

Stinner, Jona; Tyrell, Marcel

Conference Paper

Proof-Of-Work Consensus Under Exogenous Distress: Evidence from Mining Shocks in the Bitcoin Ecosystem

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Stinner, Jona; Tyrell, Marcel (2022) : Proof-Of-Work Consensus Under Exogenous Distress: Evidence from Mining Shocks in the Bitcoin Ecosystem, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/264115>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Proof-of-work consensus under exogenous distress: Evidence from mining shocks in the Bitcoin ecosystem*

Jona Stinner[†] Marcel Tyrell[†]

March 1, 2022

Abstract

We investigate the impact of exogenous shocks on network and stability parameters in proof-of-work-based permissionless blockchains with structurally concentrated mining ecosystems, such as the Bitcoin network. While plunging hashing power during a shock intuitively weakens the network by lowering the attacking costs, we show that parallel movement of important covariates (such as the coin price) can offset this effect and keep the mining process incentive compatible. Fundamental to our analysis is a theoretical framework building on [Budish \(2018\)](#) and [Capponi et al. \(2021\)](#). We then examine two exogenous shocks to the Chinese Bitcoin mining ecosystem caused by the Corona pandemic in October 2020 and grid disruptions in April 2021 to provide empirical evidence for our hypotheses. Our results illustrate i) the crucial structural parameters that affect the robustness of the consensus design against exogenous shocks, and ii) how market participants incorporate and value variation in the implied stability of the distributed ledger.

Keywords: Blockchain, distributed ledger technology, Bitcoin, proof-of-work, competition dynamic, miner.

JEL Codes: G10, D02, D41, D85

*An earlier version of this paper was entitled "Competitive momentum, incentive compatibility, and transaction security in proof-of-work consensus - Evidence from mining shocks"

[†]Contact details: Witten/Herdecke University, ISIC - The PPE Institute for Social and Institutional Change, Alfred-Herrhausen-Str. 50, 58448 Witten, Germany; corresponding author: jona.stinner@uni-wh.de.

1 Introduction

Distributed ledger technologies (DLT) promise to become the most significant innovation of the 21st century in terms of transparent, cost-effective, and automated coordination of processes in business and society. In most DLT, a blockchain replaces trusted third-party intermediaries by combining a publicly accessible, decentralized database with a permission-specific consensus algorithm. This structure facilitates interaction among a (basically) unlimited number of anonymous and arbitrarily distributed individuals at low transaction costs and independently of any existing trust relation. Central to the operation of each permissionless blockchain is a consensus mechanism designed to ensure that initiated transactions are stored in the decentralized ledger in a tamper-proof and externally verifiable format. Historically, proof-of-work consensus (PoW) has been prevalent, although several alternative concepts have emerged and are gaining traction (see [Irresberger et al., 2020](#); [Stinner & Tyrell, 2021](#)).

PoW consensus encourages some nodes in the network (also called "miners") to engage in a competitive tournament for the right to update the ledger. Miners verify and settle pending transactions in new block candidates and maintain the ledger's integrity. Hence, they form the backbone of any PoW-based network. Participation is incentivized by a reward in the form of native currency units for the miner, who first appends a valid transaction candidate to the blockchain. In this process, a miner's probability of success is proportional to the ratio of computing power employed to solve the cryptographic puzzle associated with block creation.

By construction, PoW consensus intends to maintain a symmetric distribution of hashing power among miners based on a competitive equilibrium ([Nakamoto, 2008](#)).¹ Preventing dominant actors or coalitions is critical for the stability of the consensus design since an attacker with a majority of hash capacity is able to successfully attack the ledger (see, among others, [Budish \(2018\)](#), [Eyal and Sirer \(2014\)](#), and [Nakamoto \(2008\)](#))

¹Computing power exerted to PoW mining is generally interpreted as "hash-rate" and expressed in trillions of hashes or tera-hash (TH). The hash-rate reflects the estimated number of hash computations performed to solve the cryptographic function underlying the PoW mechanism. See [Naor and Yung \(1989\)](#) or [Al-kuwari et al. \(2011\)](#) for more information on hash functions, and [Schär and Berentsen \(2020\)](#) for their application in proof-of-work-based cryptocurrencies.

for theoretical and [Shanaev et al. \(2020\)](#) for empirical work). However, looking at the Bitcoin network – the most prominent implementation of PoW consensus – we observe recurring centralization based on economic forces, technical characteristics, and strategic behavior of heterogeneous miners. These patterns pose a systemic threat to the consensus protocol and have prompted a growing body of academic research (see section 2). We extend the literature by examining the implications of exogenous shocks on the robustness and incentive compatibility of PoW consensus with structural concentrated patterns in the mining ecosystem.

Concentration patterns are central for our analysis since they substantially influence the occurrence and impact of exogenous shocks, such as cyberattacks, political and regulatory interventions, or energy-supply disruptions. This interdependence arises from two aspects: First, miners sharing a particular characteristic (e.g., location) are simultaneously affected by an exogenous disturbance that impinges on that quality. Second, a few large miners are easier for a malicious party to corrupt or compromise than a symmetrically distributed network, in order to gain control over a relevant fraction of computing capacity. Taken together, an exogenous shock is more likely to cause a severe capacity decline and develop into a systemic crisis when the mining ecosystem is clustered. Although essential for the long-term stability of decentralized PoW consensus, this interaction has received little attention in the academic literature.

Our first contribution is a theoretical framework building on [Budish \(2018\)](#) and [Caponi et al. \(2021\)](#) that captures the competition and industry dynamics of PoW consensus. The model formally describes the comparative statics of the mining market and allows to derive implications on the influencing factors and market dynamics. Moreover, the theoretical examination provides empirically verifiable hypotheses on the tension between competition, incentive compatibility, and transaction stability in the presence of exogenous shocks, which negatively influence the hash capacity provided by miners.

Our second contribution is an empirical analysis on the impact of exogenous disruptions on transaction and stability parameters in the Bitcoin network – a PoW-based permissionless blockchain ecosystem with a centralized mining industry. Based on a com-

parative natural experiment setting, we provide evidence from two exogenous shocks to the Chinese Bitcoin mining industry in autumn 2020 and spring 2021.

The first event is associated with the regional distribution of computing capacity in China, which followed a seasonal pattern over the last decade: During the rainy season (June-October), numerous mining centers were located in the southern Chinese provinces Sichuan and Yunnan to exploit cheap surplus energy from local hydropower plants. When hydro-energy becomes expensive during the winter months, mining operators relocate to the northern provinces Xinjiang and Inner Mongolia – a distance of $\approx 3,000$ kilometers – to employ coal energy. In autumn 2020, the physical transport of computing units was severely restricted by quarantine restrictions following local Coronavirus outbreaks in Xinjiang. Over 3-4 weeks, this restriction represented an unexpected negative production shock. The second event occurred in April 2021, when the flooding of a coal mine in Hutubi county, Xinjiang province, trapped 21 coal miners. Safety inspections following the accident led to the closure of the district’s commodity mines for about ten days. At this time of the year, many Bitcoin miners were still active in Northern China. However, the limited energy access due to the blackout of the coal-based energy infrastructure forced cryptocurrency miners to shut down facilities.

Both events are remarkably similar regarding their impact on the Bitcoin network: During October 2020 and April 2021, the global hash-power exerted to mining declined by 32 and 34 percent, respectively.² The observed shocks thus constitute the sharpest capacity losses in Bitcoin’s history to that date. In addition, the weekly block formation exhibits the lowest and second-lowest number of blocks ever observed between January 2012 and May 2021. Various figures demonstrate that the capacity squeeze adversely compromised transaction settlement: The average block time increased to 15.01 and 15.25 minutes (compared to 9.5 minutes o.a.), the number of unprocessed transactions peaked at 135K and 126K (compared to 17.8K o.a.), and transaction fees grew by a factor of 4.02 and 4.24. Our results show that the relation between transaction fees and settlement congestion, as studied by [Easley et al. \(2019\)](#) in a general equilibrium model,

²Each deviation was observed within a 14-day interval, i.e., within one period of global block difficulty.

holds under exogenous shocks. They further suggest that miners not affected by the shock benefit from higher fees associated with the stringency of settlement capacity.

While Bitcoin’s transactional capability was similarly impaired during the observed shocks, the robustness of the consensus protocol was not. Plunging mining capacity during a shock intuitively lowers the costs associated with a majority attack, which enhances the vulnerability of the decentralized ledger. However, whether the ledger remains immutable (i.e., the consensus mechanism incentive-compatible) depends not only on the aggregate of hashing capacity but also on the parameters affecting the economics of mining. The rationale behind this is straightforward: The opportunity costs of an attack proportionally increase with the rents extracted from honest mining. A shock naturally reduces competition and increases transaction fees. Additional movements of the BTC-USD exchange rate modify the present value of operated equipment.³ Hence, increasing gains from honest mining may offset declining attacking costs, rendering malicious actions economically unviable. Ideally, this mechanism maintains honest behavior, and the blockchain system remains incentive-compatible against a majority attack. During the described shocks, the BTC price shows an opposite trend with a trajectory of +50 percent in October 2020 and -20 percent in April 2021. Given similar trends in other parameters, we argue that Bitcoin’s fragility remained at or declined from pre-shock levels in October 2020, but significantly increased in April 2021.

We leverage this heterogeneity to study how market participants internalize information on Bitcoin’s varying stability levels. To empirically quantify this property, we estimate bid-ask spreads for 11 cryptocurrencies based on hourly trade data from Kraken.com, a major cryptocurrency exchange. Bid-ask spreads have been shown to implicitly mimic a market’s liquidity. Moreover, under distress or general uncertainty, market-makers impose higher spreads to offset risks from providing a trading venue. Hence, bid-ask spreads are a suitable instrument to reflect the implied stability of the Bitcoin network. Some cryptocurrencies in our dataset employ alternative consensus mechanisms (e.g., Proof-of-Stake), which do not require relevant amounts of electricity; Others simply appear not to

³The terms "Bitcoin price", "valuation", and "exchange rate" are used interchangeably in this paper and always refer to the BTC-USD conversion rate.

be affected by the shock. We cluster a control group from unaffected cryptocurrencies and employ a comparative regression design to isolate implications of the shocks on Bitcoin and a few other treated cryptocurrencies.

Our findings show that bid-ask spreads for trading Bitcoin in October 2020 are not significantly different from currencies in the control group. However, in April 2021, we observe significant positive spreads, as predicted by our hypotheses. Additional robustness checks confirm the result from our baseline fixed-effects regression. Consistent with our intuition, bid-ask spreads for trading Bitcoin widened considerably in April 2021, when the vulnerability of the ledger was relatively high, while we observe insignificant variation in October 2020, when price appreciation compensated shrinking attacking costs. The evidence presented indicates that sophisticated market participants, such as cryptocurrency exchanges, carefully observe Bitcoin’s stability level and internalize fluctuations during exogenous shocks accordingly. Overall, we theoretically and empirically uncover relevant fragility parameters of PoW consensus and demonstrate conditions under which the design remains robust when confronted with exogenous shocks.

The paper proceeds as follows: Section 2 provides a literature review. Section 3 describes the theoretical framework of the mining market, considering both homogeneous and heterogeneous miners. Section 4 first discusses the relevance of the Chinese mining industry to introduce the context of the empirical analysis. The section then provides descriptive results, the econometric design and calibration, as well as a discussion of the findings. Finally, section 5 concludes by shortly summarizing the paper.

2 Literature Review

Our paper joins the growing academic literature studying the implications and economics of blockchain technologies, digital currencies, and crypto-assets. Irresberger et al. (2020) and John et al. (2020) study the efficiency, while Wang et al. (2019) provide a literature review of various consensus algorithms. Prominent papers elaborating the mechanism of cryptocurrency pricing and returns include Griffin and Shams (2020), Pagnotta and

Buraschi (2018), Pagnotta (2021), Biais et al. (2021), Liu and Tsyvinski (2021), Makarov and Schoar (2020), and Li et al. (2018). A variety of authors address certain aspects of PoW consensus, such as competitive dynamics, incentive compatibility, and stability, in permissionless blockchains. Abadi and Brunnermeier (2019), Ma et al. (2019), Biais et al. (2019), and Chiu and Koepl (2017) provide equilibrium frameworks to formalize the economic mechanism of the consensus design and study its properties. Leshno and Strack (2020) use an axiomatic approach to formulate economic limits of decentralized consensus in an impossibility theorem. Budish (2018) formally establishes the incentive compatibility of mining and discusses several attack scenarios. Shanaev et al. (2020) empirically analyze the value depreciation of 13 cryptocurrencies that have been exposed to a majority attack. Prat and Walter (2021) provide a structural model of miners' entry and exit decision based on variation in the Bitcoin price. Benetton et al. (2021) point to negative externalities of organized mining activities. Garratt and van Oordt (2020) highlight the relevance of fixed costs associated with the operation of mining equipment for the robustness of the consensus protocol. Easley et al. (2019), Basu et al. (2021), Brown and Koepl (2019), Huberman et al. (2021) and Auer (2019) examine transaction fees. Lehar and Parlour (2020) suggests that colluding miners inflate fee levels based on effective price discrimination against Bitcoin users. Surveys are provided by Halaburda et al. (2021), who give an overview of the microeconomics of cryptocurrencies, and by Chen et al. (2020), who discuss several strands of the literature.

We focus on the economic robustness (i.e., incentive compatibility) underpinning PoW consensus. Related to our paper, several authors describe concentration patterns in the mining ecosystem and study potential implications/threats to the consensus design:

Eyal and Sirer (2014) elaborate centralization by so-called "selfish" miners. Slightly simplified, selfish miners form a secret coalition and pool block rewards by repeatedly creating private versions of the public blockchain. By strategically releasing blocks from the private chain, selfish miners provoke honest participants working on the public chain to waste their resources. Eventually, discouraged honest miners will leave the network (or join the selfish miners), which increases the relative share of the selfish coalition and

allows for disproportionate earnings. In practice, recurrently publishing alternative chains leads to forks in the assignment of the generally valid blockchain. This decreases users' trust in the value of the cryptocurrency, which is not in any miners' interest.

[Cong et al. \(2020\)](#) and [Savolainen and Ruiz-Ogarrio \(2020\)](#) examine the concentration of computing power in mining pools and incentives for attacks on the network arising from such constructs.⁴ Importantly, they suggest that economic barriers limit the size of pools and their incentive to execute attacks.

Moreover, the consensus design suggests a limited expansion of individual pools: Since the PoW protocol dictates proportional returns equal to the ratio of individual and global hashing power, a miner gains the same amount of rewards when working for a small or large pool. Consequently, increasing a single pool's size does not generate meaningful economic gains. In fact, honest miners will leave a pool that becomes too large, as inflating size imposes a systemic risk to the stability of the network, on which mining revenues critically depend. Pool formation by rational miners, therefore, does not treat system stability.

