

Inside Qubic’s Selfish Mining Campaign on Monero: Evidence, Tactics, and Limits

Suhyeon Lee^{*†} and Hyeongyeong Kim[‡]

^{*}Tokamak Network [†]Hashed Open Research [‡]Korea University

Abstract—We analyze Qubic’s advertised selfish mining campaign on Monero in 2025. Combining data from Monero nodes, and the Qubic pool API, we reconstruct Qubic-attributed blocks and hashrate and detect ten intervals consistent with selfish mining strategies. In these intervals, Qubic’s average hashrate share rises to the 23-34% range, yet sustained 51% control is never observed. We evaluate the campaign against the classical selfish mining model and a modified Markov-chain model that reflects Qubic’s conservative release strategy: both predict lower revenue than honest mining at the inferred parameters, and the data largely confirms this while still showing noticeable deviations from the predicted curve. We interpret this gap between model and measurements in terms of Qubic’s time-varying hashrate and coarse-grained attack segmentation.

Index Terms—Blockchain, Proof-of-Work, Selfish Mining

I. INTRODUCTION

Monero is a privacy-focused cryptocurrency whose security critically relies on the assumption that miners follow the longest-chain rule and behave approximately honestly. Among the known deviations from this assumption, selfish mining is particularly concerning because it can increase a miner’s relative revenue and degrade confirmation reliability without requiring a strict majority of the total hashrate. While selfish mining has been extensively analyzed in theory, there is limited empirical understanding of how such strategies manifest in a deployed large proof-of-work system.

In August 2025, the Monero network faced an unprecedented stress test when the Qubic mining pool aggressively expanded its hashrate, explicitly marketing its actions as a “51% takeover” and a demonstration of selfish mining capabilities [1], [2]. While selfish mining has long been theoretically established as a vulnerability of Proof-of-Work systems [3], [4], Qubic’s campaign represented a rare, publicized deviation from honest protocol adherence in a major cryptocurrency. This incident forced the community to confront a critical gap between theory and practice: determining whether Qubic’s behavior was a profitable execution of optimal strategies or merely a destabilizing signal. However, lacking granular data beyond public dashboards, the ecosystem struggled to quantify the attack’s true mechanics and impact.

Studying this incident is technically challenging. Monero’s design, including privacy-preserving transaction mechanisms and the lack of explicit pool-identifying markers in blocks, complicates reliable attribution of blocks to specific entities. Moreover, Qubic’s strategy, if selfish, would involve selectively publishing blocks and exploiting network propagation effects that are not directly observable. Any empirical evaluation therefore requires carefully combining on-chain data, pool-level statistics, and timing information to reconstruct Qubic-attributed blocks, approximate its effective hash power, and infer its deviation from honest mining behavior.

In this work, we perform such an empirical investigation of Qubic’s mining activity on Monero including the alleged attack period. By operating a Monero pruning node, collecting mining job information from the Qubic pool API, and reconstructing the timeline of Qubic-attributed blocks, we derive an empirical tie-breaking parameter and evaluate Qubic’s strategy within an adapted selfish mining framework. This allows us to quantify both the revenue implications for Qubic and the induced instability on the Monero chain.

Our contributions are as follows:

- Construction of an empirical dataset: We built a robust dataset by combining data from a local Monero node and live job notifications from the Qubic pool API. We publicly release this dataset to support future empirical research on mining attacks.
- Analysis of strategic constraints and modeling: We observed that the attacker did not follow the standard selfish mining strictly, likely to mitigate real-world network latency. We formalized this behavior into a modified selfish mining model, utilizing it to establish a theoretical lower bound for the attacker’s expected revenue.
- Empirical refutation of strategic utility: We demonstrate that the deployed strategy was economically ineffective, contrary to the context promoted by the pool. Our analysis confirms that the attacker’s observed revenue consistently underperformed the expected baseline of honest mining in the majority of the periods.

The rest of the paper proceeds as follows. Section II

details our data collection methodology, using a pruning node, a Qubic pool miner, and identification heuristic. Section III measures the pool’s mining power share and orphan-related data on the Monero network. Section IV investigates Qubic’s selfish-mining strategy and its profitability. Section V attempts to explain the discrepancy between the theoretic expectation and the observed Qubic’s revenue. Section VI discusses the mitigations and Section VII concludes the paper.

II. DATA COLLECTION

In this section, we present the data collection and processing methodologies for analysis. The dataset and the related materials are publicly available at our Github repository <https://github.com/shlee-lab/Qubic-selfish-mining-study>.

A. Block and mining information collection

For analysis, we collected two kinds of data: Monero block information by Monero nodes and Qubic’s mining job information by Qubic’s Monero mining pool.

Monero block information. We needed to operate a Monero pruning node, which is a lightweight Monero node that stores only recent blocks while still participating fully in network validation. Pruning nodes help us not only to collect block data trustfully but also to collect Monero orphan blocks that are not accessible from open Monero explorers. We operated the node during 29th September 2025 – 17th October 2025. For periods before September 29, 2025, we use publicly accessible Monero full nodes to retrieve historical block and coinbase data.

Qubic mining pool. We used the RPC API of the Qubic mining pool to retrieve mining jobs at 5-second intervals. The API has a structure similar to the Stratum protocol, which is widely used by PoW mining pools, and is openly accessible. Specifically, the `job_notify` method returns multiple values, including the mining block height and the previous block hash, enabling miners to obtain the most recent mining jobs. These data are therefore useful for analyzing how Qubic performs selfish mining on the Monero network.

B. Qubic block attribution

Since Monero provides a strong privacy at default, miners cannot be directly identified solely from the blockchain. Block rewards are paid through a special *coinbase transaction*, and while some mining pools optionally disclose a view key to prove ownership of the reward output, this practice is not mandatory. Qubic likewise does not provide view keys in real time, so it is not possible to immediately determine whether a given block was produced by Qubic through straightforward direct verification.

