Online Appendix for

Explaining Sustained Blockchain Decentralization with Quasi-Experiments: The Resource Flexibility of Consensus Mechanisms

Harang Ju[1], Madhav Kumar[2], Ehsan Valavi[3], Sinan Aral[3]
[1]Johns Hopkins Carey Business School, harang@jhu.edu
[2]Harvard Business School, mkumar@hbs.edu
[3]MIT Sloan School of Management, {evalavi, sinan}@mit.edu

A Literature Overview

Table A1 summarizes key studies on blockchain decentralization, their methodological approaches, main findings, and boundary conditions.

B Robustness Check with Consensus Layer Covariates

To test the robustness of our main difference-in-difference results of China's mining ban in Section 4.1, we include consensus layer covariates in our analysis. As explained in more detail in Section 3.1, our analyses focus on the consensus layers of Bitcoin and Ethereum. While the two blockchains differ in the types of applications they support, the consensus layer is agnostic to the application layer. The consensus layer simply manages the secure production of blocks (Nakamoto, 2008; Buterin, 2014) and indirectly the economic incentives of block production.

In our robustness tests, we include pre-ban z-scored averages of daily metrics related to block production and economic incentives as covariates to isolate the effect of resource flexibility in response to China's ban. For block production, metrics include the number of transactions, the amount of data stored in megabytes (MB), and the estimated hashrates in megahashes per second (MH/s).[1] For economic incentives, we use the average transaction fee in US dollars ($), excluding token price and block reward due to severe multicollinearity ($VIFs > 167,000$).[2] The Bitcoin metrics

[1]See Section 3.2 for an explanation of hashrates. For reference, the hashrates were 174 exahashes per second for Bitcoin and 632 terahashes per second for Ethereum at their peak before the ban.
[2]We excluded block reward and token price, as miner revenue can be derived from the transaction fee and number of transactions, avoiding redundancy.

were obtained from blockchain.com, and the Ethereum metrics were obtained from etherscan.io.[3] We remove the Exposure variable, as it is directly derived from hashrates and captured by our covariates.

In our robustness check, we find that our overall results are robust to these covariates (Table B2) with and without interactions with the treatment variable (Table B3). The treatment effects on entropy, the Nakamoto coefficient, and HHI are all qualitatively similar but reduced when interactions are added. Interestingly, the treatment effects disappear for the Nodes and Gini metrics. This attenuation is consistent with these metrics being partially driven by consensus-layer characteristics that correlate with both the treatment and participation patterns. Our primary measure, Shannon entropy, remains robust to covariate interactions, reinforcing its use as our preferred decentralization metric. This result suggests two nuanced effects and highlights the importance of using multiple decentralization metrics to understand decentralization. First, the shock had an indistinguishable effect on the number of nodes in both Bitcoin and Ethereum. Second, the inequality in block production among the remaining nodes, as measured by the Gini coefficient, remains the same.

[3]Only weekly metrics were available for Bitcoin, so the daily metrics were forward-filled from the weekly metrics.

Study	*Approach*	*Key Findings*	*Boundary Conditions*
Gencer et al. (2018)	Empirical measurement	Bitcoin and Ethereum networks are not as decentralized as assumed; mining pools dominate	Snapshot analysis; limited time period
Cong et al. (2021)	Theoretical model	Decentralized consensus can emerge despite economies of scale through appropriate mechanism design	Model assumptions; permissionless setting
Cong et al. (2023)	Causal identification	Layer-2 scaling increases decentralization in oracle data providers	Oracle subsystem only; not consensus layer
Capponi et al. (2023)	Theoretical model	Mining technology (ASICs vs. GPUs) affects decentralization through economies of scale	PoW blockchains only
Garratt and van Oordt (2023)	Theoretical model	Fixed costs in mining increase resilience to 51% attacks when reward prices drop	Shocks to rewards, not resources
Mueller-Bloch et al. (2024)	Agent-based simulation	Higher participation promotes decentralization in PoS; wealth concentration is a risk	Simulation; PoS only
Chen et al. (2021)	Empirical (blockchain platforms)	Inverted U-shaped relationship: semi-decentralization outperforms full decentralization	Platform governance, not consensus
Sai et al. (2021)	Taxonomy/Framework	Identifies multiple dimensions of decentralization across blockchain subsystems	Descriptive; no causal claims
Ju et al. (2025)	Longitudinal empirical	Crypto ecosystems show mixed trends; recent centralization in consensus, NFTs, developers	Measurement focus; limited causal identification
Halaburda and Mueller-Bloch (2020)	Conceptual commentary	Blockchain decentralization is multidimensional; centralization preferable for system development, decentralization for transaction validation	Conceptual; no empirical analysis
Hsieh and Vergne (2023)	Mixed methods (fsQCA)	Different dimensions of decentralization differently affect early-stage platform growth	Early-stage platforms; not consensus layer
This paper	Quasi-experimental	Resource flexibility of consensus mechanisms enables sustained decentralization	Three specific shocks; limited blockchains

Notes: This table summarizes key studies on blockchain decentralization, their methodological approaches, main findings, and limitations.

Table A1: Overview of literature on blockchain decentralization

Dependent variable	*Entropy*		*Nodes*		*Gini*		*Nakamoto*		*HHI*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment	-0.209***	-0.152**	5.050***	0.134	0.047***	0.005	-0.414***	-0.477***	0.012***	0.010*
	(0.049)	(0.048)	(1.229)	(1.112)	(0.010)	(0.010)	(0.094)	(0.117)	(0.003)	(0.004)
Bitcoin	0.020	0.296**	-33.980***	-33.920***	-0.327***	-0.370***	1.577***	1.776***	-0.037***	-0.055***
	(0.018)	(0.103)	(1.091)	(1.951)	(0.009)	(0.025)	(0.053)	(0.298)	(0.002)	(0.007)
After	0.020	-0.045	-6.893***	-3.501**	-0.033***	-0.011	-0.060	0.089	0.001	0.004
	(0.020)	(0.028)	(1.061)	(1.200)	(0.006)	(0.008)	(0.061)	(0.059)	(0.003)	(0.003)
Intercept	3.722***	3.672***	53.740***	60.209***	0.801***	0.848***	2.990***	2.902***	0.131***	0.140***
	(0.015)	(0.069)	(0.948)	(1.627)	(0.005)	(0.013)	(0.008)	(0.249)	(0.002)	(0.007)
Block Size (MB)		0.147		1.167		0.002		0.467*		-0.010
		(0.075)		(1.148)		(0.015)		(0.212)		(0.005)
Hashrate (MH/s)		-0.004***		-0.034***		0.000		-0.006***		0.000***
		(0.000)		(0.007)		(0.000)		(0.001)		(0.000)
Number of Tx		0.000		-0.000**		-0.000**		0.000		-0.000
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Price ($)		0.000		0.000***		0.000***		0.000		-0.000
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Fee per Tx ($)		-0.000		-0.002		-0.000		-0.003		0.000*
		(0.001)		(0.020)		(0.000)		(0.003)		(0.000)
Reward ($)		0.000		-0.000**		-0.000**		-0.000		-0.000
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Observations	1202	1202	1202	1202	1202	1202	1202	1202	1202	1202

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain and month.

Table B2: Robustness check of difference-in-difference estimation of China's mining ban in Section 4.1 with consensus layer covariates.

Dependent variable	*Entropy*		*Nodes*		*Gini*		*Nakamoto*		*HHI*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
After	-0.118***	-0.118***	-1.673	-1.673	-0.009	-0.009	-0.105	-0.105	0.013**	0.013**
	(0.030)	(0.030)	(1.727)	(1.730)	(0.014)	(0.014)	(0.101)	(0.101)	(0.004)	(0.004)
Bitcoin	0.020	0.020	-33.980***	-33.980***	-0.327***	-0.327***	1.577***	1.577***	-0.037***	-0.037***
	(0.017)	(0.017)	(1.019)	(1.021)	(0.007)	(0.007)	(0.051)	(0.051)	(0.002)	(0.002)
Intercept	3.722***	3.722***	53.740***	53.740***	0.801***	0.801***	2.990***	2.990***	0.131***	0.131***
	(0.011)	(0.011)	(0.908)	(0.909)	(0.005)	(0.005)	(0.016)	(0.016)	(0.001)	(0.001)
Treatment	-0.109	-0.022	-0.989	-0.203	-0.001	-0.000	-0.253	-0.052	0.002	0.000
	(0.063)	(0.013)	(1.554)	(0.319)	(0.013)	(0.003)	(0.145)	(0.030)	(0.005)	(0.001)
Treatment × Block Size		-0.008		-0.070		-0.000		-0.018		0.000
		(0.004)		(0.110)		(0.001)		(0.010)		(0.000)
Treatment × Fees per Tx		-0.005		-0.049		-0.000		-0.013		0.000
		(0.003)		(0.077)		(0.001)		(0.007)		(0.000)
Treatment × Hash Rate		-0.039		-0.357		-0.000		-0.091		0.001
		(0.023)		(0.561)		(0.005)		(0.052)		(0.002)
Treatment × # Txs		0.017		0.157		0.000		0.040		-0.000
		(0.010)		(0.248)		(0.002)		(0.023)		(0.001)
Block Size	-0.002	-0.002	-0.259	-0.259	-0.001	-0.001	0.017	0.017	0.000	0.000
	(0.009)	(0.009)	(0.240)	(0.240)	(0.003)	(0.003)	(0.046)	(0.047)	(0.001)	(0.001)
Fees per Tx	0.007	0.007	-0.849*	-0.849*	-0.003	-0.003	-0.045	-0.045	-0.000	-0.000
	(0.008)	(0.008)	(0.425)	(0.425)	(0.003)	(0.003)	(0.023)	(0.023)	(0.001)	(0.001)
Hash Rate	0.030**	0.030**	0.097	0.097	0.008	0.008	-0.027	-0.027	-0.002	-0.002
	(0.010)	(0.010)	(0.443)	(0.444)	(0.006)	(0.006)	(0.036)	(0.037)	(0.001)	(0.001)
# Txs	0.022*	0.022*	-1.212***	-1.212***	-0.015***	-0.015***	0.081*	0.081*	-0.003**	-0.003**
	(0.008)	(0.008)	(0.281)	(0.281)	(0.003)	(0.003)	(0.032)	(0.032)	(0.001)	(0.001)
R^2	0.306	0.306	0.962	0.962	0.953	0.953	0.727	0.727	0.687	0.687
Observations	1202	1202	1202	1202	1202	1202	1202	1202	1202	1202

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain and month. All consensus layer covariates are pre-ban averages, normalized as a Z-score for each blockchain. *Tx* refers to transactions. Block size is in megabytes, and hashrate is in megahashes per second.

