Endogenous Misbehavior in Decentralized Settlement

Pablo D. Azar* Maryam Farboodi†

March 23, 2026

Abstract

A central design goal of settlement systems is that security should not degrade when transaction values are high. We show that Proof-of-Work blockchains 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, the precise timestamps of both the winning block and its displaced competitor. Under the ideal protocol, the earlier block should always win; instead, we observe cases where a block timestamped later displaces one timestamped earlier. Instrumenting fees with major hacks and network crises, we find that exogenous fee increases raise the probability of such reversals by 12.1 percentage points on a base rate of 14.3 percent. The effect is entirely absent for forks where the earlier block wins, and survives controls for network congestion and latency. The evidence suggests that incentive-based enforcement in Proof-of-Work degrades with the value of what the chain secures.

*Federal Reserve Bank of New York; pablo.azar@ny.frb.org. The views expressed here are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

†MIT Sloan School of Management; farboodi@mit.edu.

1 Introduction

Market systems have long been understood to require the support of a government and its rule of law to enforce contracts. This legal framework provides the ultimate guarantee for the finality of settlement, ensuring that once a high-value transaction is complete, it is irreversible. Decentralized blockchain protocols represent the most ambitious attempt to depart from this model, replacing legal enforcement with a mechanism that aligns the private incentives of anonymous, decentralized validators with the public goal of maintaining an immutable transaction record. Whether decentralized enforcement is robust—and in particular, whether it remains incentive-compatible as the value of the transactions it secures increases—is an open empirical question. If, under stress, miners find it profitable to fork the chain rather than extend it honestly, the security guarantee of the system unravels.

This paper tests empirically whether decentralized settlement finality is robust to increases in the value at stake. The question is not obvious. Persistently higher fees enlarge the network's security budget, attracting hashrate and raising difficulty—a force that strengthens the ledger's resilience over time. But a specific high-fee block also creates an immediate incentive to deviate: a miner who observes a competitor's valuable block can attempt to re-mine it and capture those fees, forking the chain in the process. Whether the long-run deterrence effect or the short-run deviation incentive dominates is an empirical question.

To answer this question, we exploit a unique feature of Ethereum—the second-largest blockchain and the largest blockchain that allows for smart contracts. During Ethereum's seven-year Proof-of-Work era, the blockchain publicly recorded valid blocks that lost the race for canonical inclusion, providing precise timestamps for both the winning block and its displaced competitor in each fork. When two or more valid blocks are mined at the same blockchain height, we call them *sibling blocks*—they share a common parent and compete for the same slot. Most Proof-of-Work blockchains, including Bitcoin, discard the losing sibling, leaving no public record of competing chains. Ethereum's protocol, by contrast, retained a record of these losing sibling: a block in the next genereation (a child block) can refer to both its parent and optionally, to its uncle block for additional reward. This provides an exceptional on-chain window into the resolution of the temporary forks that are central to our analysis.

When the protocol works as intended, the winner of this race is the sibling block with the earliest timestamp—instead, we frequently observe that the winner is the block with the second or third earliest timestamp. We call these *late canonical blocks* and use the *timing gap* between them and the blocks that should have won as a measure of deviation from honest mining. When a late canonical block enters the chain, the earlier, valid block it displaced becomes an uncle: any transactions in the earlier valid block that are not in the late canonical block are reversed and must re-enter the queue for confirmation. Each such event is a fork—a branching of the ledger that is resolved only by discarding one branch and the transactions it contains. This is the blockchain analogue of a settlement failure—a completed transaction being unwound after the fact.

A theoretical literature has established that miners can profit by deviating from honest behavior, intentionally forking the chain. The deviation most relevant to our analysis is *fee-motivated re-mining*: upon observing a competitor's newly broadcast block, a miner continues mining an alternative block at the same height, racing to propagate it before the network consolidates around the original. If successful, this produces a fork in which the *later*-timestamped block displaces the earlier one—precisely the "late canonical blocks" that are the focus of our analysis. Carlsten

et al. (2016) demonstrate that when block rewards decline and transaction fees become the primary source of miner revenue, the variance in block value creates strong incentives to fork "wealthy" blocks, making such re-mining profitable even for miners with very small hashrate shares. Bahrani et al. (2025) generalize these results, proving that the incentive for strategic deviation—and hence for deliberate forking—increases with the thickness of the reward distribution's right tail, precisely the conditions created by high-fee blocks. A related but distinct strategy is *selfish mining*: Eyal and Sirer (2014) show that a miner controlling a sufficient share of hashrate can profit by *withholding* a privately discovered block and later releasing it to orphan a competitor's work—a strategy that produces forks in which the attacker's *earlier*-mined block displaces an honest block that was timestamped later. An alternative explanation is that forks are simply due to network latency and congestion: because it takes several seconds to propagate a newly mined block across the global peer-to-peer network, two miners may independently solve the puzzle at nearly the same time, creating a fork that is resolved by chance rather than strategy. Decker and Wattenhofer (2013) document that block propagation delays in Bitcoin follow a heavy-tailed distribution, with median propagation times of roughly 6.5 seconds and a long right tail driven by network topology and block size. Such delays can generate forks even under fully honest mining.

These three explanations leave different marks in the timing gap $\Delta\tau$. Under fee-motivated re-mining, the attacker races to replace a competitor's valuable block, so a successful attack produces a canonical block timestamped *after* its sibling ($\Delta\tau > 0$), and the probability of this outcome should rise with fees. Under selfish mining, the attacker withholds a block mined earlier, so the canonical block is timestamped *before* its sibling ($\Delta\tau < 0$); the strategy also predicts consecutive blocks by the same miner. Under honest latency, propagation delays are orthogonal to block value, so fees should have no systematic effect on the direction of the timing gap. Our identification strategy is designed to tell these three stories apart.

1.1 Our Results in a Nutshell

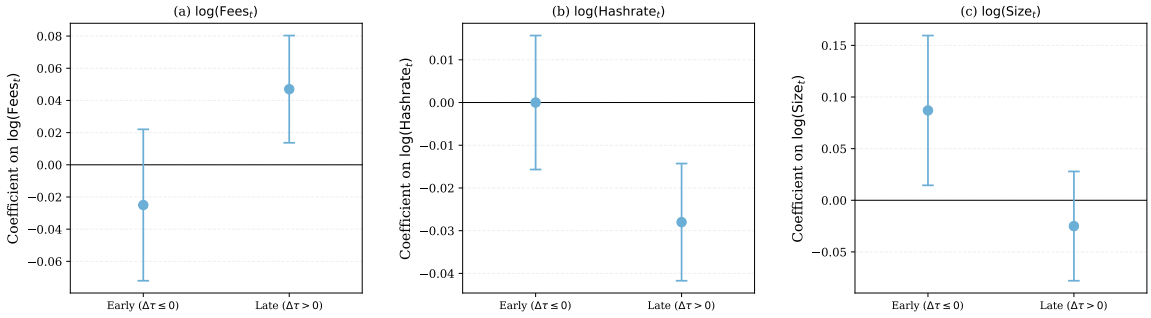


Figure 1: The central finding. Each panel plots the 2SLS coefficient (with 95% bootstrap confidence intervals) on log fees (a), log hashrate (b), and log block size (c), estimated separately for uncle events in which the canonical block is timestamped strictly earlier than its sibling (*Early*, $\Delta\tau < 0$) and events in which the canonical block is timestamped *later* than its sibling (*Late*, $\Delta\tau > 0$). All coefficients shown are treated as endogenous and instrumented as described in Section 3. A Heckman selection correction accounts for endogenous entry into the sibling sample.

Figure 1 summarises our main result by showing the IV coefficients for the probability of a late

canonical block (or an early canonical block) regressed on log fees, log hashrate, and log block size simultaneously. Using major hacks, network crises, and regulatory developments as instruments for fees, we find an asymmetric pattern across the two outcomes. Panel (a) shows the key finding: a one percent increase in fees increases the probability of a late canonical block by 0.045 percentage points, while its effect on early canonical blocks is negative and statistically indistinguishable from zero. The fee effect is one-sided: it operates on *Late* but not on *Early*. A miner attempting to displace a high-value block can succeed—producing a late canonical block—but cannot through this mechanism systematically alter the probability that an honestly-mined block arrives first.

Panels (b) and (c) sharpen the identification and rule out the main alternatives. Panel (b) shows that higher hashrate reduces the probability of late canonical blocks—consistent with deterrence, as a larger honest majority forecloses any private fork more quickly—while leaving early arrivals unaffected. Panel (c) addresses the congestion alternative: block size has opposite effects on the two outcomes, and both signs follow from the strategic re-mining mechanism. On the Early side, a large canonical block propagates slowly, so it can win the network race against a competing sibling only if it had a meaningful head start—that is, it must have been mined sufficiently earlier in absolute time to reach most nodes first despite its propagation disadvantage. Observing a large canonical block therefore selects for strictly negative timing. On the Late side, the effect is negative or negligible, because a strategic re-miner preemptively constructs a deliberately compact block around the target transaction so it propagates faster in any ensuing race; canonical blocks that result from successful re-mining are therefore systematically smaller, and a larger observed block size is evidence against strategic origin. If propagation congestion were driving the fee effect, size should push in the same direction as fees on both outcomes; it does not. The pattern—fees raising late canonical blocks but not early ones, hashrate reducing them, and block size with opposite effects on the two outcomes—is most naturally read as strategic re-mining. To our knowledge, this is the first causal evidence that the Proof-of-Work settlement guarantee weakens as the value of the transactions it secures rises.

1.2 Related Work

The theoretical case for fee-motivated re-mining was established by Carlsten et al. (2016) and extended by Bahrani et al. (2025); both predict that deviations become more profitable as the right tail of the fee distribution thickens. Biais et al. (2019) show that the blockchain protocol is a coordination game with multiple equilibria, several of which sustain persistent forks driven by information delays; their analysis implies that the timing patterns we study can arise even absent malice, which is precisely why distinguishing fee-driven from latency-driven forks is the central empirical challenge we face. Gans and Halaburda (2024) show that this problem is more severe than the standard Nakamoto condition implies: because the attacker earns block rewards while extending the attacking chain, the net cost of a majority attack can be negative, so high-fee blocks disproportionately attract attacks. Abadi and Brunnermeier (2022) show that free entry into decentralized chains drives rents toward zero, eliminating the franchise value that disciplines centralized record-keepers. A bank or clearinghouse has a long-run stake to protect; a PoW miner does not. Our findings are a consequence: without the threat of losing a franchise, miners' behavior is governed entirely by contemporaneous payoffs, which is why a single high-fee block can be enough to trigger a deviation.