Empirically, Qubic’s mining activity appears to be organized into weekly epochs. In general, an epoch ends on Wednesday, and the view key used during that period is disclosed only after the epoch has fully completed, via Qubic’s official Discord channel. A consolidated list of the publicly disclosed Qubic view keys used in our study is given in Table III in Appendix B. As a consequence of this disclosure schedule, view-key-based verification can be applied only to blocks that are already definitively included in the main chain; it cannot be used for blocks that are still within an ongoing epoch or for blocks that have already become orphaned.

To address this limitation, we leverage the `tx_extra` field of coinbase transactions as an additional attribution signal. This field contains an extra-nonce area where mining pools may freely embed auxiliary data. By examining blocks suspected to be mined by Qubic, we observe a consistent structural pattern in the extra-nonce values. This pattern is highly atypical in ordinary Monero transactions and exhibits a distinctive form that enables us to classify Qubic blocks even before the corresponding view key is published. Furthermore, for all blocks that were later confirmed to belong to Qubic through ex-post view key disclosure during our study period, we did not observe any violations of this pattern.

Table I presents representative examples of extra-nonce values observed in Qubic blocks and in blocks from other mining pools. Based on these samples, we derive a heuristic regular expression that captures the Qubic-specific structure of the extra nonce:

Listing 1: Qubic extra-nonce regex

```
([0-9a-f]{4})0{4}([0-9a-f]{8})(([0-9a-f]{8})0{10})$
```

We classify a block as mined by Qubic if the extra nonce of its coinbase transaction matches this pattern. This method enables us to attribute Qubic-mined blocks across the Monero blockchain without requiring any private keys or pool-disclosed view keys.

III. QUBIC’S MINING ON MONERO

This section characterizes Qubic’s presence on the Monero network during the observation period. We first quantify its mining power based on attributed blocks, and then examine its impact on orphan blocks and reorganizations. A focused analysis of Qubic’s selfish mining strategy follows in Section IV.

A. Qubic’s mining power share

Figure 1 shows Qubic’s mining power share in the Monero network, computed as the ratio of Qubic-attributed blocks to all main-chain blocks over weekly, daily, and hour windows. Since direct telemetry of the




Fig. 1: Qubic mining pool's mining power share on the Monero network




Fig. 2: Daily Monero chain block production




Fig. 3: Weekly distribution of orphan lengths

TABLE I: Examples of `extra_nonce` values in coinbase transactions from Qubic and non-Qubic pools.

pool's physical hashrate is unavailable, we rely on this 'effective hashrate' realized on-chain as the primary metric (α) for our selfish mining models. It was measured by the number of blocks share over the total blocks in a period. This approach prioritizes realized block production over nominal hardware performance (hash/s), as the former directly determines the success probability of chain reorganizations. Throughout the measurement period, Qubic's overall average share is approximately 22.09%, far below a sustained majority of the network hashrate.

A central point in public discussions was Qubic's claim, echoed by several media outlets, that it had achieved a 51% mining position on Monero. However, the notion of a "51% attack" is often left ambiguous, and our measurements do not support a persistent majority. In the 6-hour series, we observe several short intervals where Qubic's share approaches or briefly exceeds 50% [1]. In contrast, the corresponding daily and weekly aggregates never reach 51%, and Qubic does not maintain a stable majority at any point in our dataset. These results indicate that while Qubic temporarily concentrated substantial mining power, it did not achieve the sustained control typically associated with a practical 51% attack on the Monero network.

B. Orphan Blocks by Qubic

Although Qubic did not sustain majority mining power, its activity significantly affected the stability of the Monero chain when it engaged in selfish mining. Selfish mining naturally leads to frequent chain reorganizations, which manifest as orphan blocks on the public chain.

Figure 2 presents the daily counts of regular and orphan blocks, separated into Qubic-attributed and non-Qubic blocks, alongside the average network difficulty. We observe a pronounced increase in the number of orphan blocks during periods when Qubic is active with selfish mining, including a substantial fraction attributable to Qubic itself. This behavior indicates that Qubic’s strategy not only discarded other miners’ blocks but also caused a non-trivial number of its own blocks to be orphaned.

Figure 3 shows the distribution of orphan fork lengths over time. In periods without Qubic’s selfish mining, orphan forks are almost exclusively of length one. Once Qubic’s selfish mining begins, we observe a clear shift toward longer orphan chains, with more frequent occurrences of multi-block forks. This change reflects deeper and more disruptive reorganizations of the main chain.

IV. SELFISH MINING FOCUSED ANALYSIS

In this section, we analyze Qubic’s selfish mining behavior on Monero and evaluate the profitability of the inferred strategy. We first outline the strategic context, then justify our use of Qubic’s timestamps, identify candidate selfish mining periods, formalize the corresponding analytical model, and finally compare the theoretical revenue with Qubic’s observed mining share.

A. Selfish mining strategies

Classical selfish mining strategies, as introduced by Eyal and Sirer [3] and later optimized by Sapirshtein et al. [4], describe how a rational miner (or pool) withholds blocks and selectively publishes a private chain to gain a revenue share exceeding its relative hashrate. In practice, miners face uncertainty due to network asynchrony and incomplete information, and may adapt their withholding and release rules accordingly. Qubic's publicly visible behavior and self-reported capabilities suggest the possibility of such deviations, but its exact internal policy remains unobservable.

Our analysis is therefore necessarily indirect. We do not assume access to Qubic's internal block discovery times or private chain states. Instead, we rely on Monero main-chain and orphan blocks, Qubic-attributed blocks, and their timestamps to infer when and how Qubic may have executed selfish mining. The following subsections validate our use of timestamps for this purpose and identify periods in which Qubic's behavior is consistent with sustained strategic deviation, which we subsequently interpret through an analytical model.