Table B3: Robustness check of difference-in-difference estimation of China's mining ban in Section 4.1 with consensus layer covariates and covariate-treatment interactions.

C Identifying Ethereum Validators post Proposer-Builder Separation (PBS)

Proposer-builder separation (PBS) was designed to create an open marketplace for block building. Critically, block building involves the ordering of transactions, which can extract arbitrary opportunities by front-running, sandwiching, or other methods. This value extraction through transaction ordering is called maximal extractable value (MEV). Flashbots, who first studied MEV on Ethereum (Daian et al., 2020), developed an initial implementation of PBS in a middleware called MEV-Boost.[4] MEV-Boost was deployed on Ethereum on the same day as the Merge on block 15537940.[5] By running MEV-Boost, validators can access a competitive block-building market and sell block space to arbitragers to increase staking rewards by over 60%.[6]

MEV-Boost creates problems for data collection of validators because the block reward recipient now becomes the block builder, not the proposer (also known as the validator). However, in the context of our study on the consensus layer, we need to identify the proposer for each block, not the builder. In the block, the builder first receives the block reward, and the builder transfers the block reward plus additional fees for the proposer as a transaction in the block.[7] Thus, for a block that has an MEV builder as the block recipient, we identify the proposer within the block transactions in the following way.

First, we obtain a list of labeled addresses for MEV builders on Etherscan.[8] The list of MEV builders that we use is available at https://dune.com/queries/3665816. Because builders are in a competitive marketplace, they are incentivized to reveal their identities as builders.[9] Then, we obtain all blocks whose recipients are *not* on the list of MEV builders. These blocks were not sold to builders and have validators as the block reward recipients.

Then, we obtain a list of all blocks whose block reward recipients are in the list of MEV builders. These blocks *were* sold to builders and have the builder as a block reward recipient. In this case, we identified the transaction by which the builder sent the block reward to the proposer in a few different ways. First, if there was a transaction in an MEV block whose sender is a builder and

[4] See the GitHub repository for MEV-Boost. Accessed April 25, 2024.
[5] See deployment transaction on Etherscan. See also the blog post by Flashbots. Accessed April 25, 2024.
[6] See article by Flashbots on hackmd.io. Accessed April 25, 2024.
[7] Because the proposer ultimately signs the block, the proposer is guaranteed to receive the block reward in the same block that it validates by checking whether it receives the reward before signing the block.
[8] See etherscan.io. Accessed April 25, 2024
[9] While Ethereum allows anonymous or private builders, such builders should be rare, if they exist at all.

the receiver is a known proposer (see below for how we obtain a list of proposers), then we label the recipient of the transaction as the proposer for that block. Second, a few builders sent the block reward from another address that it owned, presumably to save money on transaction fees. This was common for larger MEV builders. We manually identify the separate addresses for MEV builders, which are available at https://dune.com/queries/3669067. Lastly, some MEV builders are also proposers. For these builders, we label the builder as also the proposer of the block; such builders are labeled at https://dune.com/queries/3665820. These methods accounted for all of the blocks that we obtain from January 1, 2022, to March 31, 2023. The SQL implementation of these methods can be found at https://dune.com/queries/3664358.

As mentioned above, we obtain a list of block reward recipients in the following ways. First, we query all addresses that deposited ETH to the Beacon staking contract on Ethereum. Now, if Ethereum did not have staking pools or smart contracts, then this list would be sufficient. Second, we obtain a list of the top proposer fee recipients from Etherscan, which may include staking pools.[10] Most, if not all, staking pools are known because they have the incentive to obtain as much stake in their pool as possible and thus market themselves to potential stakers. Third, we have manually labeled a list of MEV builders who are also recipients of block rewards. They are clearly recipients because they have received numerous fees from other labeled MEV builders. See address 0x7e2a2FA2a064F693f0a55C5639476d913Ff12D05 as an example. The final list of recipients is available at https://dune.com/queries/3665820.

[10]See list on etherscan.io. Accessed April 25, 2024.

D Post-PBS Analysis of MEV Builders

As explained in Section C, proposer-builder separation (PBS) was designed to create a competitive marketplace for block builders to ameliorate the negative externalities of maximal extractable value (MEV). Here, we report data on the MEV builders involved after the implementation of MEV-Boost, which is coincident on the same day as the Ethereum Merge. Figure D1 shows the daily fraction of blocks that have used MEV (left) and the daily fraction of proposers (*i.e.*, validators) that have produced at least one MEV block.

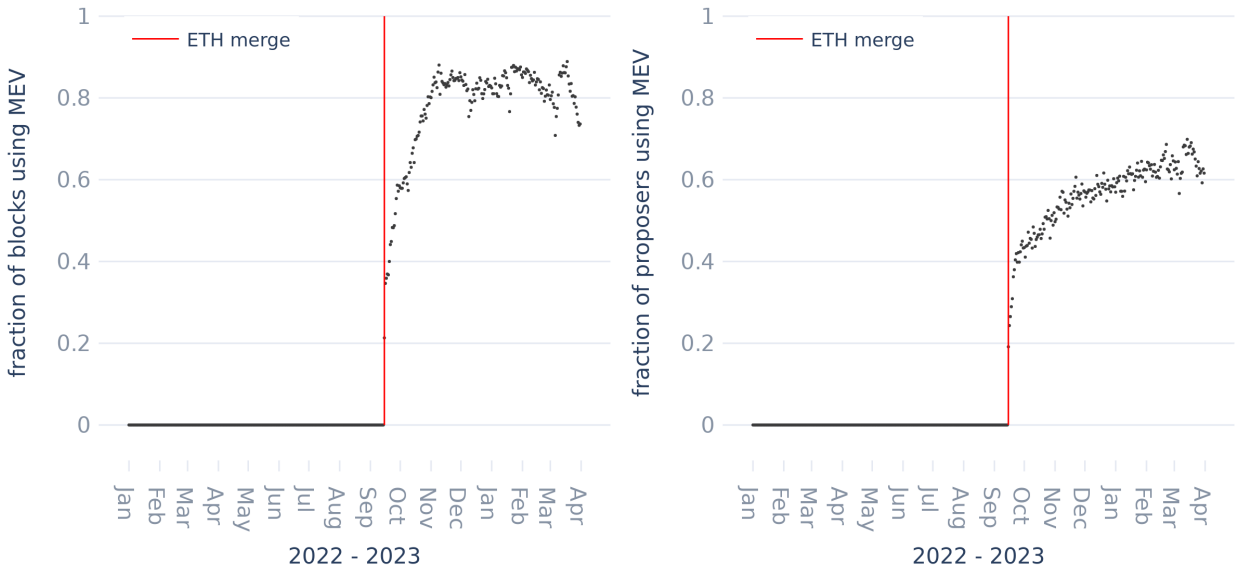

Figure D1: (Left) Fraction of blocks using MEV. (Right) fraction of proposers using MEV.

Then, to test the robustness of our analyses in Section 4.3, we replicate the same analyses without correcting for PBS. In Figure D2, we can see that entropy increases only by around 0.5 at the Merge. There is also a sharp decline in the entropy that corresponds with the increasing usage of MEV shown in Figure D1.