The relevant attackers in our setting are organized mining pools. Cong et al. (2021) model

equilibrium pool formation in Proof-of-Work networks and show that cross-pool diversification disciplines pool growth: as a pool expands it charges higher fees, which induces miners to reallocate hash power to competitors, stabilising the industry at intermediate concentration. Lehar and Parlour (2020) document that pools nonetheless exercise strategic market power—estimating that withholding-based fee extraction exceeds $300 million per year—consistent with the strategic sophistication visible in our fork record.

Our instruments exploit the DeFi ecosystem's tendency to generate sharp, discrete fee spikes. Cong and He (2019) show that smart contracts expand the contracting space by eliminating informational frictions, enabling the complex DeFi transactions that generate these thick tails. Makarov and Schoar (2022) document the systemic fragilities of the DeFi ecosystem—cascading liquidations, oracle manipulation—that produce precisely the discrete fee spikes we use as instruments. Huberman et al. (2021) show that in a permissionless payment system the only way to raise revenue is to allow congestion, so fees and network load are jointly determined in equilibrium; this structural endogeneity is the identification challenge our IV strategy addresses. On the empirical side, Easley et al. (2019) show that Bitcoin miners allocate block space strategically to maximize fee revenue and that competition for inclusion intensifies as block rewards decline—direct evidence, in a Bitcoin setting, that miners respond to fee incentives in the manner our analysis documents for Ethereum.

Saleh (2021) shows that a modest reward schedule supports consensus in Proof-of-Stake while precluding the persistent fork equilibria that arise in Proof-of-Work; the contrast underscores that the fragility we document is specific to the incentive structure of Proof-of-Work, not a property of decentralized consensus per se.[1]

2 Institutional Details: The Economics of Strategic Re-Mining

The incentive structure of Proof-of-Work (PoW) mining is the central institutional feature that enables our analysis; while mining rewards are typically modest, the distribution of these rewards has a thick right tail, creating rare, lottery-like prizes that can incentivize strategic deviations from the protocol. In a PoW system, miners engage in a continuous, winner-take-all contest to solve a computational puzzle, with the victor earning the right to append a new block of transactions to the chain and claim a reward. This economic expenditure on computation is the fundamental source of the ledger's security, as reversing a recently confirmed block requires a costly computational attack.

This reward structure, however, creates an incentive for miners to deviate from honest behavior. The deviation we study is *fee-motivated re-mining*: upon observing a competitor's newly discovered block, a rational miner chooses not to accept it as the new chain head. Instead, the miner continues mining their own version of the same block—potentially reordering or replacing its transactions—and races to propagate it before the network adopts the original. If successful, this allows the re-miner to capture the transaction fees that would otherwise have gone to the original discoverer. Carlsten et al. (2016) first showed that transaction fees create incentives for block withholding even absent block rewards, and Bahrani et al. (2025) generalise this to show that the incentive for strategic deviations increases with the thickness of the reward distribution's right tail—precisely the conditions created by high-fee blocks. Our paper provides the first causal test of how this incentive

[1] Budish and Sunderam (2024) argue that when trust is anchored in legal institutions rather than Nakamoto consensus, the competitive disruption from blockchain adoption is limited; our findings reinforce the point by showing that Proof-of-Work's trustlessness has structural limits.

varies under stress, using the umbrella term "strategic re-mining" to encompass deviations that result in observable displacement of blocks timestamped earlier than the late canonical block. A related but distinct class of deviation 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, sacrificing immediate broadcast for a probabilistic advantage. Both strategies exploit the same underlying incentive structure, but selfish mining targets the block reward whereas fee-motivated re-mining targets the transaction fees—the channel that varies exogenously in our identification strategy.

Strategic re-mining is related to, but distinct from, the majority attacks studied by Budish (2025). Budish's framework focuses on the threat of a sustained 51% attack: an adversary who controls a majority of the network's hashrate can rewrite arbitrarily many confirmed blocks, enabling double-spending at any scale. The security of the ledger in that setting depends on whether the cost of acquiring majority hashrate exceeds the value of the coins that can be double-spent. Strategic re-mining, by contrast, requires no majority: a miner needs only to solve the current puzzle faster than the honest network in a single race, not to sustain a private chain indefinitely. The motive is also narrower—capturing the transaction fees in one block rather than reversing the ledger—and the attack leaves an observable trace in the uncle record rather than erasing history. Both mechanisms share the same root fragility identified by Budish (2025): the security of Proof-of-Work is purchased at a flow cost, and when the value of what the chain secures rises, so does the incentive to subvert it.

While the incentive to strategically re-mine blocks is a general feature of PoW, Ethereum's (Buterin, 2013) DeFi ecosystem makes it a particularly powerful laboratory for studying it. DeFi protocols generate exactly the thick-tailed fee distributions that Bahrani et al. (2025) identify as the key driver of strategic deviation: a single liquidation, flash loan, or arbitrage transaction can concentrate hundreds or thousands of dollars of fees into one block, creating the high-stakes races our identification strategy exploits. Our analysis focuses on Ethereum for this reason and because its protocol provides a unique empirical window into the outcomes of these contests. Most PoW blockchains, including Bitcoin, are designed to discard blocks that lose the mining race, leaving no public record of a competing chain. Ethereum's PoW consensus mechanism, however, was explicitly designed to address this issue by retaining a record of certain recently orphaned blocks, which it termed "uncles"—and which we call *siblings*, to emphasize their competition against the canonical block in the same generation rather than a genealogical relationship to it.

We leverage this unique observability to construct our primary dependent variable: the timing gap between a canonical block and its competitor sibling. When this gap is positive—meaning the canonical block was mined *after* its competitor—a valid, earlier block has been supplanted. The transactions in the displaced block that are not replicated in the winning block become unconfirmed and must re-enter the mempool; while many may eventually be re-included, their ordering, content, and confirmation can change. This canonical reorganization is the blockchain analogue of a settlement failure, and suggests a successful re-mining attack.

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.*

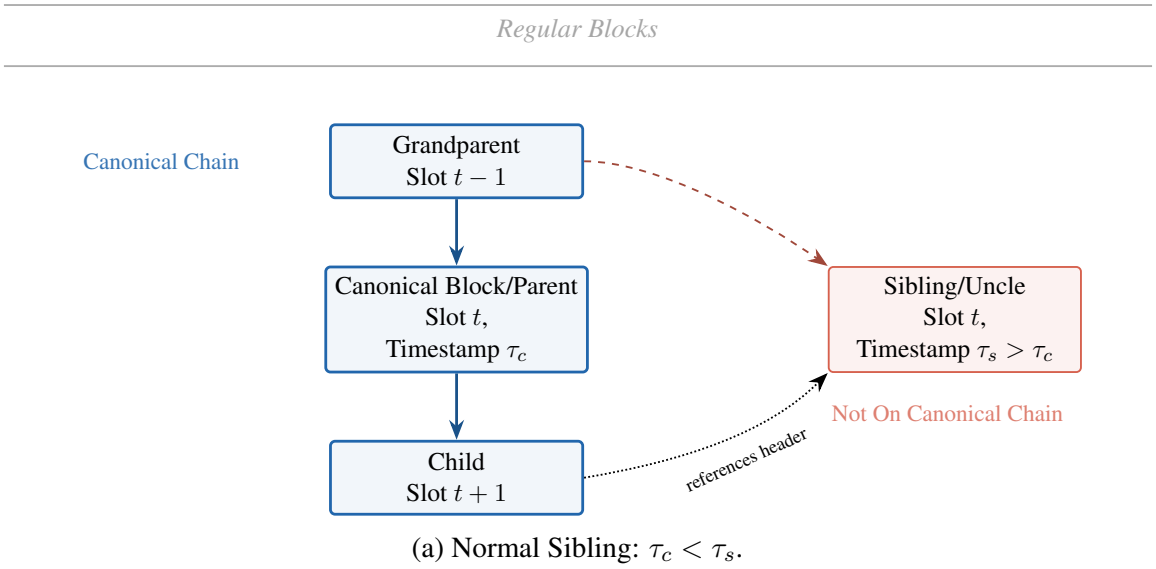
Definition 2. *A **late canonical block** is a canonical block whose timestamp is later than that of any sibling block at the same blockchain height ($\Delta\tau > 0$).*

Figure 2 provides a schematic of the mechanism that allows us to identify these settlement reversals. Panel (a) illustrates the normal inclusion of a sibling block, where the canonical block is

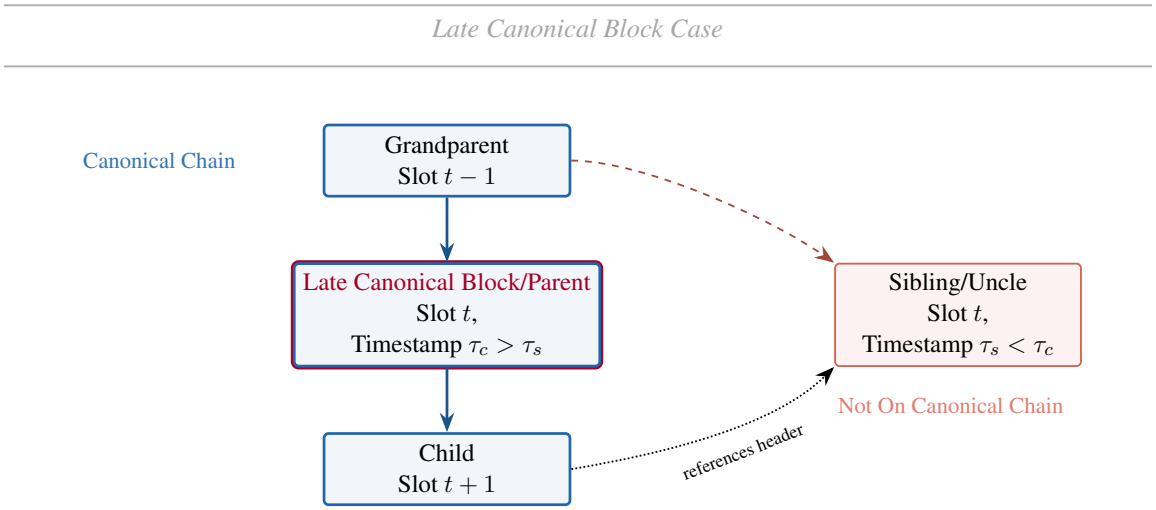
mined and recorded with an earlier timestamp (τ_c) than the competing sibling (τ_s). Panel (b) depicts the late canonical block case, which forms our primary measure of settlement reversals. In this scenario, the block that is ultimately included in the canonical chain has a later timestamp (τ_c) than its corresponding sibling (τ_s). This timestamp ordering is observable on-chain and constitutes direct evidence that an earlier, valid block was displaced by one that was timestamped later.