B. Reliability of Qubic block timestamps

To infer Qubic's strategy from on-chain artifacts, we must first assess whether Qubic's block timestamps are




Fig. 4: Qubic blocks’ timestamp difference distribution with the same height blocks including orphan blocks




Fig. 5: Time delay distribution between Qubic block timestamps and job fetch timestamps. The mean delay is 5.61 seconds, and the green line marks the near 8-second job fetch period.

usable as approximate indicators of block discovery and release behavior. We focus on heights where a Qubic block competes with at least one orphan block or where Qubic’s block itself becomes orphaned, and measure the timestamp differences between Qubic-attributed blocks and competing blocks.

As illustrated in Fig. 4, in the majority of competing blocks, Qubic’s blocks exhibit timestamps that are either antecedent to or comparable with those of competing blocks. This temporal alignment is consistent with a strategy of mining blocks early and selectively withholding them prior to broadcast. While a minor subset of Qubic blocks displays anomalously high timestamps, these outliers are predominantly associated with orphaned blocks.

Furthermore, to verify timestamp integrity, we measured the time difference between the block’s timestamp and the timestamp of the preceding job fetch response

recorded locally (Fig. 5). The mining client is configured to fetch new jobs approximately every 7 seconds, though the actual average period was recorded as 7.7 seconds due to response delays and network circumstances. Consequently, the observed mean delay of 5.61 seconds falls naturally within this effective fetch interval, representing the expected time elapsed between receiving a job and finding a valid nonce. This alignment indicates that the timestamps are a natural byproduct of the mining workflow rather than the result of deliberate falsification. Consequently, we find no compelling evidence of large-scale timestamp manipulation that would invalidate temporal reasoning regarding Qubic’s behavior.

In light of these observations, we treat Qubic’s timestamps as a sufficiently reliable proxy for inferring the relative ordering of events and for detecting patterns indicative of selfish mining. Specifically, we premise our analysis on the assumption that Qubic does not consistently alter timestamps to obfuscate private chain lead lengths or release decisions. Leveraging this assumption, we subsequently utilize orphan dynamics and timestamp data to illustrate the specific time intervals in which selfish mining was likely active.

C. Selfish mining period categorization

Qubic’s strategy may vary over time, and selfish mining, if present, is unlikely to be applied uniformly across the entire observation window. To focus our analysis on intervals with meaningful deviations, we define a heuristic as Alg. 1 to identify periods with sustained abnormal orphan activity, which serves as an indicator of potential selfish mining.

Algorithm 1 Selfish Mining Period Heuristic

Require: Blocks B with timestamps and orphan flags, thresholds τ_{\min} , d_{\min} , g_{\max}
Ensure: Valid period spans M

- 1: **Step 1:** Aggregate orphan blocks (Qubic + Other) by hour
- 2: $C[h] \leftarrow \{b \in B : b.\text{is_orphan} \wedge \lfloor b.\text{timestamp} \rfloor_{\text{hour}} = h\}$
- 3: **Step 2:** Find contiguous segments where each hour meets threshold
- 4: $S \leftarrow \{(s_i, e_i) : \forall h \in [s_i, e_i], C[h] \geq \tau_{\min} \wedge |e_i - s_i| \geq d_{\min}\}$
- 5: **Step 3:** Merge segments with gaps $\leq g_{\max}$ hours
- 6: $M \leftarrow \text{merge}(S, g_{\max})$
- 7: **return** M

The heuristic operates in three steps. First, we aggregate all orphan blocks (both Qubic and non-Qubic) into hourly bins and compute the orphan count $C[h]$ for each hour. Second, we identify contiguous segments in which




Fig. 6: Orphan block frequency and selfish mining periods

every hour satisfies a minimum orphan-count threshold τ_{\min} and the total segment length exceeds a minimum duration d_{\min} . Third, we merge neighboring segments separated by gaps shorter than g_{\max} hours to tolerate brief fluctuations in activity. The result is a set M of candidate selfish mining periods.

In our evaluation, we set $\tau_{\min} = 2$ orphan blocks per hour, $d_{\min} = 4$ hours, and $g_{\max} = 6$ hours. These parameters were chosen to capture intervals where elevated orphan activity is persistent enough to be unlikely under normal conditions, while avoiding over-fragmentation due to minor gaps. Fig. 6 shows the resulting candidate periods on the timeline. Applying this heuristic yields ten periods (P1–P10) during which Qubic’s block share and involvement in orphan blocks are noticeably higher than their global averages, suggesting that Qubic concentrated its selfish mining behavior in these windows. In particular, Qubic’s average hashrate share in these periods is 28.02%, higher than its overall share of 22.09%. This gap indicates that Qubic tends to enable selfish mining only after its hashrate reaches a comparatively high level, and to cease such behavior once its effective share falls back toward the baseline.

Using the measured $\alpha = 28.02\%$ for the selfish mining periods together with the observed low $\gamma \approx 0$, we evaluate both the classical selfish mining revenue function and the derived modified-strategy revenue.

Fig. 7 further illustrates how Qubic’s behavior varies across the identified periods by plotting the lengths of Qubic-controlled runs against the number of associated orphan blocks. In early periods (P1–P4), most points lie close to the reference blue line ($y = x - 1$) corresponding to releasing the private chain at lead 1, which is broadly consistent with the classical selfish mining strategy that maximizes the waste of the other mining power. In con-

trast, during P8, where Qubic’s activity is most intensive, we observe many runs aligned with patterns indicative of releasing at lead 2, suggesting a more conservative policy that avoids exposing the pool to tie situations when the lead is only one block. This shift supports the view that Qubic adapts its release rule depending on its effective hashrate and perceived network conditions, motivating the need for an analytical model that explicitly captures such a modified strategy.