[Capponi et al. \(2021\)](#) develop a Cournot-nash equilibrium of PoW-based mining. Akin to [Arnosti and Weinberg \(2021\)](#), the willingness (or ability) of miners to invest in innovative hardware creates heterogeneous hashing costs. While such cost variation explains structural concentration in the mining ecosystem, the authors imply that large miners do not inherently increase their advantage over small miners. In particular, capacity constraints on access to low-cost energy prevent the most efficient miners from extending their advantage indefinitely. [Capponi et al. \(2021\)](#) further show that increasing investment in hardware has two opposing effects on network immutability: On the one hand, technology investments raise the level of computing capacity exerted to mining, which intuitively makes it more costly to obtain a majority qualified for attacking the network. On the other hand, operating innovative equipment decreases the cost-per-hash, which mitigates the first effect. In equilibrium, the investment intensity is a function of mining

⁴Miners organize themselves into pools to offset idiosyncratic risks. Pool participation became more or less obligatory for participants when competition in the Bitcoin mining market intensified. Typically, miners consolidate their capacity and distribute revenues proportionally to the contributed computing power within a pool.

returns, which retains a certain level of robustness. As we will demonstrate later, this mechanic is stressed during exogenous shocks: Per definition, an exogenous disturbance pushes the total hash-rate below the equilibrium rate. Therefore, a majority attack must become cheap when capacity collapses. However, we empirically show that this cannot be taken for granted as the movement of covariates (such as coin price and transaction fees) during the shock can offset this effect.

Referring to the geographical distribution of the mining industry, [Hileman and Rauchs \(2018\)](#), [Rauchs et al. \(2019\)](#) and [Rauchs et al. \(2021\)](#) locate ≈ 65 percent of the globally operating Bitcoin mining farms in mainland China. In a recent paper, [Makarov and Schoar \(2021\)](#) exploit the semi-transparent Bitcoin blockchain to study the distribution of mining rewards within the 20 largest mining pools. Their analysis implies a large concentration in the mining industry, with a minority of 10 percent of miners controlling more than 90 percent of the total computing capacity. By tracking miners' transaction flows to local cryptocurrency exchanges, the authors reveal miners' regional composition with significant geographic capacity clusters of 60 to 80 percent in China between 2015 and 2020. On the qualitative side, [Kaiser et al. \(2018\)](#) analyze the role of the Chinese government as a looming threat to the stability of the Bitcoin network. The authors are primarily concerned with cataloging motives and potential threat scenarios that might be in the interest of the Chinese government but do not quantify the impact of these actions on the network. Eventually, [Scharnowski and Shi \(2021\)](#) investigate an exogenous shock on the energy supply in the Chinese mining market in April 2021 to highlight the effects of grid disruptions on market integration, i.e., volatility, transaction volume, and transactions costs in the network. In contrast to their work, we primarily study the fragility of the consensus design in the presence of concentrated mining capacity and exogenous shocks.

3 Theoretical framework

In this section, we outline the theoretical framework based on previous work, especially [Budish \(2018\)](#) and [Capponi et al. \(2021\)](#). The framework elaborates the properties of PoW consensus from a competition-theoretic perspective and aims to highlight its fundamental economic principles. Without loss of generality for PoW-based consensus, we primarily refer to the Bitcoin network.

3.1 Free-entry equilibrium, homogeneous miners and fragility of the blockchain

We start with the basic economic mechanisms under the assumption of free entry and homogeneous miners, following the widely established model of [Budish \(2018\)](#). The model employs the following notation: R_B denotes the expected income of a miner who succeeds in the mining competition and adds transaction block B to the chain.⁵ A miner succeeds if she is the first to bundle an unspecified number of pending transactions with the solution of a computational problem into a block candidate that is accepted by the network. From a miner’s perspective, the computational problem constitutes the major challenge as she competes with other miners for the solution that, in probabilistic terms, solely depends on the provided computational capacity. This computational tournament follows a winner-takes-it-all format, while the underlying consensus protocol rules that the winning probability always remains proportional to the provided capacity. The amount of computing power thus determines a miner’s likelihood to present the solution to the network and earn the reward R_B .

The reward consists of two components: First, the winning miner receives a freshly minted amount $S > 0$ of Bitcoins, currently 6.25 BTC. Second the miner collects any fees $\sum_{i \in B} f_i$ embedded in the transactions bundled to block B . All miners have equal access to an identical mempool – a repository that contains unsettled transactions transmitted by network participants.⁶ [Easley et al. \(2019\)](#) analyze the evolution of transaction

⁵Note that cost and revenue components are denoted per block.

⁶Unsettled transactions are stored in the mempool until they are included in a block and appended

fees in a game-theoretic model and provide empirical evidence. They show that higher transaction fees are driven by congestion, i.e., the number of unsettled transactions in the mempool. More specifically, Bitcoin users compete for the limited block size to have their transactions included in the blockchain. Because rational miners prioritize transactions according to relative fees, some users are willing to pay higher fees to reduce waiting times. It follows that increasing congestion, as indicated by a larger mempool, inflates transaction fees. The total revenues are defined as $R_B = p(S + \sum_{i \in B} f_i)$, with p as the dollar value of Bitcoin. We expect the sum of f_i to increase with the size of mempool, i.e., $\frac{\partial(\sum_{i \in B} f_i)}{\partial size_{mp}} > 0$ with $size_{mp}$ denoting the mempool size. The first reward component currently accounts for the majority of total revenues, but the ratio may change to the favor of fees in the future as the amount of fixed reward S halves in a roughly 4-year sequence.

c_B denotes the per block cost of one unit of computing capacity that miners can deploy. We assume that one unit of capacity requires both one ASIC chip and one unit of electricity as complementary production factors. This means that the cost structure can be expressed by $c_B = rC + e_B$, where C is the acquisition cost of the chip, r is the per block capital cost of the chip (including depreciation) and e_B is the per block energy cost. For simplicity, we initially assume that costs incur symmetrically to all market participants.

Let H refer to the total number of units of hash capacity in the system. Given the symmetric cost distribution, each unit has a probability of $\frac{1}{H}$ to first find a valid block in the mining competition. For example, a mining pool that controls a share of 50 percent of the total hash-rate has a 50 percent probability of success. Moreover, considering free market entry and pure competition, the following condition arises:

$$H^* c_B = R_B \tag{1}$$

Potential miners continue to invest in computing capacity until all profit opportunities are exploited. This is the typical result of a rent-seeking competition under free market

to the blockchain.

entry logic. The competition condition determines the system's equilibrium hash-rate H^* as the outcome of competitive dynamics. However, what condition assures the reliability, credibility, and stability of the decentralized architecture against attacks for a given level H^* ? To answer this question, we next study the incentive compatibility of market participants.

In simple terms, the incentive compatibility requires that a miner's expected pay-offs from 'honest' behavior (as defined by the consensus protocol) must exceed those of malicious activities (e.g., manipulation of transaction blocks). A manipulation scenario generally occurs as follows: In the event of an attack, one or more manipulated block candidates are settled and collectively accepted on the blockchain, containing transactions that exclusively benefit the attacker. Network participants collectively agree that the longest chain of blocks is considered valid. The attacker seeks to build an alternative chain that contains more aggregated blocks than the blockchain originally formed by honest miners. If the attacker combines more than 50 percent of the system's computational capacity, he is able to grow the manipulated chain faster than the honest miners. After a particular time, it is necessarily considered valid by the network. The attacker thus competes with the aggregate of honest miners. Constructing an alternative chain involves high costs for the attacker since a valid solution for the resource-intensive computational problem must be provided for each block on the alternative chain. When assessing the attacker's effort, it is essential to distinguish whether the attack originates from inside or outside the system, i.e., whether or not the computational capacity was already used for mining. Equation (1) indicates that there are H^* units under the control of honest miners. An outside attacker must employ at least $H^* + \epsilon$ to form a larger alternative chain in purely probabilistic terms. In the case of an insider attack, the attacker must control slightly more than half of the existing computational capacity, $\frac{H^*}{2} + \epsilon$.

The costs increase proportionally if an attacker controls a super-majority of hashing power: A share of $A > 1$ that delivers a super-majority of $\frac{A}{A+1}$ imposes per-block costs of $A \cdot H^* c_B$ on the attacker. $A = 2$, for example, provides a 2/3 majority. In addition, we need to consider the expected average time that an attacker needs to form a larger

alternative chain. Akin to revenues and cost parameters, the attacking time is described per block formation.⁷ If an attacker requires t blocks to establish an alternative chain with manipulated blocks, then the cost of the attack is $At \cdot H^*c_B - tR_B$, which becomes $At \cdot H^*c_B - tH^*c_B$ when competition condition (1) is included. By inserting the competition condition, the costs of an attack are calculated net of the block rewards that an attacker receives in the course of alternative chain formation. If the attack succeeds, the costs are offset by its expected payoff V_{attack} , which we discuss in more detail below. We can now set $\alpha = (A - 1)t$ and arrive at the following incentive compatibility condition:

$$\alpha \cdot H^*c_B > V_{attack} \quad (2)$$

Equation (2) implies that the costs of attacking the blockchain must exceed the expected returns of doing so (Budish, 2018). If condition (2) holds, the attack is economically unviable for a potential attacker, and system stability is maintained. It follows that the computing capacity devoted to mining is an essential variable for generating trust in a PoW-based decentralized transaction system, such as Bitcoin. Two aspects of the incentive compatibility are fundamental: First, the cost patterns derived in (2) refer to a pure flow quantity, namely the operating costs H^*c_B for maintaining the system. Influences of other variables, such as the level of confidence in the overall system and the static value of mining equipment, are not included. Second, the security of the decentralized ledger is linearly dependent on the mining costs H^*c_B . A sharp and unpredictable drop in the hash-rate, such as the exogenous shocks we will discuss later, should immediately impact the fragility of the network.

The competition condition (1) in conjunction with the incentive condition (2) dictates that the total hash-rate follows from individual investments of mutually competing miners. By investing in mining capacity, miners aim to extract rents, which is only possible if the system proves to be stable in the future. In such an environment, the equilibrium condition results from the competition condition (1) and incentive compatibility constraint

⁷A new block is added to the Bitcoin blockchain every 10 minutes on average.

(2) as follows:

$$R_B > \frac{V_{attack}}{\alpha}. \quad (3)$$

In equilibrium, the one-off rewards of an attack must be relatively small to the per-block proceeds of honest behavior to maintain stability. From a miner's perspective, employing mining power to propose blocks with valid transactions must be economically beneficial compared to the returns from a one-time attack. It follows that block revenues must be high, limiting the operating conditions and scalability of PoW consensus.

But what are the potential proceeds of an attack? The answer depends on the attacker's possible actions, which we describe below. An attacker, controlling a majority of hash-power, is able to generate an alternate chain faster than the aggregate of honest miners. The attacker can use this private chain strategically to replace the blockchain created by honest miners with his own alternative. By doing so, the attacker controls which transactions are included in the transaction ledger. More importantly, the attacker can remove transactions settled on the public blockchain by starting an alternative chain with varying transactions based on the most recent public history. The alternative chain evolves into the longest chain at a particular stage, depending on the majority of computing capacity the attacker holds, while honest miners continue to append blocks to the public blockchain. Once the alternative surpasses the public chain, honest miners accept the attacker's chain following the consensus protocol.

An attacker's proceeds first contain all block rewards from the alternative chain. However – since the attacker controls a majority of processing capacity – he could earn similar payoffs by mining on the public blockchain. The main incentive for an attack stems from the ability to select the transactions that are executed on the decentralized ledger. This opportunity is not unlimited: An attacker, for instance, cannot manipulate accounts on the blockchain or transfer Bitcoins owned by other network participants to addresses under his control. To initialize such transactions, the attacker would need to (i) successfully impair the cryptographic fundamentals or (ii) gain access to a users' private key. However, the attacker can perform a so-called "double-spending attack".

In a first step, the attacker spends his Bitcoins in exchange for (off-chain) goods,

assets, or fiat currencies. The record in the public blockchain validates the transaction, and the counterparty delivers the goods or assets. After a short lock-up period, the seller considers the payment irrevocable through the entry in the blockchain.⁸ In a second step, the attacker reverses the transactions after receiving the counter value. For this purpose, he creates an alternative blockchain that no longer contains the transactions underlying the original trading contract. By employing his majority of processing power, the alternative chain becomes the longest chain and is considered valid by the network. This approach undermines the finality of transactions on the decentralized ledger, as the attacker reverses the payment process while retaining the goods and assets. Conceptually, the attacker can use his native coins multiple times by repeatedly performing the manipulation approach. Therefore, it is not strictly a "double-spending" but rather a "multiple-spending problem".

We further follow [Budish \(2018\)](#) to specify the value of a majority attack. Assume that a typical block contains k individual transactions. The attacker can create a manipulated block by bundling k Bitcoin transactions from different addresses that he controls into one block and append it to the public blockchain. For simplicity, let Block 1 represent the manipulated bundle of transactions and the previous Block 0 the state of the blockchain with its entire history before the attack. The average value of all transactions located in Block 1 is equal to $\bar{v}_{transaction} = \frac{\sum_{i \in B} p v_i}{k}$. After an trade is initiated, the counterparties wait an lock-up period of t blocks, before delivering the goods and assets. The honest miners provide a cumulative share of H^* units of computational capacity, while the attacker controls AH^* , with $A > 1$. The attacker accumulates the block revenues of R_B for the duration of the attack, which corresponds to the period until his alternative chain surpasses the public blockchain. Further, we assume that the attack does not affect the Bitcoin value but only incurs costs and revenues in the form of stream quantities.⁹ Under these assumptions, the value of an attack is $V_{attack} = k\bar{v}_{transaction}$. If we further streamline the per-block rewards to the number of transactions in the block, i.e., $r_{transaction} = \frac{R_B}{k}$,

⁸The lock-up/escrow period considerably varies among cryptocurrencies (see [Irresberger et al., 2020](#)). Concerning Bitcoin, the escrow period is typically set to 2-3 blocks or 20-30 minutes on average.

⁹This (strict) assumption will be critically discussed and eliminated below.

equilibrium condition (3) can be specified as follows:

$$r_{transaction} > \frac{\bar{v}_{transaction}}{\alpha}. \quad (4)$$

Equation (4) illustrates that the transaction proceeds must be sufficiently high to maintain incentive compatible. The profit expectation of honest mining exceeds those of attacks on the network if condition (4) is satisfied.

