Bitcoin MEV Analysis for Stacks Mining

Draft 2.0
April 20, 2023

Commissioned by the Stacks Foundation

Jesse Soslow*

7th Avenue Group
jesse@7avenue.co
www.7avenue.co

* While Jesse is a lawyer, nothing in this paper should be construed as legal advice. The legal risks of implementing any of the solutions discussed in this report have not been analyzed.

About 7th Avenue Group

7th Avenue Group is a boutique consulting firm that helps businesses solve a wide array of problems. We have over 30 years' experience founding, building and operating technology companies, including in the web, mobile, SaaS, consumer, and crypto spaces, in both B2B and B2C models.

We've built products used by more than half of the Fortune 100, and we've founded or operated companies that have raised over $300 million from top investors, including Sequoia Capital, Union Square Ventures and Andreessen Horowitz.

If there's a new market you need to get into, an initial product idea that you want to explore, or an important strategic decision you need to analyze, we can plug into your team and help you get it done.

Table of Contents

1 Background

1.1 Factual Background

Recently, F2Pool, the third largest Bitcoin miner, devised a scheme such that the Stacks miner they run wins 100% of the Stacks blocks associated with the Bitcoin blocks they produce. This is approximately 14% of all Bitcoin blocks.

For the Bitcoin blocks that F2Pool produces, they omit the Stacks block commitments from all other Stacks miners, and only include the commitment from their Stacks miner. Because there is only F2Pool's Stacks block commitment in their Bitcoin blocks, they are not competing against any other Stacks miners for the corresponding Stacks block, and thus they win.

In running this strategy, F2Pool is also sending a very small amount of Bitcoin – significantly less than a typical Bitcoin mining commitment for a typical Stacks miner – in order to win, thus making the strategy a highly profitable one for them. Further, while typically a Stacks miner needs to consistently mine in order to mine competitively in any given block, F2Pool is able to exclusively mine their own Bitcoin blocks because they are not actually competing with any other miners in those blocks. This makes their overall strategy marginally more profitable.

1.2 Impact on Stacks

While F2Pool's strategy (the "**MEV Strategy**") is currently being run, and the Stacks network continues to operate properly, it nonetheless creates potential downstream issues:

- The expected value of Stacking will decrease, reducing the incentive to participate.
- Robust Stacking demand and participation is critical for the launch of the soon-to-be-released sBTC functionality.
- The overall profitability of miners who do not participate in this strategy (we'll call these "**Honest Miners**") will decrease.
- As the knowledge of the profitability of this strategy spreads, we might expect an increase in the number of Bitcoin miners engaging in this strategy (we'll call these "**MEV Miners**").
- Over time, with the opportunity for considerable risk-free profit to be earned by Bitcoin miners, we would expect MEV Miners to entirely crowd out Honest Miners on the Stacks blockchain.
- The above could potentially result in effectively zero value being delivered to Stackers.

Even putting aside the above issues, the MEV Strategy conflicts with certain values of the Stacks ecosystem. In particular:

- Fairness – MEV Miners are receiving something of significant value from the Stacks ecosystem for essentially no cost.

- Decentralization – MEV Miners are excluding others from being able to equally participate in the Stacks ecosystem, centralizing a significant portion of the mining power, and setting the stage for further centralization.
- Competition – The MEV Strategy conflicts with the intended functioning and design of the Stacks blockchain. The mining process is intended to be a competitive process.

Given all the above, it would be ideal for Stacks to adopt a change to its mining algorithm such that the MEV Strategy, or expected adaptations of it, no longer proved desirable to engage in by potential MEV Miners.

1.3 Structure of this Analysis

This analysis is structured in two parts. Part I examines certain potential solutions to the MEV Strategy mainly from the perspective of whether they conceptually dissuade an MEV Miner from engaging in the MEV Strategy, or adaptations thereof. The hope is that this will allow us to narrow the potential solution set to one or two solutions following discussion and consultation with core developers for the Stacks blockchain.

Part II of this paper will be devoted to further understanding the potential consequences of these one or two solutions as well as examining these solutions given empirical Stacks mining data and modeled mining data.

2 Part I: Conceptual Analysis

2.1 Evaluation Criteria

As we begin analyzing potential solutions to this MEV Strategy, it's important that we define the various criteria through which we will examine solutions. For this analysis, we are generally assuming that miners are rational actors and may engage in a strategy if the expected value of that strategy is greater than zero, and won't engage in a strategy if the expected value is zero or negative.

The current MEV Strategy can be broken down into three separate behaviors that are suboptimal for the Stacks blockchain. As initial criteria to analyze potential solutions, we should consider whether these behaviors continue to offer a positive expected value to an MEV Miner. In Part I of this paper, we consider the following three questions with respect to any potential solution.

Can an MEV Miner have a positive expected value:

- With a de minimis Bitcoin commitment?
- While not mining consistently and only mining an MEV Miner's own Bitcoin block (we'll call this behavior "opportunistic mining")?
- By excluding other miners from participation?[1]

[1] Another, perhaps more robust version of this criterion, is whether an MEV Miner can have a positive expected value by excluding other miners from participation *while* committing less Bitcoin than the aggregate amount of a "typical" Honest Block.

In Part II of this paper, having narrowed down the potential solutions to one or two in consultation with Stacks blockchain core developers, we will further consider any such solutions through an analysis of empirical blockchain data as well as considering the solution's broader potential impact on the Stacks blockchain and ecosystem. These further considerations will include:

- How fundamental or disruptive is the change to the current mining process?
- Is the change expected to alter miner behavior in significant ways?
- Does the change disproportionately affect smaller miners versus large miners?
- Is the change expected to alter stacking payout amounts or behavior?
- Will the change cause mining to further centralize?

We'll examine each viable solution through the above lenses.

2.2 Overview of Proposed Solutions

At the risk of oversimplifying, the fundamental issue with the MEV Strategy is that MEV Miners are paying a little bit in, but receiving a lot out. With that in mind, we can think of four classes of solutions to this problem:

1. If a little is paid in, only pay out a little.
2. Require that enough is paid in.
3. Exclude miners who engage in the behavior.
4. Require a certain number of miners to participate.

The third and fourth classes of solutions aren't, as far as we can tell, workable at this time.

The third class, excluding miners who engage in this behavior, requires both being able to properly detect the behavior and making it costly to be excluded as a miner. Both of these are challenges and would likely require more significant changes to the Stacks blockchain than are reasonable at the moment.[1]

The fourth class of solutions, that rely on the number of miners participating in a block, is likely not feasible, given it can be trivial for a miner to spin up an additional miner or miners.

We've therefore focused below on three solutions that are either of the first or second class:

- The first solution is of the first class and adjusts a block's coinbase payout to be commensurate with the aggregate mining commitment paid in for a block.
- The second solution is of the second class and requires a minimum total Bitcoin commitment for a payout to happen at all.
- The third solution is again of the first class and makes the probability of a miner receiving a payout always commensurate with the probability of winning a "typical" block.

---

[1] We have not analyzed it in this paper, but an example of a potential solution in this class might be one that requires a miner to put at stake some number of Stacks tokens to be a miner. This staked amount could be subject to slashing (potentially by Stackers) in the event of bad behavior. A solution such as this might be considered for the long term if some version of MEV Mining persists following an initial change.

All of these solutions eliminate the expected value to MEV Miners of running the current MEV Strategy. They all also appear to sufficiently dissuade an MEV Miner from mining with a de minimis amount of Bitcoin and dissuade them from opportunistically mining.

However, none of these strategies eliminate some positive expected value for an MEV Miner to run a strategy whereby they consistently mine some material amount of Bitcoin and exclude others from their own block.

Nevertheless, as we'll see, exploiting the Stacks blockchain following implementation of these solutions may require not-insignificant capital for relatively small returns. It may be that, despite positive expected value for running certain MEV strategies, the cost of capital may be too high for an MEV Miner to bother running such strategies.

After we narrow down the potential solutions to one or two candidates, we will analyze these issues further in Part II.

2.2.1 Proportional Coinbase Rewards

2.2.1.1 Description of Solution

Aaron B. has proposed a solution whereby the Stacks coinbase reward amount is proportioned each block such that the amount a Stacks miner receives in a given block is proportional to the aggregate Bitcoin commitment that was provided in order to win that block. We'll call this the Proportional Coinbase Reward ("**PCR**") solution.

Currently, when a miner wins a Stacks block, they receive the full block-associated coinbase for the block that they win. With the MEV Strategy, MEV Miners are therefore able to commit a very small amount of Bitcoin, but win the entire Stacks coinbase. The PCR proposal would create the concept of a "reward window" and split a block's Stacks coinbase among the winners of some number of following blocks, proportioned based on the relative total Bitcoin commitments for each block.