For months following the Merge, there is also a steady, daily decrease in Ethereum's entropy of -0.002^{***}. To understand the small but statistically significant negative slope, we investigated whether any particular group of validators had a role in the negative slope of entropy after the Merge. To do so, we simulate "knockout" experiments in which we remove particular groups of validators from the analysis after the Merge and recompute the entropy. The groups included centralized exchanges (CEX), Staking Pools such as Lido, MEV (maximal extractable value) builders, and

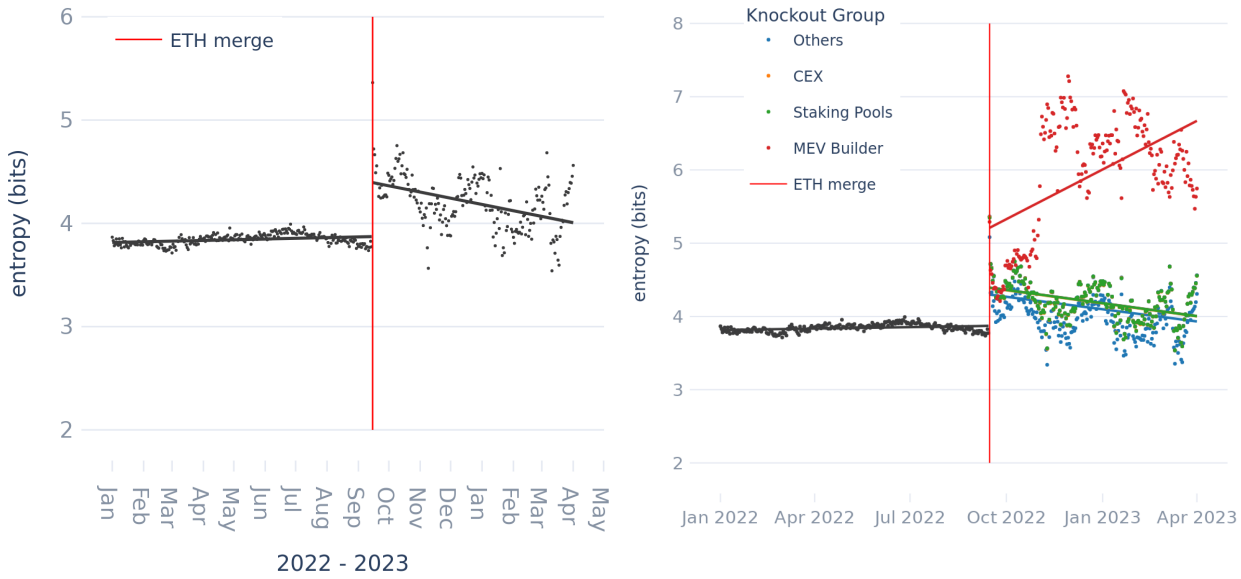

Figure D2: Knockout simulations of Ethereum's entropy after the Merge.

"Others" which include individuals or are unlabeled.[11] We see in Figure D2 that removing MEV builders from the data had a substantial positive effect on both the slope and the absolute value of entropy. This result suggests that MEV builders are a *centralizing* force on Ethereum's consensus. Removing CEX and Staking Pools validators has negligible effects while removing "Others" has a negative effect, as one can expect from removing individuals from consensus.

The results here not only validate our MEV analysis but also suggest that the composability of protocols through proposer-builder separation (PBS) or otherwise, while introducing opportunities for open protocols, market competition, and interoperability, opens up the risk of centralization in each component. This unexpected discovery, which emerged from our analysis, does not directly implicate the resource flexibility hypothesis but is an interesting element of the decentralization of blockchains that warrants further investigation. While such analysis is beyond the scope of this paper, we encourage such work in future research.

[11] See ethereum.org for a detailed explanation of MEV. List of sequencers from etherscan.io. Accessed April 23, 2024.

E Multi-Period Difference-in-Difference

To account for the rolling enforcement of China's policy change banning crypto mining, we decomposed our difference-in-difference estimation in Section 4.1 into multiple periods. Instead of simply having a single date that divides the time series into $\text{After}_t = 0$ before May 15, 2021, and $\text{After}_t = 1$ after, we create both During_t and After_t variables. The During_t variable is a binary indicator variable that equals one from May 15, 2021, to the end of July 2021 and represents the days that the ban was being enforced. The After_t variable is a binary indicator variable as before, but it now equals one starting on August 2021. We also reformulated the Exposure_{it} as the cross-term $(\text{During}_t \times \text{Exposure}_i)$, instead of $(\text{After}_t \times \text{Exposure}_i)$, to account for the exposure occurring during the rolling ban. We then specified a multi-period difference-in-difference estimation by

$$\begin{aligned} \text{Entropy}_{it} = {} & \delta_1(\text{During}_t \times \text{Bitcoin}_i) + \delta_2(\text{After}_t \times \text{Bitcoin}_i) \\ & + \beta_1 \text{Exposure}_{it} + \beta_2 \text{Bitcoin}_i + \beta_3 \text{During}_t + \beta_4 \text{After}_t + \epsilon_{it}, \end{aligned} \tag{E1}$$

where the variables are the same as explained in Equation 1.

When we estimate Equation E1, we find a similar result as in Section 4.1 of the main manuscript. In Column 1 of Table E1, we observe that there were no significant effects during the rollout, as seen in the following terms: During, (During × Bitcoin), and (During × Exposure). However, we did see a negative effect of (After × Bitcoin) after the rollout. This negative effect is similar to our results in Section 4.1 but also slightly larger because it now excludes the more minor effects during the rollout of the ban. The consistency of the negative effect observed post-rollout with our earlier findings in Section 4.1 not only validates our initial results but also suggests that the impact of the ban on Bitcoin's decentralization was more pronounced when isolating the post-rollout period, thereby reinforcing the robustness of our analysis.

While the above multi-period analysis does not reveal any significant effects during the rollout, this may be because the rollout is gradual. That is, this simple "difference-in-means" approach may not reveal the nuanced dynamics of the rollout, as we can visually see happening in Figure 2. Thus, to get an even more fine-grain understanding of the multi-period dynamics, we extended Equation E1 by adding the time variable Day_t, by itself and as cross-terms with $(\text{During}_t \times \text{Chain}_i)$ and $(\text{After}_t \times \text{Chain}_i)$. We also add the time variable as a cross-term with the control variable

Dependent variable	*Entropy*		*Nodes*		*Gini*		*Nakamoto*		*HHI*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
After	0.078*	-0.223***	-5.800***	-3.774	-0.041***	0.003	-0.050	-0.274	0.001	0.029**
	(0.033)	(0.057)	(1.411)	(3.737)	(0.007)	(0.016)	(0.084)	(0.307)	(0.004)	(0.009)
Treatment (After)	-0.497***	0.180	1.219	2.251	0.047***	-0.013	-0.832***	0.104	0.027***	-0.028*
	(0.063)	(0.143)	(1.575)	(4.272)	(0.011)	(0.030)	(0.128)	(0.350)	(0.005)	(0.012)
Bitcoin	0.020	-0.221***	-33.980***	-34.347***	-0.327***	-0.306***	1.577***	1.244***	-0.037***	-0.017***
	(0.026)	(0.057)	(1.309)	(2.064)	(0.010)	(0.012)	(0.062)	(0.118)	(0.003)	(0.004)
Day		0.002***		-0.002		-0.000**		0.003*		-0.000***
		(0.000)		(0.016)		(0.000)		(0.001)		(0.000)
During	0.010	-0.140***	-6.520***	-5.509*	-0.021*	0.001	0.010	-0.102	-0.007*	0.007
	(0.027)	(0.033)	(1.612)	(2.627)	(0.008)	(0.011)	(0.010)	(0.115)	(0.003)	(0.005)
Treatment (During)	-0.026	0.312***	6.760***	7.276*	0.035***	0.006	-0.138*	0.330*	0.010**	-0.018**
	(0.039)	(0.079)	(1.669)	(2.825)	(0.011)	(0.015)	(0.066)	(0.159)	(0.004)	(0.006)
Exposure × Day		-0.006***		-0.010		0.001*		-0.009**		0.001***
		(0.001)		(0.034)		(0.000)		(0.003)		(0.000)
Intercept	3.722***	3.828***	53.740***	53.020***	0.801***	0.785***	2.990***	3.070***	0.131***	0.121***
	(0.020)	(0.020)	(1.280)	(1.898)	(0.007)	(0.008)	(0.010)	(0.083)	(0.002)	(0.003)
Observations	1202	1202	1202	1202	1202	1202	1202	1202	1202	1202

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain and month.

Table E1: China bans crypto mining

Exposure_{it} because it is indeed the exposure that is being rolled out gradually during the ban. As such, we are even more explicitly estimating the effect of resource flexibility, or lack thereof, *in response to* the exposure while, at the same time, even controlling for the time-varying effects of the exposure itself. We specified this new time-varying multi-period difference-in-difference estimation by

$$\begin{aligned}\text{Entropy}_{it} = {} & \delta_1(\text{During}_t \times \text{Bitcoin}_i \times \text{Day}_t) + \delta_2(\text{After}_t \times \text{Bitcoin}_i \times \text{Day}_t) \\ & + \beta_1(\text{Exposure}_{it} \times \text{Day}) + \beta_2\text{Bitcoin}_i + \beta_3\text{During}_t + \beta_4\text{After}_t \\ & + \beta_5\text{Day}_t + \epsilon_{it}.\end{aligned} \tag{E2}$$

By estimating Equation E2, we first find that the exposure itself has the greatest negative effect of -0.007 bits per day. This initial result confirms the simple intuition that the direct shutdown of nodes should have a negative effect on decentralization. In fact, in this initial period, Bitcoin received this negative effect of the exposure slightly less than Ethereum did by 0.001 bits per day ($p = 0.045$), and this is visible in Figure 2, where we can see that despite Bitcoin's greater exposure, both Bitcoin and Ethereum experienced a similar decrease in entropy throughout the rollout period. However, after the rollout period, Bitcoin experienced a significant negative effect of -0.002 bits per day, consistent with our other results. Thus, although both Bitcoin and Ethereum initially saw a decline in decentralization, Bitcoin's more pronounced negative impact in the post-rollout

period highlights the differing recovery dynamics and underscores the role of resource flexibility in influencing a blockchain's resilience to external shocks. Our decomposition analysis in Section G of the Appendix reveals the behavior of individual nodes and mining pools that may further elucidate the effects of the ban even more precisely.

