

Article

Pool-Operator Concentration After the 2021 China Mining Ban: Full Block-Attribution Evidence from the Bitcoin Mining-Pool Layer, 2020–2022

Craig Steven Wright 

Business School, University of Exeter, Exeter EX4 4PU, UK; cw881@exeter.ac.uk

Abstract

The May–September 2021 cryptocurrency-mining prohibition in the People’s Republic of China removed an industry that had become heavily concentrated in Chinese jurisdictions over the years preceding the ban. Using the full population of 97,005 Bitcoin blocks mined between May 2020 and February 2022, retrieved from the mempool.space block API, with pool-operator corporate domicile and China-linkage coded from publicly documented corporate-parent records, we estimate the post-prohibition shift in pool-attributed block share. On the top-13 panel covering approximately ninety-five per cent of attributed blocks, a two-way fixed-effects projection records a differential China-linked share movement of -5.45 percentage points per pool after the 24 September 2021 NDRC and PBOC enforcement event (cluster-robust SE 2.55; CRV1 $p = 0.054$, not significant at five per cent; wild cluster bootstrap $p = 0.014$ and exact-permutation $p = 0.023$ reported as small-sample diagnostics, not relied on for significance; under a permutation conditioned on baseline share the latter rises to $p = 0.100$). This paper’s primary magnitudes are the aggregate compositional movements, which do not depend on the cluster-level test. Aggregate China-linked share falls from 78.47 to 60.01 per cent between the placebo (May–October 2020) and persistence (November 2021–February 2022) windows; non-China-linked share rises from 3.50 to 27.02 per cent, with Foundry USA absorbing approximately 15.7 percentage points (approximately 67 per cent of the non-China-linked share gain). The within-non-China-subset normalised Herfindahl–Hirschman Index falls from 10,000 to 4191, consistent with rising within-subset dispersion among non-Chinese-domiciled operators. Among the identified non-China-linked operators absorbing material share in the analytical panel, the displacement is concentrated at named legal persons, chiefly with documented United States and European corporate domicile; the residual excluded/unknown category is not used to support that claim.



Academic Editors: Juraj Fabuš and Shyan-Ming Yuan

Received: 22 May 2026

Revised: 10 July 2026

Accepted: 11 July 2026

Published: 16 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: Bitcoin mining; mining pool concentration; regulatory shock; cryptocurrency market structure; China crackdown

JEL Classification: K23; L51; O17; P14; Q48

1. Introduction

The institutional theory of permissionless consensus is built on a counterfactual: that no individual jurisdiction can extinguish the network because computational participation is geographically dispersed, the entry technology is portable, and the participants are atomistic. The May 2021 announcement of the State Council Financial Stability and

Development Committee that cryptocurrency mining and trading would be subjected to a national crackdown, and the subsequent measures during June, July, and September 2021 by Sichuan Province, the People's Bank of China, the National Development and Reform Commission, and a further nine central authorities, supplied an unusually informative empirical setting for examining that counterfactual at the pool-operator layer of a proof-of-work network. At the time of the announcement, contemporary estimates placed the People's Republic between 45 and 75 per cent of global Bitcoin mining hashrate, with the largest pool operators headquartered in Beijing, Hangzhou, and Shenzhen, and the largest hardware manufacturer, Bitmain Technologies, operating eleven mining farms across Sichuan, Xinjiang, and Inner Mongolia [1]. The prohibition was therefore not a marginal regulatory adjustment. It was a unilateral nation-state action against one of the largest documented jurisdictional concentrations of the industry of an activity that strong polycentric and atomistic accounts of permissionless consensus have often treated as difficult for any single jurisdiction to regulate effectively, given portability, open entry, and jurisdictional substitutability.

The institutional outcome of that experiment is the subject of this paper. We ask three questions. First, was the post-ban displacement of pool-attributed share diffuse across destinations, or did it concentrate at identifiable non-China-linked institutional intermediaries, chiefly in the United States and Europe? Second, was the displacement transient, with Chinese-affiliated brands recovering their pre-ban shares once the regulatory climate stabilised, or was it durable on the time horizon over which institutional adjustment is observable? Third, what does the pattern of displacement imply for the institutional theory of permissionless consensus? In particular, we ask—as an open empirical question, not a foregone conclusion—whether the evidence at the pool-operator layer bears on the claim that pool-level concentration is corrigible and that polycentric governance can resist unilateral state coercion. We are careful throughout to distinguish what the pool-attribution evidence can and cannot establish about that claim (Section 6.1), and the conclusions we draw are confined to the operator layer the data actually measure.

We answer the first two questions through a two-way fixed-effects difference-in-differences design with placebo and pre-trend windows constructed from the full population of 97,005 Bitcoin blocks mined between 1 May 2020 and 28 February 2022, retrieved from the public mempool.space block-attribution API. There is no sampling: every block in the study window is included. Our treatment group is the set of mining-pool brands whose corporate parent was domiciled in the People's Republic of China at the date of the ban. The full classification ledger is reported in Table 1 below, with the corporate-parent evidence cited in line and the underlying primary-source documents listed in the data-availability statement. The control group is the set of mining-pool brands whose corporate parent was domiciled in the United States, the European Union, the United Kingdom, Switzerland, Canada, or Japan at the same date, and which therefore lay outside the legal reach of the 24 September 2021 NDRC Notice and PBOC Circular.

What this paper tests is precise and narrow. The empirical object is the pool-attributed block share—the fraction of blocks whose coinbase output is attributed by open-source labelling rules to a named pool brand—and not the physical location or ownership of mining hardware. On that object, we test four observable implications of the view that the operator layer of permissionless consensus is organised by identifiable corporate intermediaries rather than by anonymous atomistic participants: that the China-linked operator group loses attributed share after the prohibition (H1); that aggregate network mining activity recovers even as the operator set turns over (H2); that the displaced share is absorbed in a form that is concentrated rather than broadly dispersed (H3); and that the absorbing operators are named, publicly disclosed legal persons with documentable

domicile, capital structure, and infrastructure (H4). The contribution is threefold: the use of the full block population rather than a sample to characterise the operator layer; the explicit, reproducible corporate-domicile classification of pool operators with separate treatment of the cases where formal incorporation and operational headquarters diverge; and the formal treatment of the closed-share inference problem created by using operator shares of a closed total as the outcome.

Table 1. Pool-operator classification at 21 May 2021. The classification criterion is operational headquarters and regulatory exposure, not formal place of incorporation; the two columns are shown separately so that divergent cases are visible. Full primary-source ledger in the Supplementary Materials and in the Zenodo deposit (DOI 10.5281/zenodo.20754433).

Pool Brand	Formal Incorporation	Operational HQ	Class.
China-linked (treatment group)			
AntPool	Cayman/BVI (Bitmain)	Beijing	PRC
BTC.com	BitDeer → BIT Mining Ltd. (formerly 500.com); transfer 15 April 2021	Mainland-facing (Beijing)	PRC
F2Pool	Hong Kong	Beijing & Hong Kong	PRC
ViaBTC	PRC (Bixin Group)	Shenzhen	PRC
Poolin	PRC (Blockin Inc)	Beijing	PRC
Huobi.pool	Seychelles (Huobi Global)	Beijing (to Sep 2021)	PRC
1THash	PRC	Beijing	PRC
BTC.TOP	PRC (Jiang Zhuoer)	Beijing	PRC
Non-China-linked (control group)			
Foundry USA	USA (Delaware; DCG subsidiary)	Rochester NY	US
Braiins Pool	Czech Republic	Prague	EU
MARA Pool	USA (Nevada)	Las Vegas NV	US
Luxor	USA	Seattle WA	US
SBI Crypto	Japan	Tokyo	JP
Excluded for domicile ambiguity in baseline (both assignments tested)			
Binance Pool	Disputed (multi-jurisdiction)	Disputed	Excl.

Note: “Class.” is the baseline classification. For AntPool, F2Pool, and Huobi.pool the formal incorporation jurisdiction differs from the operational headquarters; the classification follows the operational headquarters, because the binding legal reach of the 24 September 2021 NDRC Notice and PBOC Circular fell on activity conducted within the People’s Republic irrespective of incorporation. BTC.com is a transition case: its pool business was transferred from BitDeer to BIT Mining Limited (formerly 500.com) with the first closing announced on 15 April 2021 [2], before the announcement date, and it is classified China-linked on mainland-facing operational exposure rather than parentage [3]; a sensitivity specification excluding it is reported in Section 5.4. The full ledger reports legal incorporation, operational headquarters, beneficial parent/control, PRC regulatory exposure, baseline classification, and sensitivity classification in separate columns.

We find that the displacement is persistent within the study window and concentrated among identifiable substitute operators. Under the headline specification with the post-period beginning after the 24 September 2021 NDRC Notice and PBOC Circular, the difference-in-differences coefficient on the top-13 panel is -5.45 percentage points per pool (cluster-robust SE 2.55; CRV1 $p = 0.054$; wild cluster bootstrap $p = 0.014$; exact treated-label permutation $p = 0.023$ over all $C(13, 8) = 1287$ assignments). The pre-trend over the four pre-treatment months is statistically insignificant. This test has limited power: with only four pre-treatment monthly observations the slope is estimated with two residual degrees of freedom, so the parallel-trends assumption is neither established nor refuted by it. The point estimate is itself a non-trivial positive slope ($+1.161$ percentage points per month), opposite in sign to the post-period decline; a departure from parallel trends in that direction would bias the estimated decline toward zero rather than create it, so the

pre-trend does not generate the result and, if anything, understates it. The constraint is inherent in the event window—the prohibition fixes the length of the clean pre-period at four months—rather than one removable by further estimation, and the difference-in-differences coefficient is read throughout with that limitation in view. The China-linked aggregate share falls from 78.47 per cent in the placebo window to 60.01 per cent in the persistence window; the non-China-linked share rises from 3.50 to 27.02 per cent. The full window-level series, the event-study leads and lags, and the robustness analyses are reported in Section 5.

It concentrates at corporate intermediaries that publicly announced themselves, contemporaneously with the ban and afterwards, as the institutional substitute for Chinese mining-pool capacity. Foundry Digital LLC., a wholly-owned subsidiary of Digital Currency Group founded in 2019 and headquartered in Rochester, New York, launched the Foundry USA Pool in 2020 and was responsible, on the firm's contemporaneous public statements as quoted in the financial press of late 2021, for channelling more than four hundred million dollars of mining hardware into North America during the displacement window. Its share of attributed blocks rises from a pre-ban mean of 0.56% in the months May 2020 through April 2021 to a persistence-window mean of approximately 16.25%, reaching 17.62% in February 2022.

We interpret the third question as one about cryptocurrency-market structure under regulatory shock. The post-ban equilibrium is one of operator substitution rather than protocol change: the underlying ASIC capital is portable, the protocol stratum is open, and customer-side software can connect to any pool that publishes a stratum endpoint, but the legal personality of the pool operator is jurisdiction-specific and cannot be moved across regulatory boundaries. The pool-attributed shift between May 2021 and February 2022 therefore measures a change in the corporate-domicile composition of the operator set, not a change in network-level computational resources. The relevant fintech-market-structure question is whether the regulatory shock disperses pool-attributed share across a broader set of operators—the dispersive reading often associated with strong polycentric or atomistic descriptions of permissionless-consensus markets—or concentrates that share at a small set of named legal persons with substitute jurisdictional standing. The headline-share trajectory documented in this paper sits on the concentration side of that question: a single non-Chinese operator absorbs approximately 67 per cent of the non-China-linked share gain, while the within-non-Chinese subset becomes more dispersed only after accounting for entry of additional small operators. The mixed reading across the headline and within-subset measures is consistent with the parallel diagnostic results of Davidson [4] on decentralised cryptocurrency-platform governance and Murtazashvili et al. [5] on the residual role of jurisdictional legal infrastructure in nominally polycentric blockchain orders.

