STAFF
REPORTS

NO. 1205
SEPTEMBER 2026

When Higher Stakes Weaken Security

Pablo D. Azar | Maryam Farboodi

FEDERAL RESERVE BANK *of* NEW YORK

When Higher Stakes Weaken Security

Pablo D. Azar and Maryam Farboodi

Federal Reserve Bank of New York Staff Reports, no. 1205
September 2026
https://doi.org/10.59576/sr.1205

Abstract

A central design goal of settlement systems is that security should not degrade when transaction values are high. Using Ethereum's Proof-of-Work era, we show that Proof-of-Work can fail this test: higher transaction fees—the reward for a successful attack—cause miners to deviate from honest behavior and fork the chain, undermining settlement finality when it is most needed. We exploit a unique feature of Ethereum's Proof-of-Work era that records, for each fork, both the winning block and its displaced competitor's timestamp. We jointly instrument hourly mean log fees, hashrate, and block size using a crypto shock indicator constructed from Ethereum hacks, market crises, and regulatory events, together with one-hour lags of hashrate and block size. Higher fees significantly increase canonical blocks arriving exactly one second later than the siblings they displace, which is consistent with deviation from honest behavior. The converse event showing honest behavior—a canonical block arriving before its siblings—is not driven by high fees. This contrast indicates that fees affect equilibrium behavior and settlement finality in Proof-of-Work blockchains.

JEL classification: G12, G28, D82
Key words: blockchain, payments, settlement

Azar: Federal Reserve Bank of New York (email: pablo.azar@ny.frb.org). Farboodi: Cornell SC Johnson College of Business (email: mf923@cornell.edu). The authors thank Dan Aronoff, Anders Brownworth, Kevin Mei, Michael J. Lee, Michael Maurer, Peter Mueller, and Neha Narula for helpful comments, as well as participants at the MIT Digital Currency Initiative forum and the Midwestern Financial Association Conference for valuable feedback.

1 Introduction

Settlement systems are designed so that security rises with the stakes: when a bank clears a large payment, the legal framework behind it does not weaken, and courts and regulators devote more resources to enforcing high-value contracts. This property—enforcement scaling with the value at stake—is foundational to every centralized settlement system in operation today. Proof-of-Work (PoW) blockchains attempt to replace this legal infrastructure with a purely incentive-based mechanism, substituting the authority of clearinghouses and courts with costly computation. Whether decentralized enforcement preserves the same scaling property is an open empirical question—and a growing theoretical literature (Carlsten et al., 2016; Bahrani et al., 2025; Budish, 2025) suggests that PoW security can degrade as transaction values rise.

This paper provides the first causal evidence that—in a major PoW blockchain—security can degrade with the value at stake. We ask whether exogenous increases in transaction fees cause miners to deviate from honest behavior, displacing valid blocks to capture their fees. Our main specification jointly instruments hourly mean log fees, hashrate, and block size with a CryptoShock indicator—the union of major Ethereum hacks, market crises, and regulatory events—and one-hour lags of hashrate and block size.

We find that higher fees significantly increase the number of canonical blocks that arrive one second later than their siblings. In contrast, higher fees have a near-zero effect on the mirror outcome that would be expected with honest behavior: a canonical block arriving one second earlier than its sibling. This concentration exactly one second later, rather than a symmetric change around zero, is difficult to explain by network congestion.

Our empirical strategy exploits a unique institutional feature of Ethereum, the largest smart-contract blockchain. Most PoW blockchains, including Bitcoin, discard blocks that lose the mining race, leaving no record of competing chains. Ethereum's protocol, by contrast, retained a record of these losing blocks—termed *sibling blocks*—providing precise timestamps for both the winner and the displaced competitor in each fork.[1] This on-chain record of fork resolution spans Ethereum's full seven-year PoW era (2015–2022) and approximately 14 million blocks.

1.1 Our Results in a Nutshell

When the protocol works as intended, the sibling block with the earliest timestamp wins canonical inclusion. We define the timing gap as $\Delta\tau \equiv \tau_{\text{canonical}} - \tau_{\text{sibling}}$. Thus, $\Delta\tau = +1$ means that the canonical block is timestamped exactly one second later than its sibling, while $\Delta\tau = -1$ is the

[1] Ethereum's protocol specification used the term "uncle" for these blocks. We refer to them as siblings throughout to emphasize that they are blocks mined at the same blockchain height and therefore compete for the same slot in the canonical chain.

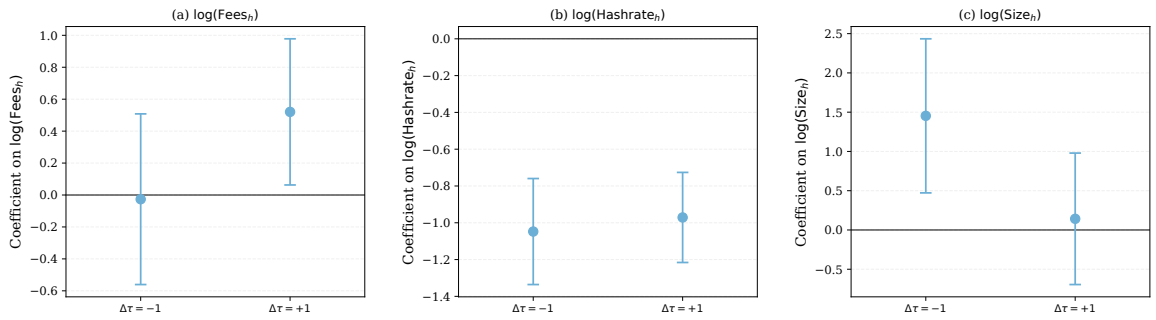

Figure 1: The central finding. Each panel compares the hourly 2SLS estimates at $\Delta\tau = -1$, where the canonical block is timestamped exactly one second earlier than its sibling, and $\Delta\tau = +1$, where it is timestamped exactly one second later. The panels report coefficients on log fees (a), log hashrate (b), and log block size (c), with 95% heteroskedasticity-robust confidence intervals. All three regressors are treated as endogenous and estimated jointly.

mirror outcome in which it is timestamped exactly one second earlier. This narrow comparison helps distinguish strategic re-mining from forces that should move close timing gaps symmetrically.

Figure 1 summarizes our main result. Panel (a) shows that higher fees significantly increase the number of canonical blocks at $\Delta\tau = +1$. The effect at the mirror outcome, $\Delta\tau = -1$, is near zero. The fee response is therefore concentrated on the side consistent with a miner displacing an already observed sibling.

Panels (b) and (c) provide useful mechanism checks. Higher hashrate reduces both exact timing outcomes. At $\Delta\tau = +1$, this is consistent with deterrence: a larger honest mining base extends the accepted chain more quickly, narrowing the window in which a competing fork can succeed. Block size, by contrast, is positive and significant at $\Delta\tau = -1$ but positive and statistically insignificant at $\Delta\tau = +1$. A larger block propagates more slowly and is therefore more likely to win only when it begins with a timing advantage. Conversely, a miner attempting to replace an already observed sibling faces a narrow window in which to assemble and propagate a competing block, making a large replacement block less likely. The absence of a statistically detectable size effect at $\Delta\tau = +1$ is consistent with this prediction.

The full timing-gap analysis sharpens this result. The positive fee response is concentrated at $\Delta\tau = +1$—the minimum positive gap—with no significant positive fee coefficients at larger positive gaps. This location is consistent with a miner observing a high-value sibling and attempting to replace it immediately: successful attempts should appear at the smallest positive gap and favor blocks that can be assembled and propagated quickly.

To our knowledge, this is the first causal evidence that PoW settlement security weakens as the value of what the chain secures rises. This finding confirms empirically what previous literature established theoretically. The observed pattern—higher fees increasing the hourly count at $\Delta\tau = +1$ but not at the mirror outcome $\Delta\tau = -1$, with no significant positive responses at larger timing

gaps—is consistent with strategic re-mining and inconsistent with the leading alternatives such as congestion which would have symmetric effects.

1.2 Implications for Settlement Infrastructure

The institutional significance of this result extends beyond blockchain. Every modern payment and settlement system—from Fedwire to CLS to central counterparties—is designed so that enforcement *strengthens* under stress. PoW inverts this relationship: the infrastructure becomes less secure precisely when security matters most. This is not a bug in a particular implementation; it is a structural feature of the incentive mechanism. As Abadi and Brunnermeier (2022) show, free entry into decentralized chains dissipates rents, eliminating the franchise value that disciplines centralized record-keepers. Without franchise value, miners' behavior is driven primarily by contemporaneous payoffs, which is why a single high-fee block can trigger a deviation.