D. Analytical model of Qubic’s selfish mining behavior

To interpret the observed periods, we compare them against analytical models of selfish mining. We first recall the classical selfish mining revenue function $R_{\text{selfish}}(\alpha, \gamma)$, which expresses the pool’s expected revenue share as a function of its relative hashrate α and the network tie-breaking parameter γ . Its revenue can be calculated based on the state machine depicted in Fig. 8. Each state presents ‘lead’ denotes the number of blocks by which the selfish miner is ahead. For example, if a selfish miner’s private chain has height 103 and the honest miners’ public chain has height 100, then the system’s state is 3. State 0’ corresponds to the situation where the selfish miner has mined a block first and the honest miners have also mined a competing block, so the network is in a tie-breaking state. The tie can be resolved in three ways: (1) an honest miner extends the honest chain with probability $(1-\alpha)(1-\gamma)$; (2) an honest miner extends the selfish miner’s chain with probability $(1-\alpha)\gamma$; or (3) the selfish miner extends the selfish miner’s own chain with probability α . Eq. 1 gives the revenue of the selfish mining [3].

$$R_{\text{selfish}}(\alpha, \gamma) = \frac{\alpha(1-\alpha)^2(4\alpha + \gamma(1-2\alpha)) - \alpha^3}{1 - \alpha(1 + (2 - \alpha)\alpha)} \quad (1)$$




Fig. 7: Distribution of orphan blocks per Qubic run length across ten selfish mining periods (a-j). The scatter plots visualize run frequency (encoded by point size and color) for each run length and orphan count pair. Reference lines include the theoretical ideal selfish mining line $y = x - 1$ (blue dashed) and the secondary threshold $y = x - 2$ (red dotted).

We then consider a slightly different selfish mining strategy. Fig. 9 represents a modified strategy motivated by Qubic's observed behavior, in which the private chain is frequently released at lead 2 and includes an additional transition pattern (e.g., from state 3 to 0) reflecting more conservative release decisions. Qubic's overall strategy is best understood as lying between these two models rather than matching either one exactly.

Using this model, we derive a closed-form expression for the expected revenue of the modified strategy as a function of α and γ .

Theorem 1. Let $R_{\text{mod}}(\alpha, \gamma)$ denote the selfish pool's long-run fraction of accepted blocks under the modified strategy encoded by the state machine in Fig. 9, where $\alpha \in (0, \frac{1}{2})$ is the pool's relative hash power and $\gamma \in [0, 1]$ is the tie-breaking parameter. Then




Fig. 8: State machine of the original selfish mining.




Fig. 9: State machine of a modified selfish mining model with a modified transition from state 3 to 0.

$[0, 1]$ is the tie-breaking parameter. Then

$$R_{\text{mod}}(\alpha, \gamma) = \frac{\alpha(\alpha^3\gamma - 3\alpha^2\gamma + \alpha^2 + 3\alpha\gamma - 2\alpha - \gamma)}{\alpha^4 - 2\alpha^3 + \alpha - 1}. \quad (2)$$




Fig. 10: Weekly γ rate and estimated state $0'$ counts

Proof. See Appendix A. \square

E. Expected revenue under observed parameters

We now apply the analytical model using parameters inferred from the identified selfish-mining periods. For each period, we estimate Qubic's relative hashrate α from its share of main-chain and orphaned blocks, and derive an empirical tie-breaking parameter γ by counting events in which Qubic and non-Qubic blocks compete at the same height and observing which block is eventually adopted.

For the relevant range of α and the near-zero γ observed in our data, both the standard and modified selfish mining models predict that Qubic's expected revenue is not superior to that of honest mining as Fig. 11 illustrates. At the average hashrate share of 28.02% during the identified selfish mining periods and $\gamma \approx 0$, the standard selfish mining model yields an expected revenue ratio of $R_{\text{selfish}} \approx 25.53\%$, while the modified model yields $R_{\text{mod}} \approx 17.82\%$. Since Qubic's actual strategy is best understood as lying between these two models, its expected revenue ratio should also fall within this interval, implying a relative loss of roughly 9% – 36% compared to honest mining at the same hashrate. These results indicate that, under realistic network conditions, Qubic's selfish mining behavior would not provide a mining reward advantage and is, from a narrow revenue perspective, irrational.

F. Comparing the theoretic expectation and observed revenue

Again in Fig. 11, we presented the observed revenue of Qubic with categorizing the global average, the selfish period mining average, and each selfish mining period's average. Table II shows the detailed numbers. Unlike our theoretic expectation, the Qubic's revenue did not approach the lower bound (the modified selfish mining).




Fig. 11: Theoretical selfish mining revenues and observed Qubic's mining revenue by periods, average of period, and the global average. The light-blue shaded area denotes the theoretically predicted revenue range when $\gamma = 0$.

TABLE II: Qubic's observed revenue overview

Period	Estimated α	Observed Revenue	Main Chain	Qubic Total	Qubic Main
P1	0.3415	0.3386	508	208	193
P2	0.3213	0.3178	440	151	143
P3	0.3312	0.3476	421	157	146
P4	0.3349	0.3314	507	215	169
P5	0.3164	0.2778	605	243	175
P6	0.3223	0.2926	769	313	237
P7	0.3102	0.2939	611	237	194
P8	0.2578	0.2248	7429	2303	1693
P9	0.3436	0.3533	617	268	212
P10	0.2346	0.1793	1053	294	199
Avg (P1-P10)	0.2802	0.2529	12960	4389	3361
Global Average	0.2209	0.2114	55937	13000	11871

Rather, the Qubic sometimes outperformed the lower bound of the normal selfish mining. And P6 and P9 outperformed the honest mining revenue even though the selfish mining average and the global average underperformed the honest mining. In the next section, we analyze this difference.

V. ANALYSIS ON DIFFERENCE BETWEEN REVENUE EXPECTATION BY THEORY AND OBSERVATION

In this section, we analyze the discrepancy between the theoretical revenue expectations derived from our models and the actual observed revenue of the Qubic pool during the identified selfish mining periods. We then provide a conjecture regarding the underlying causes of this divergence.