The remainder of this paper is organised as follows. Section 2 sets out the four testable hypotheses about the post-ban equilibrium of the mining-pool market. Section 3 documents the regulatory action through primary legal sources. Section 4 sets out the data, the pool-classification scheme, and the empirical design. Section 5 presents the difference-in-differences results, placebo tests, pre-trend tests, and persistence analysis. Section 6 interprets the results in the context of the cryptocurrency-market-structure and blockchain-governance literature, including the unresolved methodological questions identified by Sai and Vranken [6]. Section 7 closes with policy implications and limitations.

2. Pool Operators and Testable Hypotheses

The unifying empirical claim of this paper can be stated as one proposition. Models of permissionless consensus that treat the operational layer as anonymous, atomistic, and purely computational tend to abstract away from the corporate organisations that operate the pool layer. Once those organisations are made explicit, the institutional incidence of

a unilateral regulatory shock against the largest jurisdictional concentration of the industry becomes a natural object of study, and the post-2021 evidence is informative about where that incidence falls. The four hypotheses developed in the subsections below are observable implications of this view. H1 (pool brand displacement) is the implication that the operator-attributed distribution of pool-attributed block share shifts away from organisations in the prohibited jurisdiction. H2 (aggregate activity recovery) is the implication that aggregate network mining activity recovers even as the operator set turns over. H3 (form of absorption distribution) is the implication that the substitute organisations are identifiable, not anonymous, and that the form of the distribution is informative about the polycentric prediction. H4 (identifiable corporate intermediaries) is the implication that the operators absorbing the displaced share are named, publicly-disclosed organisations whose domicile, capital structure, and operational footprint can be documented from primary sources.

The empirical contribution is the documentation that the operational layer of permissionless consensus reconstitutes corporate-intermediary structure under regulatory shock, even where the protocol layer remains formally open. The blockchain-governance literature has identified operator-level concentration as the empirical question [7]; the parallel finding for decentralised autonomous organisations [4] and for nominally polycentric blockchain orders that retain off-chain institutional backstops [5] reflects the same empirical pattern observed here at the mining-pool layer.

2.1. Pool Operators as Governance Organisations

The unit of analysis in this paper is the mining-pool operator. A mining-pool operator is a legal person that solicits hashrate contributions from miners, aggregates those contributions, allocates the proceeds of successful block discoveries among contributors, retains a fee, and exercises certain operational discretions over the construction of candidate blocks. The operator is therefore not the same entity as the underlying miners, who are physically distributed and individually anonymous, and not the same entity as the protocol, whose rules are common to every pool. Pool operators are organisations in the sense of North [8] (pp. 3–5): they are purposive groups of individuals bound by a common objective, operating within institutional constraints which they themselves did not write but which they confront, contest, and occasionally amend.

We adopt this distinction directly. The distinction between an institution as a system of embedded rules and the organisation that acts within it is the one Hodgson [9] draws: on this reading the protocol is the institution—a common rule-system no operator authored—while the pool operators are organisations constituted to act within it. Why block production should be organised through such intermediary structures at all, rather than through the anonymous spot interaction the atomistic model assumes, is the question transaction-cost economics addresses: Williamson [10] argues that recurrent, non-trivial transactions are removed from spot markets and placed under specialised governance structures because organising them atomistically is itself costly, which is consistent with the emergence of identifiable pool operators as the governance layer of block production. The May 2021 prohibition was directed at organisations: it altered the legal status of mining as an industry, classifying it as an obsolete sector under the National Development and Reform Commission’s industrial-structure-adjustment catalogue [11]. It did not, and could not, alter the protocol, which is not subject to revision by a single jurisdiction. What the prohibition altered is the set of organisations that could legally participate in the protocol from within the People’s Republic. The empirical question that follows is whether the operator-level distribution of attributed share after the ban shifted in the direction the institutional reading predicts—away from China-domiciled organisations and toward organisations domiciled where mining-pool operation remained lawful—and whether that shift persisted.

The framing of pool operators as the targets of legal sanction also positions our analysis within the legal-deterrence literature initiated by Becker [12], which treats unlawful activity as the product of a rational decision-maker who weighs expected returns against the probability of detection multiplied by the expected sanction. The May 2021 prohibition raised the expected sanction for mining-pool operation in China to a level approximating the entire economic value of the activity, since continued operation became formally illegal under the September 2021 PBOC circular. Under the Becker framework, the predicted response of a rational pool operator is to exit the jurisdiction, and the empirical question is whether exit at the operator level is what the block-attribution data show.

2.2. Four Testable Hypotheses

The framework yields four testable hypotheses on the post-shock structure of the mining-pool market. The formal model derivations are set out in the Supplementary Materials and are deposited at Zenodo alongside the replication materials.

Because the hypotheses are stated as predictions of rival readings of permissionless-consensus markets, we first define those readings operationally, in terms of the measurable behaviour of pool-attributed share after a jurisdiction-specific shock. A *strong atomistic* reading holds that pool-attributed share is distributed across many small, individually immaterial, and largely unidentifiable participants, so that no named operator holds a large share and a jurisdiction-specific prohibition cannot displace a material fraction of attributed share; operationally, it predicts that the largest single operator's share remains small and that displaced share, if any, is absorbed diffusely across many entrants. A *weak atomistic* reading relaxes the identifiability requirement but retains the dispersion requirement: share may be traceable to operators, but no single operator absorbs a large fraction of any displacement. A *strong polycentric* reading holds that the network is governed by many independent centres of authority none of which can be eliminated by a single jurisdiction, and predicts, for the absorption of a displacement, that no single substitute operator dominates (the H3a concentration prediction is rejected) while the substitute subset is internally dispersed (H3b holds). A *weak polycentric* reading retains only the resilience claim—that the network as a whole survives the elimination of any one jurisdictional centre—without the prediction that absorption is internally dispersed; operationally, it is consistent with the protocol surviving a partial but material turnover of its largest jurisdictional operator concentration even if a large fraction of the displaced share is captured by a single substitute operator. These definitions are stated in terms of pool-attributed share specifically, and the tests below evaluate them at the pool-operator layer only, not at the level of the network as a whole (Section 6.1).

H1 (pool brand displacement): the ratio of pool-attributed share between Chinese-domiciled and non-Chinese-domiciled pool operators falls discontinuously after May 2021 and remains depressed in the persistence window. The mechanism is jurisdiction-specificity of operator infrastructure: pool-server hosting, low-latency stratum equipment, customer-software interfaces, and bilateral procurement contracts are tied to the mainland operating environment and lose value once the customer base relocates.

H2 (aggregate activity recovery): aggregate network mining activity exhibits a temporary shock around the ban event but returns toward its pre-ban level within the study window, consistent with redeployment rather than complete destruction of mining capacity. The Bitcoin protocol—longest-chain rule, difficulty-adjustment algorithm, block-reward structure, network-topology conventions—is unchanged across the 2021 prohibition; the directly measured Bitcoin network hashrate (obtained from Coin Metrics) is the testable proxy for aggregate proof-of-work activity, because the hypothesis concerns aggregate hashing power and not energy consumption or carbon intensity.

H3 (form of absorption distribution): distinguished into two sub-hypotheses. H3a (headline absorption concentration): the leading non-Chinese pool operator absorbs more than half of the change in pool-attributed share. H3b (within-substitute-subset dispersion): the within-substitute-subset normalised concentration index declines, indicating that the absorbing subset is not internally dominated by a single residual claimant. H3b has limited discriminating power and is reported as a characterisation rather than a discriminating test: because only one non-China-linked operator produced blocks in the placebo window, the within-subset index begins at its mechanical maximum, so any entry lowers it irrespective of the polycentric question (Section 5).

The strong polycentric reading of permissionless-consensus markets [4,7,13] predicts that H3a is rejected and H3b holds; the alternative concentration reading predicts that H3a holds and is silent on H3b. This polycentric reading traces to Ostrom's analysis of self-governance: Wilson, Ostrom, and Cox [14] generalise the design principles under which collective-action problems are resolved neither by atomistic self-interest nor by external imposition, but through identifiable institutional arrangements the participants themselves maintain—the sense in which permissionless-consensus governance is described as polycentric rather than atomistic. We test H3a against the headline share concentration of the leading substitute operator and H3b against the within-substitute-subset normalised Herfindahl–Hirschman Index.

H4 (identifiable corporate intermediaries): against the atomistic-miner alternative prominent in the theoretical literature [15]: the displaced pool-attributed share is absorbed by named, publicly disclosed corporate entities whose domicile, capital structure, and operational infrastructure can be documented from primary sources, rather than by unidentifiable atomistic participants. The atomistic model treats miners as anonymous competitors interacting only via the protocol, with participation uniform across the surface of the earth, weighted only by local cost of electricity; on this view a jurisdiction-level prohibition could not displace a meaningful fraction of pool-attributed block share to identifiable operators in other jurisdictions. If instead the operator-level corporate-structure view is correct, the displacement should be sharp, brand-specific, and persistent. H4 is a characterisation of the absorbers rather than a sharp test of the atomistic alternative: any operator large enough to absorb material share is necessarily a named entity, so H4's content is the documentation of which corporate intermediaries absorbed the displacement and their corporate structure, not a high-powered rejection of atomistic participation.

3. The May–September 2021 Prohibition

We document the regulatory chronology from primary legal sources collected and translated by the United States Library of Congress, Global Legal Research Directorate [11,16].

On 21 May 2021, the State Council Financial Stability and Development Committee, chaired by Vice Premier Liu He, held its 51st meeting and announced a crackdown on Bitcoin mining and trading [11]. This was the first central-level act and it framed the subsequent provincial and central-administrative measures.

On 18 June 2021, the Sichuan provincial government issued an order requiring that cryptocurrency mining within the province be screened, cleaned up, and terminated. Sichuan was at this date the largest mining-density region in China, accounting on Cambridge Centre for Alternative Finance estimates for approximately one-third of the country's mining-pool-attributed hashrate during the wet season [17,18]. The Cambridge data show the province's reported mining share collapsing to zero in the July–August window; the same data show similar collapses in Xinjiang, Inner Mongolia, and Yunnan.

On 24 September 2021, the People's Bank of China and nine other central authorities issued the Circular on Further Preventing and Disposing of Speculative Risks in Virtual

Currency Trading [16]. The circular declared all cryptocurrency trading transactions to be illegal financial activities. On the same day, the National Development and Reform Commission and nine other authorities issued the Notice on Regulating Virtual Currency Mining Activities, which added cryptocurrency-mining activities to the Industrial Structure Adjustment Guidance Catalogue's obsolete category, prohibited new investment, prohibited the supply of electricity to mining operations, and prohibited the operation of mining businesses under the guise of data centres [11].