F Bitcoin Hashrate Mining Maps

Here, we validate that the hashrates observed in Figure 1, which we used to estimate the Exposure covariate, do indeed represent exposure to China's mining ban. While we do not have geographical hashrates for Ethereum, we use the data available for Bitcoin as validation for our exposure measures using total hashrates. To validate our exposure measures, we use the Bitcoin Mining Map collected by The Cambridge Centre for Alternative Finance (The Cambridge Centre for Alternative Finance, 2023). The Centre has partnered with several Bitcoin mining pools and uses the IP (Internet Protocol) addresses of mining facility operators to identify the geographical location of a miner.

Figure F1 shows the monthly hashrates by country from April 2021 to September 2021, around the time of China's mining ban. First, we observe on the map that monthly hashrates decreased from 43.98% in April 2021 to 34.25% in May 2021. Then, China's hashrate disappears to 0% for the months of June and July until it recovers 22.29% in August. As an alternative visualization, Figure F2 shows a chart from Statista that displays the same data as in Figure F1 as a bar chart (Cambridge Centre for Alternative Finance, 2022). This trend accurately matches the decrease in total hashrates from May 2021 to July 2021 in Figure 1. Second, the maximum percentage of China's hashrate was 43.98%, which is close to the exposure measure of 51.14% for Bitcoin. While it is conceivable that the partnered mining pools colluded with Chinese miners to hide their IP, that seems unlikely given the similarity between the hashrate of Chinese miners in the map and the total hashrate. The close alignment between the observed hashrate percentages in China and the total hashrate data further validates our exposure measure and our analysis surrounding the impact of China's mining ban on blockchain decentralization.

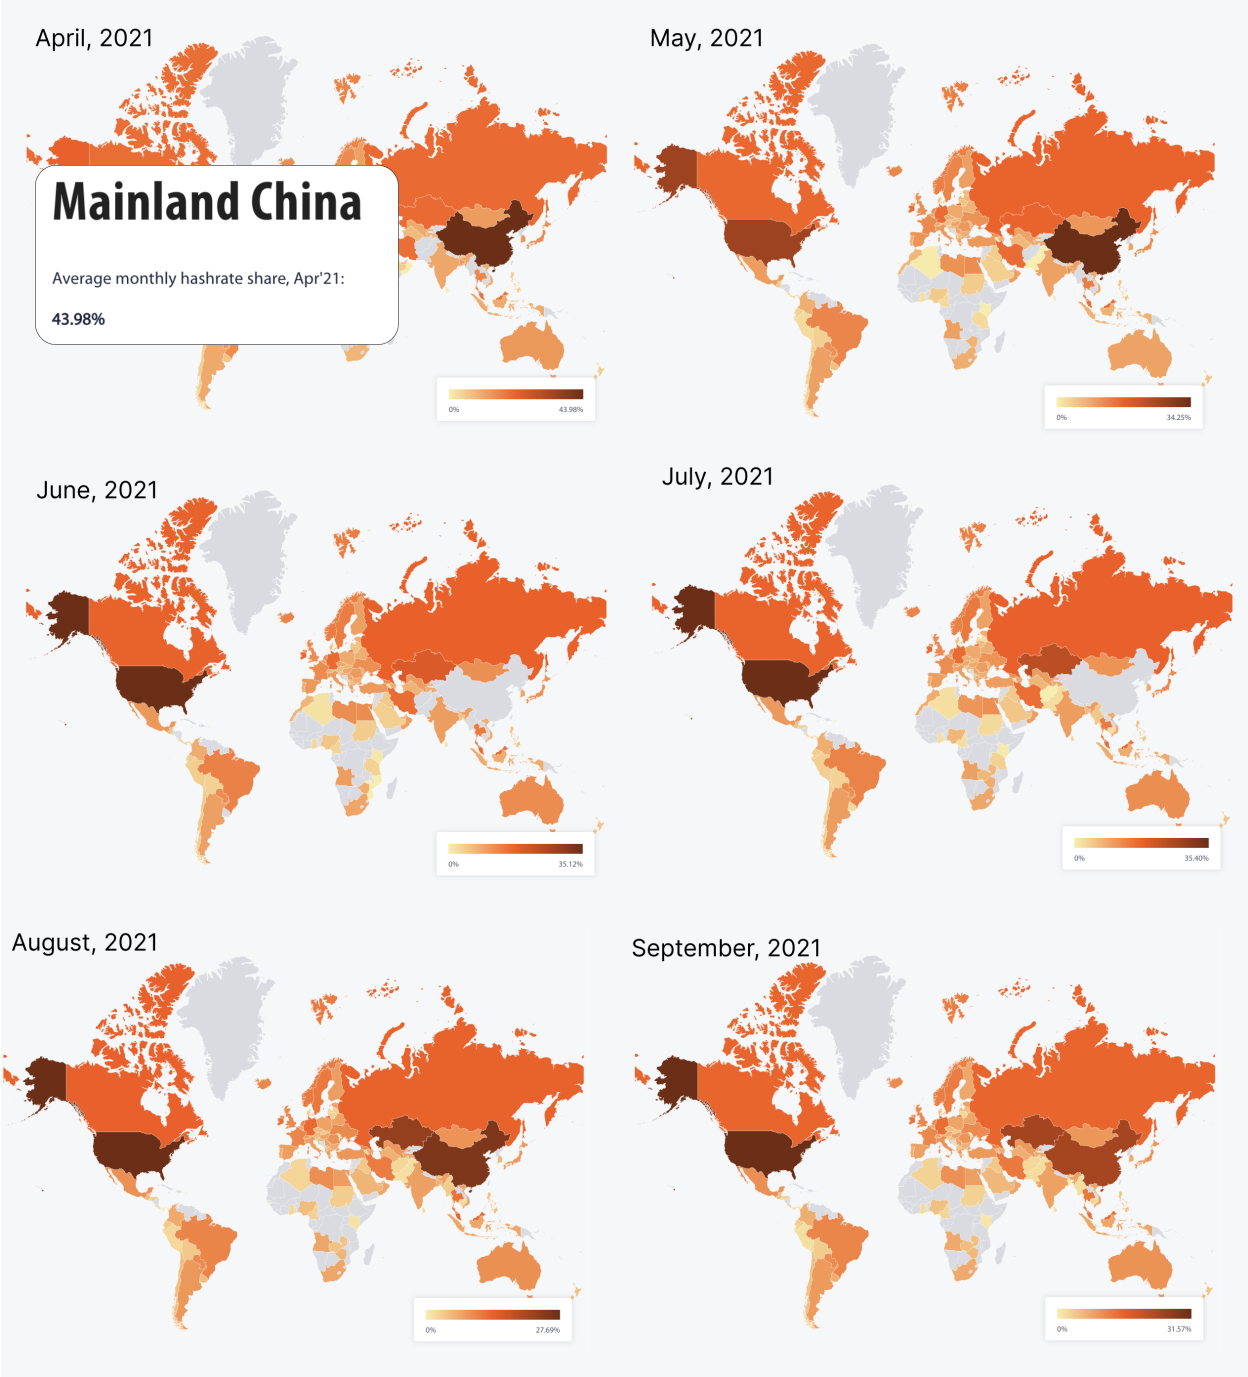

Figure F1: The Bitcoin Mining Map by The Cambridge Centre for Alternative Finance from April to September 2021.