A. Tie-breaking winning rate

To understand the deviation observed in Fig. 11, we initially hypothesized that the empirically derived tie-




Fig. 12: Qubic’s winning ratio at State 0’

breaking parameter γ might have been underestimated in our previous analysis. However, a closer inspection revealed that the primary anomaly lay not in the global γ estimation, but in the specific winning rate of Qubic at state 0’—the scenario where a selfish miner competes in single block race against the rest of the network.

As illustrated in Fig. 12, Qubic’s winning rate in these race conditions was consistently higher than its estimated mining power (α) for most periods, with the exceptions of P2 and P4. Notably, several periods exhibited a race winning rate that significantly exceeded the corresponding hashrate share. For instance, in period P1, the observed winning rate was 0.62 against an α of 0.34. Similarly, P3 showed a rate of 0.56 (vs. $\alpha = 0.33$), and P7 showed 0.49 (vs. $\alpha = 0.31$).

B. Temporal granularity of selfish mining identification

We considered several possibilities to explain why the observed behavior deviated from the theoretical models. The most plausible explanation is that the heuristic periods identified by Algorithm 1 were too rough, inadvertently capturing intervals where selfish mining was not actively employed. While our heuristic defined periods based on a rough 4-hour threshold with 6-hour gaps, it is likely that Qubic dynamically toggled its selfish mining strategy on and off with much finer granularity, responding delicately to real-time fluctuations in its mining power which Qubic can estimate precisely.

This hypothesis is supported by the data from period P8, which offers a large sample size over a long duration. In P8, the average mining power share was measured at 25.78%, yet the tie-breaking winning rate at state 0’ was approximately 35%. We conjecture that 35% represents Qubic’s true effective hashrate during the brief sub-intervals when it was actively engaged in selfish mining. If we assume an active mining power of $\alpha = 0.35$ with $\gamma \approx 0$, the theoretical revenue of selfish mining would

be bounded between 36.65% (Eq. 1, upper bound) and 28.04% (Eq. 2, lower bound).

However, the actual observed revenue for P8 was only 22.48%. This suggests that even if Qubic attempted to execute the strategy, it failed to outperform honest mining. In relative terms, this corresponds to a revenue loss of approximately 12% compared to honest mining ($\approx 25.78\%$). This significant underperformance implies that the strategy was not executed efficiently enough to overcome the inherent risks of selfish mining.

Conversely, in periods such as P1 and P3, Qubic did outperform honest mining. These periods were characterized by a high baseline mining power (30%–40%), and the exceptionally high race winning rates (Fig. 12) suggest that Qubic’s effective power was dominant during the specific moments of attack. Nevertheless, the excess revenue gained over honest mining was marginal—more or less than 5%—and these successful instances constituted only a minor fraction of the total observation timeline.

Our overall analysis indicates that Qubic’s selfish mining campaign was, on aggregate, inefficient and failed to consistently outperform honest mining. The measured revenue often fell below the theoretical profitability threshold, likely due to the difficulty of maintaining optimal strategy execution under real-world network variance. However, despite the lack of financial success for the attacker, the strategy induced severe negative externalities on the Monero network. The identified periods coincided with elevated orphan rates and significantly deeper chain reorganizations, degrading the effective reliability of transaction confirmations. This tension between the attacker’s private loss and the public harm inflicted on the network motivates the discussion of mitigation strategies in the subsequent sections.

VI. DISCUSSION

Our primary focus has been on analyzing Qubic’s selfish mining strategies. In this section, we turn to two technical defense approaches against the Qubic’s campaign in the Monero community. The first approach is to modify the main-chain selection rule so that nodes are less likely to accept blocks produced by a selfish miner. The second is *detective mining*, which mitigates deep private selfish chains without any modification to the chain-selection rules.

A. Chain-selection rule modification

A natural class of countermeasures against selfish mining is to modify the chain-selection rule so that withheld blocks become less valuable to an attacker. Existing proposals along this line include timestamp-

and freshness-based fork-choice rules and backward-compatible schemes such as Publish or Perish [5]–[7]. These mechanisms typically penalize blocks that propagate too slowly or reward chains that incorporate more timely blocks and uncles, thereby reducing the benefit of maintaining a long private chain [8], [9].

Following Qubic’s campaign, parts of the Monero community actively discussed adopting a Publish-or-Perish-style rule [10]. However, such approaches rely on relatively strong assumptions about network-wide propagation bounds and timing information, which are difficult to guarantee in a heterogeneous, permissionless environment like Monero. Moreover, deploying a new chain-selection rule at this point would mainly protect against future selfish mining attempts rather than mitigating the damage from Qubic’s already finished campaign.

More broadly, our case study reinforces a well-known limitation of proof-of-work systems: when the overall hash power is modest and concentrated, the system remains structurally exposed to selfish mining and related strategies. This tension has motivated some ecosystems to explore alternatives such as proof-of-stake or hybrid designs. While a full evaluation of such alternatives is beyond the scope of this work, our findings suggest that Monero must either secure a sufficiently large and decentralized mining power or eventually confront this design trade-off.

B. Detective mining and its limitation

In contrast to protocol-level modifications, the Monero community highlighted *detective mining* as a countermeasure that does not require any modification to the chain-selection rules [11]. One characteristic of Qubic’s selfish mining attack is that it is carried out by a public mining pool. Detective mining has been studied under the assumption that a public mining pool performs selfish mining [12]. Because the selfish mining pool is public, some information about its private chain is partially observable to miners. In particular, the previous Merkle root, which is essential for tracking the private chain, is visible. Rational miners may then prefer to mine on the leading private chain rather than on the lagging public chain. In theory, such miners obtain a higher expected reward than both the selfish miner and honest miners.