We can ease the equilibrium constraint by assuming that the system's perceived value and thus its native coin will be severely affected by a successful attack. This reflects the more realistic scenario of market participants losing trust in the stability of the decentralized system after an attack. To formally integrate this aspect into condition (4), let Δ_{attack} denote the proportional loss in the native coin's value in response to an attack. If the attacker holds native coins equal to a manipulable block ($k\bar{v}_{transaction}$) and we further assume that mining hardware can be deployed for other purposes than mining, the modified equilibrium condition is as follows:

$$r_{transaction} > \frac{1 - \Delta_{attack}}{(A - 1 + \Delta_{attack})t} \bar{v}_{transaction} \quad (5)$$

We can conclude that the loss of value in the native coin increases the cost of the attack and reduces the potential value available for double-spending. A collapse of the ecosystem results in $\Delta_{attack} = 1$, which renders a double-spending attack worthless. Thus, a higher Δ_{attack} implies lower returns of an attack that is supposed to generate additional income for the attacker.

This relationship reverses when an attack is instead intended to sabotage the blockchain, which Budish (2018) refers to as a sabotage attack. Such an attack becomes more successful as the anticipated loss in value of the native coin increases. The expected coin depreciation following the attack thus affects the miners' incentives to attack the blockchain, depending on the specific motivation. Low coin depreciation makes double-spending attacks attractive, while high depreciation entails a higher risk of sabotage attacks.

Analyzing flow quantities of costs and revenues without including other variables is

adequate if the mining technology is unspecific to the blockchain. However, this cannot be generalized to all forms of cryptocurrencies. For example, the diffusion of specialized single-purpose hardware (so-called "Application-Specific Integrated Circuits" or "ASICs"), which can be used exclusively for mining, is obligatory in the Bitcoin mining industry. When a blockchain collapses after an attack ($\Delta_{attack} = 1$), its native coin and employed specialized equipment are rendered worthless. The incentive compatibility condition then changes to

$$H^*C > V_{sabotage}. \quad (6)$$

Compared to incentive condition (2) ($\alpha \cdot H^*c_B > V_{attack}$), condition (6) is less strict, at least with respect to the left-hand side of the inequality. α is smaller than 1 provided $A < 2$, i.e., the attacker controls less than a 2/3 majority of computing capacity, then $\alpha = (A - 1)t < 1$ holds. Moreover, $c_B = rC + e_B$ is typically smaller than C . This highlights that the high degree of specialization of mining technology reduces the vulnerability to sabotage.

The theoretical framework provides initial insights into the mechanics of PoW-based blockchains. The assumption of pure competition and free market entry dictates that miners cannot extract surplus profits in the long run. Instead, profit potentials stimulate market entry and additional investment in mining technology, which escalates competition and reduces payoffs from participation. The system's stability is conserved by the same dynamics in equilibrium. However, exogenous shocks that abruptly alter one of the system's core variables may severely affect the stability and transaction capability of the network, as we will discuss in section 4. Before turning to the empirical examination, we extend our analysis to an explicit game-theoretical model of miners' strategic behavior. In subsection 3.2, we sketch a model based on [Capponi et al. \(2021\)](#) in order to analyze competition between heterogeneous miners, who differ in terms of cost-efficiency.

3.2 Cryptocurrency mining, heterogeneous miners, and fragility of the blockchain

Using the framework of [Capponi et al. \(2021\)](#) we construct a game between miners, who compete for rewards from solving the computationally costly hashing problem. $N \geq 2$ miners decide on their individual hash-rate commitment during the mining competition. c_i denotes the cost-per-hash of miner i , and we assume differences in the cost efficiency of miners because of varying individual investment levels β_i in new hardware.¹⁰ Accounting for these differences of the initial cost-per-hash across miners, we can sort the miners in order of increasing cost-per-hash, i.e., $c_i \leq c_{i+1}$. In the following analysis, we consider the investment level β_i in new hardware as given and exogenous. Since more efficient hardware lowers the cost of mining, the investment level is also a strategic variable that impacts the outcome of the mining game. h_i denotes the hash-rate of miner i exerted at the mining stage for $i = 1, \dots, N$. Of course, the individual hash-rate depends on the miner's investment profile in hardware. Similar to [Capponi et al. \(2021\)](#), we assume that miners have limited hashing capacity, which is captured by a quadratic cost term $(\gamma/2)h_i^2$. This capacity constraint originates from a bounded supply of low-cost electricity, with larger values of γ corresponding to smaller capacity. However, it is a soft constraint due to the convex cost function, i.e., the hash-rate can be raised at increasing marginal costs. $H = \sum_{j=1}^N h_j$ is the aggregate hash-rate of all miners.

$R > 0$ denotes the total revenues from mining, which are defined similar to subsection 3.1. Now the objective function of miner i is given by

$$\pi_i(\beta_i, h_i; \beta_{-i}, h_{-i}) = \frac{h_i}{H} R - c_i h_i - (\gamma/2) h_i^2. \quad (7)$$

In the game-theoretic setting, miners compete for the revenues generated from adding blocks to the blockchain by solving the computationally costly hashing problem. In the first step, following [Capponi et al. \(2021\)](#), we determine the equilibrium hash-rate and

¹⁰Cost differences also reflect the quality of the old and less efficient hardware stock. The most recently introduced hardware decreases the cost-per-hash. Typically large miners have lower costs per hash than small miners. They are able to invest more in new hardware and may receive discounts due to greater bargaining power and larger quantities.

equilibrium profits of miners. Since we treat the cost-per-hash $(c_i)_{1 \leq i \leq N}$ of all miners as exogenous, we can define $c^{(n)} = \sum_{i=1}^n c_i$ as the cumulative cost of the first n most efficient miners. [Capponi et al. \(2021\)](#) show that a Nash equilibrium hash-rate profile exists for any investment profile β of the miners $h^*(\beta) = (h_i^*(\beta))_{1 \leq i \leq N}$ with n miners active in equilibrium (see their proposition 4.1). The first-order condition of the objective function (7) provides the equilibrium hash-rates from equating marginal gains and marginal cost. This condition is given by

$$\frac{R}{H^*} \left(1 - \frac{h_i^*}{H^*}\right) = c_i + \gamma h_i^*. \quad (8)$$

solving for h_i^* provides the equilibrium hash-rate for active miner i

$$h_i^* = \frac{H^*(R - c_i H^*)}{R + \gamma (h_i^*)^2}. \quad (9)$$

The equilibrium aggregate hash-rate H^* is determined by summing over all individual equilibrium hash-rates of active miners. For the realistic case of limited hashing capacity ($\gamma > 0$), the aggregate hash-rate is given by

$$H^* = \frac{\sqrt{(c^{(n)})^2 + 4(n-1)R\gamma} - c^{(n)}}{2\gamma}. \quad (10)$$

Not surprisingly, the aggregate equilibrium hash-rate increases with rewards R and decreases with smaller hashing capacity, i.e., larger value of γ . Higher cumulative costs of the active miners $c^{(n)}$ also decrease the aggregate hash-rate H^* . Only the n most efficient miners with marginal gains at least as large as marginal costs for positive h_i are active in equilibrium. As we can see from (9), miners with lower costs c_i have higher hash-rates. The least efficient miner n controls the smallest nonzero hash-rate. We can observe from (8) that the equilibrium revenue-per-hash R/H^* is larger than the marginal gain because the marginal probability of earning the reward is decreasing in the exerted hash-rate. The marginal cost of miner i is given by $MC_i^* = c_i + \gamma h_i^*$, with $(MC_i^*)_{1 \leq i \leq N}$ as an increasing sequence. Since $MC_i^* < R/H^*$ implies that $c_i < R/H^*$, a miner is active only if its

cost-per-hash is lower than the return-per-hash.

Except for the marginal miner n , all other active miners make positive profits in equilibrium. The profit-per-hash of an active miner i is given by

$$\frac{\pi_i^*}{h_i^*} = \frac{R}{H^*} - c_i - \frac{\gamma}{2}h_i^* = \frac{R}{H^*} - (c_i + \gamma h_i^*) + \frac{\gamma}{2}h_i^*. \quad (11)$$

Because $c_i + \gamma h_i^*$ is increasing in i and h_i^* is decreasing in i , profit-per-hash must be decreasing in i . The largest miners in terms of exerted hash-rate are also most profitable. Therefore, the equilibrium mining profits $(\pi_i^*)_{1 \leq i \leq N}$ and the profits-per-hash $(\frac{\pi_i^*}{h_i^*})_{1 \leq i \leq N}$ form a decreasing sequence. The oligopolistic competitive pattern and the heterogeneity of miners are the main drivers of these results. Miners exploit their individual cost advantages to generate profits, which is the typical outcome of a Cournot-Nash equilibrium. Homogeneous cost structures would drive miners' profits toward zero since an infinite number of identical miners would operate with total costs equal to total revenues in equilibrium.

The number of active miners is given by the largest number n which satisfies the following condition:

$$c_n < \frac{c^{(n)} + R\gamma/c_n}{n - 1} \quad (12)$$

We know that miners are only active in equilibrium if the expected rewards-per-hash R/H^* are greater than the associated costs c_i . This means, $R - H^*c_i$ must be greater than zero. Inserting condition (10) for H^* we get $2\frac{R\gamma}{c_i} + c^{(i)} > \sqrt{(c^{(i)})^2 + 4(i-1)R\gamma}$. Solving for c_i and simplifying results in condition (12) for $n = i$ shows, that the equilibrium number of active miners is the largest value still satisfying this equation. Miner $i + 1$, who faces higher costs than miner i cannot be active if miner i is not active.

Inspecting condition (12) delivers some interesting insights. For instance, higher average cost of the first $n - 1$ miners $(\frac{c^{(n-1)}}{n-1})$ weakens the participation constraint for miner n . Thus, a miner's decision to become active depends on the relation of its costs to those of other miners. Increased cost heterogeneity leads to lower average costs of the first $n - 1$ miners and a smaller number of active miners in equilibrium. Full cost homogeneity

would imply that all miners are active regardless of their costs. Even more interesting is the impact of γ – the hash capacity constraint. The number of active miners, n , increases with a tighter capacity constraint. To understand this effect, it is instructive to first analyze condition (12) without imposing a capacity constraint ($\gamma = 0$). Then equation 12 becomes $c_n < \frac{n}{n-1} \frac{c^{(n)}}{n}$. The cost of the marginal miner can only be slightly higher (more precisely, by the factor $\frac{n}{n-1}$) than the average cost of all active miners, $\frac{c^{(n)}}{n}$. This means that the mining competition is highly vulnerable to centralization, which might undermine the system’s security. If miners have unbounded capacity, the most efficient one will dominate the market. In the case of $\gamma > 0$, the capacity constraint prevents the efficient miners from expanding their activities indefinitely. Their marginal cost of hashing increases with the exerted hash-rate h_i . Therefore, the number of active miners is increasing in γ . A higher mining reward R increases the number of active miners, whether it results from a higher Bitcoin price p or from higher transaction fees $\sum_{i \in B} f_i$. Higher rewards are an incentive for active miners to expand their capacity as the marginal gain increases. However, this effect is limited by the capacity constraint.

R and γ are crucial determinants of mining centralization, which, as discussed above with reference to the model of the Budish (2018), is an enormous threat to the security of the system. Since miners are only active if their cost c_i are lower than R/H^* , larger values of R and γ increase the mining decentralization. In turn, it becomes more expensive to attack the network.

3.3 Testable implications for exogenous shocks

The theoretical framework provides a number of implications, which we relate to the occurrence of exogenous shocks in this subsection. We begin by examining the incentive compatibility of homogeneous miners to exogenous shocks.

By definition, adverse exogenous shocks, such as the events described in section 4, constrain the total hash rate exercised in the system. For the course of the shock, competition condition (1) changes to $H^{(s)} c_B < R_B$, where $H^{(s)}$ denotes the exogenously constrained hash-rate during the shock with $H^{(s)} < H^*$. In accordance with Budish

(2018), a decreasing hash-rate lowers the costs associated with a majority attack and increases the system's vulnerability. However, this development might be offset by countervailing effects. While pure competition ruled out any mining profits in equilibrium, the exogenous constraint allows the remaining participants to extract rents. When taking both aspects into account and substituting R_B with $p_B(S + \sum_{i \in B} f_i)$, the adjusted (compounded) incentive compatibility condition during an exogenous shock is given by

$$\underbrace{\sum_{B=0}^t \frac{AH^{(s)}c_B - p_B(S + \sum_{i \in B} f_i)}{(1+r)^B}}_{\text{Mining costs net of rewards}} + \mathbb{E} \left[\underbrace{\sum_{B=0}^{\kappa} \frac{p_B(S + \sum_{i \in B} f_i) - H^{(s)}c_B}{(1+r)^B}}_{\text{Expected net profits}} \right] > \underbrace{\sum_{i \in B} p_B v_i}_{V_{\text{attack}}} \quad (13)$$

Where p_B corresponds to the native coin price (e.g., BTC) at the arrival of block B , t to the expected duration of a double-spending attack, and v_i to the value of transactions (in BTC) an attacker is able to double spend. Again, inequality (13) states that the expected cost of an attack (i.e., the left-hand side) must exceed the benefits V_{attack} , for the system to remain robust during an exogenous shock. Since $H^{(s)} < H^*$, the mining costs net of rewards decrease ceteris paribus compared to the equilibrium state. However, contrary to (1), the attacker must additionally weight the net present value of pending profits that arise from the difference of aggregated mining costs and rewards during the shock. Since the expected profits depend on p , which collapses after a successful offense against the network, the sum of compounded net rents increases the opportunity costs of an attack. Notably, the severity of the exogenous constraint ΔH dictates the cost decline to $H^{(s)}c_B$, which both lowers net mining costs and increases expected net rents (i.e., opportunity costs). However, rents cannot be extracted indefinitely. Rather, κ restricts profitable mining to the period until the exogenous constraint is relieved (or $\Delta p < \Delta H$ for $\Delta H < 0$) and a new equilibrium emerges. Hence, the set $B\{1, 2, 3, \dots, \kappa\}$ is primarily driven by the expected duration of the shock. Whether (13) is satisfied significantly depends on the parallel movement of reward variables (p, S and f). In fact, we next demonstrate that the covariation of p_B during the shock is crucial for the stability of the network.