The Role of Block Size. A central prediction of the strategic re-mining hypothesis concerns block size. A rational re-miner who has observed a competitor's valuable block 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 before the original. Network propagation time is roughly proportional to block size—a 1 KB block traverses the peer-to-peer network substantially faster than a 50 KB block. The re-miner therefore has a strong incentive to construct a *small* block—prioritizing the high-value transactions that motivated the attack over low-fee filler—and broadcast immediately upon finding a solution. 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 deliberately constructs a small block—including only the high-value transactions that justify the attack—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.



(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 uncles/siblings and late canonical blocks. Blue nodes form the canonical chain; the rust node is a sibling block created 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 nephew's inclusion of the uncle header. Panel (a) shows the usual ordering $\tau_c < \tau_s$. Panel (b) highlights a *late canonical block* case with an amber outline, where the canonical parent is timestamped after its sibling ($\tau_c > \tau_s$).

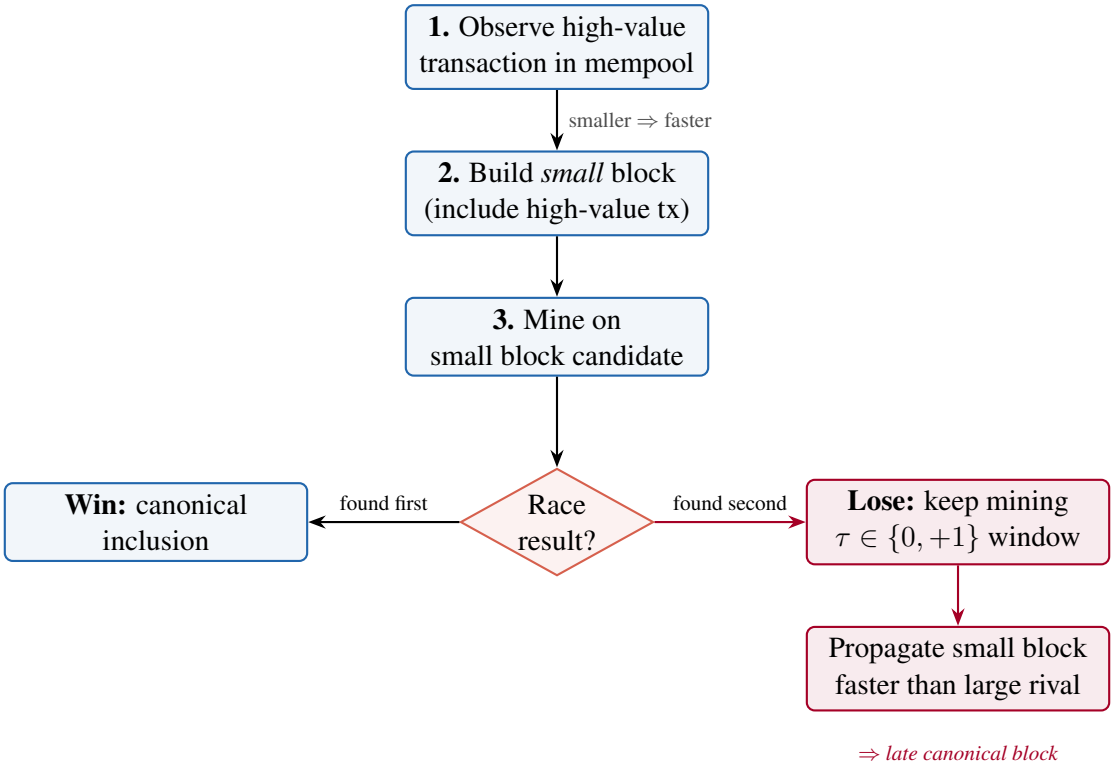


Figure 3: The strategic re-mining mechanism. A miner observing a high-value transaction in the mempool preemptively builds a deliberately small competing block (step 2) so that it propagates faster than any rival block in the event of a fork. If the race is lost initially, the miner continues mining for a few extra seconds (the $\Delta\tau \in \{0, +1\}$ window), exploiting the smaller block's propagation advantage to displace the canonical block.

3 Data

Our dataset covers the full Proof-of-Work history of the Ethereum blockchain, from its launch in August 2015 through the Proof-of-Stake transition in September 2022—spanning 2,592 days and approximately 14 million blocks. Block-level timestamps are sourced from Dune Analytics; sibling block timestamps from Blockchair. We merge these with daily network hashrate data and sibling frequency data, and construct the panel at the block level. Our key measure is the *timing gap* $\Delta\tau \equiv \tau_{\text{canonical}} - \tau_{\text{sibling}}$: the difference in seconds between the timestamp of the canonical block and that of its earliest sibling.[2] When $\Delta\tau > 0$, the canonical block arrived *after* the sibling, indicating that an earlier, valid block was displaced.

Sample Construction. The full estimation sample with where all key variables are not missing contains 13,879,763 blocks. Of these, 1,103,929 blocks (7.9%) have at least one associated sibling, forming the *sibling sample* on which we estimate the outcome equations.

Key Variables. Table 1 defines the variables used in our estimation.

Table 1: Variable Definitions

Variable	Notation	Definition
Dependent variables		
Late Canonical	$\text{Late}_t = \mathbf{1}[\Delta\tau > 0]$	Canonical block arrived after its sibling
Early Canonical	$\text{Early}_t = \mathbf{1}[\Delta\tau < 0]$	Canonical block arrived before its sibling
Timing bin	$\mathbf{1}[\Delta\tau = k]$	Timing gap equals exactly k seconds
Endogenous regressors		
Log fees	$\log(\text{Fees}_t)$	Log total block transaction fees (USD)
Log hashrate	$\log(\text{Hashrate}_t)$	Log network hashrate at block t
Log block size	$\log(\text{Size}_t)$	Log block size in bytes
Instruments		
Crypto shock indicator	CryptoShock_t	$= 1$ during major hack or network crisis
Lagged hashrate	$\text{LagHash}_{d(t)-1}$	Daily mean of $\log(\text{Hashrate})$, lagged one day
Lagged block size	$\text{LagSize}_{d(t)-1}$	Daily mean of $\log(\text{Size})$, lagged one day
Selection variables		
Lagged sibling rate	$\text{LagSiblingRate}_{d(t)-1}$	Daily sibling rate, lagged one day
Controls		
Inverse Mills ratio	IMR_t	From probit selection equation
Month FE	$\gamma_{m(t)}$	Calendar-month fixed effects

Notes: $\Delta\tau \equiv \tau_{\text{canonical}} - \tau_{\text{sibling}}$ is the timing gap in seconds between the canonical block and its sibling. Negative values indicate the canonical block arrived first; positive values indicate the sibling arrived first and was displaced.

[2] A block may have multiple siblings.

Let $d(t)$ denote the calendar day on which block t was mined. Ethereum produces approximately 7,200 blocks per day, so we define our lagged daily variables as averages over the preceding day's blocks:

$$\text{LagSiblingRate}_{d(t)-1} = \frac{1}{|\mathcal{B}_{d(t)-1}|} \sum_{i \in \mathcal{B}_{d(t)-1}} \mathbf{1}[\text{HasSibling}_i] \tag{1}$$

$$\text{LagHash}_{d(t)-1} = \frac{1}{|\mathcal{B}_{d(t)-1}|} \sum_{i \in \mathcal{B}_{d(t)-1}} \log(\text{Hashrate}_i) \tag{2}$$

$$\text{LagSize}_{d(t)-1} = \frac{1}{|\mathcal{B}_{d(t)-1}|} \sum_{i \in \mathcal{B}_{d(t)-1}} \log(\text{Size}_i) \tag{3}$$

where $\mathcal{B}_{d(t)-1}$ is the set of blocks mined on day $d(t)-1$. Table 2 reports summary statistics for both samples.

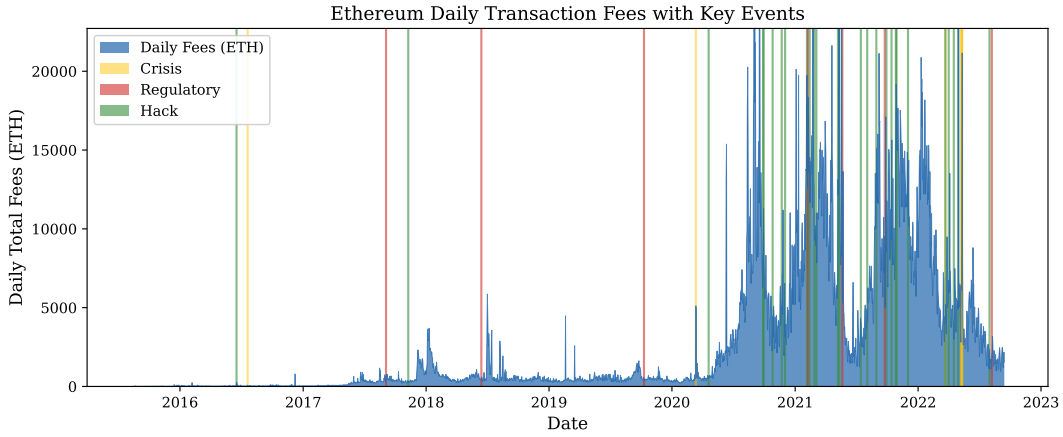


Figure 4: Daily total transaction fees (ETH) on the Ethereum blockchain, August 2015–September 2022. Yellow vertical lines indicate crisis events, red lines regulatory developments, and green lines major hacks. These events are used as instruments for transaction fee shocks in our IV strategy.