This inversion has implications for proposals to use blockchain technology for securities settlement, cross-border payments, and central bank digital currencies. Our evidence pertains specifically to PoW and to Ethereum's PoW era (2015–2022). Notably, Ethereum itself abandoned PoW in September 2022 in an event known as the Merge, transitioning to Proof-of-Stake—a decision consistent with, among other concerns, the limitations we document. The fragility we identify is specific to PoW's incentive structure, not to decentralized consensus per se: Saleh (2021) shows that Proof-of-Stake's finality mechanism precludes the class of fork equilibria that arise under PoW.

1.3 Related Work

We contribute to three literatures. First, we provide the first causal empirical evidence on the theoretical predictions of Carlsten et al. (2016), Bahrani et al. (2025), and Budish (2025) regarding strategic deviation in PoW. Biais et al. (2019) show that blockchain mining is a coordination game with multiple equilibria, several of which sustain persistent forks; our identification strategy distinguishes fee-driven from latency-driven forks. Gans and Halaburda (2024) show that the net cost of a majority attack can be negative because the attacker earns rewards while extending the attacking chain, amplifying the incentive we study.

Second, we contribute to the study of mining pool behavior and miner incentives. Cong et al. (2021) model equilibrium pool formation and show that cross-pool diversification stabilizes the industry at intermediate concentration. Lehar and Parlour (2020) document that pools exercise strategic market power, consistent with the strategic sophistication our fork record reveals. Easley et al. (2019) show that Bitcoin miners allocate block space strategically to maximize fee revenue, providing direct evidence that miners respond to fee incentives.

Third, we contribute to the literature on DeFi and financial infrastructure. Cong and He (2019) show that smart contracts expand the contracting space, enabling the complex transactions that generate the thick-tailed fee distributions central to our analysis. Makarov and Schoar (2022) document cascading liquidations and oracle manipulation that produce the discrete fee spikes whose underlying events we use as instruments. Huberman et al. (2021) show that fees and network load are jointly determined in equilibrium, motivating our IV strategy.[2]

1.4 Roadmap

Section 2 describes the institutional details of strategic re-mining and Ethereum's sibling block record. Section 3 presents the data and summary statistics. Section 4 describes our identification strategy. Section 5 presents the results. Section 6 concludes.

2 Institutional Details: The Economics of Strategic Re-Mining

In a PoW system, miners compete in a continuous, winner-take-all contest to solve a computational puzzle. The victor earns the right to append a new block of transactions to the chain and claim a reward comprising a fixed block subsidy plus the transaction fees of all included transactions. This computational expenditure is the fundamental source of the ledger's security: reversing a confirmed block requires re-doing the work that produced it.

The distribution of these rewards, however, has a thick right tail. While the block subsidy is fixed, transaction fees vary enormously. A single liquidation, flash loan, or arbitrage transaction can concentrate hundreds or thousands of dollars of fees into one block. This variance creates a tension at the heart of the mechanism: the same high-fee block that most needs secure settlement is the one that rewards strategic deviation the most.

The deviation we study is *fee-motivated re-mining*. Upon observing a competitor's newly discovered block, a rational miner may choose not to accept it as the new chain head. Instead, the miner continues mining an alternative block at the same height—potentially reordering or replacing its transactions—and races to propagate it before the network consolidates around the original. If successful, the re-miner captures the fees that would otherwise have gone to the original discoverer. Carlsten et al. (2016) first showed that transaction fees create incentives for this behavior even absent block rewards, and Bahrani et al. (2025) generalize this to show that the incentive grows with the thickness of the reward distribution's right tail.

[2] Budish and Sunderam (2024) argue that when trust is anchored in legal institutions, the competitive disruption from blockchain adoption is limited; our findings reinforce this point by showing that PoW's trustlessness has structural limits.

A related but distinct strategy is *selfish mining* (Eyal and Sirer, 2014), in which a miner withholds a privately discovered block to gain a head start on the next block. Both strategies exploit the same underlying incentive structure, but they target different payoffs and leave different traces. Selfish mining targets the block reward and produces a canonical block timestamped *before* its sibling ($\Delta\tau < 0$). Fee-motivated re-mining targets transaction fees—the channel that varies exogenously in our identification strategy—and produces a canonical block timestamped *after* its sibling ($\Delta\tau > 0$).

Strategic re-mining is related to, but distinct from, the majority attacks studied by Budish (2025). Budish's framework focuses on sustained 51% attacks in which an adversary rewrites arbitrarily many confirmed blocks. Strategic re-mining requires no majority: a miner needs only to win a single propagation race against the honest network by broadcasting a compact block before the original reaches a majority of peers. The motive is narrower—capturing fees in one block rather than reversing the ledger—and the attack leaves an observable trace in the sibling record. Both mechanisms share the same root fragility: the security of PoW is purchased at a flow cost, and when the value at stake rises, so does the incentive to subvert it.

Ethereum's Sibling Record. Most PoW blockchains, including Bitcoin, discard blocks that lose the mining race. Ethereum's protocol was designed differently. When two or more valid blocks are mined at the same blockchain height—we call these *sibling blocks*—the protocol retains a record of the losing block's header. A block in the next generation can reference its parent's sibling, earning a small additional reward for including the reference. This mechanism, originally designed to reduce wasted computation, provides an on-chain record of fork resolution spanning Ethereum's full PoW era.

We leverage this record to define the following concepts:

Definition 1. *A **canonical block** is the block at a given blockchain height that is permanently included in the main chain—the single, authoritative sequence of blocks accepted by the network.*

Definition 2. *A **late canonical block** is a canonical block whose timestamp is later than that of its earliest sibling block at the same blockchain height ($\Delta\tau > 0$).*

Figure 2 provides a schematic of this mechanism. Panel (a) illustrates the normal case, where the canonical block has an earlier timestamp ($\tau_c < \tau_s$) than the competing sibling. Panel (b) depicts the late canonical block case: the block included in the canonical chain has a later timestamp ($\tau_c > \tau_s$) than its sibling, indicating that an earlier, valid block was displaced.

The Role of Block Size. A central prediction of the strategic re-mining hypothesis concerns block size. A rational re-miner faces a propagation race: even if the re-miner finds a valid solution a fraction of a second late, the block can still win canonical status if it reaches a majority of peers

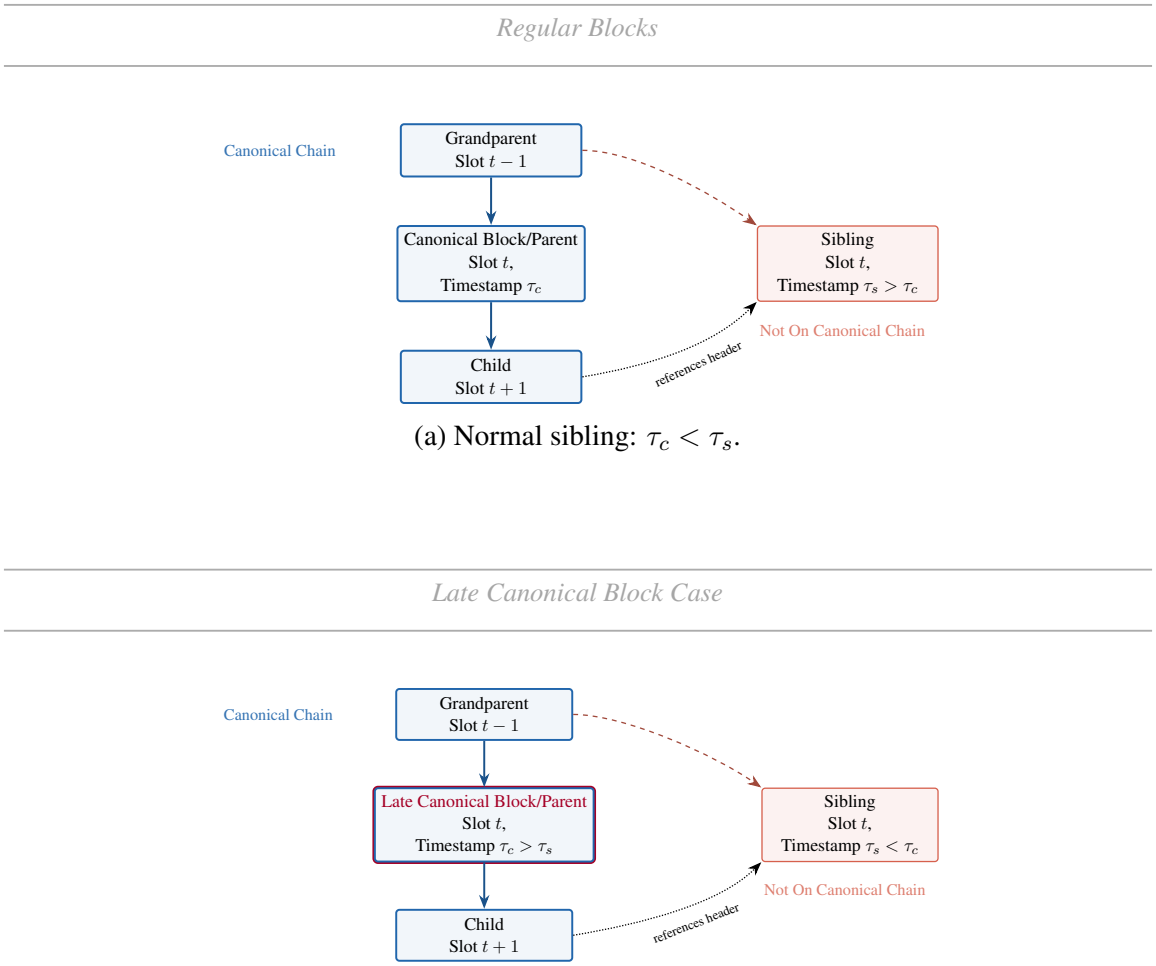

(a) Normal sibling: $\tau_c < \tau_s$.