The Bitcoin price p_B simultaneously affects costs and gains from an attack in equation (13) and is subject to significant volatility. Importantly, this price volatility has an asymmetric impact on the adjusted incentive compatibility constraint. To be more specific, consider the partial derivative of (13) with respect to p , as denoted by

$$\frac{\partial(13)}{\partial p} : - \sum_{B=0}^t \frac{S + \sum_{i \in B} f_i}{(1+r)^B} + \mathbb{E} \left[\sum_{B=0}^{\kappa} \frac{S + \sum_{i \in B} f_i}{(1+r)^B} \right] > \sum_{i \in B} v_i. \quad (14)$$

By assuming that the expected duration of the shock κ exceeds the expected time t for the attack to succeed, this can be simplified to

$$\frac{\partial(13)}{\partial p} : \mathbb{E} \left[\sum_{B=t+1}^{\kappa} \frac{S + \sum_{i \in B} f_i}{(1+r)^B} \right] > \sum_{i \in B} v_i. \quad (15)$$

Provided $\kappa > t$, equation (15) gives the net marginal costs and gains for a double-spending attack during the shock when the BTC price changes by one unit. Since the incentive compatibility condition $S + \sum_{i \in B} f_i > \sum_{i \in B} v_i$ is satisfied in equilibrium (otherwise the network was attacked prior to the shock), the change of marginal attack costs must also surpass the marginal gains. It follows that condition (13) increases for $\Delta p > 0$, and decreases for $\Delta p < 0$. This is an important observation: Adverse exogenous shocks have a divergent impact on network stability, depending on the evolution of the cryptocurrency price during the shock period. If the BTC price increases during the shock, positive marginal attack costs (partially) counteract the decreasing robustness of the network. However, if the price declines, the negative marginal attack costs lower the network robustness beyond the magnitude of the exogenous shock. Therefore, we argue that the price movement is essential for the vulnerability of PoW-based consensus during an exogenous shock.

Two properties of equation (15) require careful examination. First, the shock duration must exceed the time required for an attack ($\kappa > t$). The duration of the attack depends on the majority of computing power an attacker controls (A), and any escrow period imposed by vendors or exchanges before the asset or good is delivered. Budish (2018) estimates a simulation of all major parameter specifications (see Table 1). Estimates with

reasonable parameters for the Bitcoin network and an escrow period of six blocks show a maximum attack period of 46.02 blocks (or ≈ 7.67 hours).¹¹ The shocks elaborated in section 4.2 suggest a typical duration of at least 1-4 weeks, or the equivalent of 1008 to 4032 blocks. Consequently, $\kappa > t$ is satisfied by several orders of magnitude. Second, the coin price p_B may follow a predetermined response to the hash-rate decline (i.e., it is not truly exogenous to the shock). In line with [Biais et al. \(2021\)](#), we argue that the fundamental value of Bitcoin corresponds to its future stream of expected net transaction benefits. Since these benefits depend on the future BTC price, equilibrium prices reflect not only fundamentals but also sunspots. BTC prices thus fluctuate even when the fundamentals are constant. [Biais et al. \(2021\)](#) empirically demonstrate that a large fraction of variation in BTC returns seems to reflect extrinsic volatility. Moreover, [Liu and Tsyvinski \(2021\)](#) and [Fantazzini and Kolodin \(2020\)](#) demonstrate a unidirectional causality from the Bitcoin price to the hash-rate, i.e., miners cannot influence the price of a PoW-based cryptocurrency by choosing a certain production level. We conclude that the p_B cannot be endogenized in our framework and shows no determined relation to exogenous shocks. In addition, we provide empirical examples of the price and hash-rate development during shocks to the Bitcoin mining industry in section 4. Overall, if the price shows an increasing tendency during an exogenous shock event, we principally expect a positive impact on mining profits. As shown above, this increases the stability of the system. On the other hand, a decreasing BTC price trend lowers profits. As a result, the adjusted incentive compatibility constraint becomes tighter, which decreases the system's stability.

The mining reward R consists of three additional components that we briefly discuss: First, the per-block reward S is exogenously specified by the protocol underlying Bitcoin's consensus mechanism. Hence, S varies only about every four years in the cause of the prescribed halvings and can be considered constant in the vast majority of shocks. Second,

¹¹The observed escrow period from cryptocurrency exchanges trading Bitcoin has declined in recent years and is meanwhile well below six blocks (see [Irresberger et al. \(2020\)](#) Table A.1). Therefore, the relevant escrow period for Bitcoin is between 1 and 6 blocks. The corresponding maximum t with $A=1.05$, the smallest estimated majority, amounts to an average of 30.80 and 46.02 blocks (or 5.01 and 7.67 hours).

the evolution of transaction fees depends on the congestion in the network, as reflected by the mempool size. If transaction fee levels increase because of higher congestion during the shock, rewards from mining are positively influenced. We generally expect such a positive relation and empirically verify it in subsection 4.4. However, since transaction fees still exhibit a small proportion of total rewards, the overall effect is relatively small. Third, a negative shock on the aggregated hash-rate should expand the arrival rate of blocks for a period of up to 14 days. An increasing arrival rate temporarily decreases the total rewards until the block difficulty of solving the computational problem adjusts. Since a reduced block arrival also affects the attacker’s alternative chain (thus limiting the amount of double-spending), it does not affect the incentive compatibility condition. Drawing on the above reasoning, we expect little to no impact on the network’s vulnerability to attacks when an exogenous shock co-occurs with a significant price appreciation. This is because positive mining profits counter the negative influence of the lower aggregate hash-rate. However, when an exogenous shock and declining coin prices coincide, we expect to see signals of increasing instability.

The framework of [Capponi et al. \(2021\)](#) delivers further predictions concerning the distribution of miners’ individual hash-rates and profits which, in principle, can be tested with appropriate data. The sensitivity of a miner’s hash-rate to its own cost-per-hash parameter depends on a direct and an indirect effect. The direct effect on the individual hash-rate h_i^* is always negative when the marginal costs of mining increase with c_i . Since the indirect effect measures the sensitivity of a miner’s hash-rate to the cost-per-hash of other miners, it accounts for the strategic reactions in an oligopolistic market environment. The outcome of the indirect effect is as follows: First, the aggregate hash-rate H^* changes with h_i^* , which alters the equilibrium marginal gain of hashing and therefore affects the strategic decision of other active miners. Of course, these miners react, and the miner subject to the cost shift, considers the reactions of other miners. If the cost-affected miner controls less than a majority of the total hash-rate (i.e., $\frac{h_i^*}{H^*} < 1/2$, which is typically the case), the indirect effect is positive and (partially) alleviates the negative direct effect. However, the net effect is still a decreasing h_i^* if the miner’s cost-per-hash

increases, regardless of its hash-rate share.

Concerning the security of the system, the effects of exogenous shock events on the hash-rate allocation are particularly interesting. The vulnerability of the system depends on the mining concentration. Higher capacity concentrations typically increase the probability of an attack. If the exogenous shock affects the cost-per-hash in a heterogeneous manner, the resulting shares of the active miners are informative with respect to the stability of the entire system. The theoretical framework also provides some answers to these questions. It can be shown that miner i 's hash-rate share $\frac{h_i^*}{H^*}$ decreases if her cost-per-hash c_i increase. The sensitivity of miner i 's hash-rate share to an increasing cost-per-hash of miner j is always positive. The share of miner i increases in case of a cost-per-hash increase of miner j . Moreover, mining profits are affected similarly: Miner j becomes less competitive with increasing cost-per-hash, and when j 's profit decreases, all other miners benefit and increase their profits.

If the exogenous shock event increases the homogeneity of miners in terms of cost efficiency, decentralization increases similarly, even if the set of active miners is fixed. This mechanic enhances the system's robustness. The same intuition can be applied to the capacity constraint: An exacerbating capacity constraint due to an exogenous shock increases decentralization. Smaller miners gain market shares at the expense of larger miners. A larger mining reward R has a similar effect since it increases the marginal gain of hashing. Even though the hash-rate of each miner increases, small miners can increase their hash-rate disproportionately to large miners. This amplifies the decentralization of shares within the network if the expansion of hash capacity is not systematically constrained for exogenous reasons.

How do exogenous shocks affect the hash-rate allocation of heterogeneous miners and consensus stability? Given the relevance of cost-efficiency for the capacity and profit allocation, miners gravitate towards areas with the most cost-effective production factors. As we will describe carefully in section 4, mining clusters emerged in areas such as the Xinjiang province in Mainland China with particularly low energy fees (see Figure 1). [Makarov and Schoar \(2021\)](#) analyze the distribution of mining capacity in the context

of an exogenous shock in Xinjiang (referred to as Event [2] in chapter 4). Their results suggest that most mining companies closed between 20 and 50 percent of mining capacity during the shock, but only a small fraction lost 100 percent. The reason for this observation is that large (and most efficient) miners are diversified across multiple locations. Therefore, we argue that the exogenous shocks impaired the most efficient miners. As large miners temporarily lost the most cost-effective energy source, we expect that miner homogeneity increased with respect to cost-effectiveness. Following the mechanic of the mining game presented in section 3.2, the shock principally facilitates network stability.

However, we argue that the exogenous shock creates a temporary window of opportunity for miners with medium cost-per-hash efficiency. Temporarily, the market share of these miners increases until the exogenous constraint fades. While (previously) large miners are significantly constrained and less cost-efficient, medium-efficient miners face less competition, apply larger capacities, and generate higher profits-per-hash due to the oligopolistic nature of the mining competition. Nevertheless, they are well aware of the temporary nature of this position: Once the external restriction eases, profit potentials will drop significantly.

In such an environment, active miners may act honestly, i.e., process transactions and validate blocks adherent to the fundamentals of the consensus protocol. They temporarily earn windfall gains since the block rewards surpass hashing costs. As argued above, the magnitude of windfall gains is tied to the Bitcoin price trend and shock duration. This option becomes attractive in the event of a rising price trajectory.

Instead, the remaining miners may consider an attack when the opportunity to (collectively) control a capacity share of > 50 percent becomes feasible and economically attractive. In particular, the subset of miners whose marginal profits barely exceed marginal costs during the shock might try to coordinate for an attack. The long-term prospects and commitment of such miners are low (i.e., they do not have much "skin in the game"), as they are doomed to low (or even negative) profits once the equilibrium is restored. However, the least efficient miners control only a tiny fraction of the network capacity, and coordination costs increase with the number of coalition members. There-

fore, we expect coordination for an attack from the medium-efficient miners that gained relevant market shares due to the drop of the global hash-rate. Of course, whether these miners decide for honest or malicious behavior depends on their individual incentive compatibility condition. As demonstrated above, a decreasing Bitcoin price trajectory during the shock tightens the condition and thus increases the vulnerability of the ledger. Therefore we predict a higher fragility of the system if the exogenous shock and a decreasing Bitcoin trend coincide.

The rationale developed in this chapter provides testable predictions if the impact of an exogenous event on the cost-per-hash, capacity constraint, and rewards parameters can be accurately captured. In the next section, we analyze two events to illustrate the impact of exogenous disturbances on network and stability parameters in PoW consensus.

4 Empirical results

Section 3 provides several conclusions regarding the stability and incentive compatibility of PoW consensus in the face of exogenous shocks. This section provides empirical evidence from two shocks on the Bitcoin mining ecosystem that were exacerbated by geographic concentration.

4.1 Background

The framework described in subsection 3.2 shows that cost variation in the provision of hash capacity (c_i) determines both profits and market shares of miners. Due to the design of the PoW competition, absolute mining rewards are independent of mining capacity, as the adaptive block difficulty guarantees a constant block arrival. Economies of scale, however, are the dominant force in the organization of cost-effective mining and greatly influence the distribution of mining revenues. When the cryptocurrency mining sector matured, this dynamic inevitably led to a geographical concentration in countries with the most economical production factors. The proliferation of ASICs accelerated the concentration since it ruled out semi-professional miners and favored the organization in

data centers (Küfeoğlu & Özkuran, 2019; Song & Aste, 2020; Stinner, 2021).

With the emergence of ASICs, the People’s Republic of China developed into the dominant location for operators of mining centers over the past decade. China proved an ideal environment since it offered low-cost and rapid access to the essential production factors energy and hardware, as well as loose regulation (e.g. tax incentives).¹² Similar to large industrial customers, mining companies benefited from globally competitive energy prices in China’s subsidized energy infrastructure (Hileman & Rauchs, 2018). In addition, farm operators were able to quickly adopt innovative mining hardware from the world’s leading Chinese manufacturers (e.g., Bitmain). Based on data from three major mining pools (BTC.com, Poolin, and ViaBTC), Rauchs et al. (2021) identify an average share of 63.7 percent of Bitcoin miners in China between September 2019 and April 2021 (see Figure 1).¹³ In a recent paper, Makarov and Schoar (2021) investigate the reward distribution within the 20 largest mining pools by analyzing the transaction flow from the pool operator to individual contributors on the Bitcoin blockchain. By further tracking miners’ transaction flows to local cryptocurrency exchanges, the authors estimate miners’ regional composition between 2015 and 2020 with a significant geographic capacity cluster of 60-80 percent in mainland China. It follows that a majority of computing power dedicated to Bitcoin mining was localized in mainland China during the last decade.

[Figure 1 about here]

4.2 Event Description

We next describe two shocks on the Chinese mining ecosystem in October 2020 and April 2021. Shocks are identified endogenously from the global hash-rate (H) and defined as a decline of $\Delta H < -0.25$ within 14 days (i.e., one period of global block difficulty).

¹²By June 2021, the Chinese Government imposed a strict ban on cryptocurrency mining activities with the argument that the industry jeopardizes China’s pursuit of carbon neutrality. Thus, after June 2021, mining capacities have been shifted to Kazakhstan, the United States of America, and other countries (K Wan et al., 2021).

¹³Mining pools typically tag their block candidates with an identifier in the Coinbase transaction. Between September 2019 and April 2021, BTC.com, Poolin, and ViaBTC accounted for an aggregate of 29.7 percent of all blocks.

The first event is closely linked to the regional allocation of mining capacity within China, which exhibits a strong seasonal component. During the rainy summer months, numerous mining operators relocate their computing capacity to the southern Chinese provinces of Sichuan and Yunnan, where surplus energy from hydro-power plants is available at uniquely cheap conditions. The energy surplus is a result of structurally inadequate grid expansion, as revealed by data on the regional power generation and demand from the China National Bureau of Statistics ([NBS, 2021](#)). Between 2015 and 2020, Sichuan and Yunnan jointly supplied 42.7 percent of the nation-wide hydropower energy. Electricity generation periodically spikes by about 50 percent during the rainy summer from June to October compared to the dry winter and spring. On average, the provinces' overall electricity generation exceeded local demand by about 55.6 percent between 2011 and 2020, with the surplus increasing by 4.7 percent annually. The expansion of grid infrastructure to the energy-intensive eastern China has so far been inadequate to cope with this growth, resulting in immense excess power during the summer ([Sichuan Government, 2019](#)). Bitcoin miners have effectively monetized this structural oversupply by utilizing the overage for energy-intensive PoW mining.