Figure 4 plots daily total transaction fees (in ETH) over the full Ethereum Proof-of-Work era, from genesis in August 2015 through the Merge in September 2022. Three features of the series are immediately apparent. First, fees are low and nearly flat through 2019, then explode during the DeFi boom of 2020–2021, consistent with a dramatic thickening of the right tail of the fee distribution documented by Bahrani et al. (2025) and Carlsten et al. (2016). Second, the series is highly volatile at daily frequency, with sharp spikes that are visibly discrete rather than part of smooth trends. Third, the highlighted events—crisis days (yellow), regulatory developments (red), and major hacks (green)—cluster around the largest of these spikes, which motivates their use as instruments: each event generates an abrupt, externally-driven surge in transaction demand that shifts fees upward on the day of the event and dissipates within days, satisfying the relevance condition. Because none of these events are plausibly related to the propagation or validation mechanics that govern which block wins a fork—only to the volume and value of transactions seeking inclusion—they satisfy the exclusion restriction conditional on the controls in the outcome equation.

Our primary dependent variable is Late_t, a binary indicator equal to one when the canonical block's timestamp exceeds that of its sibling—indicating that an earlier, valid block was displaced.

Table 2: Summary Statistics

	Moments		Distribution						
	Mean	**SD**	**Min**	**p5**	**p25**	**Median**	**p75**	**p95**	**Max**
Panel A: Full sample ($N = 13{,}884{,}675$)									
HasUncle_t	0.08	0.271	0	0	0	0	0	1	1
$\log(\text{Fees}_t)$	3.98	2.81	0	0.012	1.83	3.54	6.36	8.62	17
Fees_t (USD)	1,064	7,975	1	1.01	6.23	34.3	576	5,525	23,580,512
$\log(\text{Hashrate}_t)$	32.8	2.24	22.8	28	32.1	33.2	34.2	35.6	37.3
$\log(\text{Size}_t)$	9.69	1.48	6.25	6.64	8.98	10.2	10.6	11.7	14.2
CryptoShock_t	0.017	0.131	0	0	0	0	0	0	1
$\text{LagUncleRate}_{d(t)-1}$	0.086	0.05	0.038	0.045	0.053	0.066	0.096	0.195	0.437
$\text{LagHash}_{d(t)-1}$	32.8	2.02	25.3	28.3	32.8	33.2	34.1	34.9	35
Panel B: Uncle sample ($N = 1{,}104{,}251$)									
Late_t	0.143	0.35	0	0	0	0	0	1	1
Early_t	0.236	0.424	0	0	0	0	0	1	1
$\Delta\tau_t$ (sec.)	−0.325	2.81	−197	−2	0	0	0	1	107
$\log(\text{Fees}_t)$	3.73	2.53	0	0.014	2.04	3.45	5.36	8.39	13.6
Fees_t (USD)	786	3,938	1	1.01	7.66	31.4	214	4,415	810,382
$\log(\text{Hashrate}_t)$	32.6	2.13	22.8	28	32	33	34	35.4	37.3
$\log(\text{Size}_t)$	9.42	1.61	6.25	6.29	8.38	10	10.4	11.6	14.2
CryptoShock_t	0.013	0.114	0	0	0	0	0	0	1

Notes: Panel A reports statistics for the full estimation sample used in the probit selection equation. Panel B restricts to blocks with at least one uncle, used in the IV outcome regressions. $\text{Late}_t = \mathbf{1}[\Delta\tau > 0]$; $\text{Early}_t = \mathbf{1}[\Delta\tau < 0]$. Fees_t (USD) $\equiv \exp(\log(\text{Fees}_t))$. Sample: Ethereum PoW era (August 2015–September 2022).

The $23.6M maximum in Panel A is an outlier from an erroneous gas fee paid by a Bitfinex wallet in June 2021; the miner returned the funds (*Decrypt*, Jun. 10, 2021; Etherscan tx `0x2c993...`). This block is absent from the sibling sample (Panel B).

As a contrast, we also examine $\text{Early}_t = \mathbf{1}[\Delta\tau < 0]$, which captures the opposite event: the canonical block arriving first. We remain agnostic about observations where the timing gap is exactly zero seconds, as we do not know which block arrived earlier or later in these observations.[3]

For the distributional analysis, we decompose the timing gap into integer-second bins: $\mathbf{1}[\Delta\tau = k]$ for $k \in \{-5, \ldots, +5\}$. Negative values of k correspond to the canonical block arriving first

[3]To prevent the amount of storage needed to store the blockchain from expanding unecessarily, Ethereum blocks are timestamped at a granularity of one second, resulting in about 62% of the slots having the canonical and its sibling arriving at the same timestamp.

(consistent with honest mining), while positive values indicate a later canonical block displaced an earlier sibling (consistent with strategic re-mining).

We treat $\log(\text{Fees}_t)$, $\log(\text{Hashrate}_t)$, and $\log(\text{Size}_t)$ as endogenous. Fees are directly at stake in any re-mining attempt: a miner who displaces a competitor's block captures its transaction fees, making fees a direct input to the decision to strategically re-mine.

Hashrate is endogenous for two distinct reasons, both of which are central to the paper's interpretation. First, block production is a Poisson process: the expected number of blocks found per second is proportional to hashrate. When hashrate is high, the rate of simultaneous arrivals—two independent miners solving the puzzle in the same second—rises mechanically, shifting probability mass in the timing distribution toward $\Delta\tau = 0$. Any regression that fails to control for hashrate would conflate genuine fee-driven late arrivals with Poisson simultaneity.

The second reason is that hashrate governs the intensity of the competitive environment which deters strategic re-mining. When the network is mining intensively, an attacker who attempts to extend an alternative chain faces a rapidly closing window: other honest miners advance the canonical chain, and the expected time before the attacker's fork is overtaken shrinks with γ. This result is consistent with the theoretical predictions of Eyal and Sirer (2014), Carlsten et al. (2016), and Bahrani et al. (2025): the profitability of any selfish deviation collapses as the honest network's hashrate rises, because competing miners who extend the honest chain foreclose the attacker's opportunity. Both channels predict that hashrate is negatively correlated with the probability of a late canonical block, independently of any fee effect.

Block size is endogenous because it is a direct choice variable for the miner who constructs the block. A re-miner building a competing block controls which transactions to include, and the mechanism in Section 2 predicts that a strategic re-miner will construct a smaller block—prioritizing the high-value transactions that motivated the attack over low-fee filler. Block size therefore enters the outcome equation as a third endogenous regressor, with $\text{LagSize}_{d(t)-1}$ serving as its instrument: yesterday's average block size reflects network-wide demand conditions but is predetermined with respect to any individual block's construction today.

We instrument for the three endogenous regressors using three excluded instruments: CryptoShock_t, an indicator for major hack or crisis events that generates exogenous spikes in transaction demand; $\text{LagHash}_{d(t)-1}$, the lagged daily mean of log hashrate; and $\text{LagSize}_{d(t)-1}$, the lagged daily mean of log block size. Both lagged instruments are predetermined with respect to within-day block outcomes.

The Exclusion Restriction. The causal interpretation of our IV estimates requires that the excluded instruments affect the timing gap between competing blocks *only* through their effect on the endogenous regressors, conditional on the controls in the outcome equation. For CryptoShock_t, the concern is that DeFi hacks, major crises, and regulatory changes may also alter propagation conditions, transaction composition, or mining pool behavior in ways that directly affect sibling formation and block ordering. For the lagged instruments ($\text{LagHash}_{d(t)-1}$ and $\text{LagSize}_{d(t)-1}$), the exclusion restriction is that yesterday's network-wide averages are predetermined with respect to within-day block outcomes, conditional on month fixed effects.

Several features of our design address these concerns. First, the outcome equation conditions on month fixed effects, the inverse Mills ratio (which absorbs selection on observables including network conditions), $\log(\text{Hashrate}_t)$, and $\log(\text{Size}_t)$, which capture the primary channels through

which our instrumental variables might affect block production capacity and composition. Any crisis effect on the *speed* of mining is absorbed by the hashrate control; any effect on block composition is absorbed by the size control. Second, our dependent variable is defined within the sibling sample—blocks that have already produced a fork. The relevant question is not whether our instrumental variables affect if forks occur (this is handled by the selection equation), but whether the timing gap shifts conditional on a fork existing. Third, our results provide an internal consistency check: if crises, hacks or regulatory effects affected timing gaps through non-fee channels such as congestion-driven asymmetric propagation, we would expect the fee effect to appear in the *Early* regression as well (since congestion would shift the entire timing distribution, not just the probability of a late canonical block). The null result on *Early*, reported in Table 3, is difficult to reconcile with an omitted-variable story in which crises independently reshape propagation, since such reshaping would affect both directions symmetrically.

4 Empirical Strategy

4.1 Selection into the Sibling sample

Timing gaps are only defined for blocks that have at least one sibling, so our outcome regressions condition on a selected sample. If the probability of producing an sibling correlates with the same economic forces that drive strategic re-mining—as seems plausible, since high-fee blocks may attract more mining competition—then restricting attention to blocks that have siblings introduces selection bias. To correct for this, we (adopt) a Heckman (1979) two-step procedure.

Selection Equation. We estimate a probit on the full sample of 13.9 million blocks:

$$\Pr(\text{HasSibling}_t = 1) = \Phi\big(\delta_1 \,\text{LagSiblingRate}_{d(t)-1} + \delta_2 \,\text{LagHash}_{d(t)-1} + \delta_3 \,\text{LagSize}_{d(t)-1} + \delta_4 \,\text{CryptoShock}_t + \gamma_{m(t)}\big) \tag{4}$$

where $\gamma_{m(t)}$ denotes calendar-month fixed effects. The key exclusion restriction is $\text{LagSiblingRate}_{d(t)-1}$, the network-wide sibling rate on the preceding day.