The legal architecture of the prohibition was therefore directed at operators rather than at the protocol. It targeted Chinese mining-pool operators, Chinese ASIC manufacturers, Chinese operators of mining facilities, electricity suppliers to those facilities, and the financial institutions that intermediated between Chinese miners and external markets. It did not, and could not, prohibit the protocol itself. This design feature is central to our identification: the regulatory action altered the operator set without altering the protocol layer.

4. Data and Identification

4.1. Block-Level Pool-Attribution Panel

We construct a panel of pool-operator block attributions covering the full population of 97,005 Bitcoin blocks mined between 1 May 2020 00:00 UTC and 1 March 2022 00:00 UTC. The panel runs from May 2020 through February 2022, comprising twenty-two months and an average of 4409 blocks per month. The endpoints are chosen to provide substantial pre-event coverage, a bounded post-enforcement persistence window, and deliberate truncation before later confounds: the start date precedes the announcement by twelve months, giving a placebo window and a four-month pre-trend window free of any major regulatory action against Chinese mining, and the end date follows the enforcement event by five months, the longest post-event horizon over which the pool-attribution composition can be observed before the subsequent partial return of mining capacity to China during 2022 and the mid-2022 cryptocurrency-market drawdown introduce confounding dynamics outside the scope of a single regulatory-shock design. The pre-event window is accordingly longer than the post-enforcement window; this is deliberate, prioritising a stable pre-event baseline and clean placebo period over a longer post-period that later confounds would contaminate. The window is therefore long enough to inspect pre-event stability across the available monthly data and to observe persistence, and short enough that the displacement is attributable to the 2021 prohibition rather than to longer-horizon structural change. For each block, we record the pool operator from the `extras.pool.name` attribute returned by the `/api/v1/blocks/:height` endpoint of the `mempool.space` public API, which applies `mempool.space`'s pool-identification scheme to the `coinbase` scriptsig. The complete raw API responses, the deterministic mapping from API response to pool classification, and the analysis scripts used to produce every number reported in this paper are deposited at Zenodo (DOI 10.5281/zenodo.20754433). The retrieval was completed on 11 May 2026 against the `mempool.space` production API; no sampling step is involved. Because the `mempool.space` pool-identification scheme can be updated retroactively, replication should be run against the deposited raw API responses rather than by re-querying the live endpoint, which may return revised labels for the same block heights.

No sampling of blocks is performed. All 97,005 blocks in the study window are included. The full-population design eliminates seed-dependent variability and aligns the empirical object of interest with the unit of attribution used in policy-relevant network-share statistics. The effective sample for the headline difference-in-differences estimator is 286 pool-months for the top-13 specification (13 pools across 22 months); cluster-robust standard errors are computed at the pool level with thirteen clusters.

Pool attribution is an attribution-layer measurement, not a legal-control measurement. The mempool.space labels are treated in this paper as reproducible public attribution data, not as definitive corporate-control evidence. The mempool.space pool-identification scheme applies pattern-matching rules to the coinbase scriptsig of each block, producing a deterministic pool-name label per block; the rules are open source and the labelling output is reproducible from the public block data. The scheme is consistent with the corporate-parent tracking in Stoll et al. [17] and Sun et al. [18] for the named pools that enter our analytical panel; we cross-checked the placebo-window monthly share of each top-13 pool against the corresponding monthly share reported by the BTC.com pool-statistics service where available, and confirmed agreement within one percentage point of monthly share for every top-13 pool in the placebo window. This quantitative cross-check covers the placebo window only; we have no equivalent independent series for the treatment and persistence windows, so attribution error during the migration period—including any private-label or proxy mining the scheme cannot detect (limitation (ii) below)—cannot be ruled out quantitatively, and the treatment-window grounding for the named pools rests on the corporate-parent consistency with Stoll et al. and Sun et al. noted above rather than on an independent share series. The corporate-domicile classification in Section 4.2 is then applied to the attributed pool brand, independently of the attribution-layer labelling. Three known limitations of any pool-attribution scheme apply: (i) pools may change their coinbase tag mid-period and require manual reconciliation, which we do not detect in our window for the top-13 pools but which is a generic source of attribution noise for smaller pools; (ii) some blocks attributed to a named pool may have been mined under a private-label or proxy arrangement, which the attribution scheme cannot detect; (iii) some blocks with no recognised tag enter the residual Unknown category (1.4 per cent of total in our window).

The panel covers four windows by design. The placebo window is May 2020 through October 2020, during which no major regulatory action against Chinese mining was in effect. The pre-trend window is January 2021 through April 2021, the four months immediately before the State Council announcement. The immediate-treatment window is May 2021 through October 2021, the period of active Sichuan and central-government enforcement. The persistence window is November 2021 through February 2022, when the central-government Notice and the September 24 PBOC circular had been in effect for at least one full month.

4.2. Pool Nationality Classification

The classification universe is the set of pool operators that produced at least one block in the full-population panel. The analytical panel includes those pools for which a corporate-parent domicile at the date of the 21 May 2021 State Council announcement is documentable from publicly accessible primary sources: HKEX disclosures, US SEC filings, the Hong Kong Companies Registry, the Delaware Division of Corporations, the Cayman Islands General Registry, and contemporaneous corporate disclosures of the pool-operating entity. Pools whose corporate parent is contested, ambiguous, or assertion-only at the announcement date are excluded from both groups but retained in the underlying panel. The classification rule is deterministic at the pool-name level; the rule is the same one applied by Stoll et al. [17], who partition pools into Chinese and European categories on corporate-parent grounds, and is consistent with the corporate-parent tracking in Sun et al. [18].

We classify a pool as China-linked if its corporate parent at 21 May 2021 was domiciled in or operationally headquartered in the People's Republic of China, including Hong Kong only when the corporate-parent operations were principally mainland-facing. We classify a pool as non-China-linked if its corporate parent was domiciled in the United States, the European Union, the United Kingdom, Switzerland, Japan, or Canada at the same date.

We exclude from both groups those pools whose corporate-parent domicile was ambiguous; in our sample this excludes only Binance Pool.

The criterion that governs the classification is the operational headquarters and regulatory exposure of the corporate parent, not its formal place of incorporation. The two diverge for several of the largest operators: AntPool's parent, Bitmain, carries a Cayman/BVI incorporation structure while its mining-pool operations were directed from Beijing; F2Pool's parent was Hong Kong-registered with Beijing and Hong Kong operations; Huobi's pool operated under a Seychelles-registered parent with mainland-facing operations through to September 2021. We use operational headquarters because the treatment is the binding legal reach of the 24 September 2021 NDRC Notice and PBOC Circular, which fell on mining and trading activity conducted within the People's Republic, irrespective of where the operating entity was incorporated. Formal incorporation in an offshore jurisdiction did not place a Beijing-operated pool outside the reach of the prohibition; operational presence did. Classifying on incorporation rather than operations would therefore misassign exactly those entities whose mainland operations made them the subjects of the regulatory shock. The full ledger, reported in the Supplementary Materials and deposited at Zenodo (DOI 10.5281/zenodo.20754433), records legal incorporation and operational headquarters in separate columns so that the reader may assess the classification rule; the sensitivity analysis (Section 5.4) tests the two material ambiguous cases in the baseline design, Binance Pool and BTC.com.

Two forms of classification uncertainty are present and are treated explicitly rather than assumed away. The first is genuine domicile ambiguity at the announcement date: Binance Pool's parent operated across multiple jurisdictions with a disputed principal place of business, so Binance Pool is excluded from both groups in the baseline and is then tested under both the China-linked and non-China-linked assignments in Section 5.4. The second is an ownership transition that completed shortly before the announcement: the BTC.com pool business was transferred from BitDeer Technologies Holding Company to BIT Mining Limited (formerly 500.com Limited), with the first closing announced on 15 April 2021 [2]. Because the transfer preceded the 21 May 2021 announcement, BTC.com is not a Bitmain subsidiary at the announcement date; it is classified China-linked on the basis of material mainland-facing operational exposure during the treatment window [3], and we report a sensitivity specification that excludes it entirely (Section 5.4); the result is unchanged. No top-13 pool underwent a parent-domicile change during the treatment window. We do not claim that the analyst-coded partition is free of measurement error; we claim that the error is documented, that its two material sources are tested directly, and that the headline displacement survives those tests.

4.3. Identification

The identification rests on the unilateral and exogenous character of the State Council's 21 May 2021 announcement and the subsequent provincial and central enforcement events. The chronology used throughout is as follows: 21 May 2021 (State Council announcement) is the public policy signal; 18 June 2021 (Sichuan provincial order) is the first binding sub-national enforcement; and 24 September 2021 (the NDRC Notice and the PBOC Circular) is the binding national enforcement shock. The primary post-treatment month in the headline specification is therefore October 2021, the first full month after the September enforcement event. The May 2021 announcement and the June 2021 provincial order are reported as secondary timing tests.

The announcement was not preceded by a public consultation, was not anticipated by major industry forecasts, and was not the product of bilateral negotiation with industry stakeholders. Pre-trend tests in the four months January through April 2021 show a positive but statistically insignificant drift in China-linked share: an OLS regression of the

monthly China-linked share on a linear time trend over those four months yields a slope of 1.161 percentage points per month (OLS standard error 0.565, $t = 2.054$, $p = 0.176$). The null of zero pre-trend is not rejected at conventional significance levels.

The estimating equation for the panel of pool i in month t is given in Equation (1):

$$s_{it} = \alpha + \beta \cdot \text{Treat}_i \cdot \text{Post}_t + \gamma_i + \delta_t + \varepsilon_{it}, \quad (1)$$

where s_{it} is pool i 's monthly share of attributed blocks (per cent), Treat_i takes the value one for China-linked pools and zero for non-China-linked pools, Post_t takes the value one for months from the post-period start onward and zero before, γ_i are pool fixed effects, δ_t are calendar-month fixed effects, and ε_{it} is the error term clustered at the pool level. The coefficient β is the differential change in monthly per-pool share for a China-linked pool relative to a non-China-linked pool after the post-period start. Cluster-robust inference uses the t -distribution with $G - 1$ degrees of freedom where G is the number of pool clusters. Wild cluster bootstrap [19] and exact treated-label permutation inference are reported alongside this, following Cameron and Miller [20]. Given the small number of clusters ($G = 13$), the conservative bias of CRV1 with t -distributed critical values relative to wild cluster bootstrap inference, documented by MacKinnon and Webb [21], is empirically visible: the CRV1 p -value of 0.054 is materially larger than the wild bootstrap p -value of 0.014.

A descriptive overlay of the global policy-uncertainty environment, based on the Cryptocurrency Policy Uncertainty Index (UCRY Policy) of Lucey et al. [22], is reported in the Supplementary Materials as a robustness reference. The UCRY series is not used as a regression covariate because, in a panel with calendar-month fixed effects, any aggregate time-varying covariate is collinear with the month dummies.

5. Results

5.1. Window-Level Shares

Table 2 presents the window-level shares for the China-linked treatment group, the non-China-linked control group, and the excluded/unidentified residual.

Table 2. Window-level mining-pool shares (full-population panel, $N = 97,005$ blocks).