However, this approach raises two issues. The first is the possibility of a Denial-of-Service (DoS) attack by the selfish mining pool. Detective miners do not know the transactions contained in the previous private blocks, and if the selfish mining pool refuses to release its private blocks after a detective-mining block is found, miners are unable to include transactions on the longest chain. The second issue is that some community members objected to the idea of miners intentionally building on the selfish

miner’s private chain. Moreover, detective mining further accelerates the loss incurred by honest miners.

VII. CONCLUSION

This paper presented an empirical case study of Qubic’s selfish-mining campaign on Monero. Using data from a Monero pruning node, public nodes, and the Qubic pool API, we reconstructed Qubic’s mining activity and heuristically identified ten intervals consistent with selfish mining strategies. During these periods, Qubic’s hashrate share rose only into the 23–34% and never sustained majority control, falling short of a lasting 51% takeover.

Comparing the measurements with both the classical selfish mining model and a modified Markov chain model tailored to Qubic’s conservative release policy observed in several periods, we find that the campaign was profitable in only a small subset of intervals. Nonetheless, Qubic’s activity sharply increases orphan rates and reorganization depth, degrading confirmation reliability. This illustrates that even economically fragile selfish mining campaigns can still impose meaningful harm, and that existing mitigation approaches offer limitations against such behavior in practice.

ACKNOWLEDGEMENT

This study is an extended work of the project in K-shield Junior 2025 program. We appreciate discussion from Yunkyu Lee and Huiwoo Kim.

REFERENCES

- [1] retrodrive, “Qubic performs 51% monero network takeover demonstration,” <https://qubic.org/blog-detail/historic-takeover-complete-qubic-miners-now-secure-monero-network>, Aug. 2025, accessed: 2025-10-25.
- [2] ———, “Epoch 172 recap - the monero experiment,” <https://qubic.org/blog-detail/epoch-172-recap-the-monero-experiment>, Aug. 2025, accessed: 2025-10-25.
- [3] I. Eyal and E. G. Sirer, “Majority is not enough: Bitcoin mining is vulnerable,” *Communications of the ACM*, vol. 61, no. 7, pp. 95–102, 2018.
- [4] A. Sapirshtein, Y. Sompolinsky, and A. Zohar, “Optimal selfish mining strategies in bitcoin,” in *International conference on financial cryptography and data security*. Springer, 2016, pp. 515–532.
- [5] E. Heilman, “One weird trick to stop selfish miners: Fresh bitcoins, a solution for the honest miner,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2014, pp. 161–162.
- [6] S. Solat and M. Potop-Butucaru, “Brief announcement: Zeroblock: Timestamp-free prevention of block-withholding attack in bitcoin,” in *International Symposium on Stabilization, Safety, and Security of Distributed Systems*. Springer, 2017, pp. 356–360.
- [7] R. Zhang and B. Preneel, “Publish or perish: A backward-compatible defense against selfish mining in bitcoin,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2017, pp. 277–292.
- [8] R. Pass and E. Shi, “Fruitchains: A fair blockchain,” in *Proceedings of the ACM symposium on principles of distributed computing*, 2017, pp. 315–324.

- [9] P. Szalachowski, D. Reijsergen, I. Homoliak, and S. Sun, “{StrongChain}: Transparent and collaborative {Proof-of-Work} consensus,” in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 819–836.
- [10] tevador, “(github issue #144) selfish mining mitigations (publish or perish),” <https://github.com/monero-project/research-lab/issues/144>, Aug. 2025, accessed: 2025-10-25.
- [11] R. Spagni, “(github issue #140) implement detective mining (pool software change, no protocol change required),” <https://github.com/monero-project/research-lab/issues/140>, Aug. 2025, accessed: 2025-10-25.
- [12] S. Lee and S. Kim, “Rethinking selfish mining under pooled mining,” *ICT Express*, vol. 9, no. 3, pp. 356–361, 2023.

APPENDIX

A. Proof for Theorem 1

Proof. We derive the stationary distribution of the Markov chain in Fig. 9 and then compute the expected rewards per block-finding event.

Let π_s be the stationary probability of state s . From Fig. 9 we read off the following balance equations.

a) Tie state: From state 1, an honest block (probability $1 - \alpha$) leads to $0'$, and from $0'$ the next block always resolves the fork and returns to 0. Thus

$$\pi_{0'} = (1 - \alpha)\pi_1. \quad (3)$$

b) Lead 1 and 2: From the transitions $0 \rightarrow 1$ with probability α and $2 \rightarrow 1$ with probability $1 - \alpha$, and the outgoing probability 1 from state 1, we obtain

$$\pi_1 = \alpha\pi_0 + (1 - \alpha)\pi_2. \quad (4)$$

From $1 \rightarrow 2$ with probability α and total outflow 1 from state 2,

$$\pi_2 = \alpha\pi_1. \quad (5)$$

c) Lead ≥ 3 (geometric tail): For $i \geq 3$ the figure shows: on a selfish block (α) the lead increases by one, and on an honest block ($1 - \alpha$) the selfish miner publishes its private chain and the process returns to 0. Thus for $i \geq 3$ we have pure geometric growth:

$$\pi_i = \alpha\pi_{i-1} \quad (i \geq 3), \quad (6)$$

hence $\pi_i = \alpha^{i-2}\pi_2$ for all $i \geq 2$.

d) State 0: Incoming probability flow into 0 comes from: (i) the self-loop at 0 when an honest block is found, (ii) resolution of $0'$, (iii) full publication from any $i \geq 3$ after an honest block. The balance for state 0 is therefore

$$\pi_0 = (1 - \alpha)\pi_0 + \pi_{0'} + (1 - \alpha) \sum_{i=3}^{\infty} \pi_i. \quad (7)$$

e) Solving for π_s : Using (5) and (6) we obtain

$$\pi_2 = \alpha\pi_1, \quad \pi_i = \alpha^{i-1}\pi_1 \quad (i \geq 2).$$

Substituting into (4) gives

$$\pi_1 = \alpha\pi_0 + (1 - \alpha)\alpha\pi_1 \Rightarrow \pi_1 = \frac{\alpha}{\alpha^2 - \alpha + 1} \pi_0.$$

