

On the Measurement of Carbon Footprint of Bitcoin: A Bayesian Quantile Cointegration

Analysis

Michael L. Polemis^{a,b*} and Mike G. Tsionas^c

Abstract

During recent years, there is a widespread belief among researchers and academicians that the Bitcoin usage is imposing an additional burden on the environment inducing climate change. Although several studies have focused on issues related to the energy consumption of the basic cryptocurrencies, an open question remains regarding the environmental depiction of the Bitcoin. By resorting to Bayesian analysis and quantile cointegrated vector autoregression (CQVAR) model techniques, this study seeks to disentangle the driving forces that shape the carbon footprint of the Bitcoin. The sample used in the empirical analysis consists of a daily panel dataset covering 51 developing and developed countries over the years 2016-2018. The empirical findings corroborate a causal effect between the use of the Bitcoin and its underlying carbon dioxide emissions generated by the increasing energy load. The Bayesian CQVAR is associated with positive marginal posterior means for most of the covariates of the model across all the estimated quantiles. In contrast, there is a negative and statistically significant relationship between Bitcoin miner's revenue and carbon emissions, uncovering a multimodal distribution pattern of the marginal posterior densities which is stronger at higher than in lower quantiles. This finding, suggests that the lower (higher) miner's Bitcoin revenues the more abrupt (gradual) the effect on environmental degradation. Therefore, a sustainable energy strategy focusing on the penetration of renewable energy sources along with the use of energy-efficient mining hardware will alleviate the carbon footprint of the Bitcoin.

Keywords: Finance; Bitcoin; carbon emissions; Energy load; Bayesian CQVAR

JEL codes: C11; C22; G1; F64

^{a*} Department of Economics, University of Piraeus, Piraeus, 185 34 (Corresponding author). Tel: +30 210 4142303; Fax: +30 210 4142346; Email: mpolemis@unipi.gr.

^b Hellenic Competition Commission, Athens, Greece.

^c Department of Economics, Lancaster University Management School, LA1 4YX, United Kingdom.

1. Introduction

Financial technology (FinTech) constitutes an important source of capital not only for start-up companies but also for established institutions that seek to enhance their usage of financial services provided to other companies (B2B) or final consumers (B2C).

The use of cryptocurrency through decentralized blockchain technology, the mobile and web banking applications, the crowdfunding platforms, or certain artificial intelligence tools like robotic process automation, and big data are just a few examples of breakthrough technologies aiming to make financial services more accessible to the public. While the US still dominates, Europe and Asia now attract about half of the investment flows in the FinTech industry (Leong and Sung, 2018). Fintech investments are also becoming increasingly important in developing countries in the aftermath of the global financial crisis (India, Argentina, Singapore, and Hong-Kong). Even though the invested amounts are small relative to the sizes of the economies (e.g. US, Europe), the upside potential and growth prospects of the successful start-ups are tremendous. For example, Google, FedEx, Amazon, and Intel have embraced many Fintech technologies in the US. In Europe, in the past five years, a growing number of companies (Maersk, Vodafone, AXA, BNP Paribas) have been engaged in Fintech activities, which helped them to reach new markets and create jobs.

Bitcoin (BTC) and its underlying blockchain validation process have attracted considerable attention from the FinTech industry, while regulatory authorities and policymakers are rather skeptical for its increasing use since its initial launch in 2009 (Böhme, et al, 2015). This can be justified because investing in Bitcoins or any other cryptocurrency creates a risk to the stakeholders since decentralized algorithm-trading blockchain platforms are not subject to organized financial

market regulations and this raises serious concerns of security, internal control, and market surveillance (Cheng et al, 2019).

However, a large number of consumers (e.g. investors, traders, intermediaries) have a growing interest in this kind of financial transactions. Direct evidence of consumer awareness of BTC is hard to come by. But we can obtain some indirect evidence by comparing the search intensity for these terms as a news item from Google trends.² The rationale behind using this proxy is straightforward. Attention is a scarce cognitive resource thus, the amount of attention paid to a commodity or a service should reveal changes in preferences or status. Moreover, a search is a revealed attention measure since if a term has been searched in Google, attention has been paid to it. Worldwide data for internet news searches for the number of Google searches on “*bitcoin*,” “*blockchain*,” and “*cryptocurrency*” terms are available for the post-2016 period, and the results are provided in Figure 1. As it is evident from the inspection of the relevant figure, the search intensity for the term “*bitcoin*” seems to be higher than that for “*cryptocurrency*” and “*blockchain*”, reaching its peak in December 2017 (see Panel A). Moreover, internet searches for the terms “*bitcoin*” and “*cryptocurrency*” are highest in western economies (USA, Canada, Latin America, Europe) than in developing countries such as Africa and South Asia (see Panel B and C), whereas the search for “*blockchain*” is relatively more prominent in developing economies (see Panel D). This geographic distribution further accentuates the difference in search intensities plotted in Panel A. Though indirectly, this evidence corroborates our perception that consumers have a growing interest in how BTC through a Blockchain validation process could enhance many business products and processes (Katsiampa et al, 2019; Cheng et al, 2019).

² Data were obtained by filtering “*worldwide*”, “*news item*” and “*all categories*” from the drop down menu of the website for the period 01/01/2016-11/04/2020.

<Insert Figure 1 about here>

To give but an example of its tremendous penetration in financial transactions, it is noteworthy that its market capitalization reached the level of 133 billion USD in April 2020, while the number of circulated cryptocurrencies exceeded 16.8 million BTCs with its market value currently estimated to be around 45 billion USD (Cheng et al, 2019; Atsalakis et al, 2019; Katsiampa, 2017). Moreover, the computing capacity used in this process (mining) has increased rapidly in recent years, while energy requirements have shown an enormous increase (Das and Dutta, 2020; Symitsi and Chalvatzis, 2018).

For a BTC transfer to be executed and validated, an algorithm based on computer programming must be solved through the global BTC network. Although BTC is a virtual currency, the electricity load associated with its use is vast in numbers. It is estimated that BTC's annual electricity consumption adds up to 45.8 TWh (November 2018), while at the same time it causes around 23 megatons in CO₂ emissions annually (Stoll et al, 2019). Moreover, it is estimated that BTC could produce enough carbon releases to increase the global temperature above 2°C within less than three decades jeopardizing the Paris Agreement (Mora et al, 2018).³ Figure 2 displays Bitcoin's carbon dioxide releases relative to several countries in 2018. As it is evident, countries such as the USA, Russia, Korea, Venezuela, Iran, and China constitute the most polluting ones based on BTC's electricity load used for mining.

<Insert Figure 2 about here>

³ Paris Agreement on climate change mitigation or simply "*Paris Accord*" aims to strengthen the ability of countries to deal with the impacts of climate change by keeping a global temperature rise well below 2 degrees Celsius above pre-industrial levels.

Despite the profound interest of researchers and policymakers on the possible interactions between the use of BTC and environmental degradation, little progress has been made by the academic community to unravel possible causal effects. This study is based on Bayesian quantile regression and therefore provides inferences that are conditional on the data, without reliance on asymptotic approximation as its alternative classical quantile regression firstly developed in Koenker and Bassett, (1978). By employing a CQVAR model, this work attempts to define causal effects between the main financial characteristics of the BTC (return, miner's revenues, transactions, efficiency, fees, and cost) and the probability density function (PDF) of Bitcoin's carbon releases over the quantiles. The reason for relying on quantile analysis is that it provides a more accurate description of the response distribution than the mean, since the latter usually neglects important features of the distribution, especially when we might expect a different structural relationship for the higher/lower responses than the average ones (see Taddy and Kottas, 2010).

The findings reveal the existence of positive marginal posterior effects for most of the covariates across all the quantiles. On the contrary, there is a negative correlation between miner's revenue and carbon dioxide emissions, mostly evident at higher quantiles. The empirical results remain robust when Markov Chain Monte Carlo (MCMC) techniques for numerical Bayesian inference are used organized around a Metropolis-Hastings algorithm. Moreover, we also document important differences between the CQVAR and the traditional Bayesian quantile analysis, in terms of several aspects of the data.

This study offers new insights into the existing BTC technical issues while it contributes to the emerging literature in many ways. First, we utilize Bayesian analysis to extend a CQVAR model developed in Koenker and Xiao (2006). By applying the notion of cointegration to Bayesian

analysis, we can trace structural relationships between the sample variables of our model. In this way, we compare and contrast the estimates drawn from the new (cointegrated) model with a standard Bayesian quantile regression analysis. Specifically, this study is the first -to the best of our knowledge- that uses a Bayesian approach to the CQVAR to obtain the effect of covariates across different quantiles. Second, we contribute to the broader BTC literature by dissecting the drivers of BTC carbon footprint across the quantiles, an issue that has been missing from the extant literature. Third, we provide useful policy implications to regulators and government officials from the interpretation of the empirical findings.

The rest of the paper proceeds as follows. Section 2 describes the literature in the field. In Section 3 the Bayesian CQVAR model is presented and discussed. Section 4 presents the data and the sample variables used in this study. In Section 5 the empirical findings analyzed along with the necessary sensitivity analysis to check for the robustness of the empirical findings. Finally, Section 6 concludes the paper, offering useful policy implications.

2. Literature review

The existing literature on the possible spillover effects of the use of BTC to environmental degradation is still in its infancy. The majority of related studies focus solely on the energy-mining nexus, whereas scant attention has been paid to the environmental consequences of the cryptocurrency mining (see for example Das and Dutta, 2020; Li et al, 2019; De Vries, 2018; Krause and Tolaymat, 2018; O'Dwyer and Malone, 2014).

