin{aligned} E_t^{M,\text{s}} = N^{-1}\Bigg(&\psi^{\text{s}} T_t \\ &+ \mathbb{I}^d(t)\left(\sum_{\forall n\geq 0} \delta_n^{D,\text{s}} \pi_{t-n}^{D} + \sum_{\forall n\geq 1} \rho_n^{D,\text{s}} \pi_{t-n}^{R}\right) \\ &+ \mathbb{I}^p(t)\left(\sum_{\forall n\geq 0} \delta_n^{P,\text{s}} \pi_{t-n}^{D} + \sum_{\forall n\geq 1} \rho_n^{P,\text{s}} \pi_{t-n}^{R}\right) \\ &+ \mathbb{I}^p(t) \sum_{\forall n\geq 1} \gamma_n L_{t-n} \\ &+ \text{ARMA}^{\text{s}}(p,d,q)(P,D,Q,[24])\Bigg) \end{aligned} \tag{7}$$

Here, $E_t^{M,\text{ns}}$ and $E_t^{M,\text{s}}$ are the modeled cryptocurrency miners energy consumptions during non-summer and summer months. Variable $\psi^{(\cdot)}$ is the regression coefficient for temperature. Parameters $\mathbb{I}^d(t)$, and $\mathbb{I}^p(t)$ are pre-identified binary indicators, which equal 1 when the impacts of the associated regressors are active. As discussed earlier, although the market gate closure happens earlier, facilities may still utilize the day-ahead-market-cleared data to adjust their bids in the real-time market. The goal here is not to understand how much cryptocurrency miners are bidding into each market but rather to observe how their consumption correlates with historical data, which is why, for day-ahead prices, n can be equal to zero. Here, $\delta_n^{D,\text{s}}$ $\rho_n^{D,\text{s}}$ are regression coefficients for price response, and $\delta_n^{P,\text{s}}$ $\rho_n^{P,\text{s}}$ are for peak price response. Cryptocurrency mining facilities might use different predictors, γ_n, which are combinations of historical loads based on their risk appetite during 4CP hours, also identified through $\mathbb{I}^p(t)$. Finally, the ARMA process models the variance unexplained by the regression model. Here, N is the inverse transformation used to revert the transformed dataset.

In this article, contrary to building the model in a single step, we perform multiple linear regressions to systematically extract the influence of regressors and perform regression based on the residuals from the previous step. To validate the developed regression models, at each step, we divide the data into training and testing samples and compare metrics such as mean squared error (MSE) and root mean squared error (RMSE). Here, we report only statistically significant regressors and their associated p-values for brevity. This section is divided into two subsections dedicated to modeling demand response during non-summer and summer times, respectively.

3.1. Demand response model for the non-summer months

3.1.1. Temperature effect We initially divided the datasets into training and testing sample days for regression analysis. However, we observed a significant discrepancy in the calculated Mean Squared Error (MSE). We discovered that during certain late nights, the increased temperature decreased the energy demand of cryptocurrency mining facilities. While we don't know what exactly is responsible for such discrepancy, the removal of days with higher average real-time prices based on their z-scores addressed this issue. Using this procedure, even with a 50/50 split of the dataset, the MSE for the training and testing samples remains similar. The calculated correlation coefficient ψ^{ns} is 0.14 (S.E. = 0.04, p-value = 0.00). We assumed a similar correlation holds during the daytime and removed the associated effect from the dataset to generate the residuals.

3.1.2. Price effects To capture the impact of prices as identified in (8), we regressed the energy consumption data against price data with various lag periods. We focused on the hours between 10 AM and 8 PM, setting $\mathbb{I}^d(t) = 1$ during these hours. Our analysis revealed that the strongest p-values occurred when considering day-ahead price data from 2 days prior (n=48), real-time prices from the last hour (n=1), and day-ahead prices from the previous day (n=24). The calculated values are $\delta_{48}^{D,\text{ns}} = -0.08$ (S.E. = 0.03, p-value = 0.01), $\rho_1^{D,\text{ns}} = -0.19$ (S.E. = 0.03, p-value = 0.00), and $\rho_{24}^{D,\text{ns}} = -0.11$ (S.E. = 0.03, p-value = 0.00).

Surprisingly, these results suggest that cryptocurrency mining facilities, participating in the day-ahead market, may simply observe the most recent publicly available day-ahead prices and adjust their consumption position accordingly. One-day-ahead real-time prices are not available at the time of bidding into the day-ahead market. Therefore, despite higher variability, cryptocurrency miners could be utilizing both one-day-ahead and one-hour-ahead real-time prices to adjust their energy consumption in the real-time market.