Then

$$\pi_2 = \frac{\alpha^2}{\alpha^2 - \alpha + 1} \pi_0, \quad \pi_i = \frac{\alpha^i}{\alpha^2 - \alpha + 1} \pi_0 \quad (i \geq 2),$$

and from (3),

$$\pi_{0'} = (1 - \alpha)\pi_1 = \frac{\alpha(1 - \alpha)}{\alpha^2 - \alpha + 1} \pi_0.$$

Using (7) and the geometric sum $\sum_{i=3}^{\infty} \alpha^i = \alpha^3 / (1 - \alpha)$, one checks that the balance holds with these expressions. Normalization

$$\pi_0 + \pi_{0'} + \sum_{i=1}^{\infty} \pi_i = 1$$

then yields

$$\pi_0 = 1 - 2\alpha + 2\alpha^2 - \alpha^3,$$

and thus all π_s are uniquely determined.

f) Expected rewards: We next compute the expected number of accepted blocks per block-finding event for the attacker (E_s) and for the whole network (E_t).

- In state 0, an honest block (probability $1 - \alpha$) adds one accepted block for honest miners:

$$E_s^{(0)} = 0, \quad E_t^{(0)} = (1 - \alpha)\pi_0.$$

- In state $0'$, the next block always resolves the fork and finalizes two blocks. As in the original analysis, the attacker’s expected share is $2\alpha + \gamma(1 - \alpha)$, so

$$E_s^{(0')} = \pi_{0'}(2\alpha + \gamma(1 - \alpha)), \quad E_t^{(0')} = 2\pi_{0'}.$$

- In states 1 and 2, one step does not finalize any block, only changes the lead, hence

$$E_s^{(1)} = E_s^{(2)} = 0, \quad E_t^{(1)} = E_t^{(2)} = 0.$$

- In each state $i \geq 3$, an honest block (probability $1 - \alpha$) triggers full publication of the private chain: all i blocks are accepted for the attacker and the chain returns to 0. Therefore

$$E_s^{(i)} = (1 - \alpha)i\pi_i, \quad E_t^{(i)} = (1 - \alpha)i\pi_i.$$

Summing all contributions,

$$E_s = \pi_{0'}(2\alpha + \gamma(1 - \alpha)) + (1 - \alpha) \sum_{i=3}^{\infty} i\pi_i,$$

$$E_t = (1 - \alpha)\pi_0 + 2\pi_{0'} + (1 - \alpha) \sum_{i=3}^{\infty} i\pi_i.$$

Using the geometric series derivative formula $\sum_{i=3}^{\infty} i\alpha^i = \alpha^3(3 - 2\alpha)/(1 - \alpha)^2$, we evaluate the summation terms.

Crucially, we observe that the normalization constant π_0 factorizes as:

$$\pi_0 = 1 - 2\alpha + 2\alpha^2 - \alpha^3 = (1 - \alpha)(\alpha^2 - \alpha + 1).$$

This factorization allows us to simplify the stationary probabilities:

$$\frac{\pi_0}{\alpha^2 - \alpha + 1} = 1 - \alpha.$$

Substituting this into the expressions for $\pi_{0'}$ and π_i , the common denominator cancels out:

$$\begin{aligned} \pi_{0'} &= \alpha(1 - \alpha)(1 - \alpha) = \alpha(1 - \alpha)^2, \\ \pi_i &= \alpha^i(1 - \alpha) \quad (i \geq 2). \end{aligned}$$

Now, we substitute these simplified probabilities into the reward equations. For E_s :

$$\begin{aligned} E_s &= \pi_{0'}(2\alpha + \gamma(1 - \alpha)) + (1 - \alpha) \sum_{i=3}^{\infty} i\pi_i \\ &= \alpha(1 - \alpha)^2(2\alpha + \gamma(1 - \alpha)) + (1 - \alpha)^2 \sum_{i=3}^{\infty} i\alpha^i \\ &= \alpha(1 - \alpha)^2(2\alpha + \gamma(1 - \alpha)) + (1 - \alpha)^2 \frac{\alpha^3(3 - 2\alpha)}{(1 - \alpha)^2} \\ &= \alpha(1 - \alpha)^2(2\alpha + \gamma - \alpha\gamma) + \alpha^3(3 - 2\alpha). \end{aligned}$$

Expanding and collecting terms yields the final polynomial form:

$$E_s = -\alpha(\alpha^3\gamma - 3\alpha^2\gamma + \alpha^2 + 3\alpha\gamma - 2\alpha - \gamma).$$

Similarly for E_t :

$$\begin{aligned} E_t &= (1 - \alpha)\pi_0 + 2\pi_{0'} + (1 - \alpha) \sum_{i=3}^{\infty} i\pi_i \\ &= (1 - \alpha)^2(\alpha^2 - \alpha + 1) + 2\alpha(1 - \alpha)^2 + \alpha^3(3 - 2\alpha). \end{aligned}$$

Simplifying this expression leads to:

$$E_t = -(\alpha^4 - 2\alpha^3 + \alpha^2 - 1).$$

Hence the attacker's long-run revenue ratio is

$$R_{\text{mod}}(\alpha, \gamma) = \frac{E_s}{E_t},$$

which is exactly Eq. 2. \square

B. List of Qubic's view keys

The view keys listed in Table III were obtained from the official Qubic Discord server (Accessible via <https://qubic.org/>), specifically from the disclosures made by Qubic's lead dev *dkat*. These keys are provided to enable the community to retrospectively verify the ownership and provenance of Qubic-mined blocks.