Validity of the Exclusion Restriction. The lagged daily uncle rate satisfies two conditions for a valid exclusion restriction: relevance and excludability.

Relevance. Uncle production is highly persistent: days with congested network conditions—which increase propagation delays and the likelihood of competing blocks—tend to be followed by similar conditions.

Excludability. Yesterday's aggregate sibling rate reflects network-wide propagation conditions that determine whether a block at height t produces a competing sibling, but conditional on a fork having occurred, the previous day's network-wide sibling rate has no direct channel to affect the timing gap between two specific competing blocks today. The identifying assumption is that $\text{LagSiblingRate}_{d(t)-1}$ shifts the probability of *observing* a timing gap but not the gap's direction or magnitude, conditional on the controls included in the outcome equation. This is analogous to a standard Heckman selection model in which a lagged aggregate variable predicts selection but does not directly affect the outcome—for example, using the local unemployment rate to predict labor force participation in a wage equation (Heckman, 1979), where yesterday's labor market tightness

determines who we observe working but does not directly set any individual worker's wage, given contemporaneous controls.

From the probit, we compute the inverse Mills ratio $\text{IMR}_t = \phi(\mathbf{x}_t'\hat{\boldsymbol{\delta}}) \,/\, \Phi(\mathbf{x}_t'\hat{\boldsymbol{\delta}})$ for each block in the sibling sample. Intuitively, the IMR measures how "surprising" it is that block t produced an sibling, given its observable characteristics. Blocks with a low predicted probability of sibling occurrence (high IMR) are those for which unobservable factors must have been unusually favorable to forking. Including the IMR as a control in the outcome equation absorbs any correlation between these unobservable selection factors and the timing gap, preventing selection into the sibling sample from biasing the fee and hashrate coefficients.

4.2 Outcome IV Regression

On the sibling sample ($N = 1{,}103{,}929$), we estimate:

$$Y_t = \beta_1 \log(\text{Fees}_t) + \beta_2 \log(\text{Hashrate}_t) + \beta_3 \log(\text{Size}_t) + \beta_4 \,\text{IMR}_t + \gamma_{m(t)} + \varepsilon_t \tag{5}$$

where $\log(\text{Fees}_t)$, $\log(\text{Hashrate}_t)$, and $\log(\text{Size}_t)$ are endogenous, instrumented by CryptoShock_t, $\text{LagHash}_{d(t)-1}$, and $\text{LagSize}_{d(t)-1}$. The model is exactly identified with three endogenous regressors and three excluded instruments.

We also estimate the model on two alternative specifications for robustness. The first alternative specification only uses $Log(Fees_t)$ as an endogenous regressor, and does not include hashrate or size as regressors. This is to ensure that the effect on fees is not driven by some confounding coming from the other regressors. The second alternative specification recognizes that the early days of Ethereum had high volatility in hashrate and fees, with many outliers happening on these days. We do a trimmed regression that starts our observations on January 1 2017.

First-Stage Relevance. The joint partial robust F-statistic for the three excluded instruments (CryptoShock_t, $\text{LagHash}_{d(t)-1}$, and $\text{LagSize}_{d(t)-1}$) in the first-stage regressions far exceeds conventional weak-instrument thresholds for all three endogenous regressors.

Bootstrap Standard Errors. Because the IMR is a generated regressor from the first-step probit, conventional standard errors from the two-step procedure understate uncertainty by treating the selection correction as known. We address this with a nonparametric bootstrap that re-estimates the *entire* two-step procedure—probit, IMR construction, and IV regression—in each of 1,000 replications. This fully accounts for first-stage estimation uncertainty in the selection equation. All 1,000 replications converge successfully.

5 Results

Table 5 in the appendix reports the Heckman selection equations for all three specifications. The lagged daily sibling rate is a powerful predictor of sibling occurrence across all specifications, with the lagged hashrate and block size contributing significantly in the main and trimmed specifications.

Table 3 reports the main Heckman-IV estimates across all three specifications. The fee coefficient on *Late* is positive and significant across specifications. The corresponding coefficient on *Early*

Table 3: Heckman-IV Estimates

	Main		Fees-Only		Trimmed	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
	Late_t	Early_t	Late_t	Early_t	Late_t	Early_t
$\log(\text{Fees}_t)$	0.047^{***}	-0.025	0.045^{***}	-0.017	0.052^{**}	-0.003
	(0.017)	(0.024)	(0.015)	(0.020)	(0.022)	(0.045)
$\log(\text{Hashrate}_t)$	-0.028^{***}	-0.000	—	—	-0.012	0.007
	(0.007)	(0.008)			(0.020)	(0.037)
$\log(\text{Size}_t)$	-0.025	0.087^{**}	—	—	-0.040	0.015
	(0.027)	(0.037)			(0.028)	(0.058)
IMR_t	-0.030	-0.429^{***}	-0.035	-0.466^{***}	-0.013	-0.598^{***}
	(0.032)	(0.042)	(0.040)	(0.055)	(0.089)	(0.173)
Asymmetry test (Late − Early): $\log(\textit{Fees}_t)$						
$\Delta\hat{\beta}$	0.072^{**}		0.062^{**}		0.055	
	(0.031)		(0.027)		(0.050)	
Bootstrap reps			1,000			
Month FE			Yes			
Uncle N	1,104,251		1,104,251		962,000	

Notes: Two-step Heckman-IV estimates. Main and Trimmed: $\log(\text{Fees}_t)$, $\log(\text{Hashrate}_t)$, and $\log(\text{Size}_t)$ are endogenous, instrumented by CryptoShock_t, $\text{LagHash}_{d(t)-1}$, and $\text{LagSize}_{d(t)-1}$. Fees-Only: $\log(\text{Fees}_t)$ is endogenous, instrumented by CryptoShock_t; $\log(\text{Hashrate}_t)$ and $\log(\text{Size}_t)$ enter as controls. Trimmed restricts to 2017+. Bootstrap SEs from 1,000 replications.
$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

is negative and statistically insignificant in the fees-only and main specifications, confirming that fees shift probability mass asymmetrically toward late canonical blocks. The coefficient on $\log(\text{Hashrate}_t)$ is negative and significant for *Late*, consistent with the Poisson simultaneity channel and the competitive deterrence mechanism: higher hashrate compresses the timing distribution toward zero and makes strategic re-mining less profitable. The coefficient on $\log(\text{Size}_t)$ tells the clearest story on the *Early* side: it is positive and significant at 0.087 in the main specification, indicating that heavier canonical blocks are significantly more likely to have been found first—consistent with the mechanism that a large, slow-propagating block can only win the network race if it had a meaningful head start in mining time. Equivalently, heavy blocks are less likely to be the product of strategic re-mining, because a re-miner's compact block is precisely optimised to win the propagation race; the positive Early coefficient reflects selection against strategic origin. On the *Late* side the coefficient is -0.025, directionally consistent with the same mechanism but not statistically significant at conventional levels. The distributional analysis sharpens this picture: the size coefficient at $\Delta\tau = +1$ is -0.073 (significant at the 1% level) in the main specification, confirming that larger blocks are significantly less likely to be the late-arriving canonical block in a tight fork race—heavy blocks simply do not win the propagation sprint. In the Trimmed specification the coefficient at $\Delta\tau = +1$ is -0.038, negative and directionally consistent but not significant, providing weaker but corroborating evidence. The asymmetry test confirms the difference between Late and Early fee coefficients is statistically significant in the Main and Fees-Only specifications (differences of 0.072 and 0.062, each significant at the 5% level); in the Trimmed specification the difference is 0.055 (SE 0.050), which is directionally consistent but does not reach conventional significance levels.

Table 4: Distributional Event Study: Heckman-IV

	-5	-4	-3	-2	-1	0	$+1$	$+2$	$+3$	$+4$	$+5$
MAIN ($N = 1{,}104{,}251$)											
$\log(\text{Fees}_t)$	-0.003	-0.013^{***}	-0.007	-0.007	0.011	-0.023	0.061^{***}	-0.003	-0.003	-0.003^{*}	-0.002
	(0.003)	(0.004)	(0.006)	(0.009)	(0.019)	(0.027)	(0.016)	(0.007)	(0.004)	(0.002)	(0.001)
$\log(\text{Hashrate}_t)$	-0.002	-0.009^{***}	-0.020^{***}	-0.029^{***}	0.039^{***}	0.028^{***}	0.015^{**}	-0.015^{***}	-0.015^{***}	-0.006^{***}	-0.002^{***}
	(0.001)	(0.002)	(0.002)	(0.003)	(0.007)	(0.009)	(0.006)	(0.003)	(0.001)	(0.001)	(0.001)
$\log(\text{Size}_t)$	0.004	0.021^{***}	0.020^{**}	0.036^{**}	-0.001	-0.061	-0.073^{***}	0.030^{***}	0.007	0.007^{**}	0.003
	(0.005)	(0.006)	(0.009)	(0.014)	(0.030)	(0.042)	(0.025)	(0.010)	(0.006)	(0.003)	(0.002)
IMR_t	-0.036^{***}	-0.075^{***}	-0.113^{***}	-0.115^{***}	-0.000	0.459^{***}	0.115^{***}	-0.049^{***}	-0.046^{***}	-0.023^{***}	-0.011^{***}
	(0.006)	(0.008)	(0.011)	(0.017)	(0.035)	(0.048)	(0.029)	(0.012)	(0.007)	(0.004)	(0.003)
FEES-ONLY ($N = 1{,}104{,}251$)											
$\log(\text{Fees}_t)$	-0.002	-0.010^{***}	-0.006	-0.003	0.010	-0.028	0.054^{***}	-0.000	-0.002	-0.003^{*}	-0.002
	(0.003)	(0.004)	(0.005)	(0.008)	(0.016)	(0.023)	(0.014)	(0.006)	(0.003)	(0.002)	(0.001)
$\log(\text{Hashrate}_t)$	—	—	—	—	—	—	—	—	—	—	—
$\log(\text{Size}_t)$	—	—	—	—	—	—	—	—	—	—	—
IMR_t	-0.039^{***}	-0.089^{***}	-0.132^{***}	-0.146^{***}	0.021	0.501^{***}	0.152^{***}	-0.069^{***}	-0.057^{***}	-0.029^{***}	-0.014^{***}
	(0.008)	(0.010)	(0.014)	(0.022)	(0.044)	(0.063)	(0.038)	(0.015)	(0.009)	(0.005)	(0.003)
TRIMMED ($N = 962{,}000$)											
$\log(\text{Fees}_t)$	-0.004	-0.007	-0.010	-0.007	0.032	-0.048	0.048^{**}	0.008	-0.003	-0.000	-0.000
	(0.005)	(0.007)	(0.011)	(0.017)	(0.025)	(0.051)	(0.022)	(0.010)	(0.005)	(0.003)	(0.002)
$\log(\text{Hashrate}_t)$	0.000	0.000	-0.000	0.000	0.004	0.005	-0.006	-0.002	-0.002	-0.000	-0.000
	(0.004)	(0.006)	(0.009)	(0.014)	(0.022)	(0.043)	(0.019)	(0.009)	(0.005)	(0.002)	(0.002)
$\log(\text{Size}_t)$	0.003	0.006	0.009	0.008	-0.017	0.026	-0.038	-0.005	0.002	0.000	0.000
	(0.006)	(0.009)	(0.014)	(0.021)	(0.032)	(0.065)	(0.028)	(0.013)	(0.007)	(0.003)	(0.002)
IMR_t	-0.057^{***}	-0.110^{***}	-0.205^{***}	-0.266^{***}	0.141	0.611^{***}	0.217^{**}	-0.101^{**}	-0.083^{***}	-0.029^{***}	-0.013^{*}
	(0.019)	(0.029)	(0.043)	(0.066)	(0.098)	(0.196)	(0.086)	(0.041)	(0.021)	(0.011)	(0.007)