Das and Dutta, (2020), which is the most related study to our work, investigate the relationship between Bitcoin's energy consumption measured by electricity load for mining BTC and its revenue level. They rely on quantile and Markov regime-switching regression analysis on a daily dataset comprising of 774 observations (February 2017 to March 2019). The empirical

findings indicate that there is a negative relationship between energy load and miner's revenue, which is highly significant up to the 30th quantile of the miner's revenue PDF. In other words, Bitcoin's higher energy consumption may decrease their profitability level suffering even losses since energy costs contribute substantially to total costs when the revenues are low and volatile. On the contrary, all other independent variables (transaction cost, transaction fees, total mining, number of unique BTC address used and BTC returns) are positively correlated with the miner's revenue PDF across the quantiles.

In a similar vein, Li et al, (2019) ran experiments on mining efficiency of nine kinds of cryptocurrencies and ten algorithms. The empirical results indicate a positive relationship between global electricity consumption and mining activity, while the hashing algorithm mainly determines the mining efficiency. In an earlier work De Vries (2018) assesses alternative methods in determining the current and future electricity consumption of the BTC network. This study argues, that the BTC network consumes at least 2.55 GW of electricity currently, and that it could reach a consumption of 7.67GW in the future. Krause and Tolaymat, (2018) demonstrate a methodology for calculating the minimum power requirements of four cryptocurrencies (Bitcoin, Ethereum, Litecoin, and Monero) and the subsequent energy consumed to generate one US dollar's worth of digital assets. They claim that all four encrypted cryptocurrencies consume more energy than mineral mining to produce an equivalent market value.

As it is evident from the above analysis, the BTC mining constitutes beyond any doubt a "*power-hungry*" process (see also Foteinis, 2018; De Vries, 2018). However, the direct link of BTC with the carbon emissions released by its mining is nearly overlooked by the existing literature. Recent works by Mora et al, (2018), Krause and Tolaymat (2018) and the most recent one by Stoll et al. (2019), attempt to shed light on this relationship. Mora et al. (2018) estimate the

carbon emissions of BTC by assuming the adoption trends of 40 different used technologies for which data were available. For each year, they calculate the average as well as lower and upper quantiles of percent population using all technologies and apply a logistic model to project these trends. The empirical findings of the model, concur that BTC mining might create an electricity demand capable of producing enough emissions to exceed 2 °C of global warming in less than the next three decades. The study also supports a BTC mining mobility from geographic regions where electricity is expensive to regions with low electricity costs (e.g. where the cost of electricity from renewables is cheaper than fossil fuels) due to its decentralized technological encrypted nature.

In another study, Foteinis (2018), estimates that the annual carbon footprint for bitcoin and ethereum mining exceeds 43.9 million tonnes of carbon dioxide equivalent, while Krause and Tolaymat (2018) argue that during January 2016 to June 2018, mining for four cryptocurrencies (Bitcoin, Ethereum, Litecoin, and Monero) is responsible for 3–15 million tonnes of carbon dioxide emissions. Subsequent work by Stoll et al. (2019) develops a novel methodology for estimating the carbon dioxide releases associated with BTC mining. This technique uses data from various sources including *inter alia* IPO filings of major hardware manufacturers, mining facility operations, mining pool compositions, and localization of IP addresses. The empirical results postulate Bitcoin's annual electricity consumption up to 45.8 TWh, unleashing approximately 23 million tones of CO₂. The study argues that this level can be compared to the annual carbon emissions of Jordan and Sri Lanka.

Based on the above, it is evident that the underlying literature is still incomplete in at least two aspects. First, the majority of the existing studies focus on the measurement of electricity consumption generated by BTC mining. Even though some studies, employ quantile regression techniques, and in this sense are more flexible than least squares which focus exclusively on the

conditional mean (Tsonas, 2020a), they fail to unravel structural relationships between the sample variables. This study accounts for this limitation, by relying on cointegration. Second, a small number of emerging studies have stressed their attention to the measurement of the carbon footprint of BTC mining. Although many of them use advanced techniques, they miss a causal connection between BTC market characteristics and carbon emissions associated with its mining. This happens since they simply measure the amount of carbon dioxide emitted in the atmosphere, not its generating process.

3. Model

For a regression model

$$y_t = \mathbf{x}'_t \boldsymbol{\beta} + u_t, Q_\tau(u_t | \mathbf{x}_t) = 0, t = 1, \dots, T, \quad (1)$$

where $\mathbf{x}_t \in \mathbb{R}^k$ is a vector of covariates with coefficients $\boldsymbol{\beta} \in \mathbb{R}^k$, error term u_t , and $Q_\tau(u_t | \mathbf{x}_t)$ representing the conditional τ -quantile of u_t given \mathbf{x}_t , Koenker and Bassett (1978) showed that the τ -regression quantile is any $\hat{\boldsymbol{\beta}}$ that solves

$$\min_{\boldsymbol{\beta}} T^{-1} \sum_{t=1}^T [\tau - \mathbb{I}(y_t < \mathbf{x}'_t \boldsymbol{\beta})] [y_t - \mathbf{x}'_t \boldsymbol{\beta}], \quad (2)$$

where $\mathbb{I}(\cdot)$ is the indicator function, see also Koenker and Bassett (1982). Koenker and Xiao (2006) considered quantile autoregressive models (QAR) of the following form.

$$y_t = \theta_0(U_t) + \theta_1(U_t)y_{t-1} + \dots + \theta_L(U_t)y_{t-L}, \quad (3)$$

where $\{U_t\}$ is a sequence of iid standard uniform random variables, and $\theta_j: [0,1] \rightarrow \mathbb{R}$ are unknown functions ($j = 0, 1, \dots, L$). The right-hand-side of (3) should be a monotonic function of U_t in the relevant domain of y_t . The model can also be expressed as follows.

$$y_t = \mu_0 + \alpha_{1,t}y_{t-1} + \dots + \alpha_{L,t}y_{t-L} + u_t, \quad (4)$$

where $\mu_0 = \mathbb{E}\theta_0(U_t)$, $u_t = \theta_0(U_t) - \mu_0$, $\alpha_{l,t} = \theta_l(U_t) - \mu_0$ ($l = 1, \dots, L$). Therefore $\{u_t\}$ is an iid

sequence of random variables with distribution function $F(\cdot) = \theta_0^{-1}(\cdot + \mu_0)$.

Clearly, one can extend (3) as follows:

$$y_t = \theta_0(U_t) + \theta_1(U_t)y_{t-1} + \cdots + \theta_L(U_t)y_{t-L} + v_t, \quad (5)$$

where $v_t \sim^{iid} \mathcal{N}(0, \sigma_v^2)$.

In this paper, we are interested in Cointegrated Quantile Vector Autoregressive (CQVAR) Models to examine the possibility of cointegration across different quantiles. Therefore, suppose $\mathbf{y}_t = [y_{t,1}, \dots, y_{t,n}]'$ represents an $n \times 1$ multivariate time series which has the Vector Error Correction Model (VECM) representation:

$$\Delta \mathbf{y}_t = B \Delta \mathbf{y}_{t-1} + \Lambda \mathbf{y}_{t-1} + \mathbf{u}_t, t = 1, \dots, T, \quad (6)$$

where $\Delta \mathbf{y}_t$ represents the first differences (i.e. $\Delta \mathbf{y}_t = \mathbf{y}_t - \mathbf{y}_{t-1}$), B and Λ are $n \times n$ matrix of unknown coefficients, and $\mathbf{u}_t \sim^{iid} \mathcal{N}(\mathbf{0}, \Omega)$. This form is fully general as it can accommodate arbitrary short-run dynamics of the form: $\Delta \mathbf{y}_t = \sum_{l=1}^L B_l \Delta \mathbf{y}_{t-l} + \Lambda \mathbf{y}_{t-1} + \mathbf{u}_t$, where B_l is $n \times n$. We use a triangular representation for the VECM.

For our purposes, we have the following general form:

$$\Delta y_{t,i} = \mu_i + \alpha_t + \sum_{l=1}^L \beta_{i,l}' \Delta \mathbf{y}'_{t-l} + \lambda_i (\boldsymbol{\gamma}_i' \mathbf{y}_{t-1}) + e_{t,i}, i \in \{1, \dots, n\}, t = 1, \dots, T, \quad (7)$$

where μ_i, α_t represent, respectively, country and time effects, $\boldsymbol{\gamma}_i' \mathbf{y}_{t-1}$ represents errors from the long-run or steady-state cointegrating equation, λ_i s are coefficients which should be between -1 and 0 if there is, indeed, cointegration. Our prior assumptions are as follows:

$$\mu_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_\mu^2), i = 1, \dots, n, \alpha_t \stackrel{iid}{\sim} \mathcal{N}(0, \sigma_\alpha^2), t = 1, \dots, T, \quad (8)$$

$$\boldsymbol{\beta}_i \equiv [\boldsymbol{\beta}_{i,l}, l = 1, \dots, L] \sim \mathcal{N}_{nL}(\bar{\boldsymbol{\beta}}, \Sigma_{\boldsymbol{\beta}}), i = 1, \dots, n, \quad (9)$$

$$\lambda = [\lambda_i, i = 1, \dots, n]' \sim \mathcal{N}_n(\bar{\lambda}, \sigma_{\lambda}^2). \quad (10)$$

Let us define

$$\bar{\boldsymbol{\theta}} = [\bar{\boldsymbol{\beta}}', \bar{\boldsymbol{\gamma}}', \bar{\lambda}]'. \quad (11)$$