To describe this solution more precisely, let's assume that Bitcoin and Stacks blocks match, and that block B_x is an example block that we're focused on. If the reward window is n blocks, then the following series of blocks will be relevant to determining the total coinbase reward for block B_x and the total rewards for the miner of block B_x.

$$\ldots \boxed{B_{x-n+1}} \rightarrow \ldots \rightarrow \boxed{B_x} \rightarrow \ldots \rightarrow \boxed{B_{x+n-1}} \ldots$$

In general, the coinbase reward distributed for any given block B_x is below, where commit(B_x) is the total amount of Bitcoin committed by all miners for mining that Stacks block, and coinbase(B_x) is the block-associated coinbase issued by the Stacks blockchain for that block:

The miner of block B_x's share of the coinbase for block B_x is:

$$\frac{commit(B_x)}{\sum_{i=x}^{x+n-1} commit(B_i)} * coinbase(B_x)$$

The miner of block B_{x+1}'s share of the coinbase for block B_x is:

$$\frac{commit(B_{x+1})}{\sum_{i=x}^{x+n-1} commit(B_i)} * coinbase(B_x)$$

And so on, until the miner of B_{x+n-1}, who is the last miner to receive a portion of the coinbase from block B_x.

The total coinbase reward that the miner of block B_x would expect to receive (from all blocks B_{x-n+1} through B_x), is as follows:

$$\frac{commit(B_x)}{\sum_{i=x-n+1}^{x} commit(B_i)} * coinbase(B_{x-n+1}) + ... + \frac{commit(B_x)}{\sum_{i=x}^{x+n-1} commit(B_i)} * coinbase(B_x)$$

Because we can generally assume that the block-associated coinbase for each block is equal, we can simplify this a little bit:

$$commit(B_x) * coinbase(B_x) * \left(\frac{1}{\sum_{i=x-n+1}^{x} commit(B_i)} + ... + \frac{1}{\sum_{i=x}^{x+n-1} commit(B_i)}\right)$$

With these totals established, we can examine the value an MEV Miner might expect to receive and their expected behavior as a result.

2.2.1.2 <u>Impact on MEV Miner Expected Value</u>

We need to now understand this PCR strategy through the lens of the initial three criteria we set forth in Section 2.1, above. As a reminder, we want to understand whether, with PCR in place, can an MEV Miner have a positive expected value:

- With a de minimis Bitcoin commitment?
- While opportunistic mining?
- By excluding other miners from participation?

To examine these questions, we'll make some initial simplifying assumptions:

- Block B_x will be an MEV Mined block, and all other blocks between B_{x-n+1} and B_{x+n-1} (inclusive) will be Honest Mined blocks.
- The value of Bitcoin commitments that the MEV Miner sends for the MEV Block we'll keep as a variable called C_s. Note that C_s must be greater than 0.
- The value of Bitcoin commitments that the protocol recognizes as receiving from the MEV Miner for the MEV Block we'll call C_r.
- The MEV Miner is mining opportunistically, and therefore commits no Bitcoin for any block in the series except B_x, and therefore $C_r = 0$.[2]

[2] The Stacks blockchain calculates a miner's commitment in each block as the lesser of the miner's commitment for that block (i.e., what the miner sends for that block) or the median of the miner's commitments over the last five blocks. Because we're assuming that the MEV Miner is opportunistically mining, we can assume that their median commitment over the five blocks before B_x is 0. And therefore the Stacks blockchain recognizes the miner's commitment as 0.

- The commitments in the Honest Mined blocks are approximately equal, and we'll call these commit(H).

With these assumptions, it's relatively clear that under the current version of the MEV Strategy, a miner would have a negative expected value. We can see this by examining the last equation in Section 2.2.1.1, above.

For that equation, commit(B_x) = C_r = 0. Therefore, the cumulative reward (for all blocks) that an MEV Miner would expect to receive for "winning" B_x with PCR in place would be 0. To determine an MEV Miner's expected value ("**EV**"), we would subtract away the miner's costs[3] from their expected rewards, and because C_s > 0, we see:

$$EV = 0 - C_s < 0$$

With this, we can be confident that PCR would be an effective fix to dissuade a rational miner from engaging in the MEV Strategy, as it's currently performed.

Next, though, we need to consider how a miner might adapt their strategy as a result of PCR being implemented. In particular, so long as an MEV Miner is an opportunistic miner, commit(B_x) = C_r = 0, and therefore expected rewards will always equal zero. We must consider whether an MEV Miner will then be incentivized to become a consistent miner (i.e., mining in non-MEV Mined blocks) and, if they do, will they have positive expected value by continuing to run the MEV Strategy during the blocks they mine.

To analyze this adapted strategy, let's make some new assumptions. In this case, so that we can simplify equations down to intuitive levels, we'll make some assumptions around block windows and various values:

- The reward window will be 5 blocks, so n = 5.
- Block B_x will be an MEV Mined block, and all other blocks between B_{x-4} and B_{x+4} (inclusive) will be Honest Mined blocks.
- The value of the block-associated coinbase is 1000 STX with STX worth $1.
- The total value of Bitcoin commitments for each Honest Block are all approximately equal, and we'll call this value commit(H).
- The MEV Miner is consistently mining and is sending the same commitment in all blocks in the series from B_{x-4} through B_{x+4}.
- The value of Bitcoin commitments that the MEV Miner sends for the MEV Block we'll keep as a variable called C.
- The value of Bitcoin commitments that the protocol recognizes as receiving from the MEV Miner for all blocks will also be equal to C.
- An MEV Miner can expect to break-even, in the long run, for the honest blocks that they are participating in.

[3] Additional costs for an MEV Miner not shown here are (a) the opportunity cost of the difference in Bitcoin fees between the fees offered by other Stacks miners and the lowest-fee transactions that the MEV Miner includes in their block, and (b) the opportunity cost of the Bitcoin fee of the highest-fee transaction excluded from the MEV Miner's Bitcoin block as a result of including their own Stacks mining transaction. We're assuming these costs are negligible for the purposes of this analysis.

With these assumptions, we can simplify the last equation in 2.2.1.1, above. The total rewards that an MEV Miner would expect to earn for winning block B_x would be:

$$TR = \frac{5 * C * coinbase(B_x)}{4 * commit(H) + C}$$

If, for the moment, we assume the during Honest Mined blocks, the total value committed approximately equals the value of the coinbase, in other words commit(H) ≈ coinbase(B_x), then the expected total rewards that the MEV Miner would receive for block B_x would be:

$$TR = \frac{C * 5000}{4000 + C}$$

If we subtract out the Bitcoin commitment C that the MEV Miner spends to earn these total rewards, we see the MEV Miner's total expected value for block B_x is:

$$EV = TR - C = \frac{C * 5000}{4000 + C} - C$$

If we then graph this EV function, we see that for any Bitcoin commitment C spent by an MEV Miner between $0 and $1000, the expected value is positive for block B_x.

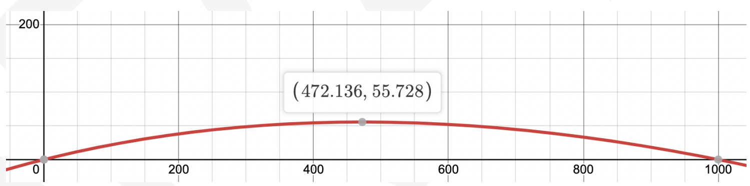

Given we assume that the MEV Miner is at least breaking even with respect to the Honest Blocks, we can see that under the assumed constraints, an MEV Miner would still have a positive expected value by consistently mining Honest Blocks while excluding other Stacks miners from the Bitcoin block that they mine.

If we adjust the assumption for commit(H) to something more realistic, we'll see this curve becomes steeper. As an example, assume that for Honest Blocks, the value of cumulative Bitcoin commitments *plus* cumulative Bitcoin transaction fees should approximately equal the value of the coinbase.

Let's assume that a typical Honest Miner spends 5 satoshis for every 95 satoshis they commit in mining.[4] We would expect that the typical cumulative commitment for the an Honest Block would be .95*coinbase(B_x).