3.1.3. Peak price effect Here, $\mathbb{I}^p(t) = 1$ between 3 PM and 7 PM. Considering a mix of regressors, we found that day-ahead prices from the past hour and real-time prices three hours prior have the strongest correlation. Notably, both sets of data are publicly available for real-time adjustment. The calculated $\delta_1^{P,\text{ns}}$ is -0.16 (S.E. = 0.05, p-value = 0.00), and $\rho_3^{P,\text{ns}}$ is -0.29 (S.E. = 0.05, p-value = 0.00). The negative sign indicates that, as cryptocurrency miners observe increasing prices leading to peak hours, they reduce their energy consumption monotonically. Note that this adjustment can only be carried out in the real-time market.

3.1.4. Autoregressive component In the ACF and partial ACF (PACF) factors calculated using the residuals, we observe spikes at lag 1 in the PACF plots without seasonal differencing, implying the strong presence of an AR(1) component. We still observe spikes appearing at around a lag of 24 in both ACF and PACF plots, suggesting the seasonality of the data. Spikes near the lag of 24 in the seasonally differenced PACF data suggest the presence of seasonal autoregressive order. We observe that the ARIMA(1,0,0)(1,1,0)[24] model fits reasonably well based on the Akaike Information Criterion (AIC). The model parameters are given as: ϕ_1 = 0.83 (S.E. = 0.02, p-value = 0.00), Φ_1 = -0.43 (S.E. = 0.02, p-value = 0.00), σ = 0.58 (S.E. = 0.02, p-value = 0.00).

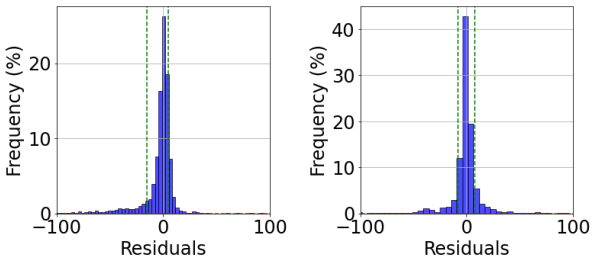

(a.) Correlation analysis alone (b) Correlation + autoregressive

Figure 10. Distribution of residuals. Green lines show the region with 75% quantile.

The Ljung-Box test shows the lack of autocorrelation in the residuals (p-value = 0.82). The ADF test indicates that the residual is stationary (p-value = 0.82), and the BP test shows that the dataset is weakly heteroskedastic (statistic = 77.5). With a 35/65 split between training and testing samples, we observe an MSE of 1.37, implying that while the ARIMA model captures the variability in the dataset well. Note that these calculations are based on transformed data, and these figures improve when computed in the original space.

3.1.5. Accuracy of non-summer model The empirical equation representing cryptocurrency miners' demand response during the non-summer months is given as:

$$\begin{aligned} E_t^{M,\text{ns}} = N^{-1}\,(&0.14T_t \\ &+ \mathbb{I}^d(t)\left(-0.08\pi_{t-48}^D - 0.19\pi_{t-1}^R - 0.11\pi_{t-24}^R\right) \\ &+ \mathbb{I}^p(t)\left(-0.16\pi_{t-1}^D - 0.29\pi_{t-3}^R\right) \\ &+ \text{ARMA}^{\text{ns}}(1,0,0)(1,1,0,[24])) \end{aligned} \tag{8}$$

To compute the overall accuracy of the model, we need to compare how much of the variability is explained using correlation analysis alone versus the additional use of an autoregressive model. From the residuals in Fig. 10, it can be seen that the correlation analysis captures a significant amount of variability, which is further enhanced by the autoregressive model. However, the lines indicating the 75% inter-quantile range show a significant amount of variance that the model could not explain, which, based on the raw data, is due to the magnitude of peaks.

The mean squared error (MSE) and mean absolute percentage error (MAPE) of the correlation analysis-only model are 25.10 and 3.27%, respectively. These values change to 32.06 and 3.55% when using the combined correlation and autoregressive model. However, the true value of the combined model is reflected in the coefficient of determination, which, considering errors only up to the 75% inter-quantile range, improves from 0.32 to 0.77.