Notes: Each column is a separate Heckman-IV regression of $\mathbf{1}[\Delta\tau = k]$ on all regressors. Main and Trimmed: $\log(\text{Fees}_t)$, $\log(\text{Hashrate}_t)$, and $\log(\text{Size}_t)$ are endogenous, instrumented by CryptoShock_t, $\text{LagHash}_{d(t)-1}$, and $\text{LagSize}_{d(t)-1}$. Fees-Only: only $\log(\text{Fees}_t)$ endogenous; $\log(\text{Hashrate}_t)$ and $\log(\text{Size}_t)$ enter as controls (— = excluded from endogenous set). Bootstrap SEs from 1,000 replications. Green cell: positive and significant at $\Delta\tau = +1$. Red cell: negative and significant at $\Delta\tau = +1$.

$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

Table 4 reports the full distributional event study across all three specifications. The fee effect is sharply concentrated at $\Delta\tau = +1$ (green column), significant at the 1% level in Main and Fees-Only. At $\Delta\tau = 0$, the fee coefficient is negative across all specs, consistent with fee pressure shifting probability mass from simultaneous arrivals toward $\Delta\tau = +1$. No other positive-τ bin is significant for fees.

Magnitudes. To assess economic significance, note that fees are volatile, with a standard deviation of $\log(\text{Fees}_t)$ in the sibling sample of 2.53. A one-standard-deviation increase in log fees thus raises the probability of a late canonical block by $0.048 \times 2.53 \approx 12.1$ percentage points—a substantial effect on a base rate of 14.3%. In dollar terms, the median block in the sibling sample generates \$31 in fees; a move from the median to the 95th percentile (\$4,415) corresponds to a 4.9 log-point increase, implying a $0.048 \times 4.9 \approx 23.5$ percentage-point increase in the probability of a late canonical block. These are economically large effects, confirming that the security of the ledger degrades meaningfully when transaction values rise.

The coefficient on $\log(\text{Hashrate}_t)$ is negative and significant for *Late*, while the corresponding coefficient for *Early* is essentially zero (-0.000, SE 0.008) and statistically insignificant—consistent with the asymmetric prediction that deterrence and simultaneity compress *Late* arrivals without mechanically affecting honest early arrivals. To interpret this coefficient, note that our hashrate measure is constructed at the block level as $\log(\text{Hashrate}_t) = \log(\text{Difficulty}_t) - \log(\text{InterBlockTime}_t)$, where InterBlockTime_t is the elapsed seconds since the previous block. Because Ethereum's difficulty adjustment algorithm recalibrates difficulty every block to target a constant average block time, an increase in aggregate mining power does not raise the global block production rate; instead, it raises Difficulty_t, which the protocol adjusts upward precisely to keep block times stable. The two components of our hashrate measure therefore capture two economically distinct channels.

The first channel operates through difficulty. Higher $\log(\text{Difficulty}_t)$ reflects greater accumulated network hashing power, as recorded by the protocol's difficulty adjustment. This captures a *deterrence* effect: when difficulty is high, it signals that the honest network commands substantial computational resources, so an attacker attempting to extend a private fork faces a larger honest majority whose aggregate power is difficult to overcome.

The second channel operates through inter-block time. Because $\log(\text{InterBlockTime}_t)$ enters our hashrate measure with a minus sign, longer inter-block times—indicating a *decreased Poisson arrival rate*—reduce $\log(\text{Hashrate}_t)$. When the realized inter-block time is long relative to the difficulty-implied expectation, blocks arrive infrequently, lowering the probability of near-simultaneous arrivals and spreading the timing distribution outward. Conversely, periods of unusually short inter-block times concentrate probability mass at $\Delta\tau = 0$, as two independent miners are more likely to finish within the same second. Both channels predict the negative and significant coefficient in column (1) and the positive hashrate effect at $\Delta\tau = 0$ visible in Table **??**. The null in column (2) is consistent with both channels: deterrence and Poisson simultaneity act symmetrically on *whether* a fork occurs, not on the direction of the timing gap conditional on a fork having already occurred. The IMR coefficient is small and insignificant in both regressions, suggesting that selection into the sibling sample does not materially bias the outcome estimates. Bootstrap standard errors are close to the robust SEs throughout, confirming that the generated-regressor correction has a negligible impact on inference.

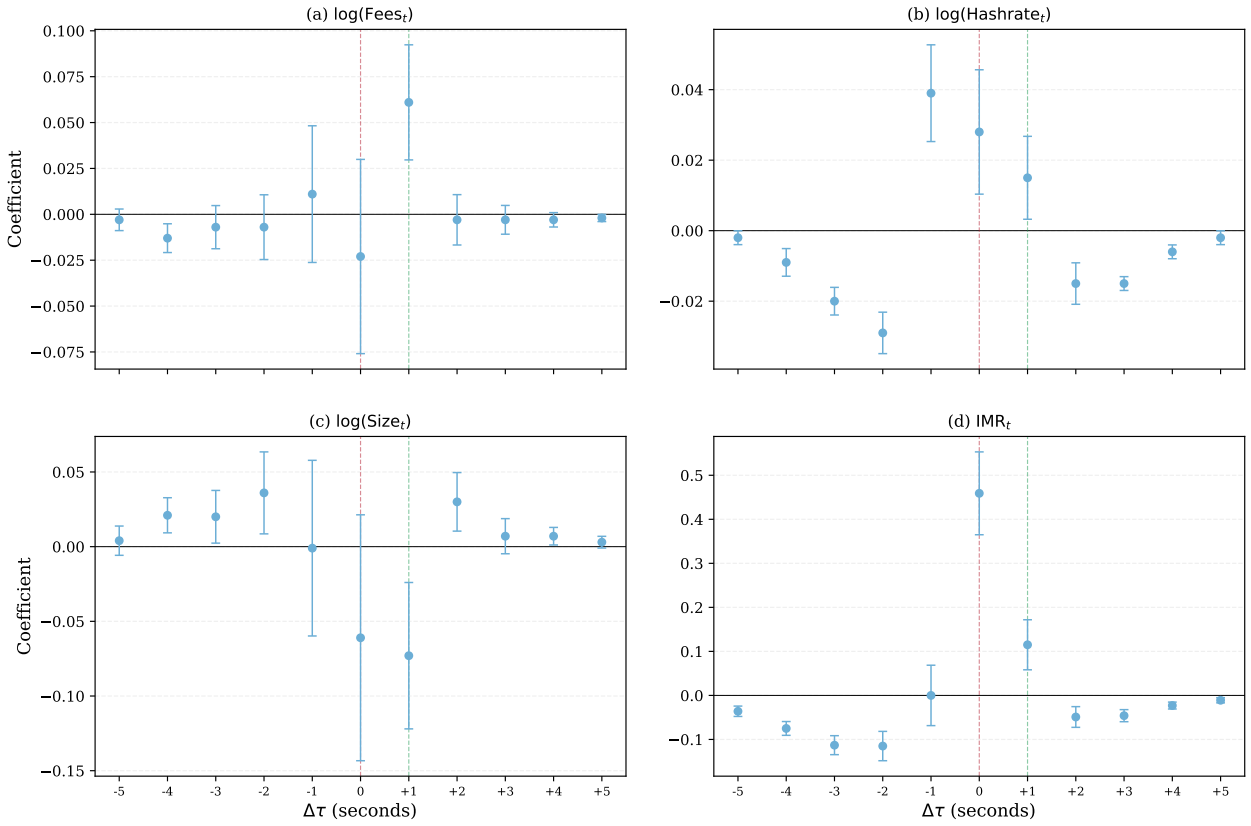


Figure 5: Heckman-IV distributional event study: all regressors. Each point is the 2SLS coefficient from a separate regression of $\mathbf{1}[\Delta\tau = k]$, $k \in \{-5, \ldots, +5\}$, on the indicated endogenous variable, instrumented by CryptoShock$_t$ and LagHash$_{d(t)-1}$, controlling for IMR$_t$ and month fixed effects. Error bars show 95% bootstrap confidence intervals (1,000 replications). Blue: positive and significant ($p < 0.05$). Red: negative and significant. Grey: insignificant. Dashed red line: $\Delta\tau = 0$; dashed green line: $\Delta\tau = +1$. Panel (a): $\log(\text{Fees}_t)$. Panel (b): $\log(\text{Hashrate}_t)$. Panel (c): $\log(\text{Size}_t)$. Panel (d): IMR$_t$.