(b) Late canonical case: $\tau_c > \tau_s$; the canonical block (amber outline) is timestamped after its sibling.

Figure 2: Ethereum sibling blocks and late canonical blocks. Blue nodes form the canonical chain; the rust node is a sibling block mined at the same height as the canonical parent. Solid arrows denote parent–child links; the dashed arc is the competing fork; the dotted arc indicates the child block's inclusion of the sibling header. Panel (a) shows the usual ordering $\tau_c < \tau_s$. Panel (b) highlights a *late canonical block* (amber outline), where the canonical parent is timestamped after its sibling ($\tau_c > \tau_s$).

before the original. Network propagation time is roughly proportional to block size—a 1 KB block traverses the peer-to-peer network faster than a 50 KB block. The re-miner therefore has a strong incentive to construct a *small* block, including only the high-value transactions that motivated the attack, and broadcast immediately. This produces a testable joint prediction: late canonical blocks should have *smaller size* (in bytes) and *larger fees* (in dollars), consistent with deliberate block-trimming for propagation advantage.

Figure 3 schematizes the strategic re-mining decision. The key insight is step 2: the re-miner

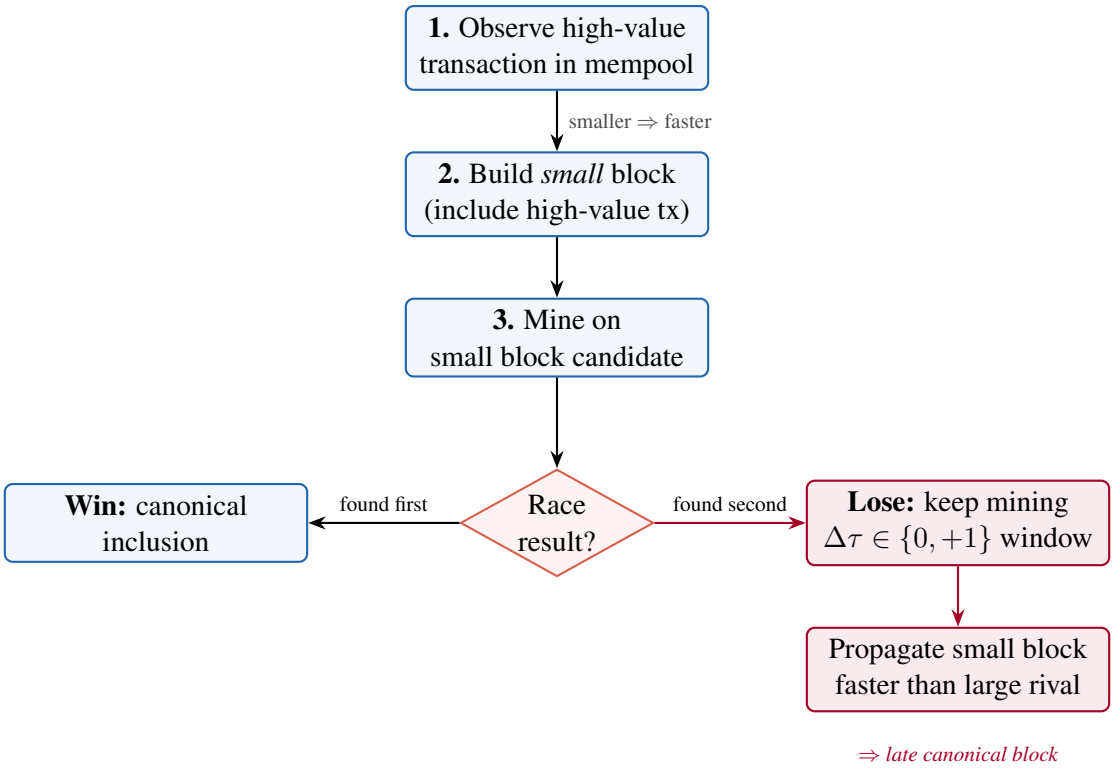

Figure 3: The strategic re-mining mechanism. A miner observing a high-value transaction preemptively builds a small competing block (step 2) so that it propagates faster than any rival. If the miner finds a solution after the competitor (step 4, right path), the smaller block's propagation advantage can still displace the rival, producing a late canonical block ($\Delta\tau > 0$).

deliberately constructs a small block to gain a propagation speed advantage. If the re-miner finds a solution even one second after the honest miner, the smaller block can still reach a majority of peers first, producing a late canonical block at $\Delta\tau = +1$. This is the empirical signature we test below.

3 Data

3.1 Blockchain and Sibling Records

Our dataset covers Ethereum's complete Proof-of-Work history, from the network's launch in August 2015 through the Merge in September 2022. The sample spans 2,592 days and approximately 14 million canonical blocks. Canonical-block timestamps, transaction fees, and block characteristics come from Dune Analytics; sibling-block timestamps come from Etherscan. We merge these records with network hashrate and with a chronology of Ethereum-mainnet exploits, market crises, and regulatory events.

The sibling record lets us observe both sides of a temporary fork. A sibling is a valid block that

competes with the canonical block at the same blockchain height but is not ultimately included in the canonical chain. For each canonical block t with at least one sibling, we define the *timing gap*

$$\Delta\tau_t \equiv \tau_{\text{canonical},t} - \tau_{\text{sibling},t}. \tag{1}$$

When a canonical block has multiple siblings, we compare its timestamp with that of the earliest sibling. A positive gap means that the eventual canonical block is timestamped later than a valid competitor it displaced; a negative gap means that the canonical block is timestamped earlier. Ethereum records timestamps at one-second resolution. Thus $\Delta\tau = +1$ is the smallest observable positive gap, while $\Delta\tau = -1$ is its closest mirror outcome. We do not assign an ordering to $\Delta\tau = 0$, which accounts for roughly 62 percent of sibling observations.

3.2 The Hourly Panel

Let $\mathcal{B}_h$ denote all canonical blocks mined in UTC clock hour h. For each integer timing gap k, we construct the hourly count

$$C_h(k) = \sum_{t \in \mathcal{B}_h} \mathbf{1}[\Delta\tau_t = k]. \tag{2}$$

Our primary outcomes are $C_h(+1)$ and $C_h(-1)$. We also construct the full timing profile for $k \in \{-5, \ldots, +5\}$ and the broader counts

$$\text{Late}_h = \sum_{t \in \mathcal{B}_h} \mathbf{1}[\Delta\tau_t > 0], \qquad \text{Early}_h = \sum_{t \in \mathcal{B}_h} \mathbf{1}[\Delta\tau_t < 0]. \tag{3}$$

Hours in which no block falls into a category receive a count of zero. Blocks without a sibling therefore do not enter a timing category, but they remain in the data used to measure hourly fees and network conditions. The outcome is defined for every clock hour rather than only for blocks selected into a sibling sample.

For the right-hand-side variables, we first take logs at the block level and then average within the hour. The resulting measures are mean log transaction fees in USD, mean log hashrate, and mean log block size. The hourly aggregation aligns the outcomes with the intraday timing of the exploit shocks and avoids treating the hundreds of blocks mined during the same shock window as independent observations.

Table 1 defines the variables in the hourly analysis.