3.2. Demand response model for the summer months

3.2.1. Temperature effect Like non-summer months, we observed similar discrepancies, where during late nights, higher temperatures are shown to lead to lower energy consumption. However, compared to non-summer times, the impact is less prominent here, which could be due to temperatures remaining high through the summer, thereby masking the relation between temperature and consumption. The calculated regression coefficient ψ^{s} is given as 0.12 (S.E. = 0.04, p-value = 0.01).

3.2.2. Price effects As in the non-summer model, we focused on the hours between 10 AM and 8 PM for all four summer months and regressed the energy consumption data against price data. Here, we observed that the strongest p-values occurred when considering real-time price data from 3 days prior ($n = 72$) and the current day-ahead prices ($n = 0$). The calculated values are $\delta_{72}^{R,\mathrm{s}} = 0.09$ (S.E. = 0.04, p-value = 0.03) and $\rho_0^{D,\mathrm{s}} = -0.40$ (S.E. = 0.04, p-value = 0.00).

This behavior essentially implies that there is a negative correlation between the real-time energy prices from 3 days prior and the current day-ahead energy prices. Specifically, if the day-ahead prices are significantly high, the real-time prices will also be higher during the same period, leading facilities, such as cryptocurrency mining operations, to significantly reduce the intensity of their operations.

3.2.3. Peak price effect Prices peak during the summer months, especially in the afternoon hours. These are the same hours when the demands peak as well, and ERCOT calculates 4CP charges based on consumption during these hours. Therefore, it is of interest to isolate how much the cryptocurrency miners are responding because of peak prices from the hedging to avoid 4CP charges. Here, we focus on the former, where we want to investigate, like the non-summer months, how increased prices contribute to cryptocurrency miners' response. Therefore, we focus on July and August datasets, the months with higher price volatility, specifically between 3 PM and 7 PM. We observe that the decision to reduce electricity consumption due to peak electricity prices is based on the recently cleared real-time prices. The coefficient showing the relationship $\rho_1^{P,\mathrm{s}}$ is -0.13 (S.E. = 0.06, p-value = 0.033).

3.2.4. 4CP effect Of the 2.9 GW of installed capacity in ERCOT, if some cryptocurrency miners try to avoid the critical peak, the peak demand could shift to later in the same day or even to the next day. To examine these effects, which may occur primarily to avoid 4CP charges, we focus on June and September, specifically between 3 PM and 7 PM. Interestingly, we find that the energy consumption of cryptocurrency miners is a weighted average of their consumption over the past two days during similar hours. The correlation coefficients are given as $\gamma_{24} = -0.89$ (S.E. = 0.11, p-value = 0.00) and $\gamma_{48} = 0.39$ (S.E. = 0.114, p-value = 0.00). This suggests that miners might be basing their behavior on ERCOT's system-wide demand from the previous day. If the demand two days ago was not too high, but the demand yesterday was high, it is likely that today's demand will also be high. This behavior appears to be completely rational.

3.2.5. Autoregressive component Like non-summer months, we observed that the ARMA model (1,0,0)(1,1,1,[24]) could explain a significant part of the variability in the residual dataset. The model parameters are given as: ϕ_1 = 0.84 (S.E. = 0.01, p-value = 0.00), Φ_1 = -0.09 (S.E. = 0.03, p-value = 0.00), Θ_1 = -0.93 (S.E. = 0.02, p-value = 0.00) and σ = 0.7 (S.E. = 0.01, p-value = 0.00). The Ljung-Box test shows the lack of autocorrelation in the residuals (p-value = 0.88). The ADF test indicates that the residual is stationary (p-value = 0.00), and the BP test shows that the dataset is weakly heteroskedastic (statistic = 94.9). With a 35/65 split between training and testing samples, we observe an MSE of 1.09, implying that while the ARIMA model captures most of the variability in the dataset.

3.2.6. Accuracy of summer model The empirical equation representing cryptocurrency miners' demand response during the summer months is given in (9). The residuals suffer from similar issues as was discussed in the earlier model (with RMSE and MAPE of 83.14 and 90.96% with the correlation-only model to RMSE and MAPE of 60.86 and 64.24% with the incorporation of autocorrelation); however, the efficacy of the model is further evidenced through the increased coefficient of determination of 0.93 to 0.99, implying that the heuristic-based correlation model itself can explain a significant portion of cryptocurrency miners' behavior, and the model get strengthened with inclusion of ARIMA model.