TABLE III: Qubic View Key List

Date	Address	View Key	Note
May 27, 2025	47hhGMKbWpKfxDiqcejWGicVvQHEYd45AEaUyKVjcZ ywL8c8mtjn3oACGfdrsLrPGP2r49gvTBnBiTVQcEkf BNFEKCDy7ME	577fd4a7278f55d2a9230d32823b81497b2e854d4a 8702b1256a17cda42a760d	
Aug. 6, 2025	43oMtdwB5aaCuM9vVaiY6u7XgxCGLwA563C7b5V3oS TSjDdhBkWeGxeZzSuD4wAydMzbvNWrF9iRGmwoMnh YnMTMcZjBrv	e935552c5665117a6ecc9fbbfd4156595c75774014 606130a01003720e063201	epoch 172
Aug. 13, 2025	49upGQgCYzxMKfBU9hYe8QH3fQQMSdReiAJx3vo9bq 6K4Yegbqp39rVKNNGh9tA3VobYmkyGxvDc1J9FnVFw 8f4UT7BsDhf	0b21ef509769c6d95899cca7cccd86b89333ca4ce0d fa3fd5aa304059aed0f903	epoch 173
Aug. 20, 2025	4AqzG7scWP19yNsFuJNpQ2CNF7LGxJNtJaEennAA48 KbLX6a7PTazW4c3FTwBPFjJ4TFq3xpZhvGvgygyVct uXxSLWrAk3S	05ada241eea8b262241762cb6be291be3aaaf123756 0a0ddd1fd4ea5cf502120f	epoch 174
Aug. 28, 2025	45hzuq7TBR73J89EXkAmZuqY9Dgpcd6xFSq4U2DMY5b o2RUSJ4V7Y3QWpMAWGF7CQ8955U6XR2F4PUVbJhU5D nzs3c16dM9Dy	9267d1762b0f3262029be73e30f5158159c2f38e86 b9d745231e57141afcc0a	epoch 175
Sep. 3, 2025	49heVghSznN9eoatkooNjLHHsRV6XiitDeW4J92cgn ey8BFtuacZGmzSA3fRKEHooC7X9xzCP9VXN6uK7Xrp oXF35FXfQCP	d761a707408f9693f9a453501dc04df6c92b82309a 90f23884564dabfed70106	epoch 176
Sep. 10, 2025	42Vt47oLyRT7C1Ch3BbapKFZgs5Hip5m3RrRVT2dbj DT8Nws76gc7NngfZvzXpZnPYGgVZFF79T5TSKWSjFx Ywk4A77WGa6	91c313b9cb0cc45e03e2f6f97e9d61566f8d0636b4 bb3bf59c59022972caad09	epoch 177
Sep. 17, 2025	47GwPhLcnWshcbekVshrzxJZwXXfDRUrjb7T6CfR1H aaiokeBwxAsQnFp779bF6rW43giviWwYbsoT1KehsG nP5L7v1vuF6	c9f5d5027465f4ff51538210e4fa110756e956c064 897f969c4a60863e227f0c	epoch 178
Sep. 24, 2025	44UsmtPAE5GC8U8vnLp7FqUFAYkWL5YYZJLNFDQrb4 6ePGpsH58ydJ2QtfmEgR834AQphJYwsLvnJrrEluFh T38bQnTebXm	a5e32d32ea8d1aac9ed47b7679ecf2cc4884dc5b38 8d9b539fede7ae5389f603	epoch 179
Oct. 1, 2025	48PSv1UrxcrQY6m1ZaDMXSM38kEv2NhjWNgiSmJoU L7BFjm5A4XkibKt2ApF5ydqsDtaMfZK8WBT7PtabvK GMfZUdZqjUe	5b78ba1a936efe94acb8e13fce72ee3581267ef668 a9c0f8967883ab12394602	epoch 180
Oct. 8, 2025	48PSv1UrxcrQY6m1ZaDMXSM38kEv2NhjWN5qTB5D9 6DN4K1EbBRM8WBTFOW3JP4AW288WTDP7VhStLu3D8F N3VZb1w9z1H	1fdbd6085b25183aadd0c241e94adb4379fb145cd5 1b968c7bf068d171275902	epoch 181
Oct. 15, 2025	48BZD46hnnGJmh5kn4py9oEupM8uL3ZobK59GDMWkZ BdHyB2ALiSD6rGa2u9in2gMtegfKciaanDYNDFew8o GHKmAcYEQTo	938e19ee3f4e0fc025e1f2b5d2eafb9bb7db899cce 00e2367d395e935a470209	epoch 182
Oct. 15, 2025	42y5h2KPKhKTW82xqf7XHMFUz32Hs8ubgCowd8Q4y5 RY8XqvNCEfzX2cCexyfsLtdD1BjT5mRMDHiBrC1t8C aT1RDeDsMrg	173ad08b17faf1a672ee8314c79e4d2e931ed85210 c925d722cb400a2c7c3405	epoch 182
Oct. 23, 2025	45w3hfgjzJjHiDSsVKh4nKdHawTGrWzr4WPT5L14qe Nx4fyyRQ73cNiLGGrrJ5pDjP8LHeDJfTMCs1UN7eBp WTPy1YX7YCV	726d15c6d9963e41965d87e311f1e51ff5722badfa 824336e893fb01a11acb00	epoch 183
Oct. 19, 2025	48CzXZ9YKcTTKAP8qc2cMfFiXpLQpxDDVEDNn5mTr F5aX9tHAY7eqfUzJkwqkXtaZhb7Ggv9rjCJYcYRzF HkPJF2BbDtx	7b8f7098b3892b7d9fb9b9ab3f96559dde693e052a c60d034e5b4115f548e119	epoch 184
Nov. 6, 2025	4A6mfADoDhNEAodQChmXRTSGsK4gheMna73TebsmKw 4rRggore7U1p8No32pJytUPTfoS5xk11aDh93BzgqZ phbkUykELdm	d696c09c6a95b91fb707911f027797c9bce4a3d7 3d86e0361e95126655fe0e	epoch 185