Table 1: Variable Definitions

Variable	Notation	Definition
Primary dependent variables		
Exact $+1$ count	$C_h(+1)$	Canonical blocks in hour h timestamped exactly one second after their earliest sibling.
Exact -1 count	$C_h(-1)$	Canonical blocks in hour h timestamped exactly one second before their earliest sibling.
Timing-bin count	$C_h(k)$	Canonical blocks in hour h with $\Delta\tau_t = k$, for $k \in \{-5, \ldots, +5\}$.
Supporting dependent variables		
Late count	Late_h	Canonical blocks in hour h with $\Delta\tau_t > 0$.
Early count	Early_h	Canonical blocks in hour h with $\Delta\tau_t < 0$.
Sibling-bearing count	$\mathrm{Sibling}_h$	Canonical blocks in hour h with at least one recorded sibling.
Endogenous regressors		
Mean log fees	$\overline{\log(\mathrm{Fees})}_h$	Hourly mean of block-level log transaction fees in USD.
Mean log hashrate	$\overline{\log(\mathrm{Hashrate})}_h$	Hourly mean of block-level log network hashrate.
Mean log block size	$\overline{\log(\mathrm{Size})}_h$	Hourly mean of block-level log size in bytes.
Excluded instruments		
Crypto shock	$\mathrm{CryptoShock}_h$	Union of exact-timestamp exploit windows and dated crisis and regulatory-event windows.
Lagged hashrate	$L.\overline{\log(\mathrm{Hashrate})}_h$	Mean log hashrate in the preceding clock hour.
Lagged block size	$L.\overline{\log(\mathrm{Size})}_h$	Mean log block size in the preceding clock hour.
Fixed effects		
Month FE	$\gamma_{m(h)}$	Calendar-month fixed effects.

Notes: $\Delta\tau_t$ is the canonical timestamp minus the earliest-sibling timestamp. Count outcomes are summed within the UTC hour; the three endogenous regressors are block-level logs averaged within the hour. Every hour remains in the panel, including hours with zero timing-bin counts.

3.3 External Shocks and Their Timing

The fee-shock data combine three event classes: Ethereum-mainnet exploits, major market crises, and regulatory developments. Exploit times are anchored to the timestamp of the exploit transaction itself rather than to the date of a subsequent announcement or news story. We round each transaction timestamp down to its UTC clock hour and activate the exploit indicator in that hour and the following two hours. Thirteen exact-timestamp exploits fall within the Proof-of-Work sample, producing 39 treated exploit hours.

The three-hour window is intentionally short. It allows the immediate transaction-demand and market response to reach the chain without converting an intraday event into a daily treatment. Crisis and regulatory events retain dated windows because a comparably precise intraday timestamp is not consistently available for those event classes. $\mathrm{CryptoShock}_h$ is the union of the exact exploit

windows and these dated events. It is active in 375 of the 61,967 estimation hours.

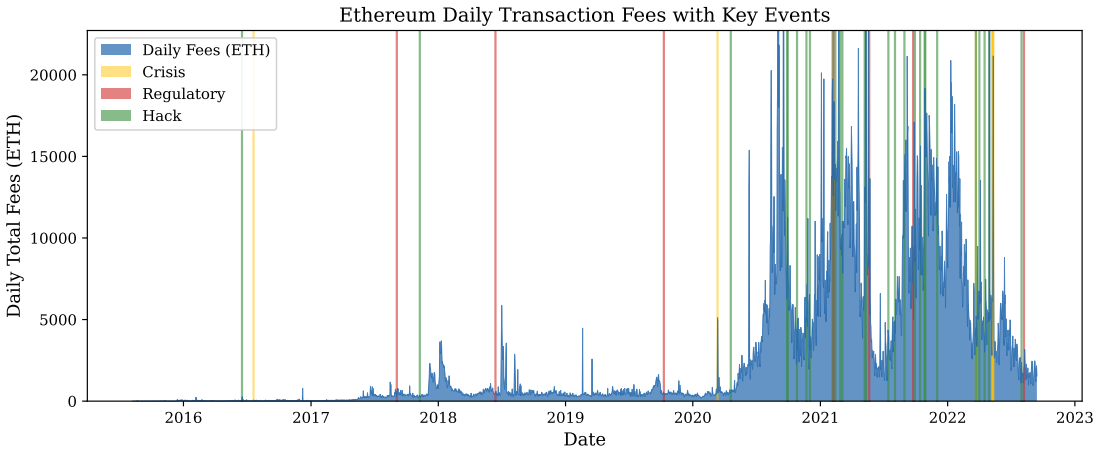

Figure 4: Daily total transaction fees (ETH) over Ethereum's Proof-of-Work era. Yellow vertical lines mark market crises, red lines regulatory developments, and green lines Ethereum exploits. The regression panel is hourly: exploit timestamps are matched to the event hour and the following two hours rather than assigned to the entire event date.

Figure 4 shows that the events composing our instrument coincide with visible disruptions to transaction fees. Fees remain low through the early sample and rise sharply with the 2020–2021 expansion of decentralized finance; against this slow-moving background, market crises (yellow), Ethereum exploits (green), and regulatory events (red) pick out discrete, sharply timed fee spikes. This visual pattern anticipates the formal relevance result: CryptoShock strongly predicts hourly mean log fees in the first stage, with a Kleibergen–Paap F-statistic of 87.12 in the Fees-Only specification.

3.4 Sample and Descriptive Statistics

The block-level record contains 13,879,763 observations with nonmissing key variables. Of these, 1,103,929 blocks (7.9 percent) have at least one associated sibling. Aggregating the complete block record produces 61,968 consecutive UTC hours. Constructing the one-hour lags of hashrate and block size removes the first hour and leaves 61,967 observations in every baseline regression.

Table 2 shows that exact one-second timing outcomes are common enough to support hourly analysis. An average hour contains 17.82 canonical blocks with a recorded sibling; the exact $+1$ and -1 outcomes average 1.982 and 2.802 blocks per hour, with medians of one and two, and the broader Late and Early counts average 2.551 and 4.198. That the -1 and Early outcomes are unconditionally more frequent is not evidence against strategic re-mining: honest mining generates early canonical blocks by default, and the empirical design asks not which outcome is more common but how each responds to exogenous changes in fees.

Table 2: Summary Statistics: Hourly Estimation Sample

Variable	N	Mean	SD	p25	Median	p75	Min	Max
$C_h(-1)$: exact -1 count	61,967	2.802	3.425	1.000	2.000	3.000	0.000	28.000
$C_h(+1)$: exact $+1$ count	61,967	1.982	2.462	0.000	1.000	3.000	0.000	20.000
Early_h: all negative gaps	61,967	4.198	5.738	1.000	2.000	4.000	0.000	53.000
Late_h: all positive gaps	61,967	2.551	3.297	1.000	2.000	3.000	0.000	33.000
Canonical blocks with a sibling	61,967	17.820	12.079	11.000	14.000	20.000	0.000	92.000
Hourly mean log fees	61,967	2.564	4.186	1.006	2.807	5.969	-7.114	12.275
Hourly mean log hashrate	61,967	32.298	2.399	31.346	33.132	33.644	24.310	35.146
Hourly mean log block size	61,967	9.455	1.412	8.360	9.877	10.633	6.476	11.671
$\mathrm{CryptoShock}_h$	61,967	0.006	0.078	0.000	0.000	0.000	0.000	1.000
Exact-timestamp hack window	61,967	0.001	0.025	0.000	0.000	0.000	0.000	1.000

Notes: Statistics are computed after the one-hour lag requirement, matching the 2SLS estimation sample. Count variables are sums within a UTC clock hour; log fees, log hashrate, and log size are block-level values averaged within the hour. $\mathrm{CryptoShock}_h$ is the union of exact-timestamp hack windows, crisis dates, and regulatory-event dates. The exact-timestamp hack indicator covers the event hour and following two hours.

4 Empirical Strategy

4.1 Hourly Instrumental-Variables Specification

Let $Y_h^{(j)}$ denote an hourly fork outcome, where j indexes either a particular timing bin or one of the broader Early and Late counts. For each outcome, we estimate

$$Y_h^{(j)} = \beta_j \, \overline{\log(\mathrm{Fees})}_h + \theta_{1j} \, \overline{\log(\mathrm{Hashrate})}_h + \theta_{2j} \, \overline{\log(\mathrm{Size})}_h + \gamma_{j,m(h)} + \varepsilon_{jh}. \tag{4}$$

Our primary estimands are β_{+1} and β_{-1}, obtained from separate regressions for $C_h(+1)$ and $C_h(-1)$. We apply the same specification to each bin from -5 through $+5$ seconds to trace the complete timing profile. The Early and Late estimates aggregate the negative and positive bins and provide a broader comparison.

Hourly mean log fees, log hashrate, and log block size are jointly determined by activity on the network, so we treat all three as endogenous. The excluded instruments are $\mathrm{CryptoShock}_h$, the one-hour lag of mean log hashrate, and the one-hour lag of mean log block size. In compact form, the first stage for each endogenous regressor X_{qh} is

$$\begin{aligned} X_{qh} = \pi_{1q} \, \mathrm{CryptoShock}_h + \pi_{2q} \, \overline{\log(\mathrm{Hashrate})}_{h-1} \\ + \pi_{3q} \, \overline{\log(\mathrm{Size})}_{h-1} + \mu_{q,m(h)} + \nu_{qh}. \end{aligned} \tag{5}$$

Here $q \in \{\overline{\log \mathrm{Fees}}, \overline{\log \mathrm{Hashrate}}, \overline{\log \mathrm{Size}}\}$. The model has three endogenous regressors and three excluded instruments and is therefore exactly identified. Calendar-month fixed effects absorb protocol changes, secular network growth, and broad market regimes while preserving the intraday variation generated by precisely timed shocks. In the baseline heteroskedasticity-robust specification, the Kleibergen–Paap rk Wald F-statistic is 22.76.