$$\begin{aligned} E_t^{M,\mathrm{s}} = N^{-1} \big(&0.12 T_t \\ &+ \mathbb{I}^d(t) \left(-0.40 \pi_t^D + 0.09 \pi_{t-72}^R\right) \\ &+ \mathbb{I}^p(t) \left(-0.13 \pi_{t-1}^R\right) \\ &+ \mathbb{I}^p(t) \left(-0.89 L_{t-24} + 0.39 L_{t-48}\right) \\ &+ \mathrm{ARMA}^{\mathrm{ns}}(1,0,0)(1,1,1,[24])\big) \end{aligned} \tag{9}$$

4. Conclusion

We present an econometric model that provides a robust framework for understanding the behavior of

large flexible cryptocurrency mining loads in the Texas power grid. By incorporating internal factors through the SARIMA process and external factors via selective external correlations, our model achieves reasonable accuracy. The quantile transformation used captures some of the nonlinearities among the variables. Our analysis reveals that cryptocurrency miners' energy consumption is mostly influenced by temperature, energy prices, and demand response strategies rather than by short-term fluctuations in cryptocurrency prices. This insight challenges common misconceptions about the drivers of currency mining activities and highlights the importance of considering multiple factors in predictive modeling.

The practical utility of our model lies in its ability to generate synthetic datasets that can simulate various grid conditions and mining behaviors. This capability is crucial for power system simulations and for developing strategies to enhance grid reliability and efficiency. Also, this study can help power grid operators better anticipate and manage the impact of these emerging technologies on the energy landscape.

References

[1] Electric Reliability Council of Texas (ERCOT), *Large flexible load taskforce interconnection and planning working session*, Jun. 2022. [Online]. Available: https://www.ercot.com/files/docs/2022/06/23/LFLTF%20working%20session_06242022.pptx.

[2] Electric Reliability Council of Texas, *Large load interconnection status & analytics update*, Accessed: 2024-05-22, Apr. 2024. [Online]. Available: https://www.ercot.com/files/docs/2024/03/26/LLI%20Queue%20Status%20Update%20-%202024-4-1.pdf.

[3] Securities and Exchange Commission, *Riot blockchain, inc. 2022 form 10-k annual report*, Online, Mar. 2023. [Online]. Available: https://www.sec.gov/edgar/browse/?CIK=0001167419.

[4] Oncor Electric Delivery Company LLC, *Tariff for retail delivery service*, Accessed: 2024-06-12, 2023. [Online]. Available: https://www.oncor.com/content/dam/oncorwww/documents/about-us/regulatory/tariff-and-rate-schedules/Tariff%20for%20Retail%20Delivery%20Service.pdf.coredownload.pdf.

[5] H. Golmohamadi, "Demand-side management in industrial sector: A review of heavy industries," *Renewable and Sustainable Energy Reviews*, vol. 156, p. 111 963, 2022.

[6] N. B. Depree, D. P. Thomas, and D. S. Wong, "The contribution and economics of demand side response towards decarbonizing the aluminium smelting industry," in *Light Metals 2022*, Springer, 2022, pp. 560–570.

[7] J. D. Rhodes, T. Deetjen, and C. Smith, "Impacts of large, flexible data center operations on the future of ercot," *Idea Smiths LLC.: Austin, TX, USA*, 2021.

[8] A. Menati, K. Lee, and L. Xie, "Modeling and analysis of utilizing cryptocurrency mining for demand flexibility in electric energy systems: A synthetic texas grid case study," *IEEE Transactions on Energy Markets, Policy and Regulation*, vol. 1, no. 1, pp. 1–10, 2023.

[9] A. Menati, Y. Cai, R. El Helou, C. Tian, and L. Xie, "Optimization of cryptocurrency mining demand for ancillary services in electricity markets," in *Hawaii International Conference on System Sciences*, 2024.

[10] J. Wooldridge, *Introductory econometrics: A modern approach*, 2020. [Online]. Available: https://books.google.com/books?id=tXqAzwEACAAJ.

[11] R. Krzysztofowicz, "Transformation and normalization of variates with specified distributions," *Journal of Hydrology*, vol. 197, no. 1, pp. 286–292, 1997, ISSN: 0022-1694.

[12] R. A. Levy, "Relative strength as a criterion for investment selection," *The Journal of finance*, vol. 22, no. 4, pp. 595–610, 1967.

[13] U.S. Department of Energy, "Electrolyzer installations in the united states," Hydrogen and Fuel Cell Technologies Office, Tech. Rep., 2024. [Online]. Available: https://www.hydrogen.energy.gov/docs/hydrogenprogramlibraries/pdfs/24001-electrolyzer-installations-united-states.pdf.

[14] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.