The components in (8) - (10) define a local prior. In turn, the global prior (which is used to “borrow strength” from the availability of panel data) is as follows:

$$\bar{\boldsymbol{\theta}} \sim \mathcal{N}_d(\boldsymbol{\theta}^*, \Omega_{\boldsymbol{\theta}}^*), \quad (12)$$

where the dimensionality $d = nL + 2n$, $\boldsymbol{\theta}^*$ represents the global mean and $\Omega_{\boldsymbol{\theta}}^*$ represents the global prior covariance matrix. To simplify prior selection we set $\boldsymbol{\theta}^* = c\mathbf{1}_d$, where $\mathbf{1}_d$ is a vector of ones in \mathbb{R}^d , c is a constant that we set to zero, $\Omega_{\boldsymbol{\theta}}^* = h\mathbf{I}_d$, where \mathbf{I}_d is the $d \times d$ identity matrix and $h > 0$ is a scalar that we set to 100 so that the local-global prior is diffuse. For the scale parameters, we have the (standard) priors

$$p(\sigma_j) \propto \sigma_j^{-(\bar{N}+1)} e^{-\bar{Q}/(2\sigma_j^2)}, j \in \{\mu, \alpha, \gamma, \lambda\}, \quad (13)$$

where $\bar{N}, \bar{Q} \geq 0$ are parameters of the prior. We set $\bar{N} = 1$ and $\bar{Q} = 0.001$ to represent “knowing little”. For $\Sigma_{\boldsymbol{\beta}}$ the prior is

$$p(\Sigma_{\boldsymbol{\beta}}) \propto |\Sigma_{\boldsymbol{\beta}}|^{-(\bar{N}+nL+1)/2} e^{-(1/2)\text{tr}\bar{\mathbf{A}}\Sigma_{\boldsymbol{\beta}}^{-1}}, \quad (14)$$

see Zellner (1971, p. 227, equation 8.15) where $\bar{\mathbf{A}}$ is a symmetric $nL \times nL$ matrix containing prior parameters, that we set to $\bar{\mathbf{A}} = c_A \mathbf{I}_{nL}$, c_A is a scalar that we set to 0.001, and $\text{tr}(\cdot)$ represents the trace operator.

We intend to perform extensive sensitivity analysis with respect to c, c_A, \bar{N}, \bar{Q} and h .

The only component of the prior that we intend to change at this stage is the following. When $\lambda_i = 0$ then $\boldsymbol{\gamma}_i$ is not identified as it does not appear in (7). Therefore, we assume

$$\boldsymbol{\gamma} = [\boldsymbol{\gamma}_i, i = 1, \dots, n]' | \boldsymbol{\lambda} \sim \mathcal{N}_n(\bar{\boldsymbol{\gamma}}, \Sigma_{\boldsymbol{\gamma}} \text{diag}(\boldsymbol{\lambda})), i = 1, \dots, n, \quad (15)$$

where we set $\bar{\boldsymbol{\gamma}} = c_{\boldsymbol{\gamma}} \mathbf{1}_n$ with $c_{\boldsymbol{\gamma}} = 0.1$, $\Sigma_{\boldsymbol{\gamma}} = h_{\boldsymbol{\gamma}} \mathbf{I}_n$, and $h_{\boldsymbol{\gamma}} = 0.1$. Moreover, $\text{diag}(\boldsymbol{\lambda})$ represents an $n \times n$ diagonal matrix containing the elements of $\boldsymbol{\lambda}$ along the main diagonal.

We intend to perform extensive sensitivity analysis with respect to $c, c_A, c_{\boldsymbol{\gamma}}, \bar{N}, \bar{Q}, h, h_{\boldsymbol{\gamma}}$.

Using as point of departure the Koenker and Xiao (2006) formulation in (5), we reformulate (7) as follows.

$$\begin{aligned} \Delta y_{t,i} &= \mu_i(U_t) + \alpha_t(U_t) + \sum_{l=1}^L \boldsymbol{\beta}_{i,l}'(U_t) \Delta \mathbf{y}'_{t-l} \\ &+ \lambda_i(U_t) (\boldsymbol{\gamma}_i'(U_t) \mathbf{y}_{t-1}) + e_{t,i}, i \in \{1, \dots, n\}, t \\ &= 1, \text{ots}, T, \end{aligned} \quad (16)$$

where $\{U_t\}$ is a sequence of iid standard uniform random variables, and $\mu_i, \alpha_t, \lambda_i: [0,1] \rightarrow \mathbb{R}$ and $\boldsymbol{\beta}_{i,l}: [0,1] \rightarrow \mathbb{R}^L$ are unknown functions. In simplified notation we can express (16) as

$$y_t = \mathbf{x}_t' \boldsymbol{\beta}(U_t) + e_t, t = 1, \dots, T, \quad (17)$$

where $\mathbf{x}_t \in \mathbb{R}^{nL+2+n}$ represents the vector containing the right-hand-side variables in (16), and $e_t \sim^{iid} \mathcal{N}(0, \sigma_v^2)$, where the prior of σ_v is the same as in (13). To specify the functions $\boldsymbol{\beta}(\cdot)$ we assume that they also depend on \mathbf{x}_t and they are monotonic functions of the form

$$\beta_j(U, \mathbf{x}_t) = \sum_{g=1}^G \delta_{jg} \Phi(\mathbf{x}_t' \boldsymbol{\gamma}_{jg} + U \eta_{jg}), j = 1, \dots, \text{dim}(\boldsymbol{\beta}), \quad (18)$$

where $\boldsymbol{\gamma}_{jg} \in \mathbb{R}^{\text{dim}(\mathbf{x}_t)}$, $\text{dim}(\mathbf{x}_t) = \text{dim}(\boldsymbol{\beta}) = nL + 2 + n$, η_{jg} are unknown parameters, parameters $\delta_{jg} > 0$ are ordered from low to high, and $\Phi(z) = \frac{1}{1+e^{-z}}$, $z \in \mathbb{R}$ is the sigmoid (logistic distribution) activation function. This formulation is a neural network (NN) whose

approximation properties as the number of nodes (G) increases, are known to be excellent (Gallant and White, 1988, Stinchcombe and White, 1989,1990, White, 1989). In turn, the no-quantile-crossing-property (NQCP) can be easily enforced using the procedure of Chernozhukov et al. (2010), see also footnote 5 in Koenker and Xiao (2006).

We define our structural parameter vector as

$$\begin{aligned} \boldsymbol{\theta} &= [\{\delta_{jg}\}, \{\gamma_{jg}\}, \{\eta_{jg}\}, j = 1, \dots, \dim(\mathbf{x}_t), g = 1, \dots, G], \boldsymbol{\theta} \in \Theta \\ &\subset \mathbb{R}^d \end{aligned} \tag{19}$$

where $d = \dim(\boldsymbol{\theta}) = G \cdot \dim(\mathbf{x}_t)$. To select the number of nodes (G) in the NN, we use the marginal likelihood or evidence defined as

$$\mathfrak{M}_G(\mathbf{y}) = \int_{\Theta} L_G(\boldsymbol{\theta}; \mathbf{y}) p_G(\boldsymbol{\theta}) d\boldsymbol{\theta}, \tag{20}$$

where $L(\boldsymbol{\theta}; \mathbf{y}), p(\boldsymbol{\theta})$ denote, respectively, the likelihood function and prior, and \mathbf{y} denotes the entire data set. All quantities are indexed by $G \in \{1, \dots, \bar{G}\}$ where \bar{G} is an upper bound that we set to $\bar{G} = 10$. In turn, we select the value of G with the maximum value of $\mathfrak{M}_G(\mathbf{y})$. The marginal likelihood in our context is estimated using the procedure of Perrakis et al, (2014). Alternatively, but equivalently, we can select the value of G that maximizes the value of the Bayes factor in favor of G nodes and against a single node:

$$\mathfrak{B}_G(\mathbf{y}) = \frac{\mathfrak{M}_G(\mathbf{y})}{\mathfrak{M}_1(\mathbf{y})}, G = 2, \dots, \bar{G}. \tag{21}$$

The values of the Bayes factor are reported in panel (d) of Figure B.1 in Appendix B from which it turns out that the optimal value of G is three. To provide statistical inferences we use a Gibbs sampler with data augmentation in (17) and (18) using 150,000 iterations the first 50,000 of which are discarded to mitigate possible start-up effects. MCMC behavior and sensitivity analysis of the benchmark results are undertaken in Appendix B.

4. Data and variables

The sample consists of a data set of 51 countries ($N = 51$) over July 02, 2016 to November 30, 2018 yielding a balanced panel of 44,982 observations.⁴ The sample period was strictly dictated by data availability of energy consumption levels for the sample countries as appeared in Stoll et al, (2019). Though the time-span is relatively short, this period has witnessed significant volatility since a steep BTC price hike (December 2017) was followed by a declining (corrective) trend throughout 2018. The description of variables used in the analysis along with their derivation is exhibited in Table 1. Similarly to Das and Dutta (2020), the variables are logarithmically differenced.

<Insert Table 1 about here>

Table 2 presents the summary statistics and the correlation matrix of the sample variables. As it is evident, from the inspection of the relevant table, we observe that the daily BTC return (RET) is equal to 0.2% showing the lowest standard deviation among the sample variables (0.041). The relevant variable is negatively skewed (-0.332) and the excess kurtosis suggests a leptokurtic distribution ($7.781 > 3$). The same findings apply to the BTC transaction cost variable (COST), while we notice rejection of normality in all variables as in many other similar studies (Das and Dutta, 2020; Koutmos, 2018; Katsiampa, 2017). The rest of the variables accounting for BTC's microstructure display the standard properties of daily asset returns as they are heavy-tailed though positively skewed (Gerlach, et Al., 2011), while the dependent variable (CO2) follows also a leptokurtic distribution. Finally, the correlations matrix presented in Panel B of Table 2, indicates that none of the independent variables are highly correlated avoiding possible problems of multicollinearity.

⁴ The countries included in the analysis are provided in the Appendix A.

<Insert Table 2 about here>

To examine the relationship between Bitcoin's market microstructure and its carbon dioxide releases, we estimate by using Bayesian quantile regression analysis the following model expressed in its algebraic form:

$$CO2_{j,t} = f((RET_t), (REV_t), (FEES_t), (TRAN_t), (EFF_t), (COST_t)) + e_{j,t} \quad (22)$$

where $CO2_{j,t}$ is the carbon dioxide emissions of BTC in country j and period t , RET is the BTC daily returns (common for every country), REV is the miner's revenue (common for every country), FEES denotes the total BTC value of transaction fees miners earn per day (common for every country), TRAN shows the historical total number of BTCs which have been mined daily (common for every country), EFF stands for the mining efficiency (common for every country), COST denotes miners revenue (common for every country), a_j is a set of country dummy variables, and a_s is a set of seasonal dummy variables. Finally, $e_{j,t}$ is the i.i.d. error term.