4.2 Identification

$\mathrm{CryptoShock}_h$ supplies discrete variation in transaction demand from events whose timing is external to the outcome of any particular fork. The exclusion restriction is that, conditional on contemporaneous hashrate, block size, and month fixed effects, these events affect hourly timing-gap counts through transaction fees rather than by independently determining which member of a sibling pair becomes canonical. The lagged network variables provide first-stage variation in their corresponding contemporaneous measures; their exclusion requires that any effect of the preceding hour's hashrate and block size on current fork resolution operates through current network conditions.

The principal alternative is that an event changes congestion or propagation generally, producing more dispersed timestamps without inducing strategic re-mining. The design addresses this concern in three complementary ways. First, hashrate and block size enter Equation (4) and are instrumented jointly with fees. Second, month fixed effects remove slow-moving changes in network capacity and usage. Third, the outcome comparison is directional. Generic congestion or timestamp noise could widen timing gaps, but it should not selectively raise $C_h(+1)$ while leaving the adjacent mirror outcome $C_h(-1)$ unchanged. The full -5 to $+5$ profile provides a further diagnostic: a diffuse propagation shock should affect several neighboring bins rather than concentrate at the minimum positive gap.

4.3 Inference and Serial Dependence

The baseline estimates use heteroskedasticity-robust standard errors at the hourly level. Adjacent hours may nevertheless share persistent network conditions. We therefore examine serial dependence in the robustness analysis using Newey–West/Bartlett HAC estimators with bandwidths of 6 and 12 hours and, separately, standard errors clustered in consecutive, nonoverlapping 6- and 12-hour blocks. HAC permits covariance to decay with the distance between observations within the chosen bandwidth; block clustering allows unrestricted dependence among observations within each intraday block. These alternatives leave the specification, point estimates, and fixed effects unchanged and vary only the estimated covariance matrix.

5 Results

Table 3 reports the paper's primary one-second comparison in the main and fees-only specifications.

The main specification shows the central result: higher fees significantly increase the hourly number of canonical blocks arriving exactly one second after their siblings, while the coefficient for the mirror -1 outcome is economically near zero. The same pattern appears when fees are

Table 3: Exact One-Second Outcomes: Hourly 2SLS

	Main		Fees-Only	
	(1)	**(2)**	**(3)**	**(4)**
	$C_h(+1)$	$C_h(-1)$	$C_h(+1)$	$C_h(-1)$
$\overline{\log(\text{Fees})}_h$	0.520^{**}	-0.026	0.534^{***}	0.121
	(0.233)	(0.273)	(0.189)	(0.221)
$\overline{\log(\text{Hashrate})}_h$	-0.971^{***}	-1.048^{***}		
	(0.125)	(0.147)		
$\overline{\log(\text{Size})}_h$	0.142	1.453^{***}		
	(0.427)	(0.500)		
Dependent-variable mean	1.982	2.802	1.982	2.802
Month fixed effects	Yes			
Observations	61,967		61,967	
Kleibergen–Paap rk Wald F	22.76		87.12	

Notes: Each column reports a separate hourly 2SLS regression. The dependent variable is the number of canonical blocks in hour h timestamped exactly one second after or before their earliest sibling. Main treats hourly mean log fees, log hashrate, and log block size as endogenous; excluded instruments are CryptoShock_h and the one-hour lags of mean log hashrate and log block size. Fees-Only treats mean log fees as endogenous, instruments it with CryptoShock_h, and excludes hashrate and size. Heteroskedasticity-robust standard errors are in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

instrumented alone. Hashrate loads negatively on both one-second outcomes. Block size, by contrast, is positive and significant at -1 but not at $+1$: as we argued in Section 2, a large canonical block is more likely to have been mined honestly, whereas a smaller block may either be honestly mined or strategically mined.

Magnitudes. The coefficient has a direct count interpretation. A one-log-point increase in hourly mean fees—approximately a 2.7-fold increase—raises the expected number of canonical blocks at $\Delta\tau = +1$ by 0.520 per hour. Relative to the observed mean of 1.982 such blocks per hour, this is an increase of approximately 26 percent.

Distributional Analysis. Figure 5 plots the complete hourly timing profile for fees, hashrate, and block size; Appendix Table 5 reports the corresponding estimates. The fee effect is sharply asymmetric, concentrated at a single timing gap.

Higher fees significantly increase the hourly number of canonical blocks timestamped exactly one second after their sibling competitors ($\Delta\tau = +1$). The fee coefficient at this bin is positive and significant at the 5% level, while the coefficients at every larger positive gap are statistically insignificant. Because timestamps are recorded with one-second resolution, $\Delta\tau = +1$ is the

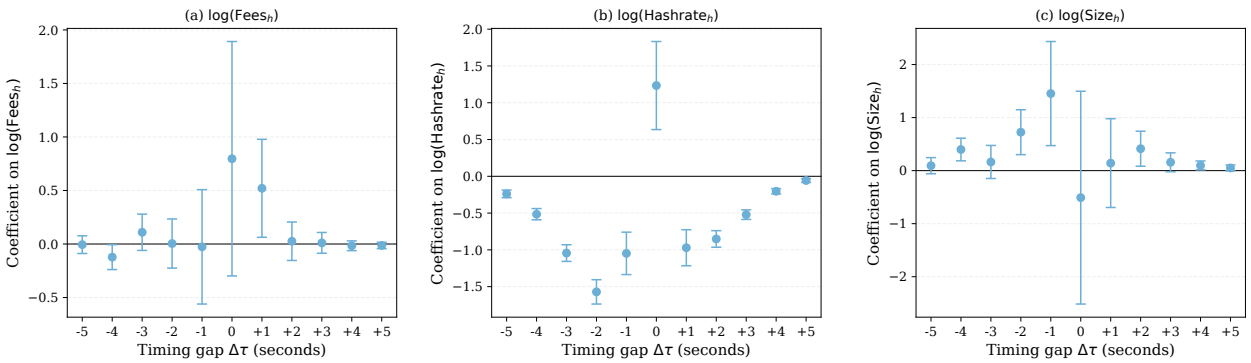

Figure 5: Hourly distributional event study. Each point is the 2SLS coefficient from a separate regression of $C_h(k)$, the number of canonical blocks in hour h with timing gap $k \in \{-5, \ldots, +5\}$ seconds. Hourly mean log fees, log hashrate, and log block size are treated as endogenous and instrumented with $\mathrm{CryptoShock}_h$ and the one-hour lags of mean log hashrate and log block size. All regressions include month fixed effects; error bars show heteroskedasticity-robust 95% confidence intervals.

minimum positive gap—the first bin in which a late canonical block can appear. The concentration at this minimum gap is consistent with a miner who, upon observing a competitor's high-value block, continues mining rather than accepting the rival block as the new chain head.

The mirror estimate at $\Delta\tau = -1$ is -0.026 (SE 0.273), economically close to zero and statistically insignificant. The coefficient at $\Delta\tau = 0$ is positive but imprecisely estimated, and the only other significant fee coefficient occurs at $\Delta\tau = -4$ with the opposite sign. The profile therefore does not resemble a diffuse widening of timestamp gaps: the economically relevant positive response is concentrated at the minimum late gap.

The hashrate coefficients are consistent with the two channels identified above. Higher hashrate increases the hourly count at $\Delta\tau = 0$ and reduces counts at nearly every nonzero timing gap. This inward compression is consistent with the Poisson simultaneity channel: blocks found during periods of high hashrates are more likely to produce near-simultaneous arrivals. The negative coefficient at $\Delta\tau = +1$ is also consistent with deterrence, as a larger honest network makes strategic re-mining less profitable.

Block size has its largest positive coefficients on the negative side of the timing distribution, including at $\Delta\tau = -1$. A large canonical block propagates more slowly and is therefore most likely to survive a close race when it was mined first and entered the network with a head start. Conversely, the size coefficient at $\Delta\tau = +1$ is small and statistically insignificant, consistent with strategically re-mined blocks being constructed compactly to improve their propagation advantage.

The fee result is stable across the main and fees-only specifications. In the hourly estimates, the coefficient at $\Delta\tau = +1$ is 0.520 when fees, hashrate, and block size are jointly instrumented with the full instrument set. It is 0.534 in a Fees-Only, specification which excludes hashrate and block

size and instruments fees solely with $\mathrm{CryptoShock}_h$. The close agreement between these estimates indicates that the result is not driven by the treatment of hashrate or block size.