Once the rainy season ends, miners return to Xinjiang and Inner Mongolia to use the relatively cheap coal power during the winter months. Based on data from [Rauchs et al. \(2021\)](#) and [Stinner \(2021\)](#), the logistics between the two most relevant regions Xinjiang and Sichuan comprised at minimum 400,000 units, or 46 percent of the domestic capacity, between April and September 2020 (see Figure 1). Although the distance between the provinces is about 3,000 kilometers, transporting the specialized computer equipment appears economically viable in the light of the industry's enormous energy demand.

The cyclical capacity transition in the Bitcoin mining industry was unexpectedly interrupted in autumn 2020 when the northern regions Xinjiang and Inner Mongolia were classified as high-risk areas following regional outbreaks of the Coronavirus. Such an intervention was not anticipated during the summer since the official figures from Chinese authorities and the World Health Organization suggested deficient infection levels. Weekly new infections declined from 31,300 in February 2020 to a moderate level of 120 to

250 cases in September and October, following drastic infection control measures ([WHO, 2021](#); [Zhong & Mozur, 2020](#)). However, the infection activity in late summer 2020 was particularly concentrated in northern provinces. To prevent a resurgence of infections, the Chinese authorities imposed extensive social restrictions, even when a few infected individuals were identified ([Reuters, 2020](#); [Xinjiang Government, 2020](#)). For example, on October 25, 2020, 138 new infections were detected in the city of Kashgar, Xinjiang, which triggered extensive testing of all 4.7 million residents in the region and massive restrictions ([BBC, 2020](#); [Hernández, 2020](#)). Comments from mining facility and pool operators on the Chinese platform Weibo show that the unexpected situation severely disrupted transportation and operational processes ([Zhao, 2020](#)). As a result, mining centers in Xinjiang could not operate efficiently for about 3-4 weeks. In addition, a considerable amount of mining equipment was stalled on its transit to the northern provinces.

The impact of the Corona pandemic exhibits an unanticipated negative production shock to the Chinese mining industry in fall 2020, manifested by a 34 percent plunge of the global hash-rate.¹⁴ Figure 2 illustrates the trajectory of the global hash-rate and block difficulty from September 2020 to May 2021. We label this shock as "Event [1]" and define its period from October 16 to November 23, 2020.

[Figure 2 about here]

The second exogenous shock occurred in April 2021, when the flooding of a major coal mine in the Xinjiang province disrupted the power supply to local Bitcoin miners. On April 10, 2021, 21 workers were trapped in the Fengyuan coal mine in Hutubi county after flooding cut off parts of the facility and disrupted communications ([CNN, 2021](#)). In response to the accident, the local government ordered to shut down the region's coal mines (and thus its power supply) in the context of safety inspections starting from April 16. As described earlier, miners migrate to Xinjiang to avoid high electricity tariffs associated with the dry season in the southern provinces. At this time of year, Xinjiang, and Hutubi County, in particular, were a magnet for Bitcoin miners, who took advantage

¹⁴This is among the most severe shocks on the mining industry to date, along with downturns in March 2020 (28.9 percent), and April 2021 (32.1 percent).

of the abundant and comparably cheap coal-fired electricity. The drastic 32 percent drop in global Bitcoin mining capacity during the post-accident shutdown is a testament to the importance of the region (see Figure 2). Once production resumed and electricity was made available, a rapid return to pre-incident capacity levels can be observed. We refer to the exogenous shock caused by the Xinjiang coal mine accident as "Event [2]" and specify its duration from April 16 to April 25, 2021.

4.3 Data and Summary Statistics

We construct our data set by examining various sources with an observation horizon from July 2020 to May 2021. First, we exploit all 46,637 blocks appended in the period from the publicly accessible Bitcoin blockchain to generate relevant block-level information, such as block-time, difficulty, transaction fees, transfer volume, and miner remuneration. Second, we derive the daily BTC-USD conversion ratio from Coindesk.com, which provides a market-representative value based on the average across multiple major cryptocurrency exchanges.¹⁵ Third, we use estimates from Blockchain.com and Bitinfocharts.com to specify the global hash-rate for any of the involved PoW cryptocurrencies.¹⁶ The provided hash data constitute retrospective estimates based on the respective historical block difficulty and observed arrival rate. In some illustrations and analyses, we thus calculate the 3-day moving average of the daily figure to balance the probabilistic element. Fourth, we use median confirmation times for a transaction to be settled on the public blockchain from Blockchain.com.¹⁷ As an additional measure of transaction capacity and demand, we obtained the number of unconfirmed transactions (mempool-transactions) from BTC.com, using web-scraping techniques.¹⁸ Table 1 contains descriptive statistics of the daily aggregated parameters of the variables listed above.

[Table 1 about here]

¹⁵See <https://www.coindesk.com/price/bitcoin>

¹⁶See <https://www.blockchain.com/de/charts/hash-rate> and <https://bitinfocharts.com/de/comparison/bitcoin-hashrate.html3y> for more information.

¹⁷See <https://www.blockchain.com/charts/median-confirmation-time>

¹⁸See <https://btc.com/stats/unconfirmed-tx>

Eventually, to derive a proxy parameter for cryptocurrency fragility, we exploit hourly data on open, high, low, and close market prices from Kraken.com for a total of 11 cryptocurrencies with varying consensus algorithms.¹⁹ For each cryptocurrency, we collected data on both the USD and EUR conversion rate and weighted figures derived on this data according to the respective trading volume. As we will explain in the next section, this approach eliminates potential confoundings in our estimates from covariation in a single fiat currency. Kraken, a major crypto exchange headquartered and regulated in the US, is considered among the most liquid and well-established exchanges in the crypto universe (Dimpfl & Peter, 2021). Regarding trading volume, Kraken ranks among the top 10 of ≈ 310 listed crypto exchanges throughout the observation period and is thus adequate to supply representative metrics.²⁰ Table 3 presents some summary statistics for the variables obtained from Kraken.com. We conclude the data sourcing by adding the overall trading volume and market capitalization for each cryptocurrency from Coinmarketcap.com.²¹

[Table 3 about here]

We begin our analysis by exploring the determinants making the combination of our observed shocks unique in the history of the Bitcoin network. Figure 2 relates the capacity decrease during the events to the endogenous block difficulty, which adjusts at a rate of roughly 14 days. Given this rigidity, declines that occur within a shorter time are expected to impact block formation considerably. Figure 3 (a) illustrates the weekly number of blocks registered on the Bitcoin blockchain for an 8-week window centered upon each event. In addition, the graph depicts the population average from January 2012 to June 2021²² of 1080 blocks and the relevant difficulty adjustments on November 3, 2020, and May 01, 2021, respectively. Both events show a substantial drop in block attachment with weekly minimum values of 787 and 817 – the lowest values ever observed to this date. This circumstance also becomes evident from Figure 3 (b), which plots the density

¹⁹See <https://docs.kraken.com/rest/> for a description of the Kraken REST API.

²⁰See <https://coinmarketcap.com/de/rankings/exchanges/>

²¹See <https://coinmarketcap.com/api/> for information about the Coinmarketcap API.

²²The period before 2012 and earlier shocks (e.g., price bubbles or reward halving) have been removed from the calculation of the population average to provide a figure of stable periods.

function for the equivalent interval (with all blocks included). Again, the observed shocks are clear outliers from the (fairly) normally distributed block count.

[Figure 3 about here]

To formally specify our event periods, we conduct weekly one-sample t-tests on the difference in observed block number against the two-tailed alternative of the population average and the observed hash-rate against its mean over the three-week interval preceding the interventions. The results are reported in Table 4. Event [1] and Event [2] deviate on a statistically highly significant level from the expected block number and hash-rate. Therefore, we consider weeks with significant divergence in either variable to specify our event periods quantitatively. Interestingly, hashing capacity was restored fast when electricity became available during the second event. Thus, the difficulty adjustment following the abrupt capacity drop caused a statistically significant above-average block arrival after May 01, 2021.

[Table 4 about here]

While a comparable decline of mining capacity characterizes both events, they differ substantially in the parallel trajectory of underlying determinants. [Prat and Walter \(2021\)](#) and [Garratt and van Oordt \(2020\)](#) demonstrate that the supply of mining capacity can be modeled as a function of the Bitcoin price. Under sufficient competition and for a given level of short-term production costs, price volatility encourages market entry or exit. Figure 4 contrasts the relative evolution of the hash-rate and Bitcoin price clustered for 20 days around each of the observed minimum values during the four most significant shocks between January 2012 and May 2021 (with $t_0 = 1$). The key finding is that the variables show the expected coherent structure in the first two shocks (i.e., the hash-rate decline is endogenous to the price), a weak co-movement in April 2021, and opposite development in October 2020.²³ The shock in October 2020 differs substantially as the endogenous determinants display an opposite trajectory: In November 2018 and March

²³We decided not to include the halving period in May 2020 into our estimation since it was anticipated by market participants and is not related to concentration patterns.

2020, the Bitcoin price dropped by about 50 percent, driven by a waning cryptocurrency hype and the global dispersion of the Corona pandemic. During the shock in April 2021, we observe a moderate decline of 20 percent. In contrast, Bitcoin appreciated by about 50 percent in October 2020, parallel to the falling hash-rate.

[Figure 4 about here]

It follows that endogenous market conditions cannot explain the declining supply of mining capacity especially during Event [1], which corroborates our argumentation of exogeneity.²⁴ Moreover, we leverage this quasi-natural experiment in our further analysis: As discussed in subsection 3.3, the network fragility during shocks depends not only on the magnitude of affected mining capacity but also on the evolution of reward parameters. Event [1] and [2] allow us to study Bitcoin’s fragility under two circumstances that are widely similar in the shock’s magnitude but substantially deviate in the price development. Since the BTC price markedly affects the economic viability of mining (and thus miners’ incentive compatibility), attacking the network in October 2020 remained expensive, given that the net present value of engines increased. In contrast, the interaction of a capacity shock and BTC price depreciation severely decreased the costs of a majority attack in April 2021. A graphical representation of the relative change in mining capacity and miners’ gross revenue during the events is given by Figure 5. We take advantage of this heterogeneity to test the hypotheses formulated in subsection 3.3 and investigate how market participants incorporate information on varying security levels of the Bitcoin blockchain. We next study the market and fragility dynamics during the two events.

[Figure 5 about here]

4.4 Blocktime, Congestion and Transaction Fees

This subsection examines the evolution of transaction capacity and fees during the described exogenous interferences. We limit the examination to a descriptive analysis in this version of the paper.

²⁴Note that the recession in April 2021 is disproportional compared to the co-movement of earlier shocks, i.e., the decrease in capacity is much larger than the price decline.

Several theoretical and empirical articles demonstrate that transaction fees in the Bitcoin network depend fundamentally on impatient users, interested in a fast settlement, rather than determinants associated with miners' revenues (e.g., block rewards) (Auer, 2019; Easley et al., 2019; Huberman et al., 2021; Möser & Böhme, 2015). Since block size exogenously dictates settlement capacity, fees typically increase with demand, which becomes transparent in the number mempool transactions and the median confirmation time.²⁵ As described earlier, the difficulty of the cryptographic function underlying the PoW consensus mechanism endogenously adapts to the exerted processing power in a bi-weekly interval. Because of this rigidity, block settlement must decrease significantly when mining capacity suddenly plunges, but block difficulty remains unchanged. Hence, we expect sharply decreasing transaction settlement, paired with rising mempool transactions and fee levels during the described events.

Figure 6 depicts daily figures for the average block time (in seconds), number of unprocessed transactions, and transaction fees per block (in BTC), each centered for ± 20 days around the observed shocks. Since both periods show a widely similar magnitude, we use the average across Event [1] and Event [2] in most of the following calculations (as not stated otherwise) and compare them to the population average based on all non-shock intervals between January 2012 and May 2021.²⁶ As expected, the arrival rate of blocks increased to > 15 minutes during the events, or by 59.5 percent compared to the population average of 9.5 minutes. Following the bottleneck in block creation, the number of transactions settled on the blockchain decreased by about 75,000 per day. In turn, the number of unprocessed transactions accumulated to 130,000 – an increase by factor 7.3 and a value only surpassed by the Bitcoin hype in December 2017. Eventually, the median confirmation time increased to 21.6 minutes for a transaction with average fees to be settled on the blockchain. This corresponds to an increase of 122 percent compared to the population average of 9.69 minutes.

²⁵In the history of Bitcoin, various technical improvements gradually increased the block capacity or established second-layer solutions to expand the settlement capacity (see Divakaruni and Zimmerman (2020) and Brown and Koepl (2019) among others). However, none of these adjustments were applied during our observation period, and therefore we treat block capacity as constant.

²⁶Non-shock intervals exclude the Bitcoin hype in December 2017, the price shock in December 2018, and Event [1] and [2].

Since transaction fees are subject to a generally increasing trend (see [Easley et al. \(2019\)](#)), we compare the movement during our shocks to the average in stable periods from January 2019 to May 2021. As predicted by our model, per block transaction fees jumped from an average of 0.59 BTC to a peak of 2.46 BTC during the event periods, corresponding to a factor of 4.16. Again, comparable fee levels are observed when the BTC price reaches a new historic all-time high, resulting in a massive increase in transaction volume. This interaction signifies that a short-term drop in computing capacity leads to a similar increase in transaction fees as we would expect from a short-term multiplication of the Bitcoin price. Event [1] and [2] show a widely similar development of executed hashing power, settlement capacity, and transaction fees. However, we observe a countervailing price trend between the shocks, which we exploit in the following subsection to examine the network vulnerability.

[Figure 6 about here]

4.5 Blockchain Fragility and Mining Shocks

Our primary interest in this subsection is to investigate the impact of temporary restrictions in the Chinese mining ecosystem on the fragility of the Bitcoin network.

As argued in subsection 3.3, PoW consensus remains incentive-compatible (i.e., robust) during mining shocks when parallel price movements offset declining attacking costs. Plunging mining capacity during the shock intuitively lowers the costs associated with a majority attack but lower competition, increasing transaction fees, and rising BTC prices may raise the net present value of operated equipment. As a result, an attacker’s increasing (opportunity) costs from more valuable equipment potentially render malicious actions economically unviable. This mechanism stimulates the economic incentives of honest behavior; the decentralized network remains robust and mining incentive compatible. Considering the evolution of relevant parameters in our observation period, we argue that the network’s fragility remained comparably equal during Event [1] but significantly increased during Event [2].

Unfortunately, the robustness of a decentralized ledger technology cannot be directly observed by any standard continuous metric. [Irresberger et al. \(2020\)](#) develop a measure of implied security across multiple cryptocurrencies from the escrow period installed by exchanges before considering a payment in a given currency irreversible. This figure reflects the perceived risk for a particular cryptocurrency of being compromised. However, it is not available in a continuous format that would allow us to decompose its variation. We, therefore, propose a proxy in the form of bid-ask spreads to empirically examine network stability.