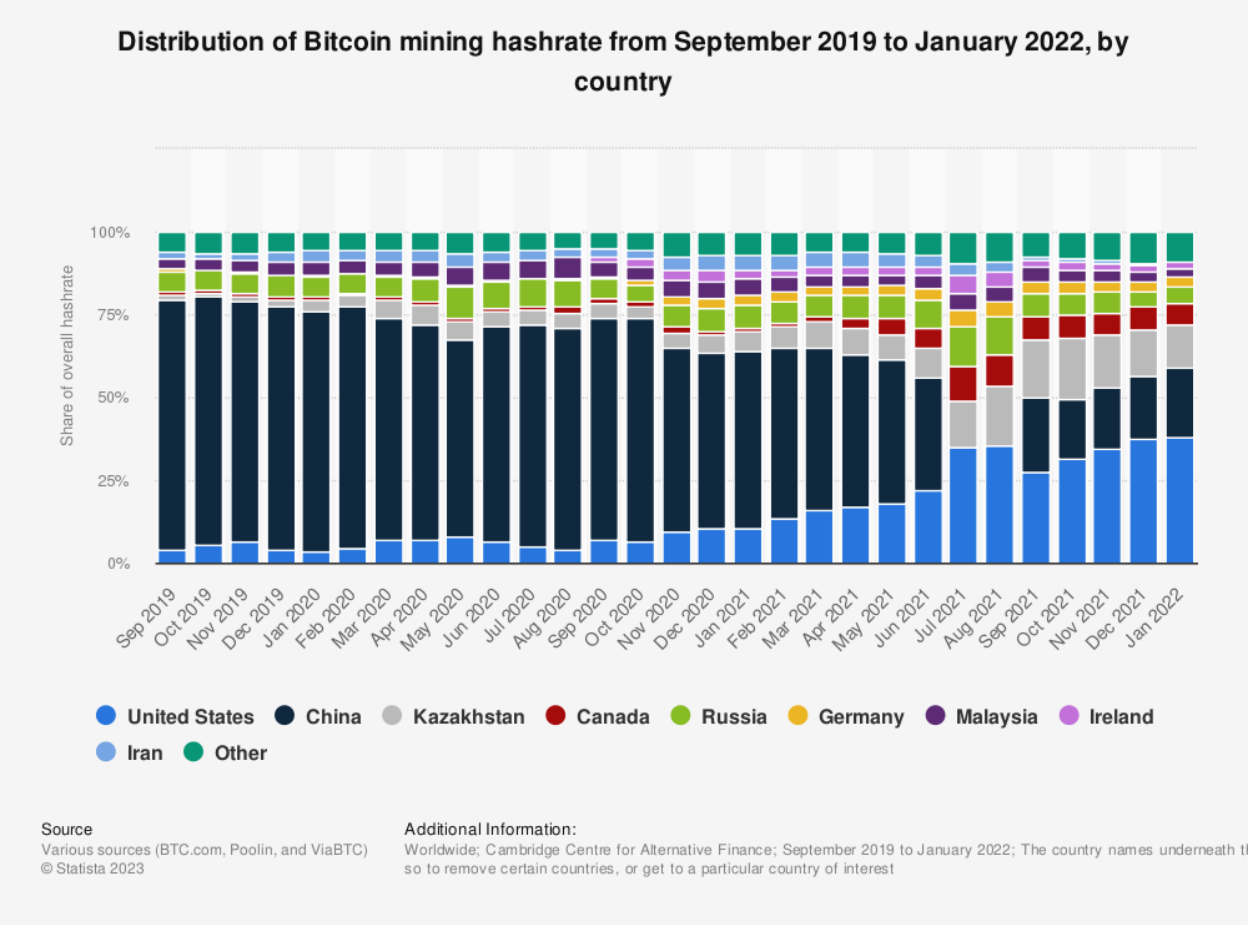

Figure F2: Distribution of Bitcoin mining hashrate from September 2019 to January 2022, by country.

G Decomposing China's Mining Ban

To better understand how China's mining ban may have affected Bitcoin, we decompose the effects of the ban to the level of individual nodes. This analysis does not take the rolling nature of the ban into account and is certainly not causal. Indeed, because of Bitcoin's pseudonymity, it is difficult to determine the identity or geographic location of nodes. Moreover, mining pools can obfuscate the distribution of rewards to their participants and make it even more difficult to find definitive, causal observations at the node level. However, by exploring node-level data, we can check for the robustness of our interpretations of our results and also investigate more nuanced effects of the ban, such as those to mining pools versus individual miners.

First, we examine the nodes that may have individually been affected by the ban. To construct the number of nodes lost after the shock shown in Figure G1, we determine the unique node addresses that received block rewards from 30 days prior to the shock on May 15, 2021, up to the day of the shock to see which nodes are no longer receiving any block rewards. If we conservatively assume that the initially high number of 18 nodes in Figure G1A simply did not yet have a chance to receive rewards, we can see that there are around 7 nodes that no longer receive rewards at the steady state. Those nodes represent, however, only 0.0147 percent of the estimated hash power of the network: a far cry from the 51.1% decrease in total hashrate we observe in Section 3.3.

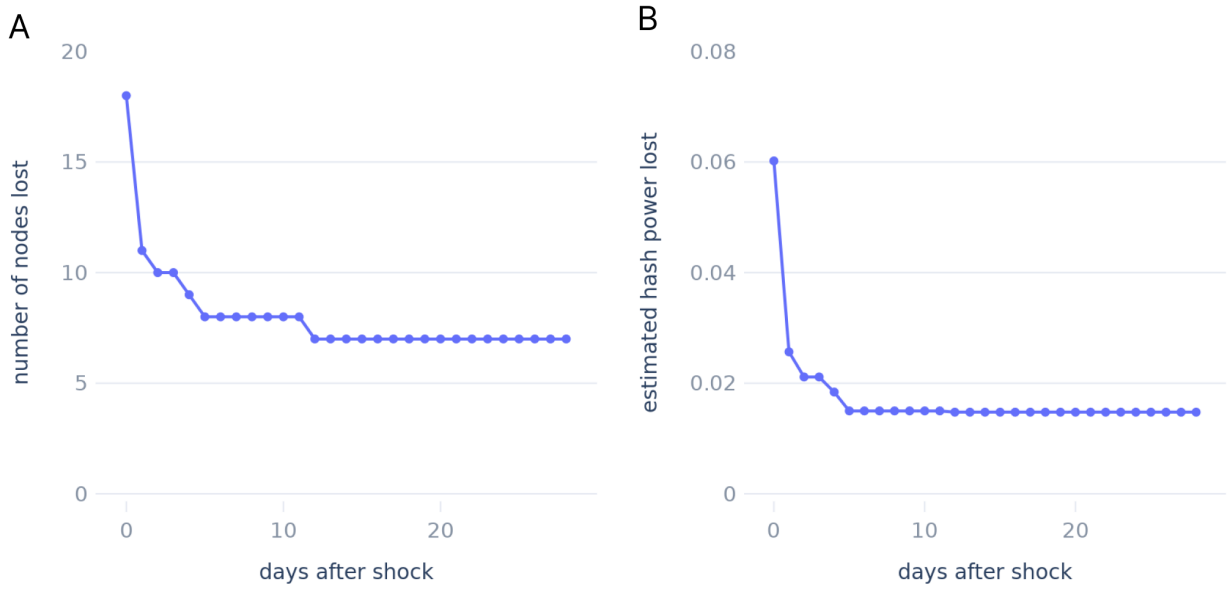

Figure G1: (A) The number of nodes lost after the initial shock. (B) Estimated hash power of the nodes lost after the initial shock.

Given that individual miners do not seem to be much affected by our initial estimation, we estimate whether mining pools may have experienced changes in their mining throughout the ban. In Figure G2A, we display the daily number of blocks mined per node address around the time of the ban with lowess trend lines. Though we cannot make any causal claims, we can clearly observe that some nodes are producing more and some fewer blocks around the time of the ban. We also show summary statistics for 30 days prior and 30 days after the ban in Figure G2B. These varying trends in block production around the time of the ban suggest that mining pools may have experienced differential impacts, underscoring the impact of the ban on mining pools versus individual miners.

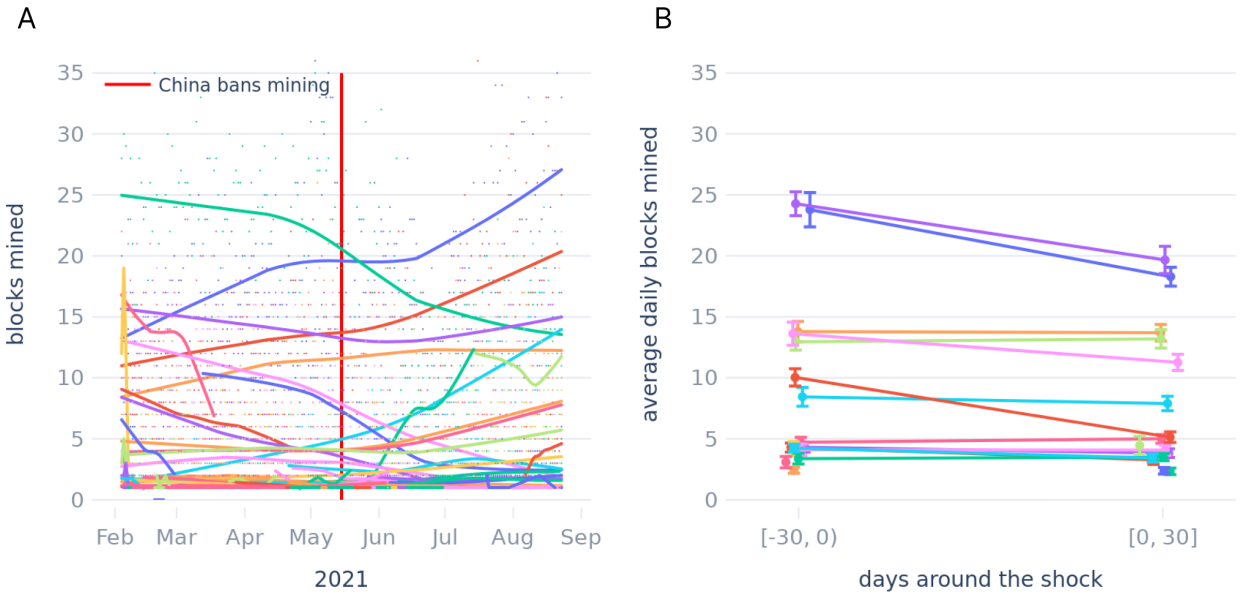

Figure G2: Number of blocks mined by Bitcoin miners around the time of the shock. (A) Daily number of blocks mined with Lowess trend lines. (B) The average number of blocks mined 30 days before and after the shock with standard errors.

While many individual nodes cannot be identified due to Bitcoin's pseudonymity, some larger nodes, such as mining pools, are labeled, allowing us to more accurately assess the impact of the ban on these specific entities within the blockchain network. We examine three examples of mining pools that have experienced a reduction in the daily nodes mined from our above analysis: the Binance pool, the Huobi pool, and an unknown pool.[12] While we cannot definitively identify the location of these pools, and even more so for their individual participants, we do know that the cryptocurrency exchanges Binance and Huobi were both initially based in China but moved their

[12]The charts for these exact nodes can be found for the Binance, Huobi, and the unlabeled pools. Accessed October 30, 2023.

headquarters abroad later on.