It is important to note that the fees that induce misbehavior are relatively small.. The 75th percentile of total block fees is only $213, the 90th is $1,908, and the 95th exceeds $4,400. If PoW were to become the dominant settlement infrastructure for finance, per-block fees would rise substantially and these incentives would become dramatically stronger.

5.1 Robustness

Table 4: Serial-Correlation-Robust Inference for the $\pm\mathbf{1}$ Result

	$C_h(+1)$			$C_h(-1)$		
Inference	Coef.	SE	p	Coef.	SE	p
HC robust	0.520	0.233	0.026	−0.026	0.273	0.923
HAC(6)	0.520	0.270	0.054	−0.026	0.343	0.939
HAC(12)	0.520	0.294	0.077	−0.026	0.385	0.946
Cluster: 6-hour blocks	0.520	0.299	0.082	−0.026	0.253	0.917
Cluster: 12-hour blocks	0.520	0.308	0.092	−0.026	0.277	0.924
KP F	22.76 / 17.59 / 10.87 / 4.62 / 2.63			same		
Observations	61,967			61,967		

Notes: Point estimates and specification are identical across rows; only the covariance estimator changes. HAC rows use a Bartlett kernel with 6- and 12-hour bandwidths. Cluster rows use consecutive, nonoverlapping blocks of 6 or 12 hours. The KP F entries follow the row order.

Table 4 allows for short-run serial correlation in two ways. With Newey–West/Bartlett HAC standard errors, the $+1$ coefficient remains significant at 10 percent using either a 6-hour bandwidth ($p = 0.054$) or a 12-hour bandwidth ($p = 0.077$). Clustering in consecutive, nonoverlapping 6- and 12-hour blocks produces p-values of 0.082 and 0.092. The -1 coefficient remains essentially zero under every correction, with p-values between 0.917 and 0.946.

The clustered rows are best understood as inference-sensitivity checks rather than equally strong IV specifications. The baseline system-level Kleibergen–Paap statistic is 22.76 and remains 17.59 and 10.87 under HAC(6) and HAC(12), respectively. Under 6- and 12-hour block clustering, however, the reported statistics fall to 4.62 and 2.63. We therefore place greater weight on the HAC estimates while reporting the clustered results as conservative checks on within-block dependence.

Aggregate Early and Late outcomes. The broader split between all positive timing gaps (*Late*) and all negative timing gaps (*Early*) provides a supporting check. With baseline heteroskedasticity-robust inference, the fee coefficient is positive for *Late* (0.555, $p = 0.047$) and near zero for *Early* (-0.146, $p = 0.691$). Figure 6 illustrates these results.

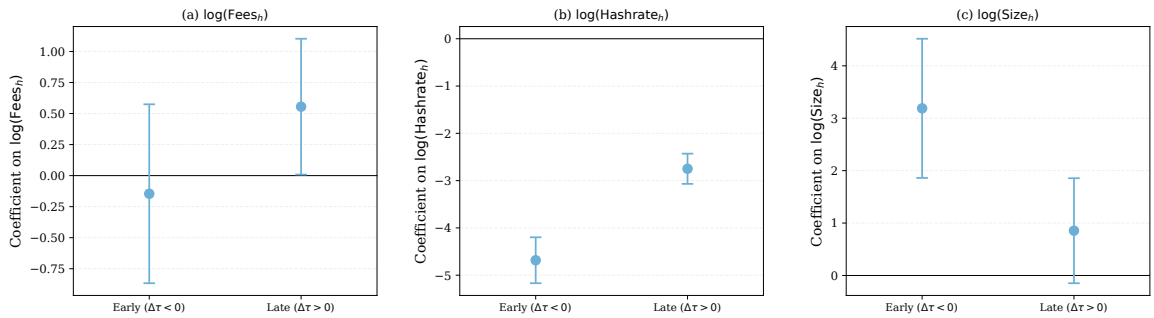

Figure 6: Aggregate Early and Late outcomes. Each panel compares the hourly 2SLS estimates for all negative timing gaps (*Early*, $\Delta\tau < 0$) and all positive timing gaps (*Late*, $\Delta\tau > 0$). The panels report coefficients on log fees (a), log hashrate (b), and log block size (c), with 95% heteroskedasticity-robust confidence intervals. All three regressors are treated as endogenous and estimated jointly.

5.2 Supporting Evidence in the Appendices

The four appendices collect the evidence supporting the main one-second result. Appendix A reports the complete timing profile underlying Figure 5: the positive fee response is concentrated at $\Delta\tau = +1$, while the mirror estimate at -1 and all larger positive gaps remain statistically insignificant. Appendix B reports the first stages and shows that CryptoShock strongly predicts hourly fees and that the baseline instrument set is jointly relevant. Appendix C provides descriptive OLS benchmarks for completeness. Appendix D then subjects the IV result to the most demanding checks. The $+1$ effect remains statistically detectable under HAC and block-cluster inference, appears directly in the reduced form, survives alternative constructions of CryptoShock, and is unchanged by mining-pool composition controls. Across these exercises, the -1 estimate remains near zero.

6 Conclusion

Every centralized settlement system in operation today shares a common design principle: enforcement scales with the value at stake. Courts devote more resources to high-value disputes; clearinghouses raise margin requirements during periods of stress; regulators intervene when systemic risk rises. Our evidence from PoW Ethereum shows that Proof-of-Work can violate this principle. Higher transaction fees—which should signal that secure settlement is most needed—instead cause miners to deviate from honest behavior, displacing valid blocks to capture their fees. In this setting, the security guarantee of the infrastructure weakens as the value at stake rises.

Our evidence speaks directly to the theoretical predictions of Budish (2025), who showed that the economic limits of PoW are binding: the security of the ledger is purchased at a flow cost, and when

the value at stake rises, so does the incentive to subvert it. We confirm this prediction causally. The effect is concentrated at the minimum possible timing gap ($\Delta\tau = +1$), consistent with fee-motivated re-mining rather than latency or selfish mining. The $+1$ result remains statistically detectable at the 10-percent level with 6- and 12-hour Bartlett HAC corrections and in the corresponding 6- and 12-hour block-cluster sensitivity checks, while the -1 estimate remains null throughout.

These findings have implications beyond blockchain. Proponents have framed PoW blockchains as a new institutional technology capable of trustless enforcement, substituting economic incentives for the authority of legal systems. Our results suggest that this form of incentive-based enforcement can face a structural limit: its performance may degrade precisely when the stakes are high. The Ethereum evidence cautions that Proof-of-Work does not by itself resolve the problem of securing high-value settlement—the problem that courts, regulators, and clearinghouses were built to solve. The economic tension between enforcement costs and attack incentives, identified theoretically by Carlsten et al. (2016), Bahrani et al. (2025), and Budish (2025), manifests empirically in the blockchain record.

References

ABADI, J., AND M. BRUNNERMEIER (2022): "Blockchain Economics," NBER Working Paper No. 25407, National Bureau of Economic Research.

BAHRANI, M., M. NEUDER, AND S. M. WEINBERG (2025): "Selfish Mining under General Stochastic Rewards," arXiv preprint arXiv:2502.20360.

BIAIS, B., C. BISIÈRE, M. BOUVARD, AND C. CASAMATTA (2019): "The Blockchain Folk Theorem," *Review of Financial Studies*, 32(5), 1662–1715.

BUDISH, E. (2025): "Trust at Scale: The Economic Limits of Cryptocurrencies and Blockchains," *The Quarterly Journal of Economics*, 140(1), 1–62.

BUDISH, E., AND A. SUNDERAM (2024): "Blockchain Technology in Traditional Finance," Working Paper, University of Chicago Booth School of Business and Harvard Business School.

CARLSTEN, M., H. KALODNER, S. M. WEINBERG, AND A. NARAYANAN (2016): "On the Instability of Bitcoin Without the Block Reward," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 154–167. ACM.

CONG, L. W., AND Z. HE (2019): "Blockchain Disruption and Smart Contracts," *Review of Financial Studies*, 32(5), 1754–1797.

CONG, L. W., Z. HE, AND J. LI (2021): "Decentralized Mining in Centralized Pools," *Review of Financial Studies*, 34(3), 1191–1235.

EASLEY, D., M. O'HARA, AND S. BASU (2019): "From Mining to Markets: The Evolution of Bitcoin Transaction Fees," *Journal of Financial Economics*, 134(1), 91–109.

EYAL, I., AND E. G. SIRER (2014): "Majority is Not Enough: Bitcoin Mining Is Vulnerable," in *Financial Cryptography and Data Security*, pp. 436–454. Springer.

GANS, J. S., AND H. HALABURDA (2024): "'Zero Cost' Majority Attacks on Permissionless Proof of Work Blockchains," *Management Science*, 70(6), 4155–4165.

HUBERMAN, G., J. D. LESHNO, AND C. C. MOALLEMI (2021): "Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System," *Review of Economic Studies*, 88(6), 3011–3040.