Distributional Analysis Figure 5 plots the full distributional event study across all regressors. Our distributional analysis—a collection of 11 regressions where the dependent variable is $\mathbf{1}[\Delta\tau = k]$ for $k \in \{-5, \ldots, +5\}$, reported in Table **??**—reveals a sharply asymmetric response concentrated at a single timing gap. Higher fees significantly increase the probability that canonical blocks are timestamped exactly one second *after* their uncle competitors ($\Delta\tau = +1$), while the coefficient at $\Delta\tau = 0$ is -0.023 (SE 0.027), negative and directionally consistent with probability mass draining away from simultaneous timestamps but not statistically significant at conventional levels. The fee coefficient at $\Delta\tau = +1$ is the *only* positive coefficient significant at the 1% level; all other positive fee coefficients are insignificant even at the 10% level. Because timestamps are recorded with one-second resolution, $\Delta\tau = +1$ is the 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, chooses to continue mining rather than accepting the rival's block as the new chain head.

At the negative end, the fee coefficients at $\Delta\tau \in \{-5, -4\}$ are small and negative, significant at

the 5% level, while those at $\Delta\tau \in \{-2, -1\}$ are small and positive but insignificant. These effects offset, leaving the *aggregate* probability of the canonical block being timestamped first unaffected by fees and explaining the null on the Early contrast. The positive but insignificant coefficient at $\Delta\tau = -1$, equal to 0.014, is suggestive of heightened competitive intensity on both sides of the fork when fees are high, consistent with miners aware of a race in progress. Since the fee effect is asymmetric—with uniquely significant weight at $\Delta\tau = +1$—and we control for network congestion through hashrate, this pattern is evidence that high fees trigger strategic re-mining, which can succeed (producing $\Delta\tau = +1$) or fail (producing $\Delta\tau = -1$).

The hashrate coefficients in the distributional analysis are consistent with the two channels identified above. Higher hashrate significantly increases $P(\Delta\tau = 0)$, with a coefficient of 0.028, significant at the 1% level. This is consistent with the Poisson simultaneity channel: blocks found in periods of short inter-block time—indicating a fast local arrival rate relative to difficulty—are more likely to produce pairs of honest miners finishing within the same second, generating uncles with no timing gap. Higher hashrate also significantly reduces probability mass at $|\Delta\tau| \geq 2$, consistent with the deterrence channel: when difficulty is high, signaling substantial honest hashing power, the timing distribution compresses inward because both strategic deviation and large propagation-driven gaps become less likely. The hashrate coefficient at $\Delta\tau = +1$ is positive and significant at 0.015 in the main specification, indicating that periods of elevated honest hashing activity also raise the probability of late canonical blocks at the minimum gap. Critically, however, this coefficient is unstable across robustness checks: in the Trimmed specification (2017 onward) it reverses to -0.006 and becomes statistically indistinguishable from zero. The fee coefficient, by contrast, remains positive, significant, and of similar magnitude across all three specifications (0.061, 0.054, and 0.048, each significant at conventional levels), never reversing sign. Because hashrate is treated as endogenous and separately instrumented throughout, its coefficient is identified orthogonally to fees; the sign instability of the hashrate effect, relative to the robustness of the fee effect, confirms that fees—not correlated hashrate movements—are the operative force at $\Delta\tau = +1$.

Even at institutionally small scales, these incentives are already large enough to sustain misbehavior. In our data, the 75th percentile of total block fees is only $213, the 90th is $1,908, and the 95th exceeds $4,400. Even this level of upside is enough to make strategic re-mining a rational response. If Proof-of-Work were to become the dominant settlement infrastructure for finance, per-block fees would likely rise substantially, and these incentives would become dramatically stronger, potentially threatening the stability of the chain.

Ruling Out Non-Strategic Explanations. Several features of the evidence jointly rule out the leading non-strategic alternatives. *Congestion-driven propagation delays* would affect timing gaps symmetrically—larger blocks during high-fee periods might propagate more slowly, but this would shift probability mass toward both positive and negative timing gaps, not exclusively toward $\Delta\tau = +1$. The null on Early rules out this channel. *Timestamp manipulation unrelated to re-mining*—such as NTP drift or rounding conventions—would produce a level effect on $\mathbf{1}[\Delta\tau = 1]$ that is constant across fee levels, not a reduced-form response to crisis instruments. *Hashrate-mediated confounding*—in which crises change mining capacity and thereby timing gaps—is controlled for by treating $\log(\text{Hashrate}_t)$ as endogenous and separately instrumenting it with lagged network difficulty. The hashrate coefficient at $\Delta\tau = +1$ is 0.015 in the main specification but flips to -0.006 (insignificant) in the Trimmed robustness check. The fee coefficient, by contrast,

is the only coefficient at $\Delta\tau = +1$ that is consistently positive and statistically significant across all three specifications, never reversing sign. This combination—endogenous control for hashrate, sign instability of hashrate's own coefficient, and robustness of the fee effect—rules out hashrate-mediated confounding as an alternative explanation.

6 Conclusion

Our findings make several contributions. First, we provide a direct empirical test of fee-motivated strategic re-mining—the class of deviations in which miners observe and displace a competitor's high-value block to capture its transaction fees. This behavior is closely related to the selfish mining vulnerability identified by Eyal and Sirer (2014) and is predicted by Bahrani et al. (2025), who show that the incentive for strategic deviation increases with the thickness of the reward distribution's right tail. By demonstrating that miners respond to high-value blocks by attempting to re-mine them, our results move this foundational debate from theory to causal evidence.

Second, we contribute to the literature on financial market infrastructure and systemic risk. Modern payment and settlement systems are designed to be bastions of stability, engineered to prevent—not encourage—failures during periods of market stress. Our findings show that PoW operates on the opposite principle, introducing a novel channel for systemic risk where the integrity of the infrastructure is inversely related to its usage under duress.

Finally, our results speak to a long-standing question in institutional economics: what are the ultimate sources of trust and enforcement in market economies? Proponents have framed decentralized blockchains as a new institutional technology capable of trustless enforcement, substituting economic incentives for the role of courts and legal systems. We show empirical evidence justifying the classic view that state-backed legal institutions are crucial for enforcing high-value contracts and property rights. The incentive structure of PoW, we find, can perversely incentivize strategic deviation precisely when the rewards are highest, suggesting that the problem of securing high-value settlement has not been resolved by the technology—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.

BÖHME, R., N. CHRISTIN, B. EDELMAN, AND T. MOORE (2015): "Bitcoin: Economics, Technology, and Governance," *Journal of Economic Perspectives*, 29(2), 213–238.

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.

BUTERIN, V. (2013): "A Next-Generation Smart Contract and Decentralized Application Platform," Ethereum Whitepaper.

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.

DECKER, C., AND R. WATTENHOFER (2013): "Information Propagation in the Bitcoin Network," in *IEEE P2P 2013 Proceedings*, pp. 1–10. IEEE.

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.

HALABURDA, H., G. HAERINGER, J. GANS, AND N. GANDAL (2022): "The Microeconomics of Cryptocurrencies," *Journal of Economic Literature*, 60(3), 971–1013.

HALABURDA, H., G. HAERINGER, AND D. OBERMEIER (2025): "Blockchain Toolbox for Digital Platforms," Forthcoming in *Cambridge University Press Handbook on Digital Platforms*, SSRN Working Paper No. 4999550.

HECKMAN, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47(1), 153–161.

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.

LIU, Y., AND A. TSYVINSKI (2021): "Risks and Returns of Cryptocurrency," *Review of Financial Studies*, 34(6), 2689–2727.

LIU, Y., A. TSYVINSKI, AND X. WU (2022): "Common Risk Factors in Cryptocurrency," *Journal of Finance*, 77(2), 1133–1177.

MAKAROV, I., AND A. SCHOAR (2020): "Trading and Arbitrage in Cryptocurrency Markets," *Journal of Financial Economics*, 135(2), 293–319.

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.

SCHILLING, L., AND H. UHLIG (2019): "Some Simple Bitcoin Economics," *Journal of Monetary Economics*, 106, 16–26.

STOCK, J. H., AND M. YOGO (2005): "Testing for Weak Instruments in Linear IV Regression," in *Identification and Inference for Econometric Models*, ed. by D. W. K. Andrews and J. H. Stock, pp. 80–108. Cambridge University Press.

Appendix

A Heckman Selection Equations

Table 5: Heckman Selection Equations

	(1) Main	**(2)** Fees-Only	**(3)** Trimmed
$\text{LagUncleRate}_{d(t)-1}$	2.511^{***}	2.470^{***}	3.094^{***}
	(0.031)	(0.026)	(0.046)
CryptoShock_t	0.022^{***}	0.022^{***}	0.008^{*}
	(0.004)	(0.004)	(0.004)
$\text{LagHash}_{d(t)-1}$	0.069^{***}		0.022^{**}
	(0.008)		(0.010)
$\log(\text{LagSize})_{d(t)-1}$	0.045^{***}		0.011^{**}
	(0.004)		(0.005)
N (Full)	13,896,162	13,896,162	12,094,132

Notes: Selection equations estimated by Heckman two-step probit with default SEs. Dependent variable: HasUncle_t. All specifications include month FE. Column (1): fees-only specification (instruments: CryptoShock_t). Column (2): main specification (instruments: CryptoShock_t and $\text{LagHash}_{d(t)-1}$). Column (3): trimmed specification (sample from 2017 onward).
$^{***}\ p < 0.01$, $^{**}\ p < 0.05$, $^{*}\ p < 0.10$.

Table 5 reports the probit selection equations for all three specifications. The lagged daily uncle rate is the primary exclusion restriction and is strongly significant across all columns.

B IV First-Stage Regressions

Table 6 reports the first-stage regressions underlying our 2SLS estimates. Each column regresses one endogenous variable on the excluded instruments and all exogenous controls (IMR and month fixed effects), estimated by OLS on the uncle sample with robust standard errors. No bootstrap is required.