	Placebo May–October 2020	Pre-trend January–April 2021	Treatment May–October 2021	Persistence November 2021– February 2022
China-linked share (%)	78.47	76.35	68.50	60.01
Non-China-linked share (%)	3.50	5.34	14.72	27.02
Excluded/unidentified (%)	18.03	18.31	16.78	12.97
N blocks (window)	26,583	17,383	26,338	17,721

The placebo and pre-trend windows reflect a stable pre-ban configuration in which Chinese-domiciled pool operators dominate the network. The immediate-treatment window shows a fall in China-linked share of approximately 10 percentage points (from 78.47% to 68.50%) and a corresponding rise in non-China-linked share from 3.50% to 14.72%. The persistence window shows the gap widening further: China-linked share falls to 60.01%, while non-China-linked share rises to 27.02%. The decrease in China-linked share between placebo and persistence is 18.5 percentage points; the increase in non-China-linked share is 23.5 percentage points; the residual excluded/unidentified share also falls, from 18.0 to 13.0 per cent.

Figure 1 shows the trajectory of the two group-level shares across the four windows. The pre-treatment trajectories do not show a large visible break, although formal power to detect a pre-trend is limited by the small number of monthly pre-treatment observations;

the post-treatment trajectories diverge. The divergence does not reverse within the available data window.




Figure 1. Group-level mining-pool shares, May 2020–February 2022. The upper (red) series is the China-linked group’s monthly share of attributed blocks; the lower (blue) series is the non-China-linked group’s share. The dotted vertical line marks the 21 May 2021 State Council announcement; the dashed vertical line marks the 24 September 2021 NDRC Notice and PBOC Circular enforcement event. Shaded backgrounds identify, left to right, the placebo (May–October 2020), pre-trend (January–April 2021), treatment (May–October 2021), and persistence (November 2021–February 2022) windows. Source: monthly block-attribution panel constructed from the full population of 97,005 mempool.space-attributed blocks.

Figure 2 disaggregates the same trajectory by individual operator. It shows the monthly share composition across the named China-linked operators, Foundry USA, the other non-China-linked operators, and the excluded and unidentified residual. The compression of the China-linked operators (the red bands) after the September 2021 enforcement event, and the simultaneous expansion of Foundry USA (the green band), are visible directly in the operator-level composition rather than only in the group aggregates.

5.2. Difference-in-Differences Estimates

Table 3 reports the two-way fixed-effects difference-in-differences estimator on three specifications. The estimator is not a conventional individual-level treatment-effect estimator: the outcome is a share of a closed-share system, so treatment-group share-loss and control-group share-gain are mechanically linked. We therefore interpret the coefficient as a linear projection of the within-cell compositional shift, not as an average treatment effect on the treated. The exact treated-label permutation inference (over all $C(13, 8) = 1287$ assignments of the eight treated labels) is reported as a diagnostic check on the unusualness of the observed treatment-label pattern relative to arbitrary relabellings of the top-13 pool panel, not as decisive causal inference. Conditional on the September 2021 enforcement event being the binding regulatory shock, the magnitude of the coefficient describes the size of the post-event differential compositional movement of attributed-block share between exposed (China-linked) and substitute (non-China-linked) pool brands. The per-pool coefficient is an average over the pools in the panel and is therefore sensitive to panel breadth—it is -5.45 on the top-13 panel and -3.23 on the 25-pool panel (Table 3)—so the panel-invariant aggregate group-share movements (the

–18.46 pp fall in China-linked share, Table 4) are treated as the primary magnitude and the per-pool coefficient as a secondary, panel-specific summary.




Figure 2. Monthly share composition of attributed blocks by operator, May 2020–February 2022. Bands are stacked to sum to one hundred per cent of attributed blocks in each month. The China-linked operators are shown individually (red shades) above the Foundry USA band (green) and the other non-China-linked operators (light green); the Binance/excluded and unidentified residual are shown at the top. The dashed vertical line marks the 24 September 2021 enforcement event. Source: monthly block-attribution panel, full population.

The primary panel-set is the top-13 panel consisting of the eight China-linked pools and the five non-China-linked pools identified in the classification table. We make this the primary panel for three reasons. First, the top-13 panel covers approximately ninety-five per cent of attributed network blocks. Second, the empirical argument concerns identifiable corporate operators with documented domicile; the smallest pools have less complete primary-source documentation and add measurement noise. Third, the top-13 panel admits exact treated-label permutation inference over all $C(13, 8) = 1287$ assignments; the extended 25-pool panel does not.

The Treat $\times$ Post coefficient is negative across all three specifications. On the primary top-13 specification with the post-period beginning at the 24 September 2021 NDRC Notice and PBOC Circular, the coefficient is –5.45 percentage points per pool, with CRV1 p -value 0.054, wild cluster bootstrap p -value 0.014, and exact treated-label permutation p -value 0.023. Under the alternative dating with the post-period beginning at June 2021 and May 2021 excluded, the top-13 coefficient is –5.06 percentage points per pool (CRV1 $p = 0.070$, wild cluster bootstrap $p = 0.045$, exact permutation $p = 0.066$). On the extended panel, under the September post-period, the coefficient is –3.23 percentage points per pool (CRV1 $p = 0.031$, wild cluster bootstrap $p = 0.001$). Inferential status is reported transparently rather than asserted. Under CRV1 with $t(G - 1)$ critical values—the only procedure here that does not require an auxiliary small-sample assumption—the headline coefficient has $p = 0.054$ and is not significant at the five percent level (it is significant at ten percent); under the June-2021 dating $p = 0.070$, and only the extended 25-pool panel reaches $p = 0.031$. The wild cluster bootstrap returns markedly smaller p -values (0.014, 0.045, 0.001) than CRV1 in every specification. That ordering is the reverse of the usual small-cluster pattern, in which the wild bootstrap is the more conservative correction to CRV1 over-rejection; with thirteen clusters and an eight-to-five treated-to-control split the bootstrap’s calibration is not established here, so the wild-bootstrap and exact-permutation p -values are reported as small-sample diagnostics rather than as the basis for a significance

claim. The paper does not rest on a conventionally sized significant treatment effect: its inferential content is the transparently reported projection together with the descriptive decomposition (Table 4) and the aggregate share movements, which do not depend on the cluster-level test. Figure 3 reports the corresponding event-study leads and lags.

Table 3. Difference-in-differences estimates on the full-population block-attribution panel.

	Primary Spec. Top-13 Panel Post = October 2021	June 2021 Post Top-13, May Excl. Post = June 2021	Full Panel (rob.) Extended 25-Pool Post = October 2021
Treat $\times$ Post (pp/pool)	−5.45	−5.06	−3.23
CRV1 std. error	(2.55)	(2.55)	(1.41)
t-statistic	−2.14	−1.99	−2.30
p-value (CRV1, t-dist)	0.054	0.070	0.031
p-value (wild cluster, B = 999)	0.014	0.045	0.001
p-value (exact permutation)	0.023	0.066	—
N pool-months	286	273	550
Number of pool clusters G	13	13	25
Pool fixed effects	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes

Note: Outcome: monthly per-pool share of attributed blocks (%). The primary specification (column 1) uses the top-13 panel; the extended 25-pool panel result (column 3) is reported as robustness. CRV1 uses the t-distribution with $G - 1$ degrees of freedom. The extended specification cannot use exact permutation because the number of assignments exceeds feasibility; the top-13 panel admits exact enumeration over $C(13, 8) = 1287$ assignments.




Figure 3. Event-study coefficients on differential China-linked share from the two-way fixed-effects DiD on the top-13 panel, full population. Relative period $k = 0$ is September 2021; $k = -1$ (August 2021) is the omitted reference. Markers left of the dashed line (blue) are pre-September national-enforcement / transition months; markers at and right of it (red) are post-enforcement lags. Error bars are 95% cluster-robust confidence intervals (CRV1, t-distribution with $G - 1 = 12$ degrees of freedom). The months between the 21 May 2021 State Council announcement and the 24 September 2021 national enforcement are not untreated pre-period months; they are interpreted as transition/anticipation months, and the corresponding coefficients are individually indistinguishable from zero. The lags decline monotonically. The headline pooled DiD coefficient of -5.45 pp/pool (Table 3; CRV1 $p = 0.054$, with the wild-bootstrap and exact-permutation values reported as small-sample diagnostics) summarises the pooled post-event movement rather than the individual leads and lags. Source: two-way fixed-effects estimation on the top-13 full-population panel.

Table 4. Pre-post share-change decomposition. Group aggregate shares between the placebo and persistence windows. The closure constraint requires the three group changes to sum to zero.

Group	Placebo (%)	Persistence (%)	Δ (pp)
Exposed (China-linked, 8 pools)	78.47	60.01	−18.46
Substitute (non-China-linked, 5 pools)	3.50	27.02	+23.52
Excluded + Unknown residual	18.03	12.97	−5.06
Sum (closure check)	100.00	100.00	0.00

Note: The three group changes sum to zero by the closure constraint. The substitute-group gain of +23.52 pp exceeds the exposed-group loss of 18.46 pp because the excluded/Unknown residual is also contracted by 5.06 pp. Roughly one-fifth of the substitute-group gain (5.06 of 23.52 pp) is therefore reallocation from the excluded/Unknown residual rather than a direct transfer from China-linked operators; the clean, residual-free magnitude of the displacement is the −18.46 pp fall in China-linked aggregate share, which is treated as the primary quantity. Part of the residual contraction is the Unknown-tag category (1.4 per cent of blocks), so a component of the 5.06 pp may reflect attribution coverage rather than a real share movement.

Closed-Share Inference and the SUTVA Assumption

The Stable Unit Treatment Value Assumption (SUTVA) does not hold by construction in the panel, because the outcome is each pool's share of the network and the network shares sum to one hundred percent at every t . Any share lost by China-linked pools is mechanically reallocated to non-China-linked pools and to the excluded/Unknown residual, so the cross-cluster independence required for classical CRV1 inference is violated. We address this in three ways. First, we report the headline DiD coefficient as a within-cell linear projection of the compositional shift rather than as a classical average treatment effect on the treated. Second, we report inference by three procedures: CRV1 with $t(G - 1)$ critical values, wild cluster bootstrap with $B = 999$ Rademacher cluster weights, and exact treated-label permutation enumerating all $C(13, 8) = 1287$ assignments. Third, we report a pre-post share-change decomposition (Table 4) that exposes the closed-share structure directly. The Supplementary Materials set out a simplifying spillover model under which the DiD coefficient is amplified relative to an underlying per-pool operator-specific displacement by a factor of approximately $(1 + \varphi_C \cdot |T|/|C|)$, with $\varphi_C \approx 0.16$ on the full-population panel (defined in the Appendix) and an implied amplification factor of approximately 1.26. This model assumes proportional reallocation of displaced share and its parameter is not separately identified from the block-attribution data; under it the headline $\beta = -5.45$ pp/pool would correspond to a per-pool operator-specific displacement of approximately -4.32 pp/pool. That figure is reported as an illustrative, model-dependent decomposition rather than an estimated quantity, and no conclusion in the paper depends on it. Estimated directly from the panel without the spillover model, the mean per-pool change between the placebo and persistence windows is -2.31 pp for the China-linked pools (95% bootstrap confidence interval $[-5.28, +0.91]$) and $+4.71$ pp for the non-China-linked pools ($[+1.48, +10.52]$); these per-pool means reproduce the aggregate windows of Table 4 ($8 \times (-2.31) \approx -18.5$ pp and $5 \times (+4.71) \approx +23.5$ pp). The model-free China-linked per-pool change is smaller in magnitude than the model-dependent -4.32 figure and its interval includes zero, whereas the non-China-linked per-pool gain is robustly positive; the displacement is therefore more firmly evidenced on the absorption side—the substitute-group gain, chiefly Foundry USA—than on the loss side, which locates the statistical strength without altering the aggregate account. Under the stated projection and spillover framework, the coefficient carries a zero null benchmark; inference is reported using CRV1 with $t(G - 1)$ critical values as a reference, and the wild cluster bootstrap and exact treated-label permutation procedures as small-cluster diagnostics, rather than as a claim that the closed-share estimator delivers a classically correctly sized test of an average treatment effect on the treated. Two further checks, computed on the deposited block data, bound these diagnostics. First, a placebo-in-time calibration of the wild cluster