5. Results and discussion

Table 3, displays the posterior means and standard deviation of the distributions per quantile drawn from the two estimated models (standard Bayesian quantile vs Bayesian CQVAR). The means of the variables RET, REV, FEES, and EFF appear to be much lower in the classical Bayesian regression analysis for nearly all the quantiles, whereas the opposite holds for the rest of the variables.

<Insert Table 3 about here>

To draw sharp inferences, we proceed with the discussion of empirical findings drawn from the standard Bayesian quantile regression without cointegration. Figure 3, depicts the marginal posterior distributions at various quantiles (τ). From the careful inspection of the relevant figure,

some general results emerge. First, it is evident that in most of the quantiles, the distributions of the variables are bimodal and, in few cases, appear to be multimodal. Second, in many variables (RET, REV, TRAN and COST) the marginal posterior densities do not differ substantially by quantile. This outcome seems to contradict the subsequent results from the Bayesian CQVAR, where we trace significant heterogeneity between the quantiles of the posterior distributions. In terms of daily returns (RET), the standard Bayesian quantile model has a much larger concentration of posterior probability around 0.8 for nearly all of the quantiles, except for the extreme ones ($\tau = 0.05$ and $\tau = 0.90$), where BTC returns average near 0.71 and 1.2 respectively (see Table 3). Similar results apply by the careful consideration of posterior densities for the BTC miner's revenues (REV). In this case, we observe that for most of the quantiles, there is a significant concentration of posterior probability around the value of -0.26, while the shape of the distributions is bimodal. However, for the 5% quantile, there is strong evidence of a multimodal PDF ranging from about -0.37 to -0.3, with a posterior mean equal to -0.29. For the rest of the variables, the standard Bayesian quantile regression analysis yields systematically positive values compared to REV.

<Insert Figure 3 about here>

Having reported the marginal posterior densities from the standard Bayesian quantile regression analysis, we proceed with the CQVAR analysis. As it is evident from Figure 4, the marginal posterior densities for most of the covariates show a positive effect with the carbon releases in nearly all the estimated quantiles. Specifically, the distributions of BTC's daily returns (RET) are mostly multimodal (see upper left part of the figure) in nearly all of the quantiles, except for some countries in lower quantiles (5% and 10%). We also notice that the PDFs in all of the

variables, depart from normality, revealing that asymptotic-based inferences do not apply in this instance (Tsionas, 2020b).

The marginal posterior distribution of the coefficient of RET in the lowest quantile ($\tau=0.05$) seems to exhibit the highest spread, while the posterior mean ranges from 0.8 to 1.2 approximately. This means that when carbon dioxide releases are low a one-unit increase (decrease) in the BTC returns, increases (decreases) on average the carbon releases stemmed from BTC mining at the same rate. A different picture emerges when we examine the 50th quantile (median). A careful look shows a multimodal distribution characterized by a smaller standard deviation since its posterior mean fluctuates roughly from 0.9 to 1.05.

On the contrary, the effect of REV seems to be negative in all of the examined quantiles, while for the 50%, 75%, 90%, and 95% quantile the marginal posteriors are multimodal. This effect is more pronounced for countries in the 75% quantile since the posterior mean ranges within the relatively broad interval [-0.45 to -0.1]. The negative effect of miner's revenues on carbon footprint can be compared with the main finding of Das and Dutta (2020), where energy consumption is negatively correlated with miner's revenues especially in low quantiles of the distribution. As it is obvious, from the upper right part of the figure, the range of the posterior mean is broader in higher than in lower quantiles of the PDF of the dependent variable, suggesting that the lower (higher) miner's Bitcoin revenues the more abrupt (gradual) the effect on environmental degradation.

The marginal posterior distributions of parameters of BTC transaction fees (FEES) are also multimodal, and there is evidence of a positive effect on the BTC carbon releases for all the sample countries across the quantiles. Similar findings apply to the parameters of the number of BTC transactions (TRAN), where the posterior mean at the lowest quantile ($\tau=0.05$) of the distribution,

ranges from 0.25 to 0.4 and averages around 0.20. This finding indicates, that when environmental degradation is at low levels (e.g. earlier stages of the BTC market development), a 10% increase in the number of mined BTCs, results in an almost 2% increase of the emitted carbon releases. The level of mining efficiency (EFF) contributes positively to BTC carbon releases across all quantiles, though this effect is stronger at the 90% quantile since the posterior mean of the distributions, which is equal to 0.9741 is maximized (see also Table 3). We notice also, a positive impact of transaction cost (COST) to BTC carbon emissions since the posterior mean is positive, though less strong at higher quantiles ($\tau = 0.95$).

From the above analysis, all measures drawn from the estimation of the Bayesian CQVAR differ substantially by quantile, revealing a significant degree of heterogeneity between the quantiles of the posterior distributions. This outcome is less pronounced for the marginal posterior densities of the standard Bayesian quantile regression model. In this sense, estimating a Bayesian cointegrated vector error correction model makes a significant difference as opposed to standard quantile regression.

<Insert Figure 4 about here>

Evidence on the existence of cointegration is provided by the marginal posterior densities of the estimated error correction terms (λ_i s) drawn from the Bayesian CQVAR. As can be observed from Figure 5, the posterior densities are concentrated around negative values as expected. From this evidence, it turns out that we have cointegration. Moreover, the PDFs are unimodal except for the 10th, 25th and 90th quantile where the distributions are bimodal. In these cases, the posterior mean ranges from -0.222 to -0.218, -0.222 to -0.220, and -0.222 to -0.217 respectively. The relevant densities express the speed of adjustment to the long-run equilibrium. For instance, for the 50% quantile (median), which averages around -0.22, we argue that in the case we are off the

long-run equilibrium, carbon releases from BTC mining, adjust towards their long-run level with about 22% of this adjustment taking place within the first day. It is noteworthy that the error-correction term lies within a narrow interval ranging from 21.5% to 22.5%. This finding suggests a relatively slow adjustment process toward the steady-state equilibrium.

<Insert Figure 5 about here>

Given that we have cointegration, two questions arise i) Is there a single cointegrating vector or not? and ii) What is / are the (long-run or steady-state) cointegration vector(s)? Evidence on the number of cointegrating vectors is provided by the Bayes factor in favor of $c > 1$ vectors relative to $c = 1$. The posterior probability of $c > 1$ vectors against $c = 1$ can be computed easily using the Bayes factor:

$$P_c(\mathbf{y}) = \frac{\mathfrak{B}_c(\mathbf{y})/\mathfrak{B}_1(\mathbf{y})}{\sum_{c'=1}^C \mathfrak{B}_{c'}(\mathbf{y})/\mathfrak{B}_1(\mathbf{y})}, c = 1, \dots, C, \quad (23)$$

where, in our application, $C = 7$. The posterior probabilities and the marginal posterior densities of the single cointegrating vector are reported in panels (a) and (b)–(g) of Figure 6.⁵ In the upper part of Figure 6, we present the posterior probability of cointegrating rank, while marginal posterior densities for the cointegrating coefficients (“*long-run*” estimates) by different quantile, are reported in the lower panel of the figure. As it is evident, there is a single cointegrating vector since the posterior probability exceeds 0.9 when the cointegrating rank equals unity ($c=1$) and (almost) zero in all of the other cases ($c=2,3,4,5,6$ and 7). Finally, the marginal posterior densities of the long-run coefficients thought different compared to their short-run counterparts, follow similar trends, indicating substantial heterogeneity across the quantiles.

⁵ To compute the Bayes factors and, therefore, the posterior probabilities of cointegrating rank $c > 1$ we use the same procedure as in the Bayesian panel cointegration study of Koop, Leon-Gonzalez and Strachan (2006) based on the Savage - Dickey density ratio (Verdinelli and Wasserman, 1995).

<Insert Figure 6 about here>

6. Concluding remarks

Despite the strong interdependence between energy, carbon dioxide emissions, and blockchain technology, their dynamics, and economic interlinkages have been scarcely investigated by previous works. This study stresses the importance of using BTC microstructure variables to explain the carbon emissions associated with its mining. We develop a CQVAR model organized around MCMC techniques for numerical Bayesian inference extending Koenker and Xiao (2006).

The empirical results portray some important differences between the traditional Bayesian quantile model and the Bayesian CQVAR, in terms of several aspects of the data. Specifically, we document considerable heterogeneity among different quantiles justifying the validity of Bayesian CQVAR. Moreover, by relying on cointegration analysis developed within a Bayesian quantile framework, we distinguish between short-run and long-run responses of the BTC financial variables across the different quantiles of the PDFs. We argue that most of the BTC microstructure variables exert a positive relationship across the quantiles with the level of carbon emitted in the atmosphere from the BTC mining. However, there is a negative relationship between miner's revenues and BTC carbon footprint, which is more pronounced at higher quantiles. We also find that the error-correction term of the cointegrating vector does not exceed the value of 22.5% at the 95% quantile, indicating a slow speed of adjustment toward the steady-state equilibrium.

Our findings will be useful for cryptocurrency energy management, in terms of pursuing suitable environmental policies associated with its usage. This study, argues that a sustainable energy strategy focusing on the penetration of renewable energy sources along with the use of energy-efficient mining hardware will alleviate the carbon footprint of the Bitcoin.

Acknowledgments

The authors would like to thank Christian Stoll, Lena Klaaßen, and Ulrich Gellersdorfer for kindly providing their data that appeared in Stoll et al, (2019). Special thanks also go to Symeoni-Eleni Soursou for providing valuable editorial assistance. All errors are the authors alone. The usual disclaimer applies.