Therefore the EV, under all other assumptions above, would be:

[4] Following drafting of this Part I, we analyzed fee data as part of our analysis in Part II. We found long-term, honest miners spent closer to 6% in fees, though we have not revised this Part I. We believe the previously-assumed 5% continues to provides results that are directionally appropriate.

$$EV = \frac{C * 5000}{.95 * 4000 + C} - C$$

We see the graph for this EV function looks similar to the one above, except with a higher expected value:

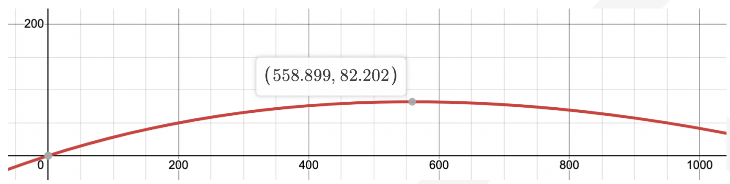

We see that the expected value increases as Bitcoin fees (as a percent of total Bitcoin committed) increases.

The next question to consider is whether this conclusion as to expected value changes if we push on any of the various of the assumptions we lay out above. Let's take each in turn:

- **Reward Window**: If you increase the reward window, the EV curve essentially flattens (and, conversely, the EV curve steepens if you decrease the reward window). As a result, adjusting the award window doesn't change the nature of the conclusion, just the degree of profits.[5]
- **One MEV Block**: If you have more than one MEV Block in the reward window, this would decrease the denominator in the reward calculation and thus lead to a higher expected return, again not impacting the conclusion.
- **Value of the Block-Associated Coinbase**: This is somewhat arbitrary and neither increasing nor decreasing significantly affects the relative outcome.
- **Value of Bitcoin Commitments**: As we show above, the lower the value of the cumulative Bitcoin commitments in Honest Blocks, the higher the expected value for MEV Miners.

All in all, we think the assumptions are conservative, and the expected value curve shown above would more likely serve as something of a floor on expected value.

One advantage of PCR, though, is that it potentially increases the rewards allocated to Honest Miners during Honest Blocks.

2.2.1.3 Key Takeaways

To go back to our original three criteria, under the assumptions above, we can reasonably conclude that the PCR approach would:

- Likely dissuade an MEV Miner from sending a de minimis Bitcoin commitment (as it would lead to a de minimis return);

[5] Note that it's not necessarily better to increase the reward window while holding all else the same, even though it appears to reduce profitability. For one thing, increasing the reward window would increase the probability of another MEV Block within the reward window which would serve to increase profits for an MEV Miner.

- Dissuade an MEV Miner from opportunistically mining; but
- Potentially not dissuade an MEV Miner from excluding other Stacks miners from their Bitcoin-mined block, depending on whether the expected value is high enough.

The last question will go to the cost of capital for the MEV Miner. While currently, the returns for the MEV Miner running the MEV Strategy as it exists today provide astronomical returns for de minimis capital outlay. For an MEV Miner to see any meaningful returns with PCR in place, they would need to outlay more significant capital, and consider the potential returns they could achieve with that capital outside of employing a version of the MEV Strategy on the Stacks blockchain.

We also need to consider the relative advantage that PCR offers by increasing the expected rewards to Honest Blocks. This may mitigate some of the downsides of MEV Miner behavior even if they continue to engage in excluding other miners during their own blocks.

If we determine that PCR warrants additional examination in Part II of this report, we will examine these questions further.

2.2.2 Minimum Bid Requirement

2.2.2.1 Description of Solution

Jude N. has proposed a minimum bid requirement ("**MBR**") solution, whereby there needs to be a minimum amount of aggregated Bitcoin funds for a particular Bitcoin block, otherwise no Stacks sortition takes place.

For a given block B_x, the MBR solution would consider the Stacks block commitments for the previous n Bitcoin blocks (so blocks B_{x-n} through block B_{x-1}) to determine a minimum threshold for the aggregate block commitment that would be required for there to be a sortition for block B_x. There would exist some function F, whose input would be the commitments over the n blocks, and whose output would be the minimum Bitcoin commitment required for a sortition for block B_x to take place.

Neither the number of blocks, n, nor the function, F, are specified in the proposal, but one potential framework that is suggested by the proposal is to identify the median of Honest Blocks, and set the threshold at or around that median.

2.2.2.2 Impact on MEV Miner Expected Value

We'll start by examining the MBR solution without specifying n or F, given that we can analyze this solution as a class. Once again, we want to understand whether, with MBR in place, can an MEV Miner have a positive expected value:

- With a de minimis Bitcoin commitment?
- While opportunistic mining?
- By excluding other miners from participation?

Given the entire class of MBR solutions is premised on a minimum Bitcoin commitment to participate in the sortition, it would seem the first criterion is satisfied by MBR. In other words, if MBR requires more than a de minimis Bitcoin commitment to participate in sortition, then an MEV Miner will not win anything with a de minimis Bitcoin commitment, and therefore an MEV Miner will have a negative expected value when making a de minimis Bitcoin commitment.

Similarly, it would seem MBR would dissuade an MEV Miner from mining opportunistically. If an MEV Miner only mined their own block, then it's likely the case that the median of the MEV Miner's bitcoin commitments for the previous five blocks would be 0. The Stacks blockchain would then consider the MEV Miner's commitment in its own block as 0, which would presumably not meet the (non-zero) minimum Bitcoin commitment required by MBR.

The third question, whether MEV Miners are still incentivized to exclude other miners from their block, gets a little bit trickier. To consider this, we'll make some assumptions and define some terms:

- The value of the Stacks coinbase for any block B_x is coinbase(B_x).
- The amount an MEV Miner commits in all blocks is consistent and equal to C.[6]
- Implementation of MBR results in a minimum required commitment in block B_x of min(B_x).

If we again assume that the MEV Miner can, in the long term, at least break-even for the Honest Blocks, then the immediate question is what does the expected value for an MEV Miner look like by implementing a strategy of consistent mining, but excluding other miners from their Bitcoin-mined block.

We see the expected value for block B_x would be:

$$EV = \begin{cases} coinbase(B_x) - C, & \text{if } C \geq min(B_x) \\ -C, & \text{if } C < min(B_x) \end{cases}$$

The graph of this function looks like the following:

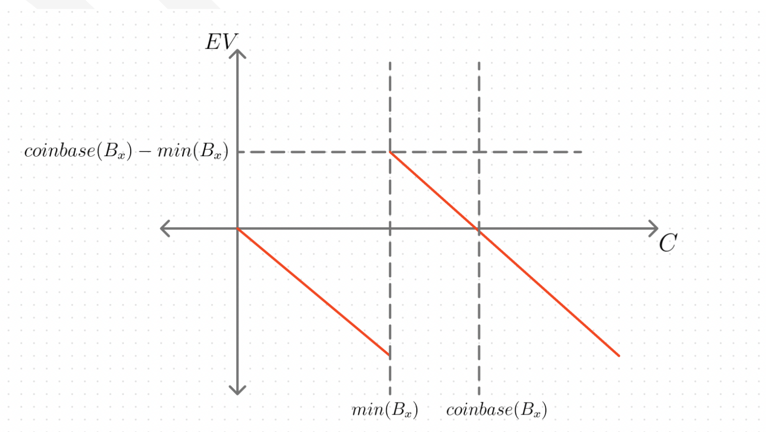

[6] As discussed above, MBR clearly dissuades opportunistic mining, so we assume consistent mining.

We immediately see that there exists a positive expected value for the miner when $\min(B_x)$ is less than $\text{coinbase}(B_x)$, and C is somewhere between those values, with expected value at its maximum at $C = \min(B_x)$.

It follows that for MBR to be effective, it would need to seek to keep the value of $\min(B_x)$ as close to the value of the coinbase as possible. The larger the spread, the greater the incentive an MEV Miner would have to exploit the gap.

With a sufficiently high $\min(B_x)$, it may be that MBR serves to effectively dissuade an MEV Miner from consistently mining at such a high bitcoin commitment relative to the value of the coinbase. Nevertheless, that may also serve as a disincentive to other miners to engage in mining as well.

Conceptually, if we think about the bitcoin committed for blocks in which $\min(B_x)$ is not met as an added cost for both MEV Miners and Honest Miners, that added cost increases as $\min(B_x)$ increases. Inevitably, increasing $\min(B_x)$ will decrease overall miner profitability.

The key question, which we anticipate considering in Part II should we decide to examine MBR further, is whether there is much room for decreasing miner profitability and whether doing so might increase centralization. If there is potential room for decreasing miner profitability, we anticipate considering whether there may be a sufficiently high $\min(B_x)$ that dissuades an MEV Miner from excluding other miners, but does not create an unreasonably high increase on the cost of mining more broadly.