bootstrap—assigning a false post-period start to each interior month of the clean pre-period and re-running the bootstrap—produced no rejections at the five percent level across nine placebo dates (every placebo p in $[0.26, 0.91]$), so there is no evidence that the bootstrap over-rejects in this thirteen-cluster, eight-to-five design, and the small headline bootstrap p is not an artefact of a liberal small-cluster procedure. Second, and more consequentially, the exact treated-label permutation treats every $C(13, 8) = 1287$ relabelling as equally likely, which is not innocuous here because the eight China-linked pools are the larger, higher-baseline-share operators (seven of the eight fall in the upper two terciles of placebo-window share). A design-based permutation that conditions on baseline share—holding the number of treated pools fixed within placebo-share terciles and permuting only within strata—raises the two-sided p -value from 0.023 to 0.100. The differential movement is thus distinguishable from zero under the wild bootstrap but not distinguishable from baseline-share-matched relabellings: the wild bootstrap addresses within-cluster dependence and small-cluster size, whereas the conditional permutation addresses the confound that the exposed group was already the concentrated one, and the latter does not reach conventional significance. This is the direct reason the coefficient is read as a descriptive compositional projection rather than an identified causal effect. As a compositional-data robustness check addressing the same concern, the Supplementary Materials re-express the three-part group composition (China-linked, non-China-linked, residual) in log-ratio coordinates, in which the unit-sum constraint does not bind: every part is strictly positive in every month, so the transform requires no zero-replacement. The non-China-to-China log-ratio rises monotonically across the four windows, by +2.31 units—a tenfold change in the ratio—between the placebo and persistence windows, and an interrupted regression of the monthly log-ratio on a post-enforcement indicator gives a post-shift of +1.72 units in the same direction. With only twenty-two monthly observations, a Newey–West t -statistic is not a reliable small-sample test, so the sign, magnitude, and timing of the shift are used rather than its nominal significance. The transform’s purpose is narrow: because every part of the composition is strictly positive, it maps the shares onto unconstrained coordinates in which no single share is mechanically bounded, providing reassurance that the displacement is not merely an artefact of bounded-share support or of conducting inference on a unit-summed outcome. It does not, and cannot, establish that the movement is anything other than a compositional reallocation—the displacement is a compositional reallocation by construction—and no such claim is made; what it shows is that the same direction and timing appear once the unit-sum support constraint is removed.

5.3. Pre-Trend Robustness Across Alternative Windows

The headline pre-trend test uses the four months January through April 2021 and is reported alongside three alternative windows for robustness. Table 5 reports the OLS slope of monthly China-linked share on a linear trend under each window. Under all four windows, the null of zero pre-trend is not rejected at conventional significance levels. The four-month window is the most adversarial, and we read it as such: it is the only window over which the slope is positive, and its point estimate (+1.161 pp per month, Table 5) is economically non-trivial. The failure to reject a zero pre-trend in this window reflects the two residual degrees of freedom rather than evidence of a flat pre-trend; with a positive pre-treatment point estimate of this size the parallel-trends assumption cannot be affirmed, and the post-event movement is interpreted accordingly as a conditional, descriptive result rather than as an identified causal effect. A positive pre-trend that reverses after the shock is also consistent with mean reversion or with anticipation following the 21 May 2021 announcement, neither of which the four-month window can exclude.

Table 5. Pre-trend slope robustness. OLS regression of monthly China-linked aggregate share (per cent) on a linear trend over each pre-treatment window.

Pre-Trend Window	L	Slope (pp/mo)	OLS SE	t	p (2-Sided)
January–April 2021 (headline)	4	+1.161	0.565	2.054	0.176
November 2020–April 2021	6	−0.089	0.519	−0.172	0.872
May 2020–April 2021	12	−0.224	0.138	−1.619	0.136
May 2020–May 2021 (announcement-month sensitivity)	13	−0.198	0.118	−1.678	0.121

Note: L is the number of monthly observations in the window. *p*-values are two-sided under the *t*-distribution with *L* − 2 degrees of freedom. The null of zero pre-trend is not rejected at the conventional five-per-cent level in any window. The fourth row (May 2020–May 2021) includes the announcement month and is therefore reported only as a boundary sensitivity, not as a pure pre-treatment trend test.

5.4. Classification Sensitivity

The classification rule produces an analyst-coded partition of pool operators into exposed (China-linked) and substitute (non-China-linked) groups, with Binance Pool excluded for corporate-domicile ambiguity. Two pools admit explicit alternative classifications: Binance Pool and BTC.com (a transition case whose pool business was transferred from BitDeer to BIT Mining Limited, formerly 500.com, with the first closing announced on 15 April 2021, before the announcement date). Table 6 reports the difference-in-differences coefficient under each alternative classification, including the Foundry-USA-excluded specification that tests whether the headline displacement is driven by a single operator.

Table 6. Classification sensitivity. Two-way fixed-effects DiD with post-period October 2021 onward. Outcome: monthly per-pool share of attributed blocks (per cent).

Classification Variant	β (pp/pool)	CRV1 SE	CRV1 p	G
Baseline (Binance excluded)	−5.45	2.55	0.054	13
Binance Pool as exposed (China-linked)	−5.08	2.50	0.063	14
Binance Pool as substitute (non-China)	−5.08	2.25	0.041	14
BTC.com excluded (transition case)	−5.16	2.64	0.076	12
Foundry USA excluded (12-pool panel)	−3.20	1.25	0.027	12
Extended 25-pool panel	−3.23	1.41	0.031	25

Note: The most informative robustness check is the Foundry USA excluded specification: removing Foundry USA from the control group yields a per-pool coefficient of −3.20 pp on the remaining 12-pool panel (CRV1 *p* = 0.027). The headline displacement is not driven by Foundry USA alone; the China-linked group's share also falls relative to the four other named non-China-linked operators. The headline finding does not depend on the Binance classification or on the BTC.com / BIT Mining transition. Across the six variants the coefficient retains its negative sign and ranges from −3.20 to −5.45 pp per pool; CRV1 significance at five per cent is attained in three (*p* between 0.027 and 0.076).

5.5. Aggregate Network-Activity Trajectory and the Protocol-Invariance Test

H2 (aggregate activity recovery) is tested by the trajectory of total network mining activity, measured by the directly observed Bitcoin network hashrate rather than by modelled electricity consumption, because the hypothesis concerns aggregate proof-of-work activity and not energy use or carbon intensity. The daily network hashrate [23] falls from a monthly mean of 161.2 EH/s in May 2021 to a monthly-mean trough of 100.3 EH/s in July 2021—with a daily trough near 58 EH/s on 27 June 2021—then recovers to a monthly mean of 173.2 EH/s in December 2021, at or above its May 2021 pre-ban monthly mean of 161.2 EH/s, and 197.8 EH/s by February 2022. The intervening monthly means trace a continuous recovery: 120.1 EH/s in June, 120.7 in August, 136.6 in September, 149.1 in October, and 160.9 EH/s in November 2021. Protocol invariance is a maintained institutional fact, observed directly: no element of the Bitcoin protocol tuple was amended in the study window. The recovery of aggregate hashing power to its pre-ban level within approximately six months of the trough, while the operator set turned over, is consistent with H2 and with capital relocation rather than capital destruction. The Cambridge electricity-consumption series [24] is retained only as a supplementary cross-check: it shows the same broad mid-2021 decline and subsequent recovery pattern, but it is not used as the primary H2 measure. It is reported in the Supplementary Materials.

5.6. Concentration Measures and Polycentric Absorption

The post-ban distribution of share gains across non-China-linked operators provides the test of H3. Of the 23.5 percentage-point increase in non-China-linked share between the placebo and persistence periods, approximately 15.7 percentage points (approximately 67 percent of the gain) accrued to Foundry USA. Foundry USA's monthly share rose from a placebo-window mean of approximately 0.56% to a persistence-window mean of approximately 16.25%, peaking at 17.62% in February 2022. On the headline market-share concentration measure, the post-ban distribution is not less concentrated than the pre-ban distribution; one operator absorbs a quantitatively material fraction of the displacement. On the within-non-China-subset normalised HHI measure, however, the post-ban distribution is more rather than less internally distributed than the pre-shock baseline.

We report two HHI measures, both plotted in Figure 4. The first is the absolute non-China-linked HHI computed from full-network percentage shares; this rises mechanically with the size of the non-China-linked subset. The second is the within-subset normalised HHI, which re-expresses non-China-linked pools' shares as percentages of the non-China-linked subset total; this is invariant to subset size and is the appropriate measure of concentration internal to the subset, with the majorisation-theoretic justification set out in Marshall, Olkin, and Arnold [25] (Ch. 1).




Figure 4. Two concentration measures, May 2020–February 2022. The solid (red, left axis) series is the full-market Herfindahl–Hirschman Index computed over all attributed pools including the unidentified residual; the dashed (purple, right axis) series is the within-non-China-linked normalised HHI computed over the five non-China-linked operators of the top-13 panel. The dashed vertical line marks the 24 September 2021 enforcement event; shaded backgrounds identify the four analysis windows. The full-market index rises after the prohibition (the attributed pool-operator market becomes more concentrated overall as a large fraction of the displaced share is captured by a single substitute operator), while the within-non-China normalised index falls (the non-China-linked subset becomes more internally dispersed as additional operators enter). The two movements are not in tension: they measure concentration at different levels of aggregation. Source: monthly block-attribution panel, full population. The normalised series is defined on the top-13 non-China-linked subset, as described in the text.

The absolute non-China-linked HHI rises from 12.3 in the placebo window to 12.4 in the pre-trend window, to 77.9 in the immediate-treatment window, to 306.0 in the persistence window. The within-subset normalised HHI, computed over the five non-China-

linked operators of the top-13 panel, falls from 10,000 in the placebo window to 4354 in the pre-trend, to 3595 in the immediate-treatment window, and stands at 4191 in the persistence window. The placebo-window normalised value of 10,000 is mechanical and definitional: in the placebo window only one of the five top-13 non-China-linked operators (Braiiins Pool) produced any blocks, so the within-subset share of that single operator was one hundred percent and the normalised index attains its theoretical maximum. Computed instead over the full population of non-China-linked operators (including pools outside the top-13 panel), the placebo-window normalised HHI is approximately 7900 rather than 10,000; the two definitions converge as the subset widens. The persistence-window normalised value of 4191 corresponds to a subset much larger in aggregate (27.02% of network share) distributed across a wider set of identifiable named operators. It should not be read as evidence of broad atomistic dispersion: a normalised HHI of 10,000 in a one-operator subset can fall mechanically as additional substitute operators enter, irrespective of any change in the concentration of underlying mining control.