Bid-ask spreads have been shown to be closely related to the stability of a market, as they implicitly reflect its liquidity. This interrelation arises as the market maker’s ability to enforce a specific spread level depends on market conditions. For example, in relatively illiquid markets, market makers impose wider spreads to offset risks associated with holding less liquid assets. Moreover, by providing a venue for buyers and sellers, liquidity providers in such markets can extract higher rents, given that fewer alternatives exist for trading the asset. In times of distress or general uncertainty, market-making becomes riskier, resulting in higher spreads and reduced exposure of liquidity providers even if markets are otherwise fairly liquid ([Anand & Venkataraman, 2016](#)). With risk-averse participants, market liquidity decreases during periods of high uncertainty ([Muranaga & Shimizu, 1999](#)). Thus, we expect bid-ask spreads to expand when sophisticated market participants, such as a major cryptocurrency exchange, evaluate market conditions as unstable (e.g., during an exogenous shock to the mining industry).

Several academic papers explore liquidity measures, such as bid-ask spreads, based on data from cryptocurrency exchanges ([Koutmos, 2018](#); [Scharnowski, 2021](#); [Scharnowski & Shi, 2021](#)). We exploit data from Kraken.com on open, close, high, and low trade prices for eight cryptocurrencies in US-Dollar and Euro and weight our results according to the trading volume in the respective fiat currency. Since the seminal work of [Roll \(1984\)](#), spread estimation from trade prices has seen considerable advances ([Abdi & Ranaldo, 2017](#); [Corwin & Schultz, 2012](#)). In this paper, we use the Efficient Discrete Generalized Estimator (EDGE) proposed by [Ardia et al. \(2021\)](#) to estimate effective

spread data from hourly open, close, high, and low trade prices. Compared to previous work, the EDGE-estimator relies on the most general conditions (e.g., includes non-frequent trade), encloses most information from discrete prices to minimize the estimation variance and therefore naturally produces fewer negative results. We aggregate daily effective spread estimates from hourly price data and ensure non-negativity by zero-setting negative results, as it is standard practice (Ardia et al., 2021).

Based on the argumentation above, we expect significant positive spreads during Event [2] compared to our baseline group of unaffected cryptocurrencies, while Event [1] exhibits no significant effect. To formally test this hypothesis, consider the unobserved structural model

$$\ln(Y_{it}) = \beta D_{it} + \gamma \ln(\mathbf{Z}_{it}) + \delta_t dt + \dots + \delta_T dT + \alpha_i + \epsilon_{it}, t = 1, \dots, T; i = 1, \dots, I, \quad (16)$$

where i identifies the cryptocurrency and t denotes each day in the observation interval t, \dots, T . D_{it} is a binary intervention indicator equal to 1 if a cryptocurrency is affected by the exogenous intervention at day t . In our baseline regression, the treatment group consists solely of Bitcoin, while seven cryptocurrencies are integrated as control groups. We examine the case of multiple affected currencies in subsection 4.6. \mathbf{Z}_{it} is a vector of control variables, including total trading volume (in all currencies), closing price in USD, and an indicator of volatility, calculated as the standard deviation of the closing price over the last three days. Each control variable has been identified as functional for the magnitude of bid-ask spreads (McInish & Wood, 1992). Standard panel unit root tests imply that closing prices are integrated by order $I(1)$ (see table 6) (Dickey & Fuller, 1979). However, first-differenced closing prices are stationary. We thus integrate closing prices using first differences and generally take the natural logarithm of all figures to interpret results as elasticities. Eventually, α_i denotes fixed effects to eliminate unobserved static heterogeneity among cryptocurrencies, and $\delta_t dt$ represents an exhaustive set of time-period dummies for each $t \in T$.

[Table 6 about here]

Although we integrate relevant control variables and time-period dummies, the observed correlation patterns may potentially be biased by the simultaneous variation of unobserved covariates. We apply a quasi-experimental (comparative) identification strategy to address this concern. To allow for a causal interpretation, we use a set of seven control cryptocurrencies with various consensus algorithms and differentiate the variation between treatment and control observations. The set of control entities includes Ethereum, Ethereum Classic, and Z-Cash (PoW), Algorand, Cardano, and Tezos (PoS), as well as EOS (DPoS). We only consider reasonably large cryptocurrencies in our control group with an average market capitalization of USD > 500M during the event periods.²⁷ Since Proof-of-Stake (PoS) and Delegated Proof-of-Stake (DPoS) consensus generally do not involve mining (and thus do not rely on large amounts of electricity), the respective currencies cannot be affected by our shocks. Moreover, the hash-rate evaluation of Ethereum, Ethereum Classic, and Z-Cash demonstrates that they were not affected by the intervention in Event [1] and [2] (see Figure 7). Notably, the currencies in our control group belong to the same asset class and trade on similar crypto-exchanges but are not affected by the shock. Therefore, we can isolate the intervention effect by comparing spreads from trading Bitcoin against US-Dollar and Euro to those of the cryptocurrencies in the control group.

We estimate equation 16 for a two-panel structure. Each panel contains the event period as specified in subsection 4.2 and the same amount of days preceding the intervention to balance shock and non-shock periods. Table 5 reports the coefficients with cluster-robust standard errors in parentheses, obtained from estimating equation 16 using currency fixed-effects regression. Regarding the first panel (Event [1]), column (1) shows the estimates without control and time-period variables, column (2) integrates the controls, and column (3) contains the estimates of the fully specified model. Results for the second panel (Event [2]) under the same reporting format are presented in columns (4), (5), and (6), respectively.

²⁷We exclude Dogecoin from the control group because the USD/DOGE exchange ratio and market capitalization multiplied by factor 10 between April 08 and May 05, 2021. This trajectory led to extensive volatility, which is significantly different from all other observed cryptocurrencies.

Consistent with our intuition, the coefficient D_{it} is insignificant on all common levels of statistical inference in the first panel. In contrast, the coefficient is of relevant magnitude and statistically significant in the second panel. This finding corroborates our hypothesis: The bid-ask spreads expanded considerably during Event [2], when the vulnerability of the Bitcoin blockchain was relatively high, whereas they showed no significant variation during Event [1] when price increases counterbalanced shrinking attacking costs. Table 5 further reveals that smaller cryptocurrencies exhibit significantly wider spreads. When considering the substantially smaller transaction volumes of those currencies as documented in Table 3, this result confirms the presumed connection between trade liquidity and spread magnitude. We conclude our analysis by performing additional robustness checks in the next subsection.

[Table 5 about here]

4.6 Robustness Checks

Our basic regression approach in subsection 4.5 employed a limited treatment group restricted solely to Bitcoin. However, Chinese miners targeting alternative PoW-based cryptocurrencies may have been simultaneously subject to constraints after the Covid-19 outbreak and coal mining accident in Xinjiang. Neglecting affected PoW-cryptocurrencies in the analysis might lead to incorrect effects estimations. Figure 7 illustrates the relative hash-rate evolution for a set of 7 PoW-based cryptocurrencies in the relevant period around each exogenous shock. In autumn of 2020, the mining activity of Bitcoin Cash (BCH), Litecoin (LTC), and Ripple (XMR) appears to be constrained in addition to Bitcoin. Interestingly, only Bitcoin Cash is subject to a similar decline in hashing power as Bitcoin in April 2021. To detect systematic spread variation in any affected PoW-currency, we utilize the entire set of PoW cryptocurrencies and adjust the treatment dummy D_{it} accordingly before repeating the estimation of the equation 16. Results are reported in Table 7 with cluster-robust standard errors in paratheses. All findings remain consistent. Overall, the robustness diagnosis lends further support to our interpretation

that the vulnerability of PoW-based consensus to exogenous shocks depends fundamentally on the concentration in the mining ecosystem and the covariation in the price of the proprietary coin.

[Figure 7 about here]

[Table 7 about here]

5 Conclusion

This paper studies the robustness of PoW-based permissionless blockchains with structurally concentrated mining ecosystems against exogenous shocks. Based on existing literature, we theoretically characterize the mining game and formally describe the economic incentive compatibility underlying PoW consensus with homogeneous and heterogeneous miners. The empirical section studies two exogenous shocks to the Chinese mining ecosystem in October 2020 and April 2021. The analysis reveals that the structural parameters, such as the hash-rate, settlement capacity, and transaction fees, were indeed exceptional. Our empirical analysis further demonstrates that the impact of exogenous shocks on the stability of PoW-based consensus substantially depends on fundamental parameters, such as covariation in the price of the cryptocurrency’s native coin. Moreover, we show that market participants incorporate and price variations in the implied stability of the distributed ledger.

References

- Abadi, J., & Brunnermeier, M. (2019). Blockchain Economics. *National Bureau of Economic Research*. <https://doi.org/10.3386/w25407>
- Abdi, F., & Ranaldo, A. (2017). A simple estimation of bid-Ask spreads from daily close, high, and low prices. *Review of Financial Studies*, 30(12), 4437–4480. <https://doi.org/10.1093/rfs/hhx084>
- Al-kuwari, S., Davenport, J. H., & Bradford, R. J. (2011). Cryptographic Hash Functions: Recent Design Trends and Security Notions. *Eprint.Iacr.Org*, 1–36. <http://eprint.iacr.org/2011/565.pdf>
- Anand, A., & Venkataraman, K. (2016). Market conditions, fragility, and the economics of market making. *Journal of Financial Economics*, 121(2). <https://doi.org/10.1016/j.jfineco.2016.03.006>
- Ardia, D., Guidotti, E., & Kroencke, T. A. (2021). Efficient Estimation of Bid-Ask Spreads from Open, High, Low, and Close Prices. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3892335>
- Arnosti, N., & Weinberg, M. (2021). Bitcoin: A Natural Oligopoly. *Management Science*, (forthcoming). <https://doi.org/10.4230/LIPIcs.ITCS.2019.5>
- Auer, R. (2019). Beyond the Doomsday Economics of “Proof-of-Work” in Cryptocurrencies. *Bank For International Settlements, BIS Working(765)*, 1–31. <https://doi.org/10.24149/gwp355>
- Basu, S., Easley, D., Hara, M. O., & Emin, G. (2021). StableFees : A Predictable Fee Market for Cryptocurrencies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3318327>
- BBC. (2020). Covid-19: China tests entire city of Kashgar in Xinjiang. <https://www.bbc.com/news/world-asia-china-54687533>
- Benetton, M., Compiani, G., & Morse, A. (2021). When Cryptomining Comes to Town: High Electricity-Use Spillovers to the Local Economy. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3779720>

- Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., & Menkveld, A. J. (2021). Equilibrium Bitcoin Pricing. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3261063>
- Biais, B., Bisière, C., Bouvard, M., & Casamatta, C. (2019). The Blockchain Folk Theorem. *Review of Financial Studies*, 32(5), 1662–1715. <https://doi.org/10.1093/rfs/hhy095>
- Brown, C., & Koepl, T. V. (2019). What Drives Bitcoin Fees? Using Segwit to Assess Bitcoin’s Long-run Sustainability. *Queen’s Economics Department Working Paper*, (No. 1423). <https://doi.org/10.1257/jep.29.2.213>
- Budish, E. B. (2018). The Economic Limits of Bitcoin and the Blockchain. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3197300>
- Capponi, A., Olafsson, S., & Alsabah, H. (2021). Proof-of-Work Cryptocurrencies: Does Mining Technology Undermine Decentralization? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3869144>
- Chen, L., Cong, L. W., & Xiao, Y. (2020). A brief introduction to blockchain economics. *Information For Efficient Decision Making: Big Data, Blockchain And Relevance*, 1–40. https://doi.org/10.1142/9789811220470_0001
- Chiu, J., & Koepl, T. V. (2017). The Economics of Cryptocurrencies Bitcoin and Beyond. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3048124>
- CNN. (2021). 21 Chinese miners trapped by underground flood in Xinjiang coal mine. Retrieved December 3, 2021, from <https://edition.cnn.com/2021/04/11/china/xinjiang-coal-mine-accident-intl-hnk/index.html>
- Cong, L. W., He, Z., & Li, J. (2020). Decentralized Mining in Centralized Pools. *The Review of Financial Studies*, 00(607), 1–40. <https://doi.org/10.1093/rfs/hhaa040>
- Corwin, S. A., & Schultz, P. (2012). A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices. *Journal of Finance*, 67(2), 719–760. <https://doi.org/10.1111/j.1540-6261.2012.01729.x>
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the Estimators for Autoregressive Time Series With a Unit Root. *Journal of the American Statistical Association*, 74(366), 427. <https://doi.org/10.2307/2286348>

- Dimpfl, T., & Peter, F. J. (2021). Nothing but noise? Price discovery across cryptocurrency exchanges. *Journal of Financial Markets*, 54, 100584. <https://doi.org/10.1016/j.finmar.2020.100584>
- Divakaruni, A., & Zimmerman, P. (2020). Ride the Lightning: Turning Bitcoin into Money. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3514125>
- Easley, D., O'Hara, M., & Basu, S. (2019). From Mining to Markets: The Evolution of Bitcoin Transaction Fees. *Journal of Financial Economics*, 134(1), 91–109. <https://doi.org/10.1016/j.jfineco.2019.03.004>
- Eyal, I., & Sirer, E. G. (2014). Majority is not enough: Bitcoin Mining is Vulnerable. In: Christin N., Safavi-Naini R. (eds) *Financial Cryptography and Data Security, FC 2014, Lecture Notes in Computer Science*, 8437, 436–454. https://doi.org/10.1007/978-3-662-45472-5_28
- Fantazzini, D., & Kolodin, N. (2020). Does the Hashrate Affect the Bitcoin Price? *Journal of Risk and Financial Management*, 13(11), 263. <https://doi.org/10.3390/jrfm13110263>
- Garratt, R., & van Oordt, M. R. (2020). Why Fixed Costs Matter for Proof-of-Work Based Cryptocurrencies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3572400>
- Griffin, J. M., & Shams, A. (2020). Is Bitcoin Really Untethered? *Journal of Finance*, 75(4), 1913–1964. <https://doi.org/10.1111/jofi.12903>
- Halaburda, H., Haeringer, G., Gans, J. S., & Gandal, N. (2021). The Microeconomics of Cryptocurrencies. *Journal of Economic Literature*, (forthcoming).
- Hernández, J. C. (2020). China Locks Down Xinjiang to Fight Covid-19, Angering Residents. <https://www.nytimes.com/2020/08/25/world/asia/china-xinjiang-covid.html>
- Hileman, G., & Rauchs, M. (2018). 2017 Global Blockchain Benchmarking Study. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3040224>