Lastly, we should consider one other strategy that implementing MBR may incentivize an MEV Miner to adopt. If MBR is implemented, it may make sense for an MEV Miner to consistently mine at some $C < \min(B_x)$, but in the MEV Miner's own Bitcoin block include other miners' transactions only to the extent that the total commitment is just greater than $\min(B_x)$, and exclude all other miner transactions. Doing so would not guarantee the MEV Miner would win, but would increase their own probability of winning, and therefore their expected value. This again points to the need for $\min(B_x)$ to be fairly close to the value of the coinbase.

2.2.2.3 Key Takeaways

MBR clearly dissuades de minimis mining and opportunistic mining, and it very much has the potential to dissuade an MEV Miner from excluding other miners, provided the minimum threshold is sufficiently high.

The primary countervailing issue, however, is that as the minimum threshold increases, the cost of mining for all miners will increase. This could lead to further miner centralization.

There may be some tolerable increase in miner cost that sufficiently disincentivizes an MEV Miner, but this would require further consideration in Part II.

2.2.3 Assumed Total Commitment

2.2.3.1 Description of Solution

We are proposing this solution, which we'll call Assumed Total Commitment ("**ATC**"). With ATC, for any block B_x, there would exist some function F whose input would be the total commitments over each of the prior n blocks and output some assumed total commitment, which we'll call assumed(B_x). As a start, we'll assume the function F and number n are the median total block commitments over the last 5 blocks.

Whereas currently the Stacks blockchain picks a miner randomly based on the percent of the total Bitcoin block commitment that a miner commits (i.e., an individual miner's commitment divided by the total commitment), with ATC the denominator would become the greater of total commitments or the median total block commitments over the last 5 blocks. If no miner is randomly chosen (which would be possible with ATC), then there will be no sortition for block B_x.

To restate the basic idea of this proposal, regardless of whether one miner (who blocks out others) or many miners participate in mining a block, there will always be a baseline assumed total commitment. The highest probability any miner could ever have of winning a block would be their commitment divided by the median of total commitments over the last five blocks.

2.2.3.2 <u>Impact on MEV Miner Expected Value</u>

We'll, yet again, start by examining this ATC proposal through the lens of the three criteria discussed in Section 2.1. With ATC in place, will an MEV Miner have a positive expected value:

- With a de minimis Bitcoin commitment?
- While opportunistic mining?
- By excluding other miners from participation?

Without getting too deep into the equations, we can develop some intuition for the first two questions. With a de minimis Bitcoin commitment, we would expect that an MEV Miner would have a de minimis probability of winning the sortition.

On the question of opportunistic mining, once again if an MEV Miner mined their own block, then the likely value of the median for the MEV Miner's commitments over the previous five blocks would be 0. As a result, the MEV Miner would likely have a 0% chance of winning their own block, and would presumably be dissuaded from opportunistically mining.

To examine the third question, and the first more rigorously, we'll again make some assumptions and define terms:

- The value of the Stacks coinbase for any block B_x is coinbase(B_x).
- The amount an MEV Miner commits in all blocks is consistent and equal to C.
- Implementation of ATC results in a minimum sortition denominator of assumed(B_x).
- The total amount of Bitcoin committed in block B_x is commit(B_x).

If we once again assume that an MEV Miner is able to at least break even during Honest Blocks, then we're focused on the MEV Miner's expected value for their own Bitcoin-mined block, B_x.

$$EV = \frac{C}{max(assumed(B_x), commit(B_x))} * coinbase(B_x) - C$$

Initially, let's assume that assumed(B_x) = coinbase(B_x), which translates to the assumption that the cumulative Bitcoin committed by all miners during recent Honest Blocks equals the value of the coinbase. While this likely isn't true, it's a place to start.

With that assumption, we see that regardless of the amount of the MEV Miner's commitment, the EV function simplifies to EV = C – C, so EV = 0.

More likely, though, for a typical block, the cumulative Bitcoin commitments *plus* cumulative Bitcoin transaction fees should approximately equal the value of the coinbase. We'll again assume an Honest Miner spends 5 satoshis for every 95 satoshis they commit in mining. We would expect that the typical cumulative commitment for the median block, B, would be .95*coinbase(B). If we insert this into our expected value equation, above, we get:

$$EV = \frac{C}{.95} - C$$

The graph of this equation is as follows:

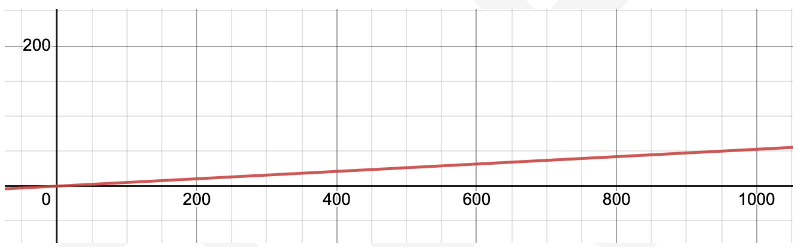

As we can see, there is a positive expected value for MEV Miners to engage in this strategy. Under these mining assumptions, the return would be approximately 5% on the capital committed in each MEV Mined block. This return increases if we assume a higher percent of miners' commitments go to fees, and decreases if we assume the opposite.

It is possible to construct a function that provides further penalty the lower C is relative to assumed(B_x) thereby creating a negative expected value, essentially up until C = assumed(B_x). The problem with a strategy such as this, however, is that much like MBR, this would create an additional cost for Honest Miners, thereby making mining overall less profitable, potentially leading to lower participation and greater centralization.

2.2.3.3 Key Takeaways

ATC, similar to PCR, likely dissuades an MEV Miner from sending de minimis Bitcoin commitments, as it will lead to de minimis expected returns. Similarly, ATC dissuades an MEV Miner from mining opportunistically.

ATC, like both other proposed solutions, can offer potential positive expected value for an MEV Miner who wishes to consistently mine some material Bitcoin commitment.

Under our various assumptions, it would appear the expected rate of return for ATC at ~5% is significantly lower than the expected rate of return estimated for PCR which, depending on the amount of Bitcoin committed by the MEV Miner, could be as high as 5x or more. We did not make sufficient assumptions to estimate a rate of return for MBR.

2.3 Core Developer Meeting and Part II Analysis

Following drafting of Part I, we discussed the above findings with Stacks core developers and which solutions made sense to examine further in Part II. We ultimately decided on examining PCR and ATC further.

We agree with this decision. While we believe MBR could serve as a powerful deterrent to the MEV behavior, were the minimum set high enough, we believe the cost to miners generally of doing so may be too high.

When we consider that, while the MEV Strategy is not ideal, it is not catastrophic for the Stacks blockchain, which has continued to run smoothly. This points to a solution that is incremental in its approach. We believe that PCR and ATC, while potentially continuing to offer some benefit to MEV Miners who wish to implement an adapted version of the MEV Strategy, are solutions that could dissuade MEV Miners entirely and are minimally disruptive to the operation of the Stacks blockchain. We'll therefore analyze these two solutions further in Part II.

3 Part II: Empirical Analysis

3.1 Overview of Part II

As discussed above, in consultation with Stacks blockchain core developers, we've decided to further analyze the Proportional Coinbase Reward and Assumed Total Commitment solutions in this Part II. As a reminder, this Part II will review historical and modeled mining data to analyze how implementation of PCR and ATC may affect miners, their expected returns, and the Stacks ecosystem more broadly.

To start, we first performed a high-level analysis of mining data to-date to set our baselines. Initially, our goal is to understand how many miners have participated in mining, how frequently they've participated, what their mining commitments and fees have looked like, and their profitability.

We'll then retroactively apply PCR and ATC to the mining data to see how these strategies might have affected miner profitability overall and on the segments we've introduced.

3.2 Historical Mining Data

3.2.1 Miner Count and Commitment Count

We received and analyzed mining data from the start of the Stacks 2.0 blockchain (when the current version of mining for the Stacks blockchain was introduced) on January 14, 2021, through to data encompassing Bitcoin block 784,957, which was mined on April 11, 2023.

During that period, 271 distinct miners mined the Stacks blockchain. They, in total, made around 680,000 mining commitments to the Stacks blockchain, with some miners making only one commitment and the maximum number of commitments coming from a miner who has made approximately 59,000 commitments.

Miners, in aggregate, have sent a total of approximately 2600 Bitcoin in mining commitments and spent an additional approximately 150 Bitcoin in fees.

When we examined the number of commitments per miner, we noticed that the distribution was heavily skewed to miners who made relatively few commitments. In other words, a large percent of the 271 miners have made a small number of the total commitments.