The Foundry case is the central corporate-intermediary anchor; its monthly trajectory is plotted in Figure 5. Foundry Digital LLC. was founded in 2019, headquartered in Rochester, New York, as a wholly-owned subsidiary of Digital Currency Group (DCG), whose own headquarters is at Stamford, Connecticut. By its own contemporaneous public statements, the firm channelled in excess of four hundred million dollars of mining hardware into North America during the displacement window. The DCG parent's existing capital structure permitted Foundry to act as the corporate intermediary that the post-ban configuration required: a US-incorporated entity with established compliance practice, capital backing through DCG's wider portfolio, and a customer-acquisition story expressly framed as the regulatory-jurisdiction substitute for the displaced Chinese pool operators.




Figure 5. Foundry USA monthly share of attributed blocks, May 2020–February 2022. The dotted vertical line marks the 21 May 2021 State Council announcement; the dashed vertical line marks the 24 September 2021 enforcement event. Foundry USA's share rises from a placebo-window mean of zero (the pool produced no blocks before November 2020) to a persistence-window mean of approximately 16 percent, peaking at 17.62 percent in February 2022. Shaded backgrounds identify the four analysis windows. Source: monthly block-attribution panel, full population.

5.7. The Institutional-Substitution Test

Hypothesis 4 (identifiable corporate intermediaries) is satisfied if the displacement is absorbed by named, publicly-disclosed corporate entities whose domicile, capital structure, and operational infrastructure can be documented from primary sources rather than by

unidentifiable atomistic participants. The data are consistent with this prediction. Foundry USA, Braiins, MARA Pool, Luxor, and SBI Crypto are each publicly registered legal persons with disclosed parent companies (Digital Currency Group; Braiins Systems; Marathon Digital Holdings; Luxor Technologies; SBI Holdings). Each has a corporate domicile, a registered office, public disclosure obligations, and a documented capital history. Among the identified non-China-linked operators absorbing material share in the analytical panel, the absorbers are named legal persons with documentable domicile, capital structure, and operational infrastructure. The residual excluded/unknown category is not used to support H4.

6. Discussion

Taken together, H1–H4 identify a coherent pattern: at the pool-operator layer, the prohibition produced operator substitution through identifiable corporate intermediaries rather than dispersion across anonymous participants. The displacement of China-linked operators (H1) and the recovery of aggregate network activity alongside turnover in the operator set (H2) are jointly informative. They indicate that what moved was the corporate-domicile composition of the operator layer, not the network's computational capacity: the aggregate-activity trajectory is consistent with capital relocating along an existing institutional path rather than being destroyed, and the substitute operators were pre-positioned before the shock (Digital Currency Group was founded in 2015; the Foundry USA Pool launched in 2020), so the absorption was not *de novo* entry by previously absent participants.

The form of that absorption (H3) is the analytically informative part, and it gives a mixed reading that resists a single label. At the level of headline shares, a large fraction of the post-ban gain is captured by one named operator, which is less consistent with broad horizontal redistribution than with identifiable operator substitution; within the non-China-linked subset, however, concentration falls as additional operators enter. The two facts are not in tension once the measures are distinguished (Section 5, and Figure 4): the attributed pool-operator distribution became more concentrated overall while the substitute subset became internally more dispersed. The same diagnostic mixed reading appears in related work on decentralised cryptocurrency-platform governance and on the residual role of jurisdictional legal infrastructure in nominally polycentric blockchain orders [4,5]. The identifiability of the absorbing operators (H4) is what makes the substitution interpretation supportable: every operator absorbing material displaced share is a named legal person with documentable domicile, capital structure, and infrastructure. The atomistic-mining model [15,26–28] is not contradicted in its formal results, which follow correctly from its premises; rather, its premise of anonymous, jurisdictionless, atomistic participation does not describe the pool-operator layer of the industry over the period studied. Hinzen, John, and Saleh [29] model a distinct structural constraint of the same mechanism: the requirement to reach consensus under network delay generates negative network effects that bound Bitcoin's adoption in equilibrium. That result concerns the protocol-level economics of consensus, whereas the present claim concerns the corporate composition of the operators who carry it out; the two are complementary descriptions of one system rather than competing ones. As set out in Section 6.1, this is a statement about the operator layer and not about the geographic distribution of underlying hashrate, which the block-attribution data do not measure.

6.1. What Pool Attribution Can and Cannot Show

The empirical object of this paper is the pool-attributed block share: the fraction of blocks in a window whose coinbase output is attributed, by the open source labelling rules of mempool.space, to a named mining-pool brand. This quantity is not the same as the share of physical hashrate owned or operated within a jurisdiction, and the distinction is central to the interpretation of every result above. Three limits follow.

First, pool attribution measures the operator that assembled and broadcast a block, not the location, ownership, or jurisdiction of the miners who supplied the underlying hashrate. A pool brand domiciled in one jurisdiction routinely aggregates hashrate from miners in many jurisdictions, and individual miners can redirect their hashrate to a different pool at low cost and at short notice. A fall in a pool brand's attributed share is therefore consistent with two distinct underlying movements that the block data cannot separate: physical relocation or shutdown of the miners formerly directing hashrate to that brand, and re-delegation of unmoved hashrate from one pool brand to another. The displacement documented here is a movement in the corporate-domicile composition of the pool-operator layer; it is not a direct measurement of where mining hardware is physically located. That a pool's attributed share reflects miners' own allocation of hashpower across pools—under the fee and risk-sharing terms each pool offers—rather than the ownership of that hashpower, is developed formally in the recent mining-pool literature discussed in Section 6.2.

Second, block-attribution counts are statistical estimates of mining power, not direct observations of it. Because block discovery is a stochastic process, the share of blocks attributed to an operator over a finite window is an estimate of that operator's contribution to the network hash function, subject to sampling variation that narrows only as the observation window lengthens; Milad et al. [30] formalise the statistical confidence attainable in mining-power estimates derived from block observations on proof-of-work chains. The window-level shares reported in this paper should be read with that estimation character in mind, which is one reason the analysis aggregates over multi-month windows rather than interpreting single-block or single-day attribution.

Third, the pool-attribution layer is one stratum of a multi-layer system, and decentralisation at one layer does not imply decentralisation at another. Ovezik, Karakostas and Kiayias [31] set out a stratified framework in which a blockchain's decentralisation must be assessed separately across distinct layers—among them hardware, software, network, consensus, economics, client API, governance, and geography—and show that these layers need not move together. The present paper observes the pool-operator stratum (closest to the consensus and geography layers as those authors define them) and is silent on the others: it does not measure the distribution of mining hardware ownership, the distribution of full nodes, client or developer concentration, or protocol-governance authority. Claims in this paper are confined to the pool-attribution layer accordingly, and the broader conclusions about permissionless consensus in Section 7 are scoped to that layer rather than to the network as a whole.

Fourth, block attribution is preferred here over IP-geolocation methods for measuring the operator layer, and the reason bears directly on what the measure can support. Methods that assign a miner's location from the IP address of the pool it connects to are unreliable, because a miner can join a pool irrespective of its own geographic location; Sai and Vranken [6], in their review of geographic-attribution methods, describe the attribution of mining activity to the jurisdiction of a pool's servers as provably inaccurate on that ground. Block attribution does not require an assumption about where the underlying miners are physically located: it records, reproducibly from public block data, the operator that assembled and broadcast each block. This is precisely why the measure speaks to operator-domicile composition and not to the geographic distribution of physical hashrate, and it is the basis on which the corporate-domicile classification, rather than a server-location classification, is used throughout.

6.2. Comparison with Concurrent Studies

The 2021 prohibition has been studied from several disciplinary angles, and it is useful to position the present pool-attribution evidence relative to that work, both the concurrent empirical studies of the same event and the theoretical literature on mining-market structure. We first locate this paper's findings against the concurrent empirical record, and then set out three economic mechanisms that have been advanced as explanations for how displaced mining activity is absorbed, together with the testable prediction each implies and what the present evidence does and does not say about it.

Concurrently with the prohibition, an energy-and-emissions literature documented the same displacement through the lens of mining location and electricity sourcing. de Vries, Gellersdörfer, Klaaßen and Stoll [32] find that, following the crackdown, miners migrated primarily to Kazakhstan and the United States, that the renewable share of the electricity mix powering the network fell from an average of 41.6 per cent in 2020 to 25.1 per cent in August 2021, and that the carbon intensity of mining may consequently have risen by approximately 17 percent; their location estimates draw, in part, on the same four large pools (including Foundry USA) that feature in the present panel. Xiao, Cui, Xiang, Liu, and Zhang [33], using provincial-level data on Chinese mining through 2021, document the seasonal migration of mining pools across Chinese provinces and the regional power-demand shocks it generated, quantifying cumulative Chinese mining emissions over 2017–2021. The earlier baseline study of Stoll, Klaaßen, and Gellersdörfer [17] established the location-and-electricity methodology those later studies extend. These studies and the present paper describe the same event from complementary measurement positions: the energy literature tracks where physical mining capacity drew its electricity and what that implied for emissions, while the present paper tracks the corporate-domicile composition of the pool operators to which blocks were attributed. The two positions agree on the central qualitative fact—that displacement out of China was substantial and was absorbed disproportionately by operators and locations in the United States—and they are mutually corroborating rather than competing, because they measure different quantities. The present paper's contribution within this record is the use of the full block population to characterise the operator layer specifically, and the formal treatment of the closed-share inference problem that the operator-share outcome creates. Beyond the change in physical location, both records describe a change in the industrial form of the activity. Xiao, Cui, Xiang, Liu, and Zhang [33] document that mining within mainland China progressively ceased after June 2021, with operations relocating abroad; the present paper identifies the corporate counterpart of that relocation at the attribution layer—the displaced share was absorbed disproportionately by institutionally organised operators, including listed firms, domiciled chiefly in the United States, rather than reconstituted under the domestically organised pools that had previously dominated. The shock thus coincided with a change in the corporate-domicile composition of the pool-operator layer, which the domicile-coded panel measures directly.

The mechanism literature offers several non-exclusive accounts of how displaced mining activity is absorbed. We state three, with their predictions, as rival *ex ante* explanations; the present design does not discriminate among them, and we treat them as complementary *ex post*.

The first is a capital-mobility and relocation mechanism. Prat and Walter [34] model the mining industry as a free-entry market in which investment in mining hardware is effectively irreversible—machines have little resale value outside mining—so that aggregate computing power tracks the expected return to mining and recovers as displaced capital is redeployed. The prediction is that aggregate hashrate recovers after a regulatory shock because physical mining capital relocates or is redeployed rather than destroyed.

The present evidence is consistent with this prediction at the aggregate level: aggregate network activity falls sharply after the prohibition and recovers toward its pre-ban level within the study window (the H2 result), even as the pool-attribution composition shifts.