Figure G3 displays the balance of BTC in the pools, obtained from blockchain.com, as a proxy for mining activity, since mining results in an increased balance of BTC. We observe a clear, lasting reduction in the BTC balance in May 2021 for the Binance pool. For the Huobi Pool, we saw a slight reduction in BTC balance in May 2021 and a near-complete reduction from June to the end of July, which directly matches our analysis of total hashrates and shock exposure in Section 3.3. For the Unknown Pool, we see a similar but less pronounced decrease in the BTC balance in May 2021 and then a larger decrease in July 2021. In contrast to these mining pools, some mining pools, such as those indicated by the prominent, rising blue and red lines in Figure G2A, see even greater mining activity after the ban. The observed variations in the trends of mining pools not only align with our analysis of the impact of resource flexibility on decentralization in response to a major policy shock, but they also reveal the complex dynamics of blockchain decentralization in response to policy shocks.

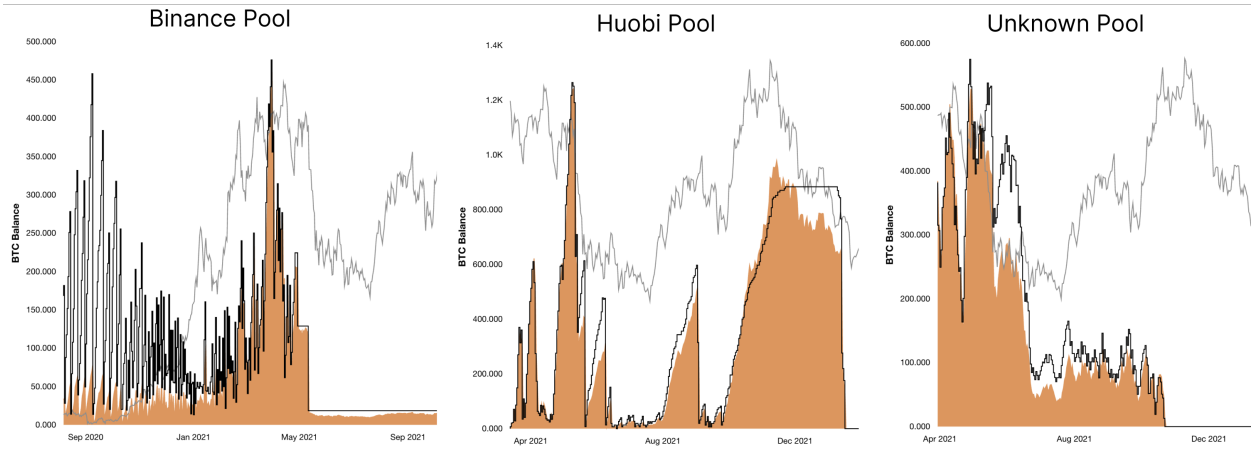

Figure G3: The BTC balance of major Bitcoin mining pools during China's mining ban.

As an interesting side note, the mining pools that experienced greater activity after the ban are labeled as those of Antpool and ViaBTC. Both companies were, and still are, headquartered in China. In contrast, Binance and Huobi, both of which experienced marked decreases in mining (Figure G3), are no longer based in China, which suggests a consolidation of Bitcoin mining among Chinese miners. Whether or not such consolidation of Bitcoin mining was deliberate or happenstance, it may further support our resource flexibility hypothesis in either case. A deliberate consolidation would require high frictions, such as the costs and logistical challenges of relocating specialized

ASIC hardware and significant electricity consumption, to prevent miners from migrating out of the jurisdiction. On the other hand, if the consolidation was happenstance, it still underscores the rigidity and lack of adaptability in less flexible systems like Bitcoin. Either way, the observed consolidation among Chinese miners lends further credence to our resource flexibility hypothesis, highlighting the role of resource flexibility in shaping the decentralization dynamics of blockchain networks.

H Lagged Effects of China's mining ban

Table H1 shows the coefficients of the lagged difference-in-difference estimation for China's mining ban according to Equation 3.

Dependent variable	*Entropy*	*Nodes*	*Gini*	*Nakamoto*	*HHI*
After	0.212***	12.114	-0.039***	0.332**	-0.018**
	(0.054)	(16.931)	(0.007)	(0.126)	(0.005)
Bitcoin	0.104*	-32.609***	-0.329***	1.765***	-0.049***
	(0.043)	(1.077)	(0.010)	(0.140)	(0.005)
Intercept	3.609***	53.090***	0.809***	2.660***	0.146***
	(0.038)	(0.851)	(0.005)	(0.112)	(0.005)
Treated$_{\tau=0}$	-0.224***	-12.614	0.027*	-0.299	0.018**
	(0.065)	(16.938)	(0.010)	(0.174)	(0.006)
Treated$_{\tau=120}$	-0.140***	-1.267**	0.015***	-0.533***	0.012***
	(0.033)	(0.471)	(0.004)	(0.110)	(0.002)
Treated$_{\tau=180}$	-0.360***	-4.251***	-0.020***	-0.438***	0.023***
	(0.054)	(0.684)	(0.006)	(0.118)	(0.004)
Treated$_{\tau=60}$	-0.041	-1.067**	-0.009	0.017	0.001
	(0.027)	(0.368)	(0.008)	(0.117)	(0.002)
Treated$_{\tau=-120}$	0.082***	1.017***	-0.001	0.017	-0.006**
	(0.022)	(0.234)	(0.003)	(0.124)	(0.002)
Treated$_{\tau=-180}$	-0.008	-0.998	-0.000	0.208	-0.000
	(0.025)	(0.677)	(0.009)	(0.149)	(0.002)
Treated$_{\tau=-60}$	-0.023	0.600	0.026*	-0.233	0.004
	(0.034)	(0.464)	(0.012)	(0.172)	(0.003)
Observations	2002				

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain and month.

Table H1: Lagged effects of China's ban on crypto mining.

We also conduct lagged analyses with a finer timescale of 50 days in Figure H1.

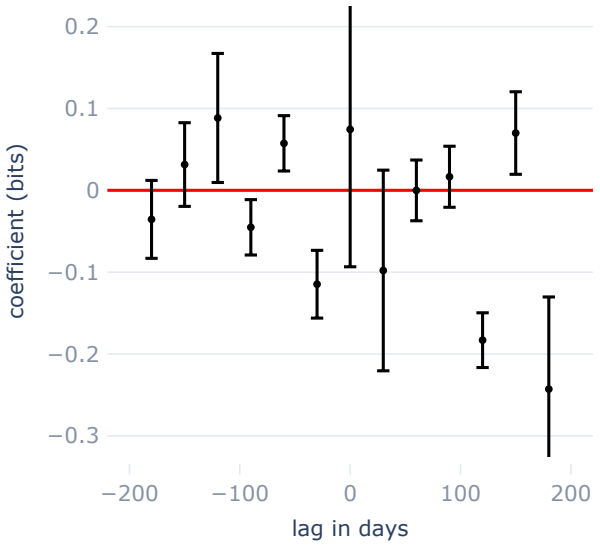

Figure H1: Coefficients for the lead and lag effects Treatment$_\lambda$ on the entropy of China's mining ban with 50-day intervals. Error bars indicate 95% confidence intervals.

I Lagged Effects of Hetzner's Shutdown

Table I1 shows the coefficients of the lagged difference-in-difference estimation for Hetzner's shutdown of Solana nodes according to Equation 3.

Dependent variable	*Entropy*	*Nodes*	*Gini*	*Nakamoto*	*HHI*
Treated$_{\tau=0}$	-0.392***	-343.680***	0.039	-6.909***	0.005
	(0.113)	(61.343)	(0.052)	(1.775)	(0.004)
Treated$_{\tau=10}$	0.194***	75.900	-0.037***	-0.300	-0.000**
	(0.038)	(42.539)	(0.005)	(0.630)	(0.000)
Treated$_{\tau=20}$	0.148***	5.600	-0.023***	6.000***	-0.000***
	(0.020)	(5.144)	(0.003)	(0.887)	(0.000)
Treated$_{\tau=30}$	0.037	11.527*	0.001	2.791**	-0.000***
	(0.021)	(5.240)	(0.004)	(0.928)	(0.000)
Treated$_{\tau=-10}$	0.011	25.300***	0.000	0.400	-0.000*
	(0.007)	(4.739)	(0.001)	(0.308)	(0.000)
Treated$_{\tau=-20}$	0.015*	31.300***	-0.002	0.400	-0.000
	(0.007)	(5.892)	(0.001)	(0.287)	(0.000)
Treated$_{\tau=-30}$	0.009	53.400***	0.002	0.300	0.000
	(0.013)	(7.932)	(0.002)	(0.341)	(0.000)
Observations	162	162	162	162	162

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain and day.

Table I1: Lagged effects of Hetzner's shutdown of Solana nodes.

J Synthetic Difference-in-Difference for the Hetzner Shutdown with Multiple Bandwidths

Table J1 shows the coefficients of the synthetic difference-in-difference estimation for Hetzner's shutdown of Solana nodes according to Equation 4 using multiple bandwidths in days before and after the shock on November 2, 2022.