- Huberman, G., Leshno, J. D., & Moallemi, C. (2021). Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. *The Review of Economic Studies*, 88(6), 3011–3040. <https://doi.org/10.1093/restud/rdab014>
- Irresberger, F., John, K., & Saleh, F. (2020). The Public Blockchain Ecosystem: An Empirical Analysis. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3592849>
- John, K., Rivera, T., & Saleh, F. (2020). Economic Implications of Scaling Blockchains : Why the Consensus Protocol Matters. *SSRN Electronic Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract%7B%5C_%7Did=3750467
- K Wan, A., Liu, L., Luo, J., Lam, E., Lee, J., & Ossinger, J. (2021). Bitcoin (\$BTC USD) News: China PBOC Says Crypto-Related Transactions Illegal. <https://www.bloomberg.com/news/articles/2021-09-24/china-deems-all-crypto-related-transactions-illegal-in-crackdown>
- Kaiser, B., Jurado, M., & Ledger, A. (2018). The looming threat of China: An analysis of Chinese influence on Bitcoin. *arXiv*, 1–22. <https://doi.org/1810.02466>
- Koutmos, D. (2018). Liquidity uncertainty and Bitcoin’s market microstructure. *Economics Letters*, 172. <https://doi.org/10.1016/j.econlet.2018.08.041>
- Küfeoğlu, S., & Özkuran, M. (2019). Energy Consumption of Bitcoin Mining. *Cambridge Working Papers in Economics*. <https://doi.org/https://doi.org/10.17863/CAM.41230>
- Lehar, A., & Parlour, C. A. (2020). Miner Collusion and the BitCoin Protocol. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3559894>
- Leshno, J. D., & Strack, P. (2020). Bitcoin: An Axiomatic Approach and an Impossibility Theorem. *American Economic Review: Insights*, 2(3), 269–286. <https://doi.org/10.1257/aeri.20190494>
- Li, T., Shin, D., & Wang, B. (2018). Cryptocurrency Pump-and-Dump Schemes. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3267041>
- Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6), 2689–2727. <https://doi.org/10.1093/rfs/hhaa113>

- Ma, J., Gans, J. S., & Tourky, R. (2019). Market Structure in Bitcoin Mining. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3103104>
- Makarov, I., & Schoar, A. (2020). Trading and Arbitrage in Cryptocurrency Markets. *Journal of Financial Economics*, 135(2), 293–319. <https://doi.org/10.1016/j.jfineco.2019.07.001>
- Makarov, I., & Schoar, A. (2021). Blockchain Analysis of the Bitcoin Market. *SSRN Electronic Journal*. <https://ssrn.com/abstract=3942181>
- McInish, T. H., & Wood, R. A. (1992). An Analysis of Intraday Patterns in Bid/Ask Spreads for NYSE Stocks. *The Journal of Finance*, 47(2), 753. <https://doi.org/10.2307/2329122>
- Möser, M., & Böhme, R. (2015). Trends, Tips, Tolls: A Longitudinal Study of Bitcoin Transaction Fees. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8976, 19–33. https://doi.org/10.1007/978-3-662-48051-9_2
- Muranaga, J., & Shimizu, T. (1999). Market Microstructure and Market Liquidity. *Committee on the Global Financial System Publications*, 11.
- Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System. *www.bitcoin.org*. <https://bitcoin.org/bitcoin.pdf>
- Naor, M., & Yung, M. (1989). Universal One-Way Hash Functions and their Cryptographic Applications. *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, 33–43. <https://doi.org/10.1145/73007.73011>
- NBS. (2021). National Bureau of Statistics China. Retrieved September 1, 2021, from <https://data.stats.gov.cn/english/index.htm>
- Pagnotta, E. (2021). Decentralizing Money: Bitcoin Prices and Blockchain Security. *The Review of Financial Studies*. <https://doi.org/10.1093/rfs/hhaa149>
- Pagnotta, E., & Buraschi, A. (2018). An Equilibrium Valuation of Bitcoin and Decentralized Network Assets. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3142022>

- Prat, J., & Walter, B. (2021). An Equilibrium Model of the Market for Bitcoin Mining. *Journal of Political Economy*, 129(8), 2415–2452. <https://doi.org/10.1086/714445>
- Rauchs, M., Blandin, A., Dek, A., & Wu, Y. (2021). Cambridge Bitcoin Electricity Consumption Index. <https://www.cbeci.org/cbeci/comparisons>
- Rauchs, M., Blandin, A., Klein, K., Pieters, G. C., Recanatini, M., & Zhang, B. Z. (2019). 2nd Global Cryptoasset Benchmarking Study. *SSRN Electronic Journal*, (December). <https://doi.org/10.2139/ssrn.3306125>
- Reuters. (2020). China’s Kashgar detects 137 new asymptomatic COVID cases. <https://www.reuters.com/article/us-health-coronavirus-china-cases-idUSKBN27A00Z>
- Roll, R. (1984). A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *The Journal of Finance*, 39(4), 1127–1139. <https://doi.org/10.1111/j.1540-6261.1984.tb03897.x>
- Savolainen, V., & Ruiz-Ogarrio, J. (2020). Too Big to Cheat: Mining Pools’ Incentives to Double Spend in Blockchain Based Cryptocurrencies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3506748>
- Schär, F., & Berentsen, A. (2020). *Bitcoin, Blockchain, and Cryptoassets - A Comprehensive Introduction*. MIT Press.
- Scharnowski, S. (2021). Understanding Bitcoin liquidity. *Finance Research Letters*, 38. <https://doi.org/10.1016/j.frl.2020.101477>
- Scharnowski, S., & Shi, Y. (2021). Bitcoin Blackout : Proof-of-Work and the Centralization of Mining. *SSRN Electronic Journal*, 1–2.
- Shanaev, S., Shuraeva, A., Vasenin, M., & Kuznetsov, M. (2020). Cryptocurrency value and 51% attacks: evidence from event studies. *SSRN Electronic Journal*. <https://doi.org/http://dx.doi.org/10.2139/ssrn.3290016>
- Sichuan Government. (2019). The approval of the Yazhong-Jiangxi UHV DC transmission project will effectively alleviate the problem of excess hydropower in Panxi (in Chinese). <http://www.sc.gov.cn/10462/12771/2019/8/28/14c0331ca0b44234a04ce11c40ca35cd.shtml>

- Song, Y.-D., & Aste, T. (2020). The Cost of Bitcoin Mining has Never Really Increased. *Front. Blockchain*, 3(October), 1–8. <https://doi.org/10.2139/ssrn.3574512>
- Stinner, J. (2021). Inside the Black Box of Bitcoin Mining : A Theoretical Framework and Simulation Evidence. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3853812>
- Stinner, J., & Tyrell, M. (2021). The New World of Blockchain Economics: Consensus Mechanism as a Core Element. In T. A. Herberger & J. J. Dötsch (Eds.), *Digital transformation and sustainability in the global economy - risks and opportunities* (pp. 9–19). Springer Proceedings in Business; Economics.
- Wang, W., Hoang, D. T., Hu, P., Xiong, Z., Niyato, D., Wang, P., Wen, Y., & Kim, D. I. (2019). A Survey on Consensus Mechanisms and Mining Strategy Management in Blockchain Networks. *IEEE Access*, 7, 22328–22370. <https://doi.org/10.1109/ACCESS.2019.2896108>
- WHO. (2021). World Health Organization - Health Emergency Dashboard. Retrieved September 3, 2021, from <https://covid19.who.int/region/wpro/country/cn>
- Xinjiang Government. (2020). Provincial Health Commission. Retrieved September 3, 2021, from <https://wjw.xinjiang.gov.cn/>
- Zhao, W. (2020). Bitcoin hash rate drops over 10% as Chinese miners migrate to fossil fuel plants. <https://www.theblockcrypto.com/post/82404/bitcoin-hashrate-miners-migrate>
- Zhong, R., & Mozur, P. (2020). To Tame Coronavirus, Mao-Style Social Control Blankets China. *The New York Times*. <https://www.nytimes.com/2020/02/15/business/%7B%5C%7D0Achina-coronavirus-lockdown.html>

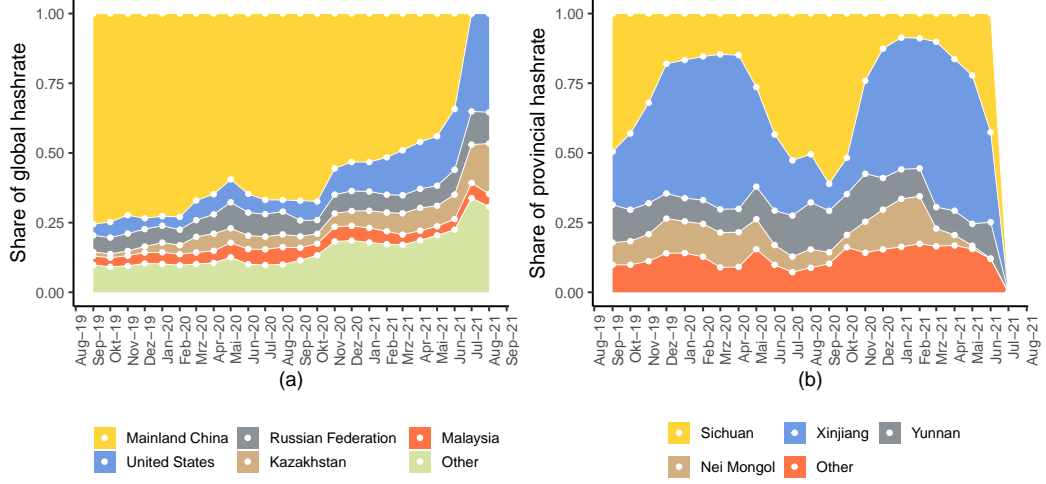


Figure 1: Global and provincial hash-rate distribution (Sep.2019-Aug.2021). Panel (a) shows the monthly percentage composition of the global hash-rate by country. "Other" summarizes countries contributing < 5 percent during the observation period (e.g., Iran, Canada, Germany, and Ireland). Panel (b) displays the percentage hash-rate distribution localized in China, classified by provinces and months. Provinces contributing < 5 percent during the observation period are summarized (e.g., Gansu, Zhejiang, and Beijing). Both plots are based on data from [Rauchs et al. \(2021\)](#).

Table 1: This table contains descriptive statistics on daily transaction and network parameters. Figures are constructed from 44,204 blocks (i.e., block 637091 to 681295) that were appended to the Bitcoin blockchain between July 2020 and May 2021. See subsection 4.3 for further details.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Hash-rate (EH/s)	304	140.5	17.5	97.9	127.2	153.6	186.4
Block Difficulty (Trillion)	304	19.4	2.2	15.8	17.3	21.4	23.6
Mean Block Time (Seconds)	304	601.0	69.4	469.6	555.6	627.2	915.2
Block Number	304	145.4	14.8	91	138	156	185
Transaction Number (100K)	304	310.4	34.1	194.6	290.4	333.8	404.8
Transaction Volume (M BTC)	304	1.91	0.66	0.59	1.45	2.28	4.19
Sum of Block Rewards (BTC)	304	909.0	92.7	568.8	862.5	975.0	1,156.2
Sum of Transaction Fees (BTC)	304	109.0	49.9	24.6	75.8	134.6	301.8
BTC Price (1K USD)	304	27.2	18.3	9.1	11.4	46.7	63.3
Mempool Transactions (1000)	417	34.7	28.0	0.4	12.7	47.8	135.9
Confir. Time (Minutes)	304	12.7	3.4	5.0	10.2	14.4	25.2




Figure 2: Global hash-rate and block difficulty (Sep.2020-May.2021). This figure shows the 3-day moving average of the estimated global hash-rate and the standard block difficulty. Event periods from October 16 to November 23, 2020 and April 16 to April 25, 2021, are highlighted by dashed vertical lines.

Table 2: This table provides daily descriptive statistics on effective bid-ask spreads, closing prices, trading volume (in all currencies), and volatility (defined as the standard deviation of the closing price in a 3-day period) for nine cryptocurrencies and an observation period from January 2020 to October 2021. See subsection 4.3 for further details.

Variable	Levels	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Bid-Ask Spread (%)	ADA	638	0.4	0.2	0.0	0.3	0.5	2.4
	ALGO	617	0.7	0.4	0.0	0.5	0.8	4.1
	BTC	638	0.1	0.1	0.0	0.1	0.1	0.8
	EOS	638	0.4	0.2	0.0	0.3	0.5	2.1
	ETC	638	0.6	0.3	0.0	0.4	0.7	3.2
	ETH	638	0.2	0.1	0.0	0.1	0.2	0.9
	XMR	638	0.4	0.2	0.0	0.3	0.5	2.2
	XTZ	638	0.5	0.2	0.0	0.3	0.6	3.0
	ZEC	638	0.6	0.3	0.0	0.4	0.7	2.4
	all	5721	0.4	0.3	0.0	0.2	0.6	4.1
Closing Price (USD)	ADA	638	0.6	0.8	0.0	0.1	1.2	3.0
	ALGO	638	0.6	0.5	0.1	0.3	1.0	2.4
	BTC	638	25389.0	17903.9	4970.8	9512.5	40621.8	63503.5
	EOS	638	3.7	1.6	1.8	2.6	4.3	14.4
	ETC	638	20.9	24.3	4.0	6.0	33.6	134.1
	ETH	638	1170.6	1122.3	110.6	236.4	2077.8	4168.7
	XMR	638	155.2	97.0	33.0	68.3	229.6	483.6
	XTZ	638	3.2	1.3	1.2	2.3	3.6	7.5
	ZEC	638	94.6	56.9	24.5	55.0	126.6	318.9
	all	5742	2982.1	9930.0	0.0	2.3	202.3	63503.5
Trading Volume (\$M)	ADA	638	2318.9	2987.4	20.8	211.9	3533.6	19142.0
	ALGO	638	222.9	358.6	17.0	60.5	255.8	4812.1
	BTC	638	40915.1	21730.5	12252.6	27186.0	49062.1	350967.9
	EOS	638	2907.8	2020.1	673.6	1682.1	3521.0	20328.7
	ETC	638	2103.3	2986.8	373.3	765.8	2233.6	42721.4
	ETH	638	20978.4	11776.3	5109.0	12743.5	25825.4	84482.9
	XMR	638	580.5	1793.0	41.8	111.8	738.2	28959.1
	XTZ	638	290.5	300.2	30.5	119.1	359.3	2721.4
	ZEC	638	609.7	993.5	90.8	283.4	620.5	12719.4
	all	5742	7880.8	15673.4	17.0	229.3	5868.1	350967.9
Volatility	ADA	638	0.0	0.0	0.0	0.0	0.0	0.3
	ALGO	638	0.0	0.0	0.0	0.0	0.0	0.5
	BTC	638	688.5	782.6	6.9	125.8	1031.6	4397.8
	EOS	638	0.2	0.3	0.0	0.0	0.2	2.7
	ETC	638	1.1	2.6	0.0	0.1	1.0	32.6
	ETH	638	44.7	63.1	0.6	5.3	66.1	505.0
	XMR	638	5.7	7.4	0.2	1.6	6.7	81.1
	XTZ	638	0.1	0.2	0.0	0.0	0.2	1.1
	ZEC	638	4.3	5.3	0.1	1.3	5.3	56.6
	all	5742	82.7	338.3	0.0	0.0	6.2	4397.8




Figure 3: Weekly block number during two shocks and overall block density. Panel (a) shows the decrease of weekly blocks during two exogenous shocks together with the population average (red line) and the block difficulty adjustments (yellow line). Panel (b) displays the overall density of weekly blocks, including the population mean (red line). The block creation during the two shocks is highlighted (red).