LEHAR, A., AND C. A. PARLOUR (2020): "Miner Collusion and the Bitcoin Protocol," Working Paper, University of Calgary and UC Berkeley.

MAKAROV, I., AND A. SCHOAR (2022): "Cryptocurrencies and Decentralized Finance (DeFi)," *Brookings Papers on Economic Activity*, 141–215.

SALEH, F. (2021): "Blockchain without Waste: Proof-of-Stake," *Review of Financial Studies*, 34(3), 1156–1190.

YAISH, A., G. STERN, AND A. ZOHAR (2023): "Uncle Maker: (Time)Stamping Out The Competition in Ethereum," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pp. 135–149. ACM.

Appendix

A Complete Hourly Timing Profile

Table 5: Distributional Event Study: Hourly 2SLS

Timing gap k	-5	-4	-3	-2	-1	0	$+1$	$+2$	$+3$	$+4$	$+5$
FULL SAMPLE ($N = 61{,}967$; Kleibergen–Paap rk Wald $F = 22.76$)											
Outcome mean	0.069	0.121	0.253	0.692	2.802	11.071	1.982	0.355	0.101	0.039	0.019
$\overline{\log(\text{Fees})}_h$	-0.006	-0.123^{**}	0.110	0.005	-0.026	0.797	0.520^{**}	0.026	0.011	-0.016	-0.014
	(0.042)	(0.059)	(0.087)	(0.117)	(0.273)	(0.559)	(0.233)	(0.092)	(0.050)	(0.023)	(0.015)
$\overline{\log(\text{Hashrate})}_h$	-0.239^{***}	-0.515^{***}	-1.044^{***}	-1.571^{***}	-1.048^{***}	1.234^{***}	-0.971^{***}	-0.852^{***}	-0.521^{***}	-0.203^{***}	-0.057^{***}
	(0.026)	(0.039)	(0.058)	(0.084)	(0.147)	(0.305)	(0.125)	(0.057)	(0.033)	(0.018)	(0.013)
$\overline{\log(\text{Size})}_h$	0.093	0.398^{***}	0.163	0.725^{***}	1.453^{***}	-0.509	0.142	0.413^{**}	0.157^{*}	0.098^{**}	0.051^{*}
	(0.078)	(0.109)	(0.159)	(0.216)	(0.500)	(1.023)	(0.427)	(0.168)	(0.091)	(0.043)	(0.028)
Month fixed effects						Yes					

Notes: Each column reports a separate hourly 2SLS regression for $C_h(k)$, the number of canonical blocks in hour h with timing gap exactly k seconds. The table uses the complete Proof-of-Work sample. Hourly mean log fees, log hashrate, and log block size are endogenous; excluded instruments are CryptoShock_h and the one-hour lags of mean log hashrate and log block size. Heteroskedasticity-robust standard errors are in parentheses. The shaded column marks the minimum positive timing gap, $\Delta\tau = +1$.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 5 reports the numerical estimates underlying Figure 5. The fee coefficient is positive and statistically significant at $\Delta\tau = +1$, while the mirror coefficient at $\Delta\tau = -1$ is close to zero. None of the larger positive-gap coefficients is statistically significant. The only other significant fee coefficient is negative and occurs at $\Delta\tau = -4$, moving in the opposite direction from the strategic re-mining prediction. Hashrate is negatively associated with nearly every nonzero timing bin, and larger blocks load most strongly on the negative-gap bins.

B IV First-Stage Regressions

Table 6 reports the hourly first stages. Fees-Only projects mean log fees on $\mathrm{CryptoShock}_h$. Main projects each of the three endogenous hourly regressors on CryptoShock, one-hour lagged mean log hashrate, and one-hour lagged mean log block size. All columns include calendar-month fixed effects. CryptoShock strongly predicts mean log fees, while the lagged network variables strongly predict their contemporaneous counterparts. The heteroskedasticity-robust Kleibergen–Paap rk Wald F-statistics are 87.12 for Fees-Only and 22.76 for Main.

Table 6: Hourly First-Stage Regressions

	Fees-Only	Main		
	Mean log fees	Mean log fees	Mean log hashrate	Mean log block size
$\mathrm{CryptoShock}_h$	0.400^{***}	0.349^{***}	-0.000	0.013^{**}
	(0.043)	(0.040)	(0.004)	(0.006)
$L.\overline{\log(\mathrm{Hashrate})}_h$		0.307^{***}	0.548^{***}	0.048^{***}
		(0.018)	(0.006)	(0.006)
$L.\overline{\log(\mathrm{Size})}_h$		1.311^{***}	0.013^{***}	0.719^{***}
		(0.013)	(0.002)	(0.005)
Joint excluded-instrument F	87.12	3548.99	2860.71	8184.40
Kleibergen–Paap rk Wald F	87.12		22.76	
Month fixed effects			Yes	
Observations			61,967	

Notes: Each column reports an hourly first-stage regression. Fees-Only instruments mean log fees with $\mathrm{CryptoShock}_h$. Main jointly instruments mean log fees, mean log hashrate, and mean log block size with $\mathrm{CryptoShock}_h$ and one-hour lags of mean log hashrate and log block size. All regressions include calendar-month fixed effects. Heteroskedasticity-robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

C OLS Main Results

This section reports OLS estimates on the same 61,967-hour sample used for the 2SLS analysis. OLS is descriptive because fees, hashrate, and block size respond jointly to network conditions. Its value here is as a contrast: the raw fee association is spread broadly across timing bins, whereas the instrumented fee effect is concentrated at $\Delta\tau = +1$.

C.1 OLS Baseline and Contrast

Table 7: Exact One-Second Outcomes: Hourly OLS

	Main		Fees-Only	
	(1)	**(2)**	**(3)**	**(4)**
	$C_h(+1)$	$C_h(-1)$	$C_h(+1)$	$C_h(-1)$
$\overline{\log(\text{Fees})}_h$	0.193***	0.242***	0.352***	0.429***
	(0.018)	(0.021)	(0.013)	(0.016)
$\overline{\log(\text{Hashrate})}_h$	0.036	0.080		
	(0.045)	(0.054)		
$\overline{\log(\text{Size})}_h$	0.671***	0.790***		
	(0.041)	(0.051)		
Dependent-variable mean	1.982	2.802	1.982	2.802
R^2	0.653	0.735	0.651	0.733
Month fixed effects	Yes			
Observations	61,967			

Notes: Each column reports a separate OLS regression on the same hourly estimation sample used for the 2SLS results. Main includes hourly mean log fees, log hashrate, and log block size; Fees-Only includes hourly mean log fees alone. The dependent variable is the number of canonical blocks in hour h timestamped exactly one second after or before their earliest sibling. Heteroskedasticity-robust standard errors are in parentheses. No selection correction or instrumentation is used.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 7 reports the exact one-second outcomes. In both specifications, mean log fees are positively associated with $C_h(+1)$ and $C_h(-1)$. Thus the raw hourly association is spread across both sides of zero, unlike the 2SLS result concentrated at $+1$.

C.2 OLS Distributional Event Study

Table 8: Distributional Event Study: Hourly OLS

Timing gap k	-5	-4	-3	-2	-1	0	$+1$	$+2$	$+3$	$+4$	$+5$
MAIN ($N = 61{,}967$)											
Outcome mean	0.069	0.121	0.253	0.692	2.802	11.071	1.982	0.355	0.101	0.039	0.019
$\overline{\log(\text{Fees})}_h$	0.043^{***}	0.083^{***}	0.133^{***}	0.214^{***}	0.242^{***}	0.051	0.193^{***}	0.120^{***}	0.061^{***}	0.025^{***}	0.012^{***}
	(0.005)	(0.007)	(0.009)	(0.013)	(0.021)	(0.040)	(0.018)	(0.008)	(0.004)	(0.003)	(0.002)
$\overline{\log(\text{Hashrate})}_h$	-0.150^{***}	-0.307^{***}	-0.538^{***}	-0.780^{***}	0.080	4.510^{***}	0.036	-0.420^{***}	-0.272^{***}	-0.099^{***}	-0.019^{***}
	(0.012)	(0.017)	(0.026)	(0.038)	(0.054)	(0.119)	(0.045)	(0.024)	(0.015)	(0.009)	(0.007)
$\overline{\log(\text{Size})}_h$	0.010	0.009	0.086^{***}	0.268^{***}	0.790^{***}	0.659^{***}	0.671^{***}	0.207^{***}	0.052^{***}	0.014^{**}	0.006
	(0.010)	(0.015)	(0.019)	(0.029)	(0.051)	(0.098)	(0.041)	(0.019)	(0.010)	(0.006)	(0.004)
FEES-ONLY ($N = 61{,}967$)											
$\overline{\log(\text{Fees})}_h$	0.043^{***}	0.080^{***}	0.146^{***}	0.266^{***}	0.429^{***}	0.273^{***}	0.352^{***}	0.163^{***}	0.069^{***}	0.027^{***}	0.014^{***}
	(0.003)	(0.005)	(0.006)	(0.010)	(0.016)	(0.031)	(0.013)	(0.006)	(0.003)	(0.002)	(0.001)
Month fixed effects						Yes					

Notes: Each column reports a separate hourly OLS regression for $C_h(k)$, the number of canonical blocks in hour h with timing gap exactly k seconds. Main includes hourly mean log fees, log hashrate, and log block size; Fees-Only includes hourly mean log fees alone. All regressions use the same sample as the hourly 2SLS estimates and include calendar-month fixed effects. Heteroskedasticity-robust standard errors are in parentheses. No selection correction or instrumentation is used. The shaded column marks $\Delta\tau = +1$.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 8 makes the contrast visible across the full timing profile. In both OLS specifications, higher fees are positively associated with counts at nearly every nonzero timing gap. The association is therefore diffuse rather than concentrated at the minimum positive gap. This broad pattern is consistent with high-fee hours also being high-activity or congested hours, conditions that can raise fork counts on both sides of zero. Instrumenting isolates a different pattern: a positive effect at $+1$, a near-zero effect at -1, and no positive effect at larger positive gaps.