Dependent variable	±10 days (1)	±20 days (2)	±30 days (3)	±40 days (4)	±50 days (5)
Panel A: Entropy					
Treatment	-0.363	-0.271	-0.180	-0.142	-0.156
	(0.734)	(0.727)	(0.696)	(0.557)	(0.505)
Panel B: Nodes					
Treatment	-332.145	-289.867	-262.205	-232.356	-222.801
	(302.324)	(298.490)	(286.137)	(229.482)	(207.284)
Panel C: Gini					
Treatment	0.040	0.025	0.009	-0.001	-0.008
	(0.352)	(0.348)	(0.333)	(0.266)	(0.243)
Panel D: Nakamoto					
Treatment	-6.433	-6.632	-4.456	-2.739	-1.977
	(11.403)	(11.284)	(10.838)	(8.874)	(8.235)
Panel E: HHI					
Treatment	0.002	0.002	0.003	0.005	0.006
	(0.025)	(0.024)	(0.024)	(0.019)	(0.018)
Observations	84	164	244	323	403

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are derived from placebo tests.

Table J1: Synthetic difference-in-difference estimation of Hetzner shutdown with multiple bandwidths.

K Price Correlation between BTC and ETH

Here, we measure the price correlation between BTC, ETH, and other financial assets to understand the level of price correlation between BTC and ETH. Using price data from Yahoo Finance, we find that BTC and ETH had a relatively high Pearson's correlation coefficient of 0.93 (Figure K1).[13] This is in comparison with correlations of 0.89 and 0.86 with the S&P 500 for BTC and ETH, respectively. Correlations are shown for other prominent financial assets in Figure K1 for comparison. TLT is the iShares 20+ Year Treasury Bond ETF, and VNQ is the Vanguard Real Estate Index Fund ETF.

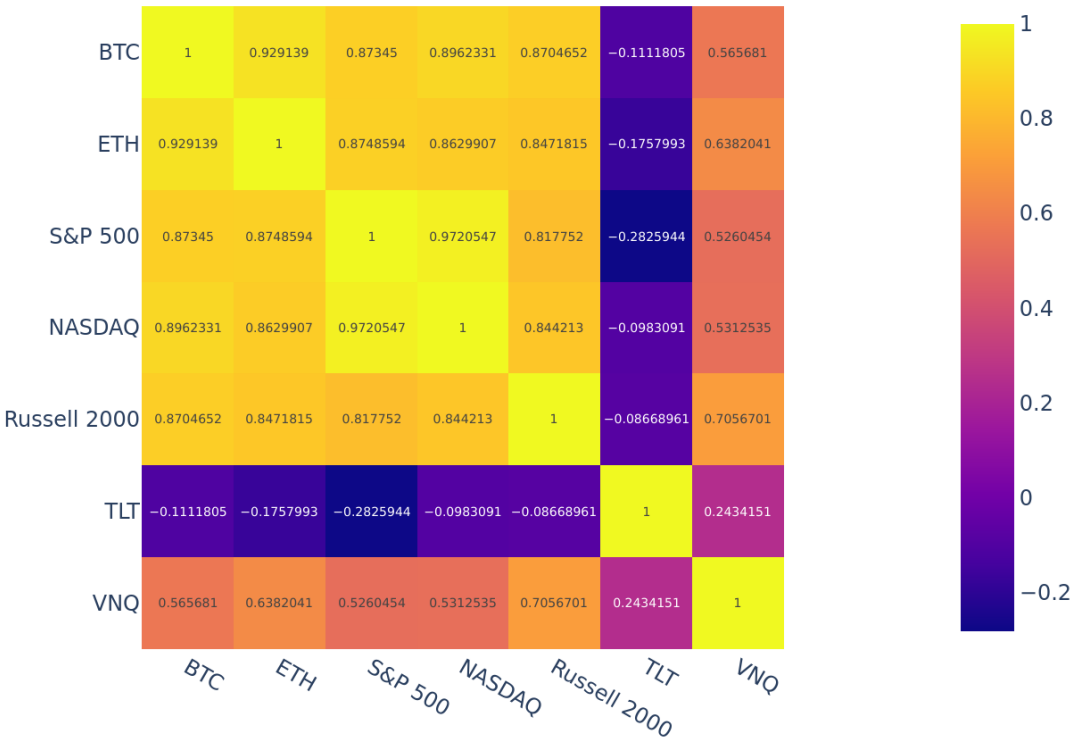

Figure K1: Pearson's correlation coefficients between BTC, ETH, the S&P 500 Index, NASDAQ Composite, Russell 2000, TLT (the iShares 20+ Year Treasury Bond ETF), and VNQ (the Vanguard Real Estate Index Fund ETF).

[13] Data obtained from Yahoo Finance for BTC, ETH, S&P Index, NASDAQ Composite, Russell 2000, TLT, and VNQ. Accessed April 5, 2024.

L Price Analysis during Shocks

Here, we test the hypothesis that market sentiment influences a blockchain's recovery from shocks. Below are estimations of Equation 5 on the daily price of tokens native to the blockchains studied here. Using price data around key events for BTC, SOL, and ETH, we find that these shocks had no positive impact on prices; some even dropped significantly.

Dependent variable	*Price*		
	BTC	SOL	ETH
Intercept	55929.964***	30.415***	1674.996***
	(770.736)	(0.488)	(28.007)
After	-14880.065***	-7.246***	-351.507***
	(1600.222)	(1.928)	(36.226)
Days	-25.735	-0.054***	0.713
	(17.386)	(0.014)	(0.841)
After × Days	-121.117**	-0.201***	0.936
	(38.166)	(0.046)	(1.314)
Observations	121	121	121

Notes: *p<0.05; **p<0.01; ***p<0.001. Heteroskedasticity-robust standard errors. Each column is a separate single-blockchain regression.

Table L1: Event studies on the token price of BTC during China's ban, SOL during Hetzner's shutdown, and ETH during the Merge.

M Clustering at the blockchain level

As a robustness check, we demonstrate that clustering only at the blockchain level, not the blockchain-month level, results in unrealistically low standard errors, which would lead to overconfident statistical significance (Cameron et al., 2008; Petersen, 2009). Table M1 reproduces the results of Table 3 with clustering at the blockchain level.

Dependent	*Entropy*		*Nodes*		*Gini*		*Nakamoto*		*HHI*	
variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
After	0.064	0.072	-3.093	-3.203	-0.033	-0.034	0.241	0.251	-0.004	-0.004
	(0.126)	(0.119)	(3.430)	(3.229)	(0.045)	(0.042)	(0.243)	(0.229)	(0.008)	(0.007)
Exposure		-0.032		0.424		0.003		-0.038		0.002
		(0.029)		(0.784)		(0.010)		(0.056)		(0.002)
Bitcoin	0.020***	0.020***	-33.980***	-33.980***	-0.327***	-0.327***	1.577***	1.577***	-0.037***	-0.037***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Treatment	-0.209***	-0.201***	5.050***	4.943***	0.047***	0.047***	-0.414***	-0.404***	0.012***	0.012***
	(0.000)	(0.007)	(0.000)	(0.198)	(0.000)	(0.003)	(0.000)	(0.014)	(0.000)	(0.000)
Intercept	3.672***	3.672***	52.990***	52.990***	0.791***	0.791***	3.069***	3.069***	0.138***	0.138***
	(0.090)	(0.090)	(0.351)	(0.351)	(0.018)	(0.018)	(0.286)	(0.286)	(0.013)	(0.013)
Observations	1202	1202	1202	1202	1202	1202	1202	1202	1202	1202

Notes: *p<0.05; **p<0.01; ***p<0.001. Standard errors are clustered by blockchain.

Table M1: China bans crypto mining

N Hetzner Shutdown Details

The reasons behind Hetzner's selective enforcement against Solana nodes, as opposed to nodes from other blockchains, remain unexplained by the company. Hetzner's public communications state a broad prohibition against the use of their services for any crypto-related activities, encompassing both PoW and PoS mechanisms, as well as indirectly related activities such as cryptocurrency trading. The absence of similar actions against other chains during the same period suggests a potentially unique issue with Solana's relatively high resource usage that triggered enforcement. For example, Solana nodes require 256GB of RAM, while Ethereum nodes require just 8 GB of RAM.[14] If indeed Solana was targeted due to its excessive resource usage, this would support the role of resource efficiency of consensus mechanisms in enabling sustained blockchain decentralization.