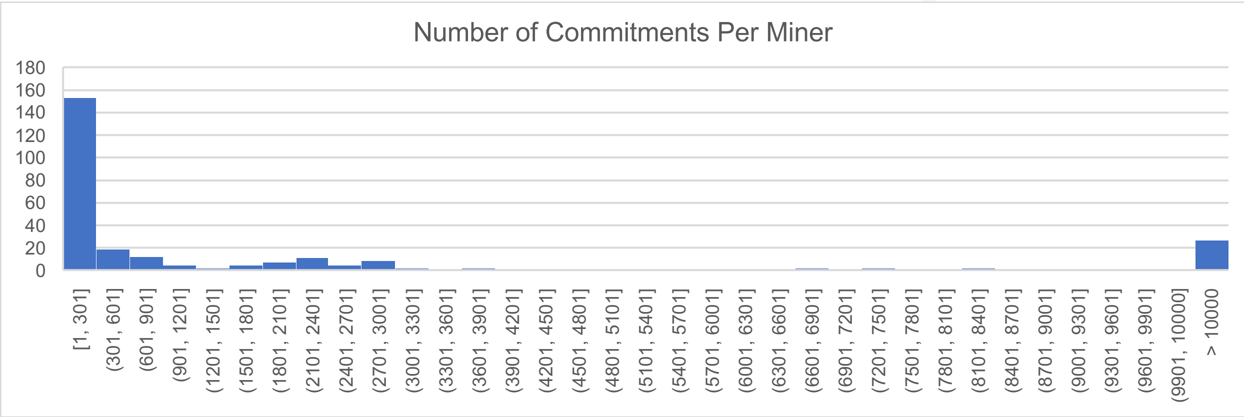

We see from this histogram that over half the miners had fewer than 300 block commitments. With a typical Stacks block time of about 11.5 minutes, 300 blocks equates to mining for less than three days. We're therefore calling call these miners "tourist miners."

These tourist miners account for 56% of all miners, but less than 1.5% of total commitments.

As a result, for the analysis going forward, we've removed these 153 tourist miners from the data and concentrated on the impact of potential solutions to the remaining 118 miners that mined more consistently throughout the period. We'll call these miners "long-term miners." We believe retaining tourist miners in the dataset would skew the results towards optimizing a solution for them, rather than the long-term miners.

Here's the distribution of miner commitments with tourist miners removed:

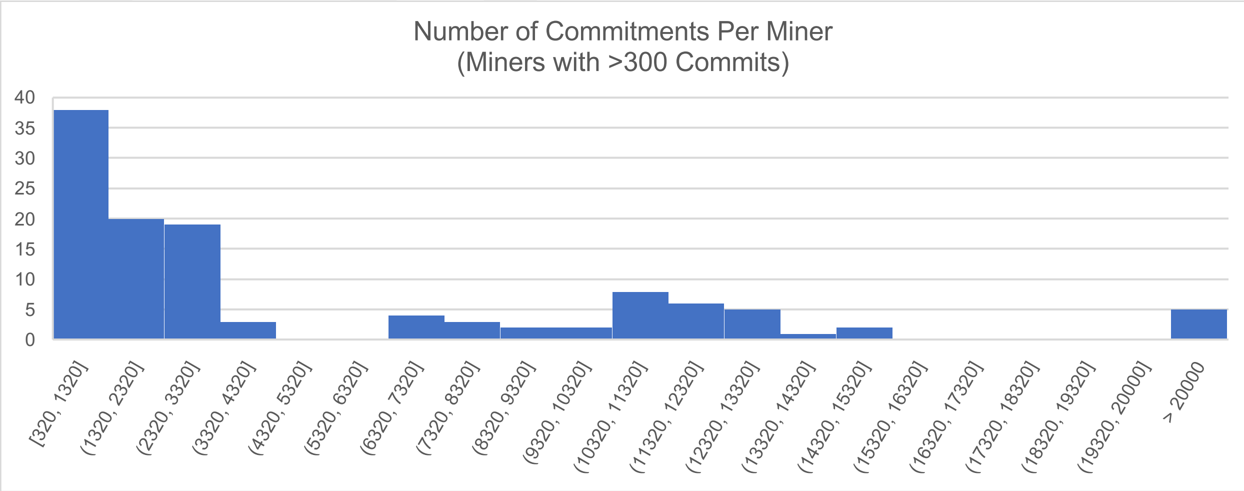

3.2.2 Historical Profitability

To examine profitability – both historical and projected for the solutions we'll analyze – we used the average daily price of Stacks and Bitcoin for each day during the period. While we'd ideally have the prices at the time miners sent commitments for each commitment, we believe that given the number of data points, the average daily price should serve as a reasonable baseline to directionally understand both historical and projected profitability.

In addition to the tourist miners, we've also excluded the one MEV Miner from our profitability analysis. First, from a baseline perspective they are not representative of the miner population. And, second, the solutions we propose will have, as designed, significant negative impact on their profitability. We did not want this negative impact to be interpreted as impacting typical miners.

Over the examined period, the long-term miners have cumulatively spent approximately $110.4 million (inclusive of fees), and earned approximately $112.8 million, resulting in total profits of approximately $2.4 million.

The long-term miners' profits ranged from approximately $1.1 million in losses for one miner to approximately $750,000 in profit for another.

To examine the long-term miner population, we broke the group down by quintiles based on the average value of each long-term miner's block commitments. If a long-term miner's average block commitment was in the bottom 20th percentile, they would be in the 1st Quintile, if the average commitment was between the 20th and 40th percentile, they would be in the 2nd Quintile, and so on.

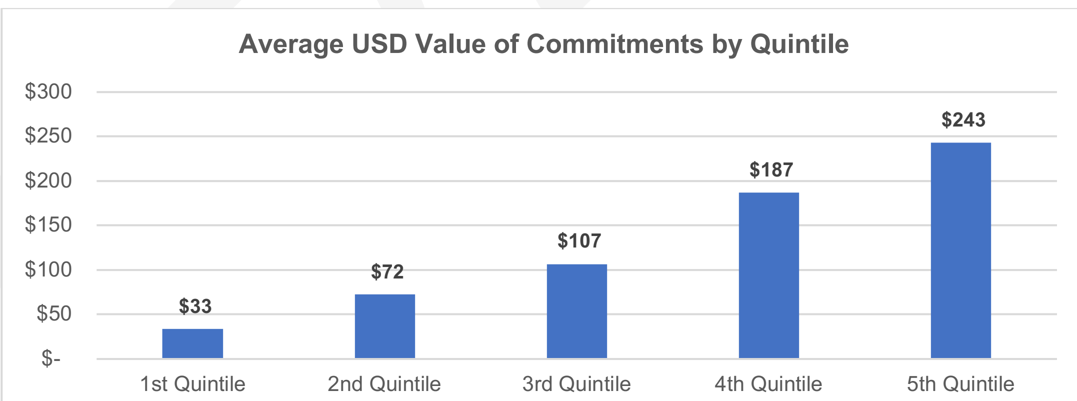

When we examine each quintile we see that the percent of miners who are profitable within each quintile range from around 48% to 67%, with no obvious correlation to the average commitment value for the miners.

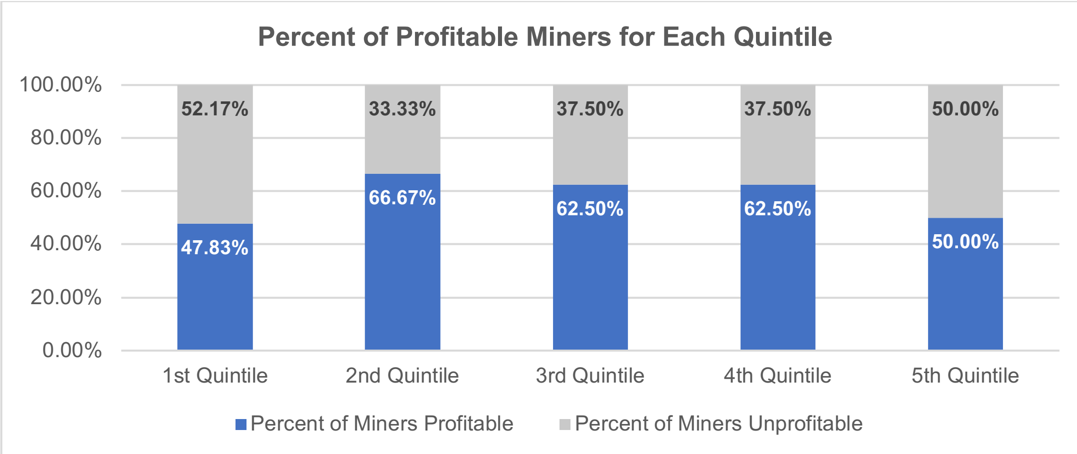

We also see that the distribution of profits does not appear related to the average commitment amount. For our baseline profitability, the 2nd and 4th quintile miners earned the majority of profits, while the 3rd quintile experienced losses.