Appendix A

Table A.1. Countries included in the sample

Abbreviation	Country
US	United States of America
VE	Venezuela
RUS	Russia Federation
KOR	Korea Republic
UA	Ukraine
CN	China
IR	Iran
RO	Romania
MY	Malaysia
ZA	South Africa
BG	Bulgaria
TH	Thailand
HU	Hungary
VN	Vietnam
TR	Turkey
KZ	Kazakhstan
ES	Spain
CZ	Czech Republic
DE	Germany
ID	Indonesia
SW	Sweden
IT	Italy
CA	Canada
UK	United Kingdom
AU	Austria
NL	Netherlands
PO	Poland
MN	Mongolia
LV	Latvia
BR	Brazil
MX	Mexico
PT	Portugal
FR	France
CS	Serbia and Montenegro
LT	Lithuania
CO	Colombia
NO	Norway
GR	Greece
PE	Peru
JP	Japan
HR	Croatia
MA	Morocco
IN	India
PK	Pakistan
SK	Slovak Republic
PH	Philippines
SD	Sudan
LY	Libyan Arab Jamahiriya
SI	Slovenia
MD	Moldova
AR	Argentina

Appendix B - MCMC performance and posterior sensitivity analysis

Our MCMC procedure for (17) and (18) is based on the following observations:

i) Given $\{\boldsymbol{\gamma}_{jg}, \eta_{jg}\}$ the model is linear in the remaining parameters so we can use a standard Gibbs sampling update.

ii) For the nonlinear parameters $\{\boldsymbol{\gamma}_{jg}, \eta_{jg}\}$ we use an Accept-Reject Metropolis-Hastings (ARMH, Chib and Jeliazkov, 2005) algorithm instead of a random-walk Metropolis-Hastings. This results in better mixing properties and better acceptance rates.

iii) For the latent variables $\{U_t\}$ we use a Metropolis-Hastings update. Specifically, given the current MCMC draw, say $U_t^{(s)}$, we draw a candidate U_t^c uniformly distributed in $(U_t^{(s)} - H_t, U_t^{(s)} + H_t)$ where H_t is adjusted during the burn-in phase so that the average acceptance rate ranges from 15 - 25%. Of course, U_t^c must be in the interval (0,1) as their prior is a standard uniform distribution. In turn, the candidate is accepted with the Metropolis-Hastings probability $\min \left\{ 1, \frac{p(U_t^c | \cdot, \mathbf{y})}{p(U_t^{(s)} | \cdot, \mathbf{y})} \right\}$ where $p(U | \cdot, \mathbf{y})$ denotes the conditional posterior density of U_t conditional on all other parameters and latent variables (and, of course, the data, \mathbf{y}).

To assess the performance of MCMC we rely on relative numerical efficiency (RNE, Geweke, 1992) which should close to one if iid sampling from the posterior were possible. We present densities of median RNE for latent variables (U_t) and structural parameters in (18) in Figure B.1, panel (a).

<Insert Figure B.1 about here>

To perform sensitivity analysis concerning prior assumptions about $c, c_A, \bar{N}, \bar{Q}, h$, we vary these parameters as follows:

- i) c and c_A are drawn randomly from a Student- t distribution with zero location, 5 degrees of freedom and scale parameter 10.
- ii) The logs of \bar{N}, \bar{Q}, h are drawn randomly from a Student- t distribution with zero location, 5 degrees of freedom and scale parameter 10.

We generate 10,000 such priors and we repeat posterior analysis using the Sampling-Importance-Resampling (SIR) algorithm, see Rubin (1987, 1988), and Smith and Gelfand (1992). The MCMC batch required for SIR is set to 20% of the final benchmark MCMC sample which is 20,000 draws. We report percentage deviations in posterior means and percentage deviations in posterior standard deviations in panel (b) of Figure A.1. As these deviations are fairly trivial relative to benchmark posterior means and posterior standard deviations we can be relatively confident that our benchmark prior is fairly representative of what would happen had we adopted different priors in a family of (very) different priors.

Tables and Figures

Table 1: Description of the sample variables

Variable	Definition	Description	Formulation	Source
CO ₂	Carbon dioxide emissions	Estimated carbon dioxide emissions for mining BTC daily in kgCO ₂ eq/kWh. This can be calculated if we multiply the electricity load (ELE) with the average emission factor (AVEF).	$CO2 = ELE * AVEF$	Own calculations based on Stoll et al, (2019)
RET	Daily returns	Market daily returns expressed in USD per BTC. The returns are calculated by taking the natural logarithm of the ratio of two consecutive prices (Urquhart, 2016; Katsiampa, 2017)	$RET = \log\left(\frac{P_t}{P_{t-1}}\right) * 100$	Stoll et al, (2019)
REV	Miners revenues	Total miners revenue per day in USD. The revenue equals the number of BTCs mined per day plus transaction fees multiplied by the market price.	$REV = (TRAN + FEES) * RET$	Stoll et al, (2019)
FEES	Transaction fees	Total BTC value of transaction fees miners earn per day in USD.	-	Stoll et al, (2019)
TRAN	Transactions	Total number of BTC mined per day in physical units.	-	Stoll et al, (2019)
EFF	Mining efficiency	Number of hashes in a second (HASH), divided by the consumed electricity load (ELE).	$EFF = HASH/ELE$	Li et al, (2019)
COST	Transaction cost	Total transaction cost of BTC mined per day in USD. It is expressed as miners revenue (REV) divided by the number of transactions (TRAN).	$COST = REV/TRAN$	Das and Dutta, (2020)

Table 2: Summary statistics and correlations (whole sample)

Panel A: Descriptive statistics							
<i>Variables</i>	<i>Mean</i>	<i>Standard Deviation</i>	<i>Median</i>	<i>Minimum</i>	<i>Maximum</i>	<i>Skewness</i>	<i>Kurtosis</i>
<i>CO₂</i>	0.003	0.137	0.000	-3.203	3.342	0.103	36.591
<i>RET</i>	0.002	0.041	0.002	-0.225	0.246	-0.332	7.781
<i>REV</i>	0.0009	0.130	-0.002	-0.502	0.639	0.195	4.473
<i>FEES</i>	-0.0005	0.244	-0.012	-2.484	2.236	0.018	25.202
<i>TRAN</i>	0.0002	0.136	-0.008	-0.561	0.509	0.429	3.834
<i>EFF</i>	0.002	0.138	-0.001	-0.556	0.722	0.185	4.418
<i>COST</i>	0.0007	0.207	0.018	-0.648	1.014	-0.108	3.927
Panel B: Correlation matrix							
<i>Variables</i>	<i>CO₂</i>	<i>RET</i>	<i>REV</i>	<i>FEES</i>	<i>TRAN</i>	<i>EFF</i>	<i>COST</i>
<i>CO₂</i>	1						
<i>RET</i>	-0.0132	1					
<i>REV</i>	-0.3934	0.0241	1				
<i>FEES</i>	0.0181	-0.0164	-0.0966	1			
<i>TRAN</i>	0.3370	0.0363	-0.2180	0.5336	1		
<i>EFF</i>	0.8835	-0.3129	-0.3717	-0.0817	0.2456	1	
<i>COST</i>	-0.4671	-0.0086	0.4700	-0.4095	-0.2905	-0.3938	1

Notes: The table reports summary statistics for the actual data. CO₂ denotes the carbon dioxide releases emitted from BTC mining, RET is the BTC daily returns, REV is the miner's revenue, FEES denotes the total BTC value of transaction fees miners earn per day, TRAN shows the historical total number of BTCs which have been mined daily, EFF stands for the mining efficiency computed as the ratio of the number of hashes in a second, divided by the power is consumed (Li et al, 2019), COST denotes miners revenue divided by the number of transactions. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020).


Table 3: Posterior means and standard deviations for the two models per quantile (τ)

Quantile	Standard Bayesian quantile model	Bayesian CQVAR model
	<i>RET</i>	
$\tau = 0.05$	0.7172 (0.0132)	0.9848 (0.0552)
$\tau = 0.10$	0.7860 (0.0059)	0.9741 (0.0430)
$\tau = 0.25$	0.7649 (0.0030)	0.9684 (0.0245)
$\tau = 0.50$	0.7669 (0.0085)	0.9718 (0.0235)
$\tau = 0.75$	0.7872 (0.0155)	0.9645 (0.0268)
$\tau = 0.90$	1.2005 (0.0202)	0.9513 (0.0381)
$\tau = 0.95$	1.3723 (0.0434)	0.9418 (0.0495)
	<i>REV</i>	
$\tau = 0.05$	-0.2945 (0.0072)	-0.2370 (0.0466)
$\tau = 0.10$	-0.3105 (0.0021)	-0.1848 (0.0356)
$\tau = 0.25$	-0.3005 (0.0037)	-0.1243 (0.0214)
$\tau = 0.50$	-0.2843 (0.0030)	-0.2160 (0.0390)
$\tau = 0.75$	-0.2780 (0.0028)	-0.1976 (0.0309)
$\tau = 0.90$	-0.3226 (0.0079)	-0.2026 (0.0325)
$\tau = 0.95$	-0.3050 (0.0064)	-0.0733 (0.0496)
	<i>FEES</i>	
$\tau = 0.05$	-0.0585 (0.0118)	0.0311 (0.0075)
$\tau = 0.10$	0.0178 (0.0059)	0.0378 (0.0063)
$\tau = 0.25$	-0.0030 (0.0091)	0.0352 (0.0055)
$\tau = 0.50$	0.0328 (0.0084)	0.0410 (0.0061)
$\tau = 0.75$	0.0097 (0.0063)	0.0423 (0.0068)
$\tau = 0.90$	-0.0875 (0.0167)	0.0445 (0.0079)
$\tau = 0.95$	-0.0397 (0.0118)	0.0379 (0.0084)
	<i>TRAN</i>	
$\tau = 0.05$	0.5258 (0.0130)	0.1982 (0.0466)
$\tau = 0.10$	0.4837 (0.0039)	0.1503 (0.0318)
$\tau = 0.25$	0.5851 (0.0030)	0.0981 (0.0210)