Column (1) reports the fees-only specification, in which CryptoShock_t is the sole excluded instrument; the partial F of 496 far exceeds the Stock and Yogo (2005) threshold. Columns (2)–(4) report the main specification with three endogenous regressors and three excluded instruments; all partial F-statistics exceed 1,300, confirming strong instrument relevance for each endogenous regressor. The instruments display economically intuitive patterns: CryptoShock_t raises fees (0.199) while slightly reducing hashrate (-0.023), consistent with crisis-driven demand surges; $\text{LagHash}_{d(t)-1}$ is the dominant predictor of contemporaneous hashrate (0.889), reflecting the high persistence of mining capacity; and $\log(\text{LagSize})_{d(t)-1}$ strongly predicts both fees (0.923) and size (0.602).

Table 6: IV First-Stage Regressions

	Fees Only	*Main Specification*		
	(1) Fees	(2) Fees	(3) Hash	(4) Size
CryptoShock_t	0.182^{***}	0.199^{***}	-0.023^{***}	0.030^{***}
	(0.008)	(0.008)	(0.008)	(0.006)
$\text{LagHash}_{d(t)-1}$		0.323^{***}	0.889^{***}	0.128^{***}
		(0.010)	(0.014)	(0.012)
$\log(\text{LagSize})_{d(t)-1}$		0.923^{***}	-0.001	0.602^{***}
		(0.007)	(0.008)	(0.006)
$\hat{\lambda}$ (IMR)	-2.605^{***}	-2.026^{***}	-0.211^{***}	-0.196^{***}
	(0.016)	(0.019)	(0.025)	(0.021)
R^2	0.907	0.909	0.811	0.731
Excl. instr. F	496	7,694	1,307	3,312
Uncle N	1,104,251		1,104,251	

Notes: Each column regresses one endogenous variable on the excluded instruments and all exogenous controls (IMR and month FE). OLS on the uncle sample with robust standard errors. Column (1): fees-only specification with CryptoShock_t as the sole excluded instrument. Columns (2)–(4): main specification with three excluded instruments (CryptoShock_t, $\text{LagHash}_{d(t)-1}$, $\log(\text{LagSize})_{d(t)-1}$). Excl. instr. F is the partial F-statistic from a joint test of the excluded instruments. $^{***}\,p < 0.01$, $^{**}\,p < 0.05$, $^{*}\,p < 0.10$.

C OLS Main Results

This section reports the Heckman-corrected OLS estimates for the main specifications. Because fees are endogenous, OLS estimates are subject to attenuation bias; the IV estimates in the body of the paper are the preferred causal specification. The OLS results are included here for completeness. As shown below, the OLS and IV distributional patterns are qualitatively distinct in a revealing way: OLS captures a broad, symmetric dispersion of probability mass away from $\Delta\tau = 0$ that is consistent with generic network congestion, while IV isolates a sharp, asymmetric concentration at $\Delta\tau = +1$ that is the causal signature of strategic re-mining. This contrast motivates the IV strategy and reinforces the interpretation of the instrumented estimates.

C.1 OLS Baseline and Contrast

Table 7 reports the Heckman-corrected OLS estimates across all three specifications. The direction of bias differs between the two outcomes, reflecting distinct omitted-variable mechanisms. For *Late*, the OLS fee coefficient (0.002 in the main specification) is attenuated toward zero relative to the IV estimate (0.047): fees are positively correlated with unobserved network congestion, which independently reduces the probability of late canonical blocks by raising honest hashing intensity, biasing OLS downward. For *Early*, the pattern is the opposite: the OLS fee coefficient is positive and significant (0.005) even though the IV estimate is negative (-0.025). This upward bias reflects positive omitted-variable bias from congestion—high-fee periods coincide with congested network

Table 7: Heckman-OLS Estimates

	Main		Fees-Only		Trimmed	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
	Late_t	Early_t	Late_t	Early_t	Late_t	Early_t
$\log(\text{Fees}_t)$	0.002^{***}	0.005^{***}	0.005^{***}	0.010^{***}	0.002^{***}	0.004^{***}
	(0.001)	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
$\log(\text{Hashrate}_t)$	-0.017^{***}	0.011^{***}			-0.018^{***}	0.008^{***}
	(0.000)	(0.000)			(0.000)	(0.000)
$\log(\text{Size}_t)$	-0.004^{***}	0.014^{***}			-0.006^{***}	0.010^{***}
	(0.001)	(0.001)			(0.001)	(0.001)
IMR_t	-0.158^{***}	-0.428^{***}	-0.151^{***}	-0.421^{***}	-0.184^{***}	-0.573^{***}
	(0.008)	(0.010)	(0.008)	(0.010)	(0.009)	(0.013)
Month FE			Yes			
Uncle N	1,208,657		1,208,657		988,572	

Notes: OLS estimates with Heckman selection correction. No instrumentation. IMR_t from first-step probit included as control. Robust SEs in parentheses. Month FE included.
$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

conditions that increase the probability of any fork, including honest early arrivals, pushing the OLS estimate above its true causal value and flipping its sign. The contrast between attenuation on *Late* and upward bias on *Early* is itself consistent with the strategic re-mining interpretation: congestion is a confounder that runs in the opposite direction to the strategic channel for each outcome, and instrumenting for fees removes it.

C.2 OLS Distributional Event Study

Table 8: Distributional Event Study: Heckman-OLS

	-5	-4	-3	-2	-1	0	$+1$	$+2$	$+3$	$+4$	$+5$
MAIN ($N = 1{,}208{,}657$)											
$\log(\text{Fees}_t)$	0.000^{***}	0.001^{***}	0.001^{***}	0.003^{***}	0.000	-0.007^{***}	0.000	0.002^{***}	0.001^{***}	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\log(\text{Hashrate}_t)$	0.001^{***}	0.001^{***}	0.001^{***}	0.002^{***}	0.001^{***}	0.007^{***}	-0.010^{***}	-0.003^{***}	-0.001^{***}	0.000^{***}	0.000^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\log(\text{Size}_t)$	0.000^{***}	0.001^{***}	0.002^{***}	0.003^{***}	0.005^{***}	-0.010^{***}	-0.003^{***}	-0.001^{***}	0.000^{**}	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
IMR_t	-0.033^{***}	-0.061^{***}	-0.119^{***}	-0.135^{***}	-0.011	0.586^{***}	0.005	-0.066^{***}	-0.051^{***}	-0.022^{***}	-0.010^{***}
	(0.001)	(0.002)	(0.003)	(0.005)	(0.008)	(0.012)	(0.007)	(0.003)	(0.002)	(0.001)	(0.001)
FEES-ONLY ($N = 1{,}208{,}657$)											
$\log(\text{Fees}_t)$	0.000^{***}	0.001^{***}	0.001^{***}	0.004^{***}	0.003^{***}	-0.015^{***}	0.001^{***}	0.003^{***}	0.001^{***}	0.000^{***}	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\log(\text{Hashrate}_t)$	—	—	—	—	—	—	—	—	—	—	—
$\log(\text{Size}_t)$	—	—	—	—	—	—	—	—	—	—	—
IMR_t	-0.034^{***}	-0.063^{***}	-0.121^{***}	-0.135^{***}	0.000	0.572^{***}	0.014^{**}	-0.064^{***}	-0.053^{***}	-0.022^{***}	-0.010^{***}
	(0.001)	(0.002)	(0.003)	(0.005)	(0.008)	(0.012)	(0.007)	(0.003)	(0.002)	(0.001)	(0.001)
TRIMMED ($N = 988{,}572$)											
$\log(\text{Fees}_t)$	0.000	0.000^{***}	0.000	0.001^{***}	0.001	-0.006^{***}	0.001	0.001^{***}	0.000^{***}	0.000	0.000^{*}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
$\log(\text{Hashrate}_t)$	0.001^{***}	0.001^{***}	0.001^{***}	0.001^{***}	0.000	0.010^{***}	-0.012^{***}	-0.003^{***}	-0.001^{***}	0.000^{***}	0.000^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\log(\text{Size}_t)$	0.000^{***}	0.001^{***}	0.002^{***}	0.002^{***}	0.004^{***}	-0.003^{***}	-0.005^{***}	-0.001^{**}	0.000^{***}	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
IMR_t	-0.044^{***}	-0.083^{***}	-0.169^{***}	-0.235^{***}	0.032^{***}	0.757^{***}	0.053^{***}	-0.123^{***}	-0.069^{***}	-0.027^{***}	-0.011^{***}
	(0.002)	(0.002)	(0.004)	(0.006)	(0.009)	(0.015)	(0.008)	(0.004)	(0.002)	(0.001)	(0.001)

Notes: Each column is a separate Heckman-OLS regression of $\mathbf{1}[\Delta\tau = k]$ on all regressors. No instrumentation. Main and Trimmed: all four regressors included. Fees-Only: only fees and IMR (— = excluded). Robust SEs in parentheses.
$^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

Table 8 reports the OLS distributional event study across specifications. The OLS pattern differs from the IV results in a diagnostically important way. In the Main specification, the fee coefficient at $\Delta\tau = 0$ is significantly negative (-0.007, $p < 0.01$), indicating that higher fees reduce the probability of simultaneous arrivals. However, the coefficient at $\Delta\tau = +1$ is precisely zero and statistically insignificant, and significant positive coefficients appear symmetrically at both positive and negative timing gaps ($\Delta\tau \in \{\pm 2, \pm 3, \pm 4, \pm 5\}$). This symmetric dispersion pattern—probability mass draining from $\Delta\tau = 0$ and spreading evenly in both directions—is the fingerprint of network congestion: high fees coincide with congested periods in which propagation delays push competing blocks apart in time, symmetrically. It is not the signature of strategic re-mining, which would load exclusively on $\Delta\tau = +1$. The IV results, by isolating fee variation driven by exogenous crisis instruments, strip away this congestion confound and recover the asymmetric $\Delta\tau = +1$ spike that is the causal evidence of strategic behavior. OLS coefficients are attenuated relative to IV throughout, as expected.