D Additional Hourly Robustness Results

This appendix reports the complete hourly estimates behind the robustness checks summarized in Section 5.1. All tables use the same 61,967-hour sample and the same calendar-month fixed effects as the main specification.

D.1 Serial-Correlation-Robust Inference for the Exact One-Second Outcomes

Table 9: Full Serial-Correlation-Robust Inference: Exact One-Second Outcomes

	$C_h(+1)$			$C_h(-1)$			
Inference	Fees	Hashrate	Size	Fees	Hashrate	Size	KP F
HC robust	0.520^{**}	-0.971^{***}	0.142	-0.026	-1.048^{***}	1.453^{***}	22.76
	(0.233)	(0.125)	(0.427)	(0.273)	(0.147)	(0.500)	
HAC(6)	0.520^{*}	-0.971^{***}	0.142	-0.026	-1.048^{***}	1.453^{**}	17.59
	(0.270)	(0.144)	(0.493)	(0.343)	(0.183)	(0.627)	
HAC(12)	0.520^{*}	-0.971^{***}	0.142	-0.026	-1.048^{***}	1.453^{**}	10.87
	(0.294)	(0.156)	(0.538)	(0.385)	(0.202)	(0.704)	
Cluster: 6-hour blocks	0.520^{*}	-0.971^{***}	0.142	-0.026	-1.048^{***}	1.453^{***}	4.62
	(0.299)	(0.156)	(0.546)	(0.253)	(0.162)	(0.468)	
Cluster: 12-hour blocks	0.520^{*}	-0.971^{***}	0.142	-0.026	-1.048^{***}	1.453^{***}	2.63
	(0.308)	(0.166)	(0.563)	(0.277)	(0.185)	(0.513)	
Month fixed effects	Yes						
Observations	61,967						

Notes: Each outcome and inference row reports the same hourly 2SLS point estimates; only the covariance estimator changes. Fees, hashrate, and size denote the coefficients on hourly mean log fees, log hashrate, and log block size. HAC uses a Bartlett kernel. Cluster rows use consecutive nonoverlapping blocks. All three regressors are endogenous; excluded instruments are $\mathrm{CryptoShock}_h$ and one-hour lags of mean log hashrate and log block size. Heteroskedasticity-robust or serial-correlation-robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 9 expands the compact inference summary in Table 4 to all three endogenous regressors. The $+1$ fee coefficient remains positive and significant at 10 percent under HAC(6), HAC(12), and 6- and 12-hour block clustering; the -1 fee coefficient remains near zero throughout. Hashrate

is negative for both outcomes. Block size is positive and precisely estimated at -1 but small and insignificant at $+1$, consistent with a propagation advantage for smaller replacement blocks.

D.2 Hourly Reduced Form

Table 10: Hourly Reduced-Form Regressions

	$C_h(+1)$	$C_h(-1)$
CryptoShock_h	0.184^{**}	0.010
	(0.074)	(0.089)
$L.\overline{\log(\text{Hashrate})}_h$	-0.366^{***}	-0.513^{***}
	(0.045)	(0.053)
$L.\overline{\log(\text{Size})}_h$	0.772^{***}	0.997^{***}
	(0.031)	(0.038)
R^2	0.650	0.733
Month fixed effects	Yes	
Observations	61,967	

Notes: Each column regresses the hourly outcome directly on the excluded instruments and calendar-month fixed effects. No endogenous regressors appear in these reduced forms. Heteroskedasticity-robust standard errors are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 10 regresses the exact one-second outcomes directly on the excluded instruments. CryptoShock raises $C_h(+1)$ by 0.184 blocks per hour ($p = 0.013$) and has a near-zero association with $C_h(-1)$ (0.010, $p = 0.913$). This directional pattern appears before mapping the instruments into the three endogenous regressors.

D.3 CryptoShock Composition

Table 11 varies the composition of CryptoShock. Using hacks and market crises produces a $+1$ fee coefficient of 0.468 ($p = 0.005$); entering hacks/crises and regulation as separate instruments produces 0.467 ($p = 0.005$); and entering all four components separately produces 0.397 ($p = 0.011$). The corresponding -1 coefficients are statistically insignificant. Regulatory events alone do not provide a viable specification: the Kleibergen–Paap statistic falls to 0.10, so that row is reported as a weak-identification diagnostic and is not interpreted.

Table 11: Sensitivity to CryptoShock Composition: Hourly 2SLS

Shock specification	$C_h(+1)$	$C_h(-1)$	KP F
Pooled CryptoShock	0.520^{**}	-0.026	22.76
	(0.233)	(0.273)	
Hacks and market crises	0.468^{***}	0.226	32.89
	(0.168)	(0.180)	
Regulatory events only	-1.678	10.579	0.10
	(7.106)	(21.188)	
Hacks/crises and regulation separate	0.467^{***}	0.233	24.74
	(0.168)	(0.180)	
Four components separate	0.397^{**}	0.198	20.67
	(0.156)	(0.161)	
Month fixed effects	Yes		
Observations	61,967		

Notes: Cells report the coefficient on hourly mean log fees. Every specification also uses one-hour lagged mean log hashrate and log block size as excluded instruments and treats all three hourly regressors as endogenous. Pooled uses the union of hacks, market crises, and regulatory events. Two Groups enters hacks/crises and regulation separately. Four Components enters exact-timestamp hacks, market crises, positive regulatory events, and negative regulatory events separately. Heteroskedasticity-robust standard errors are in parentheses. The regulatory-only specification is weakly identified (KP $F = 0.10$) and is reported but not interpreted. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

D.4 Mining-Pool Composition and F2Pool Exclusion

Table 12 addresses whether the one-second result reflects changes in which mining pools produce canonical blocks. Adding the hourly canonical-block shares of the seven named pools leaves the $+1$ fee coefficient essentially unchanged at 0.521 ($p = 0.025$), while the -1 coefficient remains near zero. We also exclude every F2Pool-mined canonical block from the outcome counts because Yaish et al. (2023) document a distinct parent–child timestamp strategy associated with that pool. The $+1$ coefficient remains 0.531 ($p = 0.017$), and remains 0.533 ($p = 0.017$) when the composition controls are added simultaneously; the corresponding -1 coefficients are statistically insignificant. The result is therefore not attributable to hourly pool composition or to F2Pool's documented timestamp practices.

Table 12: Mining-Pool Composition and F2Pool Exclusion: Hourly 2SLS

Specification	$C_h(+1)$	$C_h(-1)$	KP F
Baseline	0.520^{**}	-0.026	22.76
	(0.233)	(0.273)	
Mining-pool composition controls	0.521^{**}	-0.027	22.82
	(0.233)	(0.273)	
Exclude F2Pool outcomes	0.531^{**}	0.179	22.76
	(0.223)	(0.251)	
Exclude F2Pool + composition controls	0.533^{**}	0.183	22.82
	(0.223)	(0.250)	
Month fixed effects		Yes	
Observations		61,967	

Notes: Cells report coefficients on hourly mean log fees; heteroskedasticity-robust standard errors are in parentheses. Every row uses the Main specification, jointly instrumenting hourly mean log fees, log hashrate, and log block size with $\mathrm{CryptoShock}_h$ and one-hour lags of mean log hashrate and log block size. Mining-pool composition controls are the hourly shares of canonical blocks mined by Ethermine, SparkPool, F2Pool, Nanopool, Hiveon, Flexpool, and 2Miners; Other is omitted. F2Pool-exclusion rows remove F2Pool-mined canonical blocks from the outcome counts while retaining network-wide regressors and instruments. All 13,884,675 block records match to a pool assignment. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.