$\tau = 0.50$	0.5619 (0.0080)	0.1784 (0.0396)
$\tau = 0.75$	0.5673 (0.0047)	0.1613 (0.0318)
$\tau = 0.90$	0.6319 (0.0109)	0.1587 (0.0307)
$\tau = 0.95$	0.5436 (0.0089)	0.0358 (0.0465)
<i>EFF</i>		
$\tau = 0.05$	0.7894 (0.0134)	0.9603 (0.0139)
$\tau = 0.10$	0.9266 (0.0041)	0.9657 (0.0109)
$\tau = 0.25$	0.9034 (0.0092)	0.9727 (0.0076)
$\tau = 0.50$	0.9602 (0.0057)	0.9719 (0.0081)
$\tau = 0.75$	0.9299 (0.0149)	0.9715 (0.0091)
$\tau = 0.90$	0.7787 (0.0163)	0.9741 (0.0136)
$\tau = 0.95$	0.8547 (0.0130)	0.9641 (0.0170)
<i>COST</i>		
$\tau = 0.05$	0.5619 (0.0191)	0.1696 (0.0463)
$\tau = 0.10$	0.5228 (0.0063)	0.1190 (0.0354)
$\tau = 0.25$	0.5399 (0.0102)	0.0681 (0.0205)
$\tau = 0.50$	0.5219 (0.0096)	0.1573 (0.0384)
$\tau = 0.75$	0.5014 (0.0144)	0.1406 (0.0301)
$\tau = 0.90$	0.1468 (0.0261)	0.1359 (0.0318)
$\tau = 0.95$	-0.0196 (0.0371)	0.0092 (0.0510)

Notes: RET is the BTC daily returns, REV is the miner's revenue, FEES denotes the total BTC value of transaction fees miners earn per day, TRAN shows the historical total number of BTCs which have been mined daily, EFF stands for the mining efficiency, COST denotes miners revenue divided by the number of transactions. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020). Posterior standard errors of estimates are reported in parentheses.

Figure 1: Comparison of the relative search intensity from Google trends (2016-2020)



Panel A: Search intensity for the terms “bitcoin” (in blue), “cryptocurrency” (in red) and “blockchain” (in yellow) as a News item. The maximum value of search activity is normalized to 100, accessed April 6, 2020.



Panel B: Geographic distribution of search intensity for “bitcoin” as a News item.




Panel C: Geographic distribution of search intensity for “cryptocurrency” as a News item.



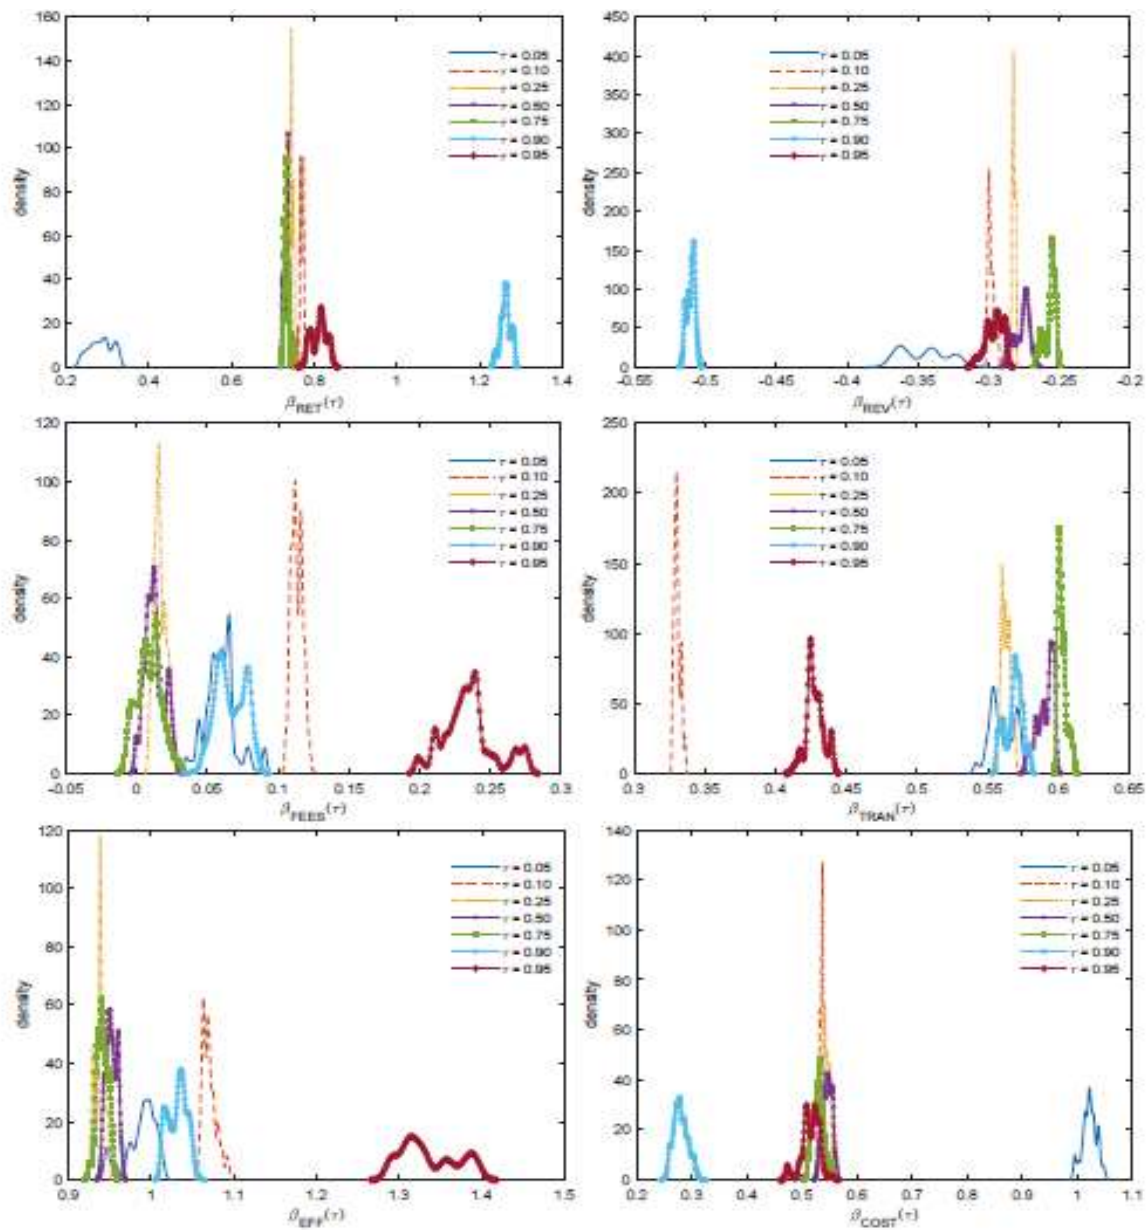
Panel D: Geographic distribution of search intensity for “blockchain” as a News item.

Figure 2: Ten most polluting countries



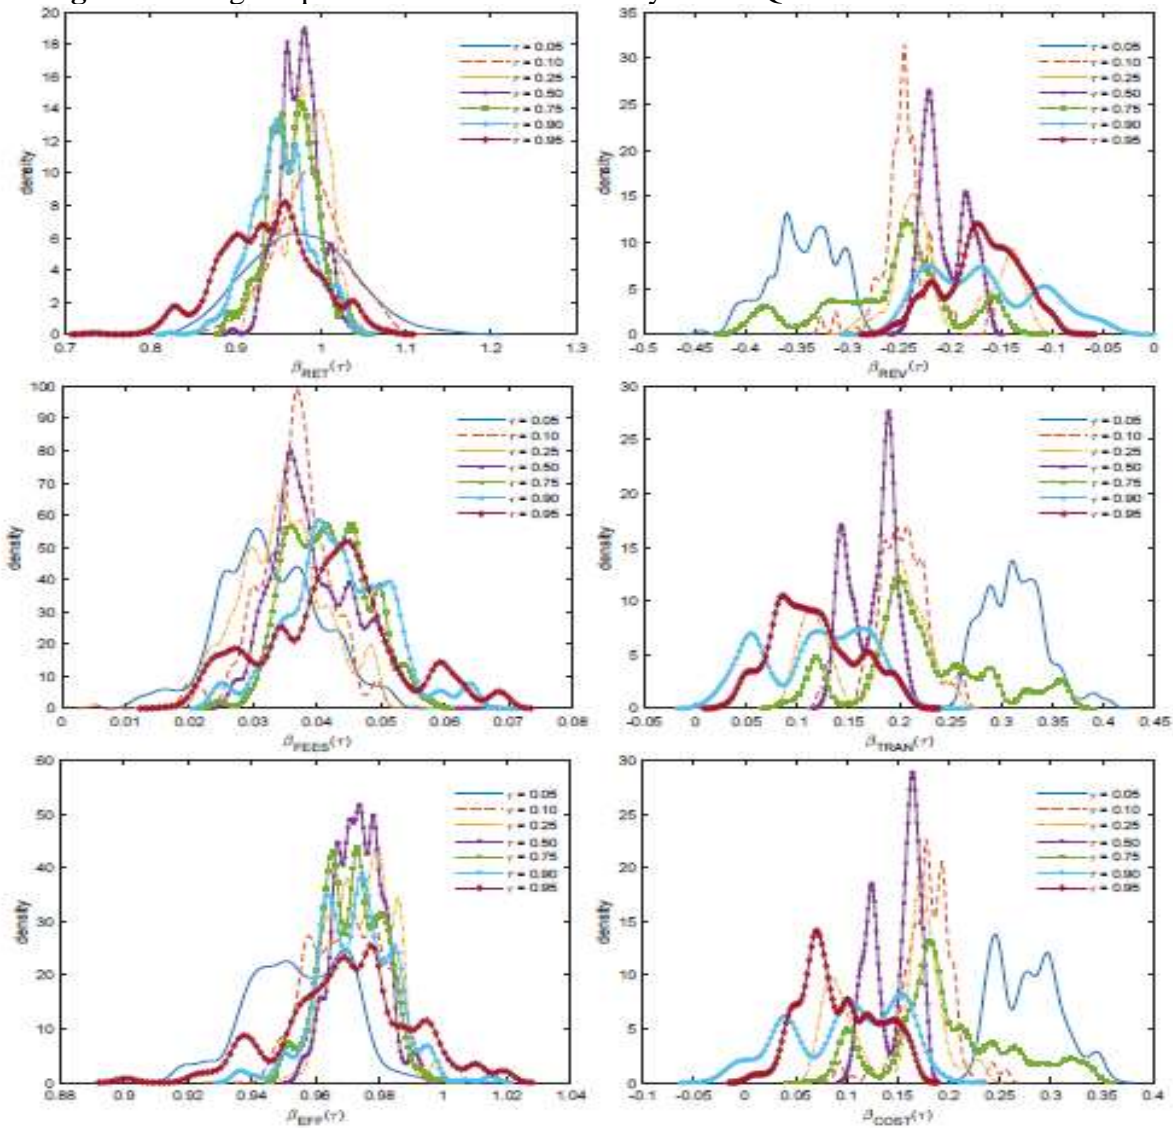
Notes: This figure shows the average carbon dioxide releases from BTC usage in 2018. The carbon releases are measured in kg CO₂ equivalent per kWh. **Source:** Own calculations