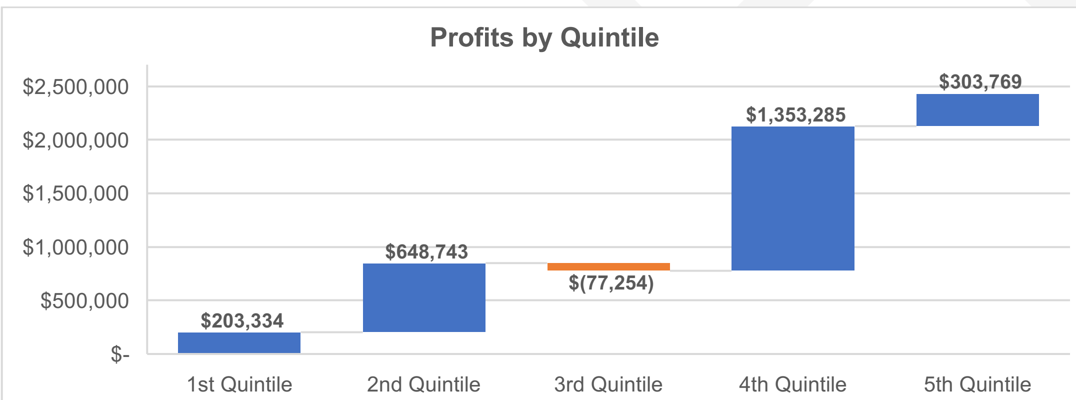

Lastly, as we examine the impact on quintiles, we should keep in mind the percent of all commitments each quintile is responsible for. We see that the top two quintiles comprise over a majority of historical commitments, with the remaining commitments fairly equally distributed among the bottom three quintiles.

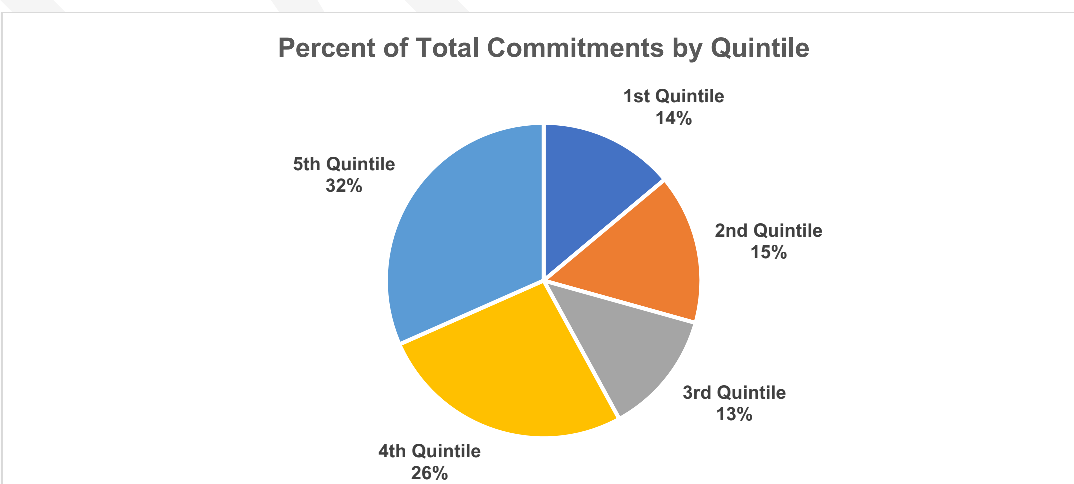

3.2.3 Historical Profitability – Trimmed Data

While we found this baseline data, presented above, to be a helpful initial framing, we noticed that the extremes of the mining data had the potential to significantly skew the overall data. There appeared to be some outlier miners on both the unprofitable and profitable end of the spectrum, that we were worried might skew our data towards a solution that could benefit them at the expense of more typical Stacks miners.

As a result, we trimmed the data set of both the top and bottom 5% of miners based on profits per commitment. In other words, both the extremely profitable and extremely unprofitable miners were removed from the trimmed data. This left a total of 105 long-term miners in the data set.

For completeness, we examined our proposed solutions from both the perspectives of the untrimmed and trimmed data sets. In general, we believe the trimmed data comparison is more useful, as it focuses on the middle 90% of miners, but we'll reference any significant deviations with the untrimmed data set, when applicable.

For the baseline historical profitability, let's examine the four charts from Section 3.2.2, above, for the trimmed data:

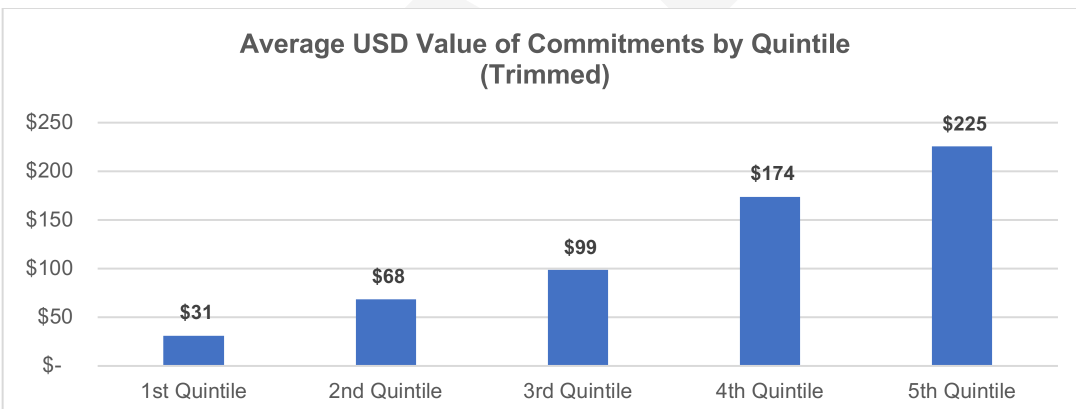

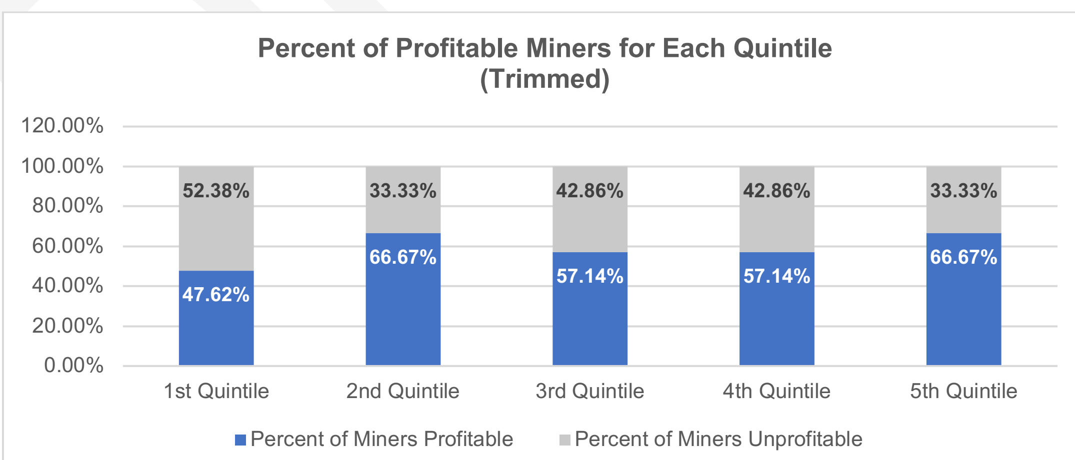

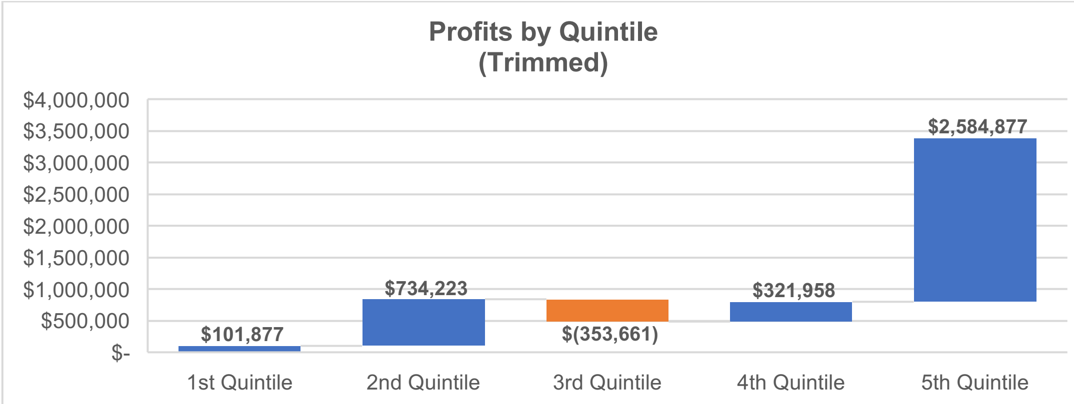

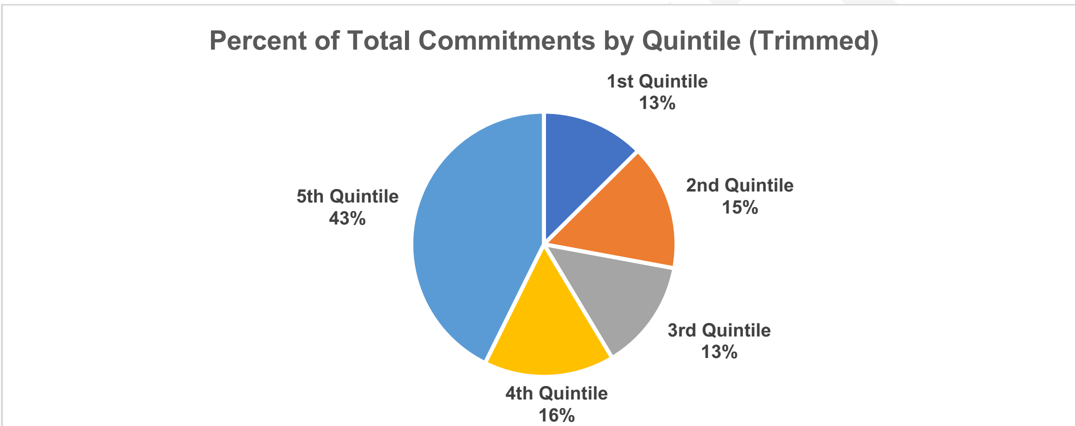

We see that, in comparison to the untrimmed data, the overall profitability increases significantly for the trimmed data. This is the result of the fact that the cumulative unprofitability of the unprofitable miners removed from the data was greater than the cumulative profitability of the removed profitable miners.

We see that the 4th and 5th quintiles are significantly impacted, with the 4th quintile decreasing in profits and the 5th increasing. Moreover, the percent of profitable miners increased to 67% from 50% for the top quintile, but otherwise remained relatively stable.

We also see that the percent of commitments attributable to the top quintile increases, while that percent for the 4th quintile decreases. For the trimmed data, the top quintile accounts for 43% of commitments, with the remainder somewhat equally distributed among the bottom four quintiles.

3.3 Projected Miner Profitability

We now want to examine the potential impact to profitability for non-MEV, long-term miners with certain of our potential solutions implemented.

3.3.1 Proportional Coinbase Rewards

We'll start by examining the Proportional Coinbase Reward ("PCR"), discussed in more depth in Section 2.2.1, above.

As a reminder, the PCR solution modifies the coinbase awarded for a given block based on the aggregate amount of Bitcoin committed by miners for that block. The higher the aggregate commitments for a block, the greater the coinbase awarded to the winner of that block. A block winner will also gets a portion of the coinbase rewards from a certain number of preceding blocks that have been mined, again proportional to the total amount of Bitcoin committed for the block that the miner won.

3.3.1.1 Methodology

We assumed for this analysis of PCR that the block window is 5 blocks. Under that assumption, we overlayed the PCR reward distribution for historical miner commitments.

We considered whether historical miner behavior was best representative of anticipated miner behavior under PCR, or whether miners may behave materially differently. We ultimately concluded that we have low confidence modeling any strategic behavior changes miners may adopt as a result of PCR, and believe the highest confidence behavior projection would be based the historical mining data.

Some potential strategies miners may adopt with PCR could include either increasing or decreasing commitments as a result of relatively small or large total commitments.

Keeping in mind that the Stacks protocol considers a miner's commitment to be the lesser of the miner's commitment in the block or the median of the miner's commitments over the previous 5 blocks, it's not obvious that miners would necessarily adopt any of these strategies with PCR in place over the relatively consistent mining strategy that miners appear to have adopted to-date.

This approach also appears to be supported by the distribution of block rewards with PCR implemented. In other words, to the extent PCR would adjust a miner's behavior, it would likely be a result of the expected rewards a miner could earn from a given block. With PCR implemented, we see that the rewards remain tightly clustered around 1000 Stacks:

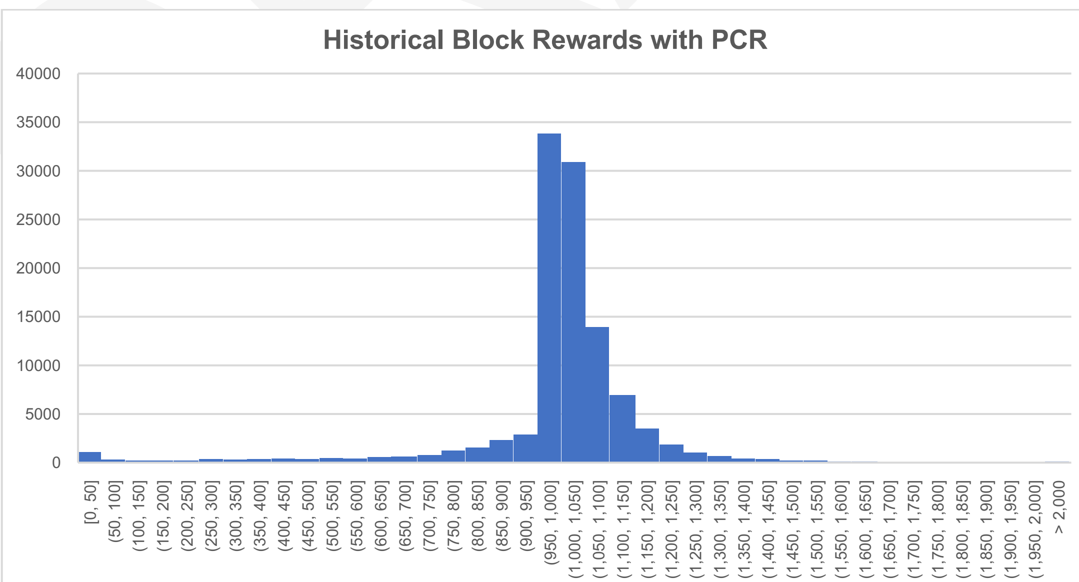

As a result of all the above, we believe it is most reasonable to rely on historical data in examining the impact of PCR.

3.3.1.2 Results

As we might expect, the overall profits remain consistent with our baseline data, with a slight uptick given, we believe, the transfer of block rewards from the MEV miner to our long-term miners.

When we compare the expected profits for PCR with the baseline historical profits for each quintile, we see that while the 4th and 5th quintile are largely unaffected, the 2nd and 3rd quintiles see significant increase in expected profits (40% and 50%, respectively). And expected profits for the bottom quintile drop 98% from about $102,000 in profits to about $2000 in expected profits. The direction and proportional magnitude of these changes are similar with the untrimmed data as well.

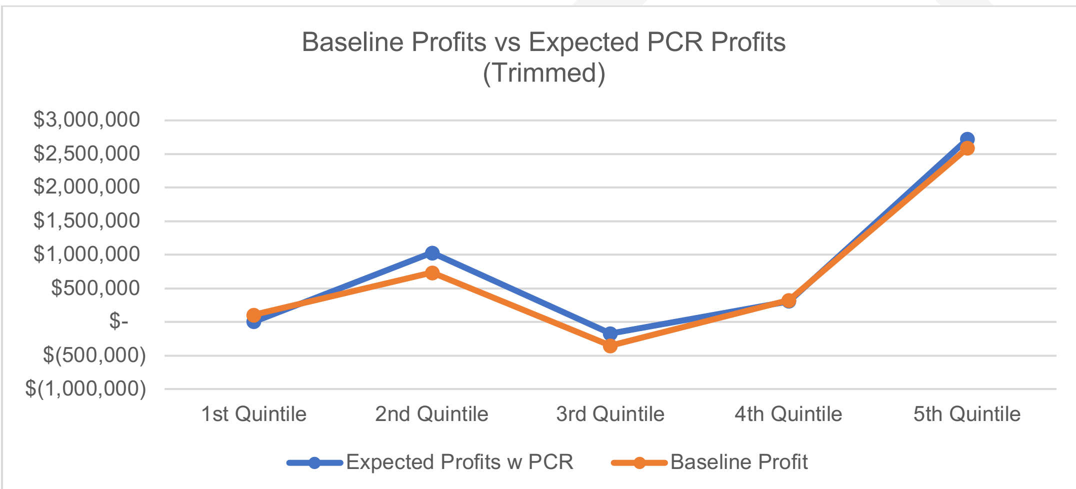

The percent of profitable miners in the bottom quintile decreases as well, from about 48% to 43%. The 3rd quintile also drops, while the remainder either stays the same or increases.