The second is an oligopoly and economies-of-scale mechanism. Arnosti and Weinberg [35] show that, under asymmetric costs or economies of scale, mining equilibrium is concentrated rather than atomistic: with economies of scale every active participant holds a market share bounded below, so that the number of viable operators is small. The economics of the mining reward points the same way: Easley, O'Hara and Basu [36] model the sector's transition from a regime governed by the block subsidy toward one in which transaction fees and the strategic conduct of large operators increasingly govern participation, so that the operators best placed to absorb displaced activity are those with the scale to remain profitable as the reward structure tightens. This concentration reading must be set against the contrasting empirical finding of Cong, He, and Li [26], who document that over the long run no single pool persistently dominates Bitcoin mining, because risk-sharing draws miners to reallocate across pools. That result is not in tension with the present one: it concerns the long-run absence of a permanent single leader, whereas the present paper measures a specific post-shock reallocation of attributed share over a bounded window; moreover, the decentralising mechanism it identifies operates through miners switching pools—a movement at the miner layer that the present paper explicitly separates from the operator-attribution layer it measures. The prediction is that displaced activity is absorbed disproportionately by operators possessing scale, financing, customer-onboarding capacity, and infrastructure, rather than dispersed across many small entrants. The present evidence is consistent with this prediction: a single large, well-capitalised operator (Foundry USA, backed by Digital Currency Group) absorbed approximately 67 per cent of the gain to the non-China-linked subset.

The third is a mining-technology and heterogeneous-efficiency mechanism. Capponi, Ólafsson, and Alsabab [37] model mining as a rent-seeking game with capacity constraints and find that improvements in hardware efficiency do not necessarily allow larger miners to increase their advantage; under their calibration, such improvements can instead allow smaller miners to expand and new miners to enter. The prediction is that hardware-efficiency and entry conditions shape whether displaced activity concentrates or disperses, in a direction that depends on the configuration of capacity constraints and is not signed a priori. The present design does not test this mechanism: the relative efficiency and entry conditions of the absorbing operators are not identifiable from pool-attribution data, so the evidence here neither supports nor contradicts it, and it is relevant only as a mechanism that could shape which operators are able to absorb displaced miners.

The present design does not adjudicate among these mechanisms as mutually exclusive causes. It documents the realised destination of pool-attributed displacement after the regulatory shock. The evidence is most consistent with operator substitution under conditions in which mobile mining capital, scale economies, financing capacity, and technological heterogeneity jointly shape the new pool-layer equilibrium, and it is not designed to identify the relative contribution of any one of them.

This literature also clarifies the measurement distinction on which the present paper depends. Studies of mining geography and electricity sourcing estimate where physical mining activity occurs; studies of block attribution estimate which pool brand assembled and broadcast blocks; and studies of corporate structure identify the legal persons operating the pool layer. The present article contributes to the third of these objects while using the first two as external context. Its claim is therefore narrower than a geography-of-hashrate claim and more institutional than an electricity-sourcing claim: it measures how a

jurisdiction-specific prohibition changed the corporate-domicile composition of the pool-attributed operator layer.

A recent strand of mining-pool research, although not concerned with the 2021 prohibition, bears on the interpretation of the operator-attribution measure used here, because it concerns how pool shares are formed and what they do and do not reveal about the miners behind them. Goffard, Albrecher, and Fouque [38] model how an individual miner optimally divides hashpower across competing Pay-Per-Share pools, trading the risk transferred to the pool manager against the management fee; in their framework, a pool's attributed share is an equilibrium outcome of miners' allocation decisions under those incentives rather than a measure of who owns the underlying hashpower. Bar-On, Kormargodski, and Weinstein [39] extend the heterogeneous-cost analysis of proof-of-work to settings with external (proof-of-useful-work) rewards, characterising how regional cost differences, hardware access, and outside computational rewards shape the equilibrium distribution of mining effort—a cost-heterogeneity mechanism complementary to the three set out above. The architecture of pooling is itself a design variable rather than a fixed feature: Sakurai and Shudo [40] document that most operating pools are centrally administered, with the bulk of network hashrate concentrated in them, and propose a distributed-pool construction—across a main-chain contract, a storage chain, and a child chain—intended to deliver stable, fair, budget-balanced rewards without a central operator; Chatterjee, Majumder, and Mazumdar [41], in a distributed useful-proof-of-work design, likewise treat the pool as a coordination layer in which individual miners collaborate on a shared task and each contribution is separately verifiable. Two implications follow for the present design. First, the operator concentration documented here is a property of the prevailing centrally administered pool architecture rather than an immutable feature of proof-of-work, since distributed-pool designs may alter how block attribution maps onto operator control. Second, because a pool's share reflects miners' fee-and-risk-driven allocation choices, the operator-attribution measure speaks to the corporate-domicile composition of the intermediary layer and not to the ownership or location of the hashpower that miners direct through it, reinforcing the distinction drawn in Section 6.1. These four are recent working papers rather than peer-reviewed articles; they are cited only for the specific mechanism points noted, while the interpretive weight rests on the peer-reviewed mining-economics literature cited above.

6.3. Economic, Social, and Financial Implications

The pool-attribution evidence carries several implications that do not depend on resolving the mechanism question. For market structure and competition, the finding that a unilateral prohibition against the dominant jurisdiction was followed by a large share gain captured by a single substitute operator, rather than by broad dispersion, indicates that the pool-operator layer of a permissionless network can exhibit the concentration dynamics of a scale-driven industry even when the underlying protocol imposes no barrier to entry; the relevant competitive constraint operates at the operator layer, not the protocol layer. For regulation and financial supervision, the displacement toward named, publicly disclosed corporate operators with documented domicile (chiefly in the United States and the European Union) means that the activity became more, not less, legible to financial regulators after the shock: the substitute operators are legal persons with registered offices, disclosure obligations, and identifiable capital structures, which has direct consequences for how mining activity can be monitored, taxed, and supervised in the jurisdictions that absorbed it. For systemic resilience, the survival of the protocol through a partial but material turnover of its largest jurisdictional concentration of operators is evidence that the network is resilient to a jurisdiction-specific shock; but the large share gain captured

by one substitute operator also relocates, rather than removes, the concentration-related risks (operational dependence, regulatory exposure, and potential single points of failure) that motivate concern about mining centralisation. These implications are confined to the pool-attribution layer and should not be read as claims about the distribution of mining hardware or of network control more broadly, for the reasons set out in Section 6.1.

6.4. *Scope Conditions and Limitations*

The empirical claims of this paper are limited in five well-defined ways. First, the dependent variable is pool-attribution of blocks, not physical hashrate; the relationship between the two is mediated by the pool-attribution scheme of mempool.space, and the displacement documented here is a movement in operator-domicile composition rather than a direct measurement of where mining hardware is physically located (Section 6.1). Second, the persistence window is four months, and we cannot rule out reversal at longer horizons; subsequent industry data indicating a partial return of mining activity to China after the study window is consistent with the displacement being durable on the institutional-adjustment horizon observed here without being permanent. Third, the classification rule combines four primary-source dimensions and is therefore an analyst-coded construct; the rule is reported in full and the deposited ledger permits replication under alternative rules, and the classification-sensitivity analysis (Table 6) shows the headline result is robust to the two genuinely ambiguous cases. Fourth, the two-way fixed-effects estimator on shares is not a classical average treatment effect on the treated; the headline coefficient is a linear projection of the within-cell mean shift under the closed-share constraint (Section 6.1 and the closed-share inference discussion above). Fifth, the strong polycentric reading is not supported on the headline-share measure but the weak within-subset reading is, so the paper's conclusion about polycentric governance is correspondingly qualified. The environmental-policy interpretation of the displacement—in particular, whether the relocation of mining activity altered aggregate carbon emissions through a change in the energy mix of the substitute jurisdictions—is outside the scope of this paper, which measures the corporate-domicile composition of the operator layer and not energy sourcing; the concurrent emissions and energy-sourcing studies discussed in Section 6.2 address that question directly.

6.5. *External Validity*

This study identifies one historical episode, and its external validity is correspondingly bounded. Four questions delimit how far the finding should be generalised. First, whether similar patterns would follow other regulatory shocks: the mechanism documented here—re-concentration of operator-attributed share at large, well-capitalised, jurisdictionally-compliant intermediaries—depends on the prohibition being jurisdiction-specific while the operator market exhibits scale economies, so the pattern should be expected to recur for jurisdiction-specific shocks against concentrated operator markets, but not necessarily for shocks that are global in reach or that fall on markets without scale economies. Second, whether the finding is specific to Bitcoin: the result concerns the mining-pool layer of a proof-of-work network with a portable hardware base and an open protocol stratum, features Bitcoin shares with other large proof-of-work chains but not with networks whose validation rights are allocated differently. Third, whether proof-of-stake systems would behave differently: in a proof-of-stake network the analogue of the pool operator is the staking pool or validator, and the analogue of portable mining hardware is transferable stake, so the relocation channel that operates through physical capital here would operate, if at all, through a different and less location-bound mechanism; the present evidence does not speak to that case. Fourth, whether the result generalises beyond mining pools to

other layers of the same network: as set out in Sections 6.1 and 7, it does not, because the layers need not move together and only the operator layer is observed here. The finding is therefore best read as evidence about how the operator layer of a concentrated proof-of-work market responds to a jurisdiction-specific prohibition, not as a general law of permissionless-consensus networks.

7. Conclusions and Policy Implications

The May–September 2021 cryptocurrency-mining prohibition by the People’s Republic of China is followed by a large documented jurisdictional shift in pool-attributed block share in a permissionless-consensus network. The pool-attributed compositional shift is sharp, brand-specific, persistent, and concentrated at the level of headline aggregate shares. Between the placebo and persistence windows, exposed (China-linked) mining-pool operators’ aggregate share fell by approximately 18.5 percentage points; substitute (non-China-linked) operators’ aggregate share rose by approximately 23.5 percentage points; Foundry USA alone absorbed approximately 15.7 percentage points, approximately 67 percent of the substitute-group share gain. The Bitcoin protocol was unchanged across the prohibition. The compositional shift documented here is consistent with operator substitution through identifiable corporate intermediaries following the regulatory shock; we do not claim a classical causal-identification result or a conventionally significant treatment effect, and the magnitude is properly interpreted as a descriptive compositional displacement rather than as an average treatment effect on the treated.

The cryptocurrency-market-structure implications are calibrated. At the pool-attribution layer specifically, the atomistic-participation premise of jurisdictionless readings of permissionless consensus [15,26–28] does not describe the observed pool-attribution pattern well during this window: the post-shock equilibrium at that layer is not the broad, atomistic re-distribution those frameworks describe, but rather a concentration of pool-attributed share at a small number of named, publicly disclosed corporate operators in substitute jurisdictions—chiefly Foundry USA. The evidence preserves a weaker resilience reading of polycentric governance: the protocol survived a partial but material turnover of its largest jurisdictional concentration of operators because operator substitution occurred through legally identifiable intermediaries in substitute jurisdictions. The headline pool-attribution measure provides limited support for the strong polycentric reading [7,42]; the weak polycentric reading is consistent with the within-subset trajectory, and is presented as an interpretation consistent with the evidence rather than as a demonstrated property of network governance.