Figure 2: Marginal posterior densities from standard Bayesian quantile regression




Notes: This figure displays the marginal posterior densities across several quantiles (τ) drawn from the standard Bayesian quantile regression analysis. RET denotes the BTC daily returns, REV is the miner's revenue, FEES denotes the total BTC value of transaction fees miners earn per day, TRAN shows the historical total number of BTCs which have been mined daily, EFF stands for the mining efficiency and COST denotes miners revenue divided by the number of transactions. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020).

Figure 3: Marginal posterior densities from Bayesian CQVAR



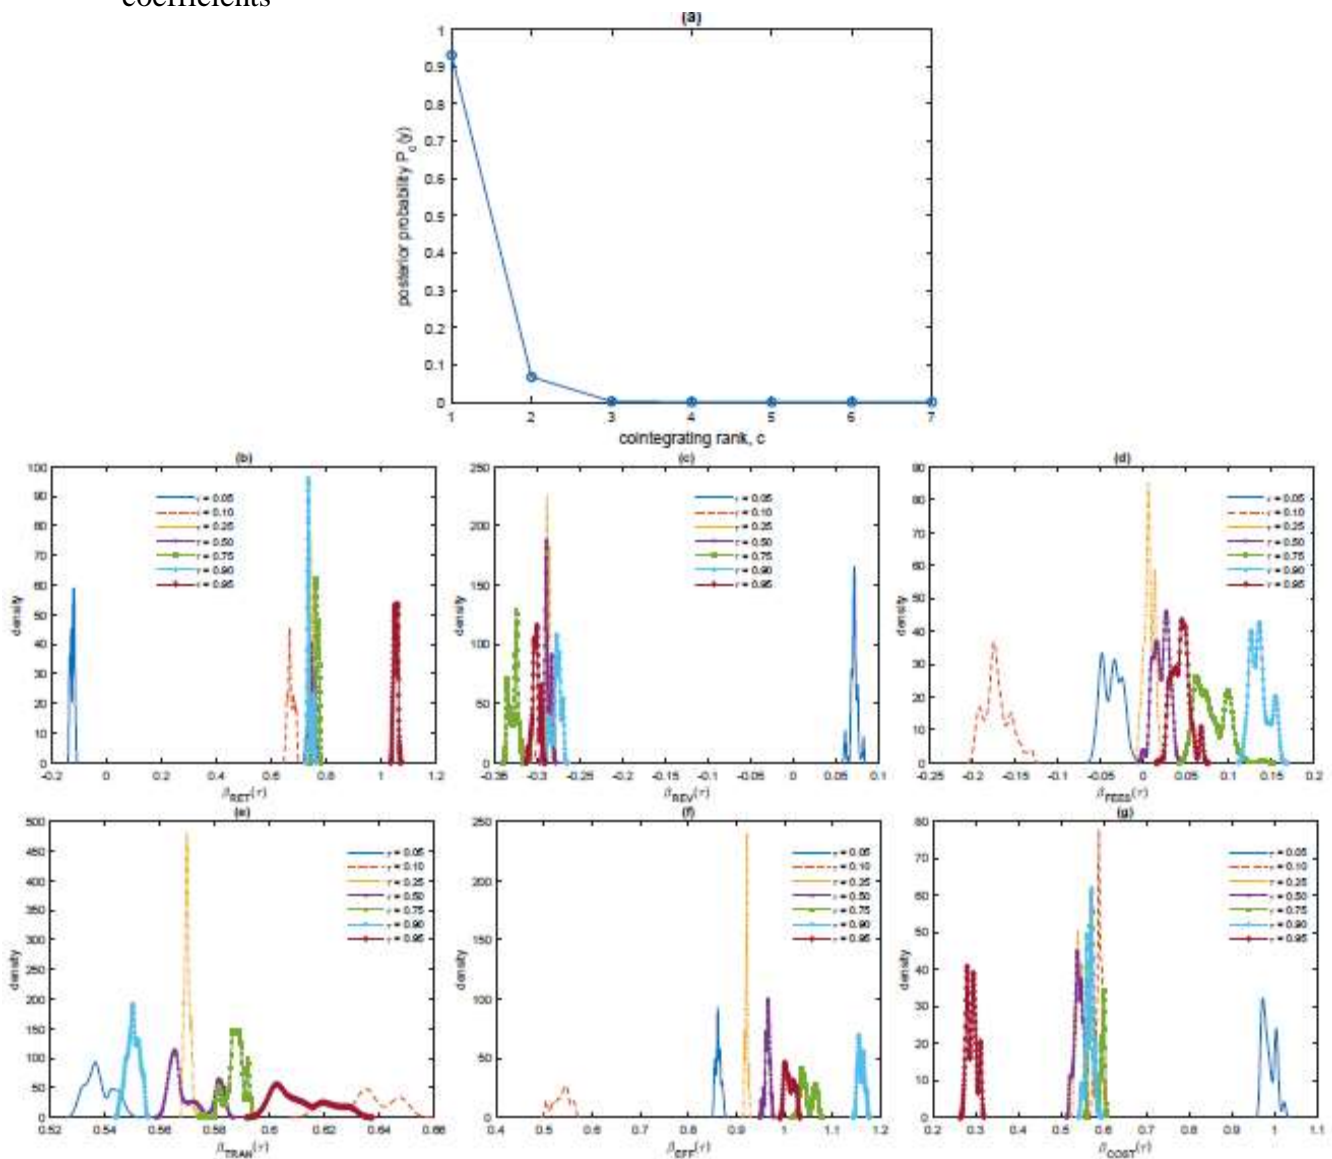
Notes: This figure displays the marginal posterior densities across several quantiles (τ) drawn from the Bayesian cointegrated quantile vector autoregression model. RET denotes the BTC daily returns, REV is the miner's revenue, FEES denotes the total BTC value of transaction fees miners earn per day, TRAN shows the historical total number of BTCs which have been mined daily, EFF stands for the mining efficiency and COST denotes miners revenue divided by the number of transactions. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020).