Table 4: This table provides results from one-sample t-test against the two sided alternative of the sample mean or population mean for the number of blocks and network hash-rate. The test calculation includes population mean and standard deviation for the block number, whereas tests on the hash-rate deploy sample figures. Test periods cover the relevant weeks around the presumed shocks (see subsection 4.2).

Panel 1

	Week	Sum of Blocks	t -Value (Blocks)	p -Value (Blocks)	Average Hash-Rate	t -Value (Hash-Rate)	p -Value (Hash-Rate)
1	2020-09-25	1021	-1.72	0.1366	140.95	-0.12	0.905
2	2020-10-02	1028	-1.52	0.1798	140.33	-0.27	0.799
3	2020-10-09	1040	-1.18	0.2842	142.62	0.47	0.655
4	2020-10-16	968	-3.23	0.0179**	138.65	-0.53	0.617
5	2020-10-23	817	-7.54	3e-04***	118.55	-3.31	0.016**
6	2020-10-30	890	-5.45	0.0016***	115.16	-6.34	0.001***
7	2020-11-06	1077	-0.12	0.9081	127.37	-5.69	0.001***
8	2020-11-13	1054	-0.78	0.467	130.31	-3.20	0.019**
9	2020-11-20	1099	0.51	0.6301	137.73	-0.81	0.447
10	2020-11-27	1011	-2.00	0.092*	133.03	-2.85	0.029**

Panel 2

	Week	Sum of Blocks	t -Value (Blocks)	p -Value (Blocks)	Average Hash-Rate	t -Value (Hash-Rate)	p -Value (Hash-Rate)
1	2021-03-26	1054	-0.78	0.467	165.26	-0.71	0.505
2	2021-04-02	1008	-2.09	0.0817*	167.70	0.18	0.866
3	2021-04-09	1047	-0.98	0.3667	168.61	0.29	0.785
4	2021-04-16	787	-8.39	2e-04***	134.58	-4.97	0.003***
5	2021-04-23	941	-4.00	0.0071***	160.65	-2.08	0.083*
6	2021-04-30	1154	2.08	0.0832*	170.61	0.82	0.444
7	2021-05-07	1214	3.79	0.0091***	182.92	4.71	0.003***

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$




Figure 4: Evolution of Bitcoin price and global hash-rate during four shocks (2018-2021). This figure shows four shocks to the Bitcoin mining industry (upper panel) and the corresponding change in the BTC price (lower panel). Shocks are defined as significant short-term recessions in the global hash-rate of > 25 percent within 14 days. All Variables are standardized with $t_0 = 1$ and a period centered ± 20 days around the observed hash-rate minimum values.




Figure 5: Evolution of miners' gross revenues and global hash-rate for two mining shocks. This figure depicts miners' gross revenues and the global hash-rate for two exogenous shocks in October 2020 and April 2021. Gross revenues are the product of block rewards, transaction fees, and daily USD/BTC conversion rates. Figures are standardized with $t_0 = 1$.

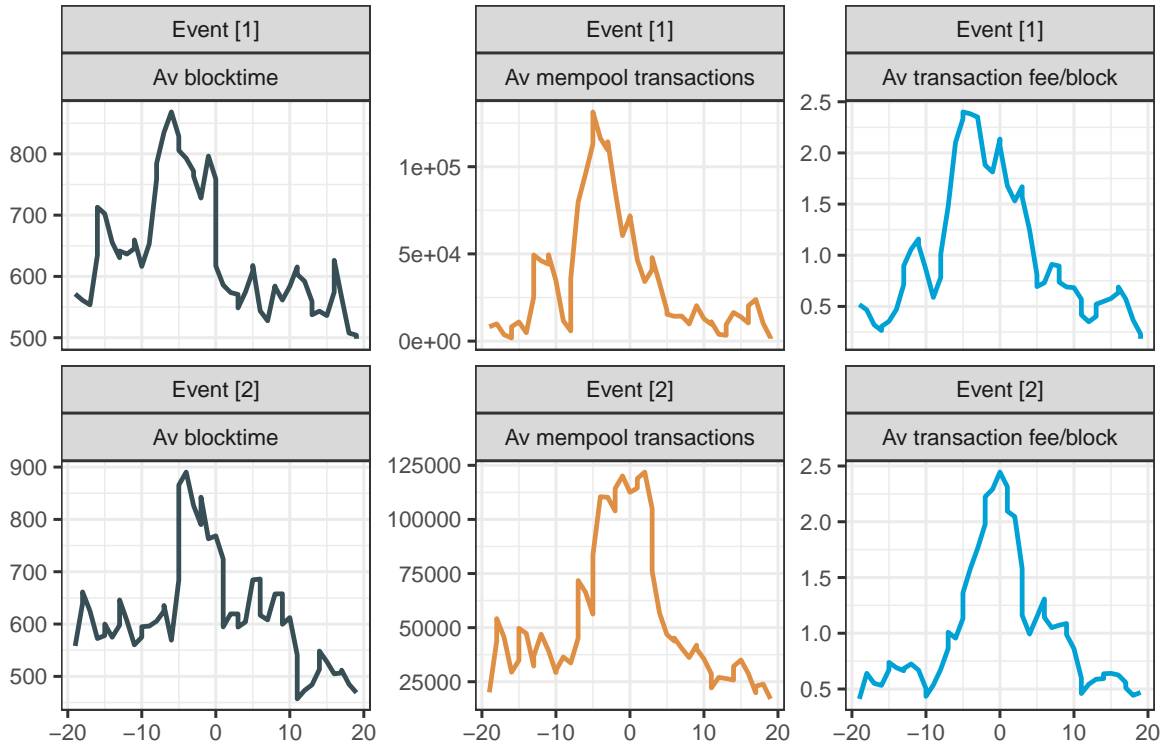


Figure 6: Effects of exogenous shocks on transaction parameters. This chart illustrates the average of block time (in seconds), number of mempool transactions, and transaction fees (in BTC per block), with $t = 0$ corresponding to the observed minimum hash-rate of each shock.




Figure 7: PoW-based cryptocurrencies and mining shocks. This chart illustrates the standardized hash-rate of seven Pow-based cryptocurrencies during the exogenous shocks in October 2020 and April 2021. All figures are standardized with $t_0 = 1$; vertical (red) lines indicate shock periods.

Table 5: Currency fixed-effects regression: Effect of shock treatment dummy (D_{it}) on $\ln(\text{Average Spread})$. This table reports regression results from estimating equation 16.

Figures are specified as follows: The log average spread denotes the estimated mean bid-ask spread from trading the respective cryptocurrency against USD and EUR (weighted by volume), log volume the overall trading volume in all currencies, the change of log closing price denotes the first difference of the daily closing price in USD, and log volatility the moving standard deviation of the closing price in a 3-day period. The shock treatment dummy (D_{it}) is equal to one if the currency at day t was affected by a shock (denoted by (1) and (2)) and zero otherwise. The sample period of Panel 1 (September 9 to November 23, 2020) and Panel 2 (April 3 to April 25, 2021) contains 600 and 184 observations, respectively, for a total of 8 cryptocurrencies.

	<i>Dependent variable:</i>					
	Ln(Average Spread)					
	<i>Panel 1</i>			<i>Panel 2</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
treatment (D_{it}^1)	0.16 (0.23)	0.19 (0.17)	0.32 (0.21)			
treatment (D_{it}^2)				0.55*** (0.00)	0.42*** (0.02)	0.19*** (0.05)
$\ln(\text{Volume})$		0.24** (0.07)	0.17** (0.03)		0.47*** (0.12)	0.37** (0.12)
$\ln(\text{Volatility})$		0.08** (0.02)	0.01 (0.05)		0.04 (0.05)	0.02 (0.07)
$\Delta \ln(\text{Closing Price})$		-0.60 (0.58)	1.20 (1.03)		-1.56*** (0.35)	-0.28 (0.38)
ALGO	0.58*** (0.00)	0.95*** (0.13)	0.92*** (0.12)	0.53*** (0.00)	1.59*** (0.29)	1.38*** (0.28)
BTC	-1.80** (0.12)	-3.61*** (0.51)	-2.66*** (0.50)	-1.67*** (0.00)	-3.29*** (0.36)	-2.75*** (0.74)
EOS	-0.19*** (0.00)	-0.66*** (0.12)	-0.39** (0.11)	0.51*** (0.00)	0.40*** (0.11)	0.45** (0.16)
ETC	0.62*** (0.00)	0.37*** (0.08)	0.60** (0.18)	0.84*** (0.00)	0.94*** (0.20)	0.93** (0.26)
ETH	-1.04*** (0.00)	-2.39*** (0.33)	-1.64*** (0.35)	-0.71*** (0.00)	-1.95*** (0.27)	-1.63** (0.53)
XTZ	0.02*** (0.00)	0.12 (0.10)	0.24 (0.18)	0.75*** (0.00)	1.56*** (0.30)	1.42*** (0.29)
ZEC	0.22 (0.12)	-0.28 (0.22)	0.09 (0.35)	0.78*** (0.00)	1.26** (0.40)	1.20** (0.46)
				(0.00)	(0.02)	(0.05)
Time Per. D.	NO	NO	YES	NO	NO	YES
R ²	0.78	0.80	0.84	0.78	0.85	0.90
Adj. R ²	0.77	0.79	0.82	0.77	0.84	0.88
Num. obs.	600	600	600	184	184	184
RMSE	0.41	0.39	0.37	0.42	0.36	0.30
N Clusters	8	8	8	8	8	8

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Table 6: This table reports results from the Augmented Dickey-Fuller test (ADF-test) for unit roots in panel data structures (Dickey & Fuller, 1979).

Variable	Unit root test	Test statistic	<i>p</i> -value
In levels:			
<i>log</i> (Average Spread)	ADF (LP=17)	-8.244	0.0000
<i>log</i> (Trading Volume)	ADF (LP=17)	-3.9455	0.0116
<i>log</i> (Closing Price)	ADF (LP=17)	-1.691	0.7092
<i>log</i> (Volatility)	ADF (LP=17)	-8.0129	0.0000
In first-differences:			
$\Delta \log$ (Closing Price)	ADF (LP=17)	-15.762	0.000

Table 7: Currency fixed-effects regression: Effect of shock treatment dummy (D_{it}) on $\ln(\text{Average Spread})$. This table reports regression results from estimating equation 16. Figures are specified as follows: The log average spread denotes the estimated mean bid-ask spread from trading the respective cryptocurrency against USD and EUR (weighted by volume), log volume the overall trading volume in all currencies, the change of log closing price denotes the first difference of the daily closing price in USD, and log volatility the moving standard deviation of the closing price in a 3-day period. The shock treatment dummy (D_{it}) is equal to one if the currency at day t was affected by a shock (denoted by (1) and (2)) and zero otherwise. The sample period of Panel 1 (September 9 to November 23, 2020) and Panel 2 (April 3 to April 25, 2021) contains 825 and 253 observations, respectively, for a total of 11 cryptocurrencies.

	<i>Dependent variable:</i>					
	Ln(Average Spread)					
	<i>Panel 1</i>			<i>Panel 2</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
treatment (D_{it}^1)	0.11 (0.08)	0.04 (0.08)	0.20 (0.11)			
treatment (D_{it}^2)				0.56*** (0.02)	0.33*** (0.09)	0.12 (0.08)
$\ln(\text{Volume})$		0.18** (0.04)	0.11** (0.04)		0.46*** (0.09)	0.39*** (0.11)
$\ln(\text{Volatility})$		0.09*** (0.02)	0.02 (0.03)		0.02 (0.03)	-0.02 (0.04)
$\Delta \ln(\text{Closing Price})$		-0.79 (0.51)	0.69 (0.82)		-1.38*** (0.29)	-0.23 (0.29)
ALGO	0.58*** (0.00)	0.82*** (0.10)	0.78*** (0.12)	0.53*** (0.00)	1.58*** (0.22)	1.45*** (0.26)
BCH	-0.48*** (0.04)	-1.33*** (0.20)	-0.81*** (0.23)	0.17** (0.01)	0.03 (0.23)	0.39 (0.32)
BTC	-1.77*** (0.04)	-3.47*** (0.31)	-2.50*** (0.28)	-1.68*** (0.01)	-3.06*** (0.27)	-2.39*** (0.49)
EOS	-0.19*** (0.00)	-0.63*** (0.08)	-0.36*** (0.07)	0.51*** (0.00)	0.44*** (0.07)	0.54*** (0.10)
ETC	0.62*** (0.00)	0.32*** (0.09)	0.56*** (0.12)	0.84*** (0.00)	0.99*** (0.14)	1.07*** (0.17)
ETH	-1.04*** (0.00)	-2.34*** (0.23)	-1.56*** (0.20)	-0.71*** (0.00)	-1.80*** (0.18)	-1.40*** (0.34)
LTC	-0.35*** (0.04)	-1.11*** (0.16)	-0.68*** (0.18)	0.18*** (0.00)	-0.16 (0.16)	0.08 (0.24)
XMR	-0.08 (0.04)	-0.74*** (0.19)	-0.32 (0.24)	0.30*** (0.00)	0.97** (0.30)	1.07** (0.34)
XTZ	0.02*** (0.00)	0.00 (0.11)	0.11 (0.15)	0.75*** (0.00)	1.57*** (0.22)	1.54*** (0.25)
ZEC	0.25*** (0.04)	-0.30 (0.18)	0.07 (0.24)	0.78*** (0.00)	1.33*** (0.29)	1.45*** (0.33)
Time Per. D.	NO	NO	YES	NO	NO	YES
R ²	0.73	0.75	0.81	0.75	0.82	0.89
Adj. R ²	0.73	0.75	0.78	0.74	0.81	0.87
Num. obs.	825	825	825	253	253	253
RMSE	0.40	0.39	0.36	0.40	0.34	0.29
N Clusters	11	11	11	11	11	11

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$