These conclusions are confined to the mining-pool attribution layer and do not extend to the network as a whole. As Ovezik, Karakostas, and Kiayias [31] argue, decentralisation must be assessed separately across distinct strata—hardware, software, network, consensus, economics, client API, governance, and geography—which need not move together. The present study observes only the pool-operator stratum. It does not measure the geographic or ownership distribution of mining hardware, the distribution of full nodes, client or developer concentration, or protocol-governance authority, and it therefore makes no claim about decentralisation at those layers. A finding of operator-layer re-concentration is consistent with the underlying hashrate being more, less, or equally geographically dispersed than before the prohibition; the block-attribution data cannot distinguish these, for the reasons set out in Section 6.1. The claim of this paper is the narrower and supportable one: that at the pool-operator layer, a unilateral prohibition against the dominant jurisdiction was followed by re-concentration at identifiable corporate intermediaries in substitute jurisdictions, rather than by dispersion across anonymous atomistic participants.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/fintech5030063/>. The Supplementary Materials contain the formal model of pool-operator capital under regulatory shock; the estimator's population projection and inference; the persistence analysis; the concentration metrics underlying the polycentric reading; the institutional-substitution material; the supplementary tables and the data-construction protocol, including the full primary-source pool-classification ledger; the protocol-invariance material; the multiple-comparison note; additional concentration measures; the heterogeneity and asymptotic-distribution results; the specification tests; the bounds on the structural parameter; the robustness and inference details; the scope conditions and limitations in full; the aggregate-hashrate and energy-consumption notes; and the supporting descriptive figures.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: Replication data (the full 97,005-block JSONL output from the mempool.space block API for heights 628,350 through 725,354), retrieval scripts, the classification ledger, and analysis scripts that reproduce every numerical result in this paper are deposited at Zenodo (DOI 10.5281/zenodo.20754433) under a CC BY 4.0 licence. The supplementary robustness re-estimations reported in Section 5—the placebo wild-bootstrap size calibration, the baseline-conditional permutation, and the model-free per-pool displacement—are reproduced by an additional deposited script, `robustness_checks.py`. The Bitcoin network-hashrate series used in the aggregate-activity analysis is obtained from Coin Metrics community data and is included in the replication materials as `coinmetrics_hashrate_v1_0_0.csv`.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Bitmain Technologies Holding Company. *Application Proof for Listing on the Stock Exchange of Hong Kong*; Application Proof Filed 26 September 2018; Stock Exchange of Hong Kong: Hong Kong, 2018.
2. BIT Mining Limited. *BIT Mining Limited Announces First Closing of BTC.com Businesses Acquisition*; Exhibit 99.1 to Form 6-K; U.S. Securities and Exchange Commission: Washington, DC, USA, 15 April 2021. Available online: https://www.sec.gov/Archives/edgar/data/1517496/000110465921050757/tm2113103d1_ex99-1.htm (accessed on 18 June 2026).
3. BIT Mining Limited. *BIT Mining's Mining Pool Subsidiary to Exit Mainland China Market*; Exhibit 99.1 to Form 6-K; U.S. Securities and Exchange Commission: Washington, DC, USA, 2021. Available online: https://www.sec.gov/Archives/edgar/data/1517496/000110465921126330/tm2130056d1_ex99-1.htm (accessed on 18 June 2026).
4. Davidson, S. The nature of the decentralised autonomous organisation. *J. Inst. Econ.* **2025**, *21*, e5. [CrossRef]
5. Murtazashvili, I.; Palida, A.; Madison, M.J. The past, present, and future of polycentric legal order: A comparative institutional analysis of lex mercatoria and blockchain. *J. Inst. Econ.* **2026**, *22*, e2. [CrossRef]
6. Sai, A.R.; Vranken, H. Promoting rigor in blockchain energy and environmental footprint research: A systematic literature review. *Blockchain Res. Appl.* **2024**, *5*, 100169. [CrossRef]
7. Alston, E.; Law, W.; Murtazashvili, I.; Weiss, M. Blockchain networks as constitutional and competitive polycentric orders. *J. Inst. Econ.* **2022**, *18*, 707–723. [CrossRef]
8. North, D.C. *Institutions, Institutional Change and Economic Performance*; Cambridge University Press: Cambridge, UK, 1990.
9. Hodgson, G.M. What are institutions? *J. Econ. Issues* **2006**, *40*, 1–25. [CrossRef]
10. Williamson, O.E. Transaction-cost economics: The governance of contractual relations. *J. Law Econ.* **1979**, *22*, 233–261. [CrossRef]
11. Zhang, L.; Wang, Y. *China: National Development and Reform Commission Issues Notice Restricting Cryptocurrency Mining*; Library of Congress, Global Legal Monitor: Washington, DC, USA, 2022. Available online: <https://www.loc.gov/item/global-legal-monitor/> (accessed on 11 May 2026).
12. Becker, G.S. Crime and punishment: An economic approach. *J. Political Econ.* **1968**, *76*, 169–217. [CrossRef] [PubMed]
13. Allen, D.W.E.; Berg, C.; Lane, A.M.; MacDonald, T.; Potts, J. The exchange theory of web3 governance. *Kyklos* **2023**, *76*, 659–675. [CrossRef]
14. Wilson, D.S.; Ostrom, E.; Cox, M.E. Generalizing the core design principles for the efficacy of groups. *J. Econ. Behav. Organ.* **2013**, *90*, S21–S32. [CrossRef]
15. Budish, E. Trust at scale: The economic limits of cryptocurrencies and blockchains. *Q. J. Econ.* **2025**, *140*, 1–62. [CrossRef]

16. Zhang, L. *China: Central Bank Issues New Regulatory Document on Cryptocurrency Trading*; Library of Congress, Global Legal Monitor: Washington, DC, USA, 2021. Available online: <https://www.loc.gov/item/global-legal-monitor/2021-10-13/china-central-bank-issues-new-regulatory-document-on-cryptocurrency-trading/> (accessed on 11 May 2026).
17. Stoll, C.; Klačaßen, L.; Gällersdörfer, U. The carbon footprint of Bitcoin. *Joule* **2019**, *3*, 1647–1661. [\[CrossRef\]](#)
18. Sun, W.; Jin, H.; Jin, F.; Kong, L.; Peng, Y.; Dai, Z. Spatial analysis of global Bitcoin mining. *Sci. Rep.* **2022**, *12*, 10694. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Cameron, A.C.; Gelbach, J.B.; Miller, D.L. Bootstrap-based improvements for inference with clustered errors. *Rev. Econ. Stat.* **2008**, *90*, 414–427. [\[CrossRef\]](#)
20. Cameron, A.C.; Miller, D.L. A practitioner’s guide to cluster-robust inference. *J. Hum. Resour.* **2015**, *50*, 317–372. [\[CrossRef\]](#)
21. MacKinnon, J.G.; Webb, M.D. Wild bootstrap inference for wildly different cluster sizes. *J. Appl. Econom.* **2017**, *32*, 233–254. [\[CrossRef\]](#)
22. Lucey, B.M.; Vigne, S.A.; Yarovaya, L.; Wang, Y. The cryptocurrency uncertainty index. *Financ. Res. Lett.* **2022**, *45*, 102147. [\[CrossRef\]](#)
23. Coin Metrics. *Community Network Data: Bitcoin (BTC) Network Hash Rate*; Coin Metrics, Inc.: New York, NY, USA, 2026. Available online: <https://coinmetrics.io/community-network-data/> (accessed on 18 June 2026).
24. Cambridge Centre for Alternative Finance. *Cambridge Bitcoin Electricity Consumption Index (CBEI): Network Hashrate and Electricity Consumption Data*; Cambridge Judge Business School, University of Cambridge: Cambridge, UK, 2024. Available online: <https://ccaf.io/cbnsi/cbeci> (accessed on 11 May 2026).
25. Marshall, A.W.; Olkin, I.; Arnold, B.C. *Inequalities: Theory of Majorization and Its Applications*, 2nd ed.; Springer: New York, NY, USA, 2011.
26. Cong, L.W.; He, Z.; Li, J. Decentralized mining in centralized pools. *Rev. Financ. Stud.* **2021**, *34*, 1191–1235. [\[CrossRef\]](#)
27. Makarov, I.; Schoar, A. *Blockchain Analysis of the Bitcoin Market*; NBER Working Paper 29396; National Bureau of Economic Research: Cambridge, MA, USA, 2021. [\[CrossRef\]](#)
28. Auer, R.; Frost, J.; Vidal Pastor, J.M. *Miners as Intermediaries: Extractable Value and Market Manipulation in Crypto and DeFi*; BIS Bulletin 58; Bank for International Settlements: Basel, Switzerland, 2022. Available online: <https://www.bis.org/publ/bisbull58.htm> (accessed on 11 May 2026).
29. Hinzen, F.J.; John, K.; Saleh, F. Bitcoin’s limited adoption problem. *J. Financ. Econ.* **2022**, *144*, 347–369. [\[CrossRef\]](#)
30. Milad, M.; Ovezik, C.; Karakostas, D.; Woods, D.W. Statistical confidence in mining power estimates for PoW blockchains. *arXiv* **2024**, arXiv:2403.13736. [\[CrossRef\]](#)
31. Ovezik, C.; Karakostas, D.; Kiayias, A. SoK: A stratified approach to blockchain decentralization. In *Financial Cryptography and Data Security 2024*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2025. [\[CrossRef\]](#)
32. de Vries, A.; Gällersdörfer, U.; Klačaßen, L.; Stoll, C. Revisiting Bitcoin’s carbon footprint. *Joule* **2022**, *6*, 498–502. [\[CrossRef\]](#)
33. Xiao, Z.; Cui, S.; Xiang, L.; Liu, P.J.; Zhang, H. The environmental cost of cryptocurrency: Assessing carbon emissions from bitcoin mining in China. *J. Digit. Econ.* **2023**, *2*, 119–136. [\[CrossRef\]](#)
34. Prat, J.; Walter, B. An equilibrium model of the market for Bitcoin mining. *J. Political Econ.* **2021**, *129*, 2415–2452. [\[CrossRef\]](#)
35. Arnosti, N.; Weinberg, S.M. Bitcoin: A natural oligopoly. *Manag. Sci.* **2022**, *68*, 4755–4771. [\[CrossRef\]](#)
36. Easley, D.; O’Hara, M.; Basu, S. From mining to markets: The evolution of Bitcoin transaction fees. *J. Financ. Econ.* **2019**, *134*, 91–109. [\[CrossRef\]](#)
37. Capponi, A.; Ólafsson, S.; Alsabah, H. Proof-of-work cryptocurrencies: Does mining technology undermine decentralization? *Manag. Sci.* **2023**, *69*, 6455–6481. [\[CrossRef\]](#)
38. Goffard, P.-O.; Albrecher, H.; Fouque, J.-P. Hashpower allocation in Pay-Per-Share blockchain mining pools. *arXiv* **2025**, arXiv:2511.13777. [\[CrossRef\]](#)
39. Bar-On, Y.; Komargodski, I.; Weinstein, O. Proof of work with external utilities. *arXiv* **2025**, arXiv:2505.21685. [\[CrossRef\]](#)
40. Sakurai, A.; Shudo, K. FiberPool: Leveraging multiple blockchains for decentralized pooled mining. *arXiv* **2025**, arXiv:2501.15459. [\[CrossRef\]](#)
41. Chatterjee, D.; Majumder, A.; Mazumdar, S. ScalWork: Useful proof-of-work with distributed pool mining. *arXiv* **2025**, arXiv:2504.14328. [\[CrossRef\]](#)
42. Bodó, B.; de Filippi, P. Trust in context: The impact of regulation on blockchain and DeFi. *Regul. Gov.* **2025**, *19*, 740–760. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.