Figure 4: Marginal posterior densities of λ_{is} from Bayesian CQVAR



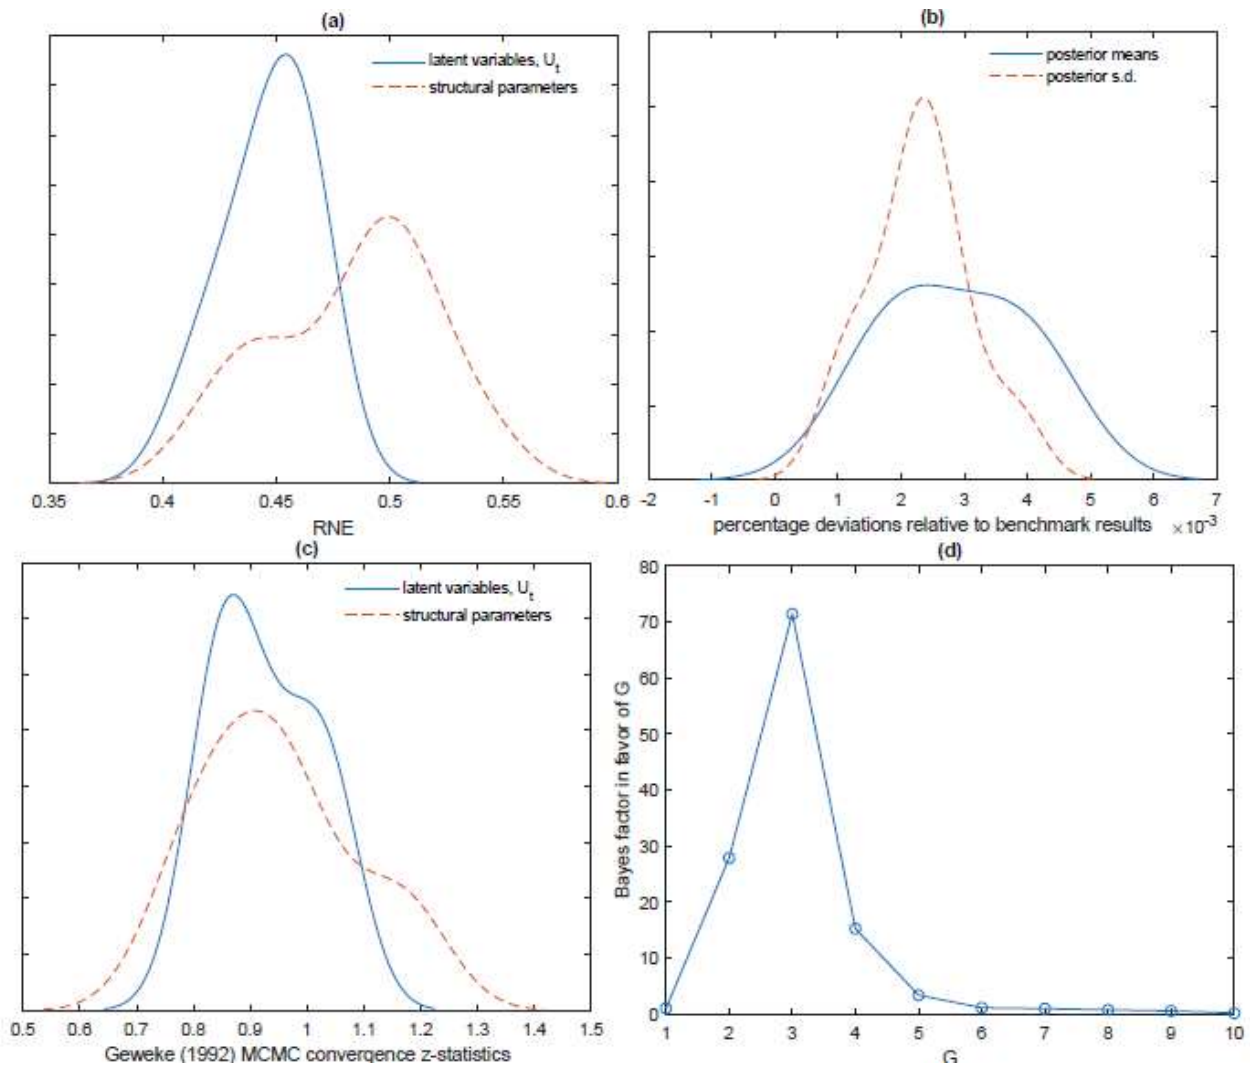
Notes: This figure displays the marginal posterior densities of the estimated quantile parameters drawn from the Bayesian cointegrated quantile vector autoregression model. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020).

Figure 5: Marginal posterior densities of cointegrating rank and cointegrating coefficients



Notes: This figure displays the marginal posterior densities of cointegrating rank and cointegrated coefficients across several quantiles (τ) drawn from the cointegrated quantile vector autoregression model. RET denotes the BTC daily returns, REV is the miner's revenue, FEES denotes the total BTC value of transaction fees miners earn per day, TRAN shows the historical total number of BTCs which have been mined daily, EFF stands for the mining efficiency and COST denotes miners revenue divided by the number of transactions. The data are daily time series. The observations are equal to 44,982. The variables are logarithmically differenced following Das and Dutta (2020).

Figure B.1: MCMC performance and posterior sensitivity analysis



Notes: MCMC convergence diagnostics (viz. Geweke (1992) absolute values of z-statistics, distributed as standard normal as the number of MCMC draws grows to infinity) are reported in panel (c) of Figure A.1.

References

- Atsalakis, G., Atsalaki, I., Pasiouras, F., and Zopounidis, C. (2019). Bitcoin price forecasting with neuro-fuzzy techniques, *European Journal of Operational Research*, 276 (2): 770-780.
- Böhme, R., Christin, N., Edelman, B., Moore, T., (2015). Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*. 29, 213–238.
- Cheng, S.F., Franco, G.D., Jiang, H., and Lin. P. (2019) Riding the Blockchain Mania: Public Firms' Speculative 8-K Disclosures. *Management Science* 65, No. 12 5901–5913.
- Chernozhukov, V., Fernandez-Val, I., Galichon, A., 2010. Quantile and probability curves without crossing, *Econometrica* 78, 1093–1125.
- Chib, S., and Jeliazkov, I. (2005). Accept–reject Metropolis–Hastings sampling and marginal likelihood estimation. *Statistica Neerlandica* 59 (1), 30–44.
- Das, D., and Dutta, A. (2020). Bitcoin's energy consumption: Is it the Achilles heel to miner's revenue? *Economics Letters*, 186: 108530
- De Vries, A. (2018). Bitcoin's growing energy problem. *Joule* 2, 801–805.
- Gallant, A.R. and White, H. (1988). There exists a neural network that does not make avoidable mistakes. Proceedings of the IEEE Second International Conference on Neural Networks (1), 657–664. San Diego: SOS Printing.
- Foteinis, S. (2018). Bitcoin's alarming carbon footprint. *Nature* 554.
- Gerlach, R., Chen, C., and Chan, N. (2011) Bayesian Time-Varying Quantile Forecasting for Value-at-Risk in Financial Markets, *Journal of Business & Economic Statistics*, 29:4, 481-492,
- Geweke, J. (1992). Evaluating the Accuracy of Sampling-Based Approaches to the

Calculation of Posterior Moments. In *Bayesian Statistics 4* (eds. J.M. Bernardo, J. Berger, A.P. Dawid and A.F.M. Smith), Oxford: Oxford University Press, 169-193.

Katsiampa, P., Moutsianas, K., Urquhart, A. (2019). Information demand and cryptocurrency market activity, *Economics Letters*, 185, 108714.

Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models, *Economics Letters*, 158: 3-6.

Koenker, R., and Bassett, G. (1978). Regression quantiles, *Econometrica*, 46, 33–49.

Koenker, R., and Bassett, G. (1982). Robust Tests for Heteroscedasticity based on regression quantiles, *Econometrica*, 50, 43–61.

Koenker, R., and Xiao, Z. (2006). Quantile Autoregression, *Journal of the American Statistical Association*, 101:475, 980–990.

Koop, R., Leon-Gonzalez, R., and R. Strachan (2008). Bayesian Inference in a Cointegrating Panel Data Model. Chib, S., Griffiths, W., Koop, G. and Terrell, D. (Ed.) *Bayesian Econometrics (Advances in Econometrics, Vol. 23)*, Emerald Group Publishing Limited, Bingley, pp. 433–469.

Koutmos, D. (2018). Bitcoin returns and transaction activity, *Economics Letters*, Volume 167: 81-85.

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

Leong, K., and Sung, A. (2018). FinTech (Financial Technology): What is It and How to Use Technologies to Create Business Value in Fintech Way? *International Journal of Innovation, Management and Technology*, 9 (2): 74–78.

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

- Li, J., Li, N., Peng, J., Cui, H., and Wu, Z. (2019). Energy consumption of cryptocurrency mining: A study of electricity consumption in mining cryptocurrencies, *Energy*, 168: 160-168.
- O'Dwyer, K. J. and Malone, D. (2014). Bitcoin mining and its energy footprint. In 25th IET Irish Signals & Systems Conference 2014 and 2014 China-Ireland International Conference on Information and Communities Technologies <http://doi.org/cvqm>.
- Perrakis, K., Ntzoufras, I. and Tsionas, M. G. (2014). On the use of marginal posteriors in marginal likelihood estimation via importance sampling, *Computational Statistics & Data Analysis*, 77 (C), 54–69.
- Rubin, D. B. (1987). A noniterative sampling-importance resampling alternative to the data augmentation algorithm for creating a few imputations when fractions of missing information are modest: The SIR algorithm, *Journal of the American Statistical Association* 82, 543–546.
- Rubin, D. B. (1988). Using the SIR Algorithm to Simulate Posterior Distributions. In *Bayesian Statistics 3*, ed. J. M. Bernardo, M. H. DeGroot, D. V. Lindley, and A. F. M. Smith, 395–402. Oxford: Oxford University Press.
- Smith, A. F. M., and Gelfand, A. (1992). Bayesian statistics without tears: A sampling–resampling perspective, *The American Statistician* 46, 84–88.
- Stinchcombe, M., and White, H. (1989), Universal approximation using feedforward networks with non-sigmoid hidden layer activation functions, in *Proceedings of the IEEE 1989 International Joint Conference on Neural Networks*, Vol. 1, IEEE, New York: 613–618.
- Stinchcombe, M., and White, H. (1990), Approximating and learning unknown mappings using multilayer feedforward networks with bounded weights, in *Proceedings of the IEEE 1990 International Joint Conference on Neural Networks*, Vol.

- 3, IEEE, New York, pp. 7–16.
- Stoll, C., Klaaßen, L., and Gellersdo, U (2019). The Carbon Footprint of Bitcoin, *Joule*, 3(7): 1647-1661.
- Taddy, M., and Kottas. A. (2010) A Bayesian Nonparametric Approach to Inference for Quantile Regression, *Journal of Business & Economic Statistics*, 28:3, 357-369,
- Tsionas, M. (2020a). Quantile Stochastic Frontiers, *European Journal of Operational Research*, 282 (3): 1177-1184.
- Tsionas, M. (2020b). On a model of environmental performance and technology gaps, *European Journal of Operational Research*, <https://doi.org/10.1016/j.ejor.2020.02.025>.
- Urquhart, A. (2016). The inefficiency of Bitcoin, *Economics Letters*, 148: 80-82.
- White, H. (1989). Learning in artificial neural networks: A statistical perspective, *Neural Computation*, 1, 425–464.
- Verdinelli, I., and Wasserman, L. (1995). Computing Bayes factors using a generalization of the Savage–Dickey density ratio, *Journal of the American Statistical Association*, 90, 614–618.
- Zellner, A., (1971). *An introduction to Bayesian inference in econometrics*. Wiley, New York.