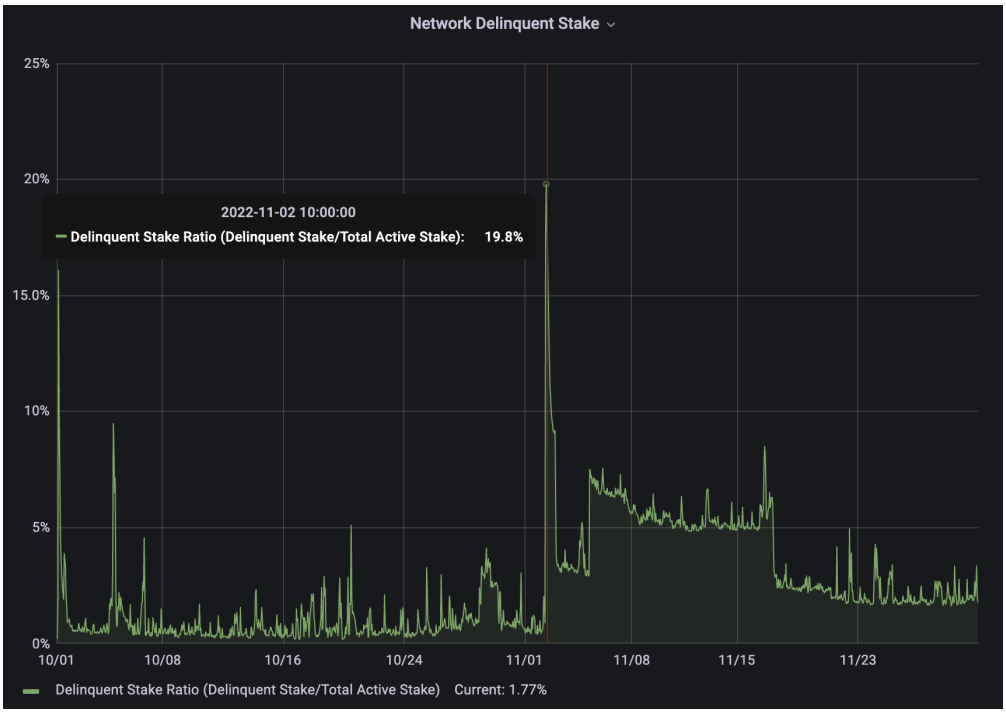

Figure N1: Delinquent stake ratio for Solana during the Hetzner shutdown.

[14] See hardware requirements for Solana and Ethereum. Accessed April 14, 2024.

O Ethereum Merge Exogeneity

We have several reasons to believe the exogeneity of the Merge as a technical shock. First, the Merge was a literal merging of Ethereum's execution layer with a new, external Proof-of-Stake consensus layer, called the Beacon chain, which has been live and rewarding validators since December 1, 2020, but without validating any transactions on Ethereum.[15] The Merge simply replaced the Proof-of-Work consensus layer with the existing PoS Beacon chain. The Merge occurred literally in the span of less than 20 seconds and did not require changes to the network before or after the Merge.[16] Second, while the Merge had been announced in advance, it has been delayed five times since 2017 across five forks: Byzantium in 2017, Muir Glacier in 2020, London in August 2021, Arrow Glacier in December 2021, and finally Gray Glacier in June 2022.[17] Third, the Merge was a technical upgrade of nearly unprecedented scale or complexity in the blockchain ecosystem. Together, these factors highlight the difficulty of anticipating any effect on the Ethereum ecosystem on September 15, 2022, by participants.

[15] See more on the Beacon chain on ethereum.org. Accessed April 22, 2024.
[16] See block 15,537,394 on etherscan.io. Accessed April 22, 2024.
[17] See the forks of Ethereum on ethereum.org. Accessed April 22, 2024.

P Data Collection Details

The hashrate is a measure of the collective computing power in a Proof-of-Work blockchain network; specifically, it is the total number of hash functions that all nodes in a network compute per second. While it cannot be directly measured, it is estimated from the number of blocks currently being mined and the current block difficulty. The computational power of an individual node can be affected by gradual hardware degradation, hardware upgrades, changes in efficiency due to electricity or cooling, or, most importantly for our context, changes in the amount of resources used by the node, which can be hash power or blockchain tokens. Most crypto miners are profit-maximizing and thus operate at maximum efficiency given the high costs of ASICs and GPUs, even in geographic regions with cheap electricity.[18]

[18] See compassmining.io.

Q Additional Decentralization Measures

The Gini coefficient Gini_{it} is a measure of inequality within a distribution and quantifies the disparity in block production across nodes in blockchain i on day t. Formally, it is defined as

$$\text{Gini}_{it} = \frac{\sum_{k=1}^{N_{it}} \sum_{l=1}^{N_{it}} |x_k - x_l|}{2N_{it}^2 \overline{x}_{it}},$$

where x_k and x_l are values from X_{it} and $\overline{x}_{it}$ is the average of all x_k in X_{it}. It is unitless and is bounded by $[0, 1]$. A higher Gini coefficient means less decentralization.[19]

The Nakamoto coefficient Nakamoto_{it} measures the minimum number of nodes required to collectively control over 51% of the block production on a given day. Formally, it is defined as

$$\text{Nakamoto}_{it} = \min\{n \in [1, ..., N_{it}] : \sum_{k=1}^{n} p_k \geq 0.51\},$$

where p_k is the fraction of blocks produced by the k^{th} node on day t (Srinivasan and Lee, 2017). The Nakamoto coefficient is a key measure for assessing the resilience of the blockchain against potential collusion or majority attacks. It has units of nodes and is bounded by $[1, N]$. A higher Nakamoto coefficient means more decentralization.

Finally, HHI_{it} indicates the concentration of block production among nodes in blockchain i on day t. Formally, it is defined as

$$\text{HHI}_{it} = \sum_{k \in K_{it}} p_k^2.$$

HHI_{it} is calculated as the sum of the squares of individual market shares of all participants. It is unitless and is bounded by $[1/N, 1]$. A higher HHI means less decentralization.

[19] The Gini coefficient can be misleading without also considering the number of nodes. For example, consider the case when there is only one node (*e.g.*, a dictatorship is one-man-one-vote). The Gini coefficient would then be zero, indicating full decentralization, but it is merely "fully equal" for only the single node. Hence, we consider it important to use multiple metrics when possible and Shannon entropy when one requires a single metric to encapsulate both the number of entities and the distribution of power over those entities.

References

Buterin, Vitalik. 2014. Ethereum: A next-generation smart contract and decentralized application platform. URL https://ethereum.org/669c9e2e2027310b6b3cdce6e1c52962/Ethereum_Whitepaper_-_Buterin_2014.pdf. Accessed on October 13, 2023.

Cambridge Centre for Alternative Finance. 2022. Distribution of bitcoin mining hashrate from september 2019 to january 2022, by country [graph]. In Statista. URL https://www.statista.com/statistics/1200477/bitcoin-mining-by-country/. Retrieved October 30, 2023.

Cameron, A. Colin, Jonah B. Gelbach, Douglas L. Miller. 2008. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* **90**(3) 414–427.

Capponi, Agostino, Sveinn Ólafsson, Humoud Alsabah. 2023. Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Management Science* **69**(11) 6455–6481.

Chen, Yan, Jack I. Richter, Pankaj C. Patel. 2021. Decentralized governance of digital platforms. *Journal of Management* **47**(5) 1305–1337. doi:10.1177/0149206320916755.

Cong, Lin William, Zhiguo He, Jiasun Li. 2021. Decentralized mining in centralized pools. *The Review of Financial Studies* **34**(3) 1191–1235. doi:10.1093/rfs/hhaa040.

Cong, Lin William, Xiang Hui, Catherine E. Tucker, Luofeng Zhou. 2023. Scaling smart contracts via layer-2 technologies: Some empirical evidence. *Management Science* **69**(12) 7306–7316. doi:10.1287/mnsc.2023.00281.

Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, Ari Juels. 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. *2020 IEEE Symposium on Security and Privacy (SP)*. 910–927. doi:10.1109/SP40000.2020.00040.

Garratt, Rodney J., Maarten R. C. van Oordt. 2023. Why fixed costs matter for proof-of-work–based cryptocurrencies. *Management Science* **69**(11) 6482–6507.

Gencer, Adem Efe, Soumya Basu, Ittay Eyal, Robbert Van Renesse, Emin Gün Sirer. 2018. Decentralization in bitcoin and ethereum networks. *Financial Cryptography and Data Security*, vol. 10957. 439–457. doi:10.1007/978-3-662-58387-6_24.

Halaburda, Hanna, Christoph Mueller-Bloch. 2020. Toward a multidimensional conceptualization of decentralization in blockchain governance. *Academy of Management Discoveries* **6**(4) 712–714.

Hsieh, Ying-Ying, Jean-Philippe Vergne. 2023. The future of the web? The coordination and early-stage growth of decentralized platforms. *Strategic Management Journal* **44**(3) 829–857. doi:10.1002/smj.3455.

Ju, Harang, Ehsan Valavi, Madhav Kumar, Sinan Aral. 2025. Are crypto ecosystems (de)centralizing? a framework for longitudinal analysis. *Commun. ACM* **68**(12) 104–111. doi:10.1145/3737446. URL https://doi.org/10.1145/3737446.

Mueller-Bloch, Christoph, Jonas Valbjørn Andersen, Jason Spasovski, Jungpil Hahn. 2024. Understanding decentralization of decision-making power in proof-of-stake blockchains: an agent-based simulation approach. *European journal of information systems* **33**(3) 267–286.

Nakamoto, Satoshi. 2008. Bitcoin: A peer-to-peer electronic cash system. URL https://bitcoin.org/bitcoin.pdf.

Petersen, Mitchell A. 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* **22**(1) 435–480.

Sai, Ashish Rajendra, Jim Buckley, Brian Fitzgerald, Andrew Le Gear. 2021. Taxonomy of centralization in public blockchain systems: A systematic literature review. *Information Processing & Management* **58**(4) 102584.

Srinivasan, Balaji S., Leland Lee. 2017. Quantifying decentralization. URL https://news.earn.com/quantifying-decentralization-e39db233c28e. Accessed: 10-08-2023.

The Cambridge Centre for Alternative Finance. 2023. Bitcoin mining map. URL https://ccaf.io/cbnsi/cbeci/mining_map. Accessed: October 30, 2023.