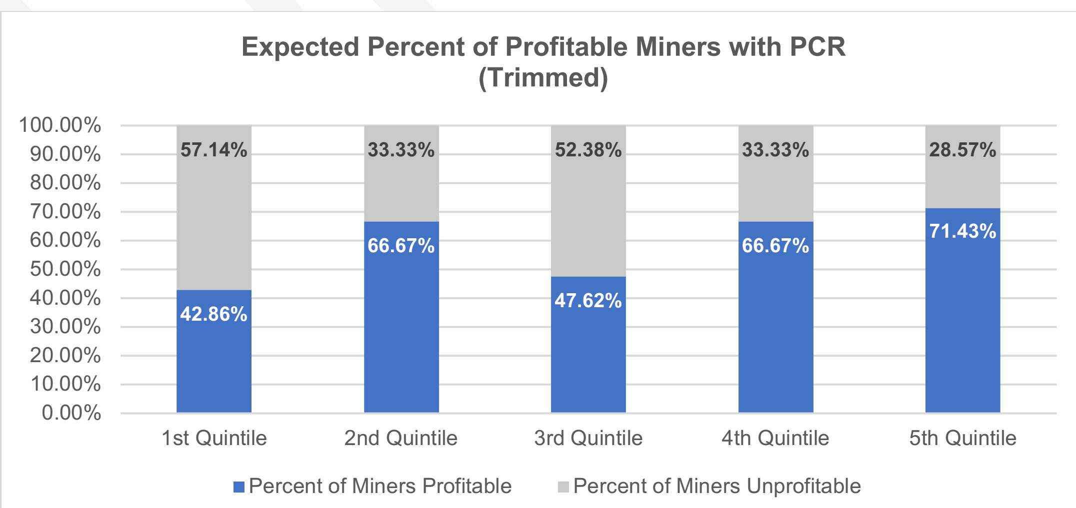

Overall, if we assume historical miner behavior remains consistent with future miner behavior once PCR is implemented, then the expected profits under PCR are largely consistent with the current mining algorithm. The exception to this is for the bottom quintile, which sees a significant percent reduction in expected profits.

One potential reason for the drop to the bottom quintile miners might be that with smaller average commitments, the bottom quintile miners would be more likely to win blocks with smaller aggregate commitments. If they're winning blocks with smaller aggregate commitments, then we would expect that their coinbase rewards would decrease as a result of the implementation of PCR. This, in turn, would be expected to decrease their total profitability as well as the percent of bottom quintile miners that are profitable. We discuss the implications of this further in Section 3.4 (Ecosystem Considerations), below.

3.3.2 Assumed Total Commitment

We'll next examine the potential impact on miner profitability of implementing the Assumed Total Commitment ("ATC") solution, described in Section 2.2.3, above. Again as a reminder, ATC assumes some minimum baseline total commitment for a block, where a miner's probability of winning that block is, at most, the value of their commitment divided by that minimum baseline. If the value of the total commitments in a block is less than that minimum baseline, then there's a probability that no miner is chosen, in which case there is no sortition for that block.

3.3.2.1 Methodology

We assumed for this analysis that the assumed total commitment for a given block is the greater of the median of the total commitment from the last 5 blocks, and the total commitment for that given block.

Again, as with PCR, we focused on analyzing the impact of ATC using historical data. Historical data is arguably even more appropriate for ATC, given that the dynamics of mining should be almost identical to the current mining process.

The primary way in which ATC would affect historical mining data is by creating the potential for miners to have not won any rewards for blocks they had historically won. We would expect ATC to reduce profitability to some extent for miners.

To model lost blocks, we first calculated the theoretical number of lost blocks based on historical commitments. We found that we should expect approximately 6.01% of blocks to be missed as a result of the implementation of ATC.

We then ran a Monte Carlo simulation for 1000 rounds of historical data to see the variation of the number of missed sortitions. We see that the number of missed sortitions stayed relatively close to the theoretical 6.01%, staying in a range of $\pm$ 0.15%.

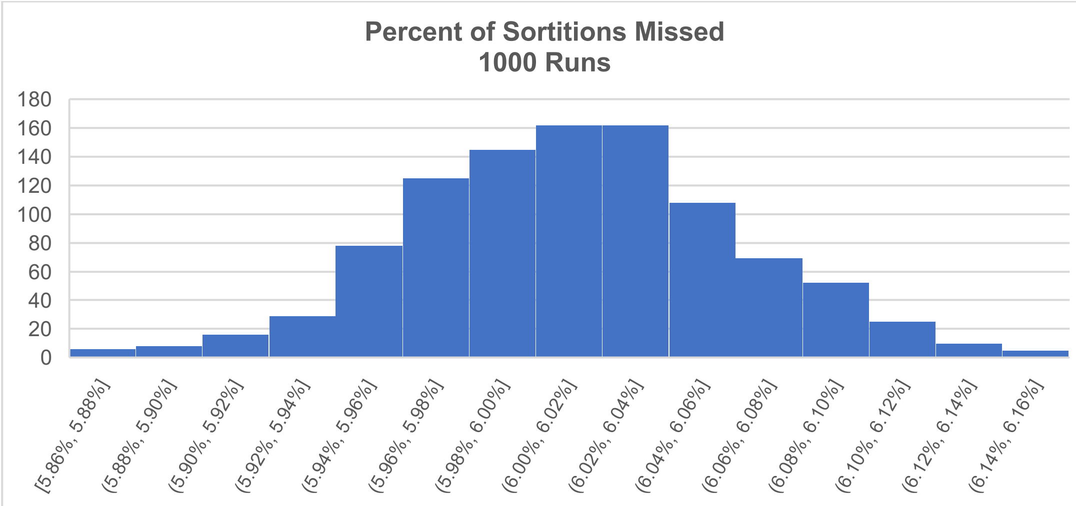

3.3.2.2 Results

We overlayed our missed sortition analysis for ATC to examine resulting miner profitability and found some dramatic results. In particular, the approximately 6% of missed blocks ended up having a profound impact on miner profitability.

Overall profitability for the trimmed miner group went from about $3.4 million in profits to about $1.4 million in losses, for a total decrease of about $4.6 million in profits. The most dramatic effect was for the top quintile.

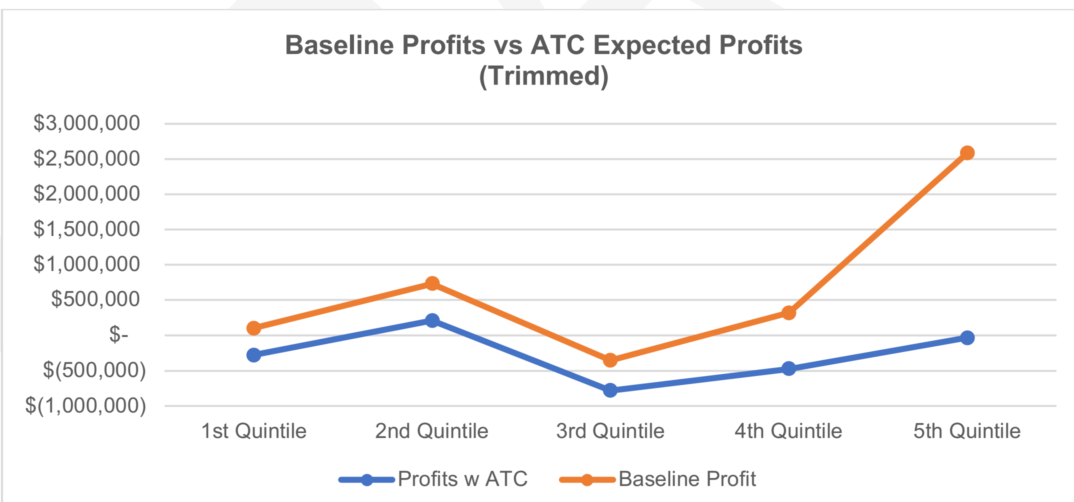

As one might expect given the above, the percent of profitable miners for each quintile dropped as well. Each quintile, except the 2nd, saw a dramatic decrease in the percent of profitable miners. The 2nd quintile saw a relatively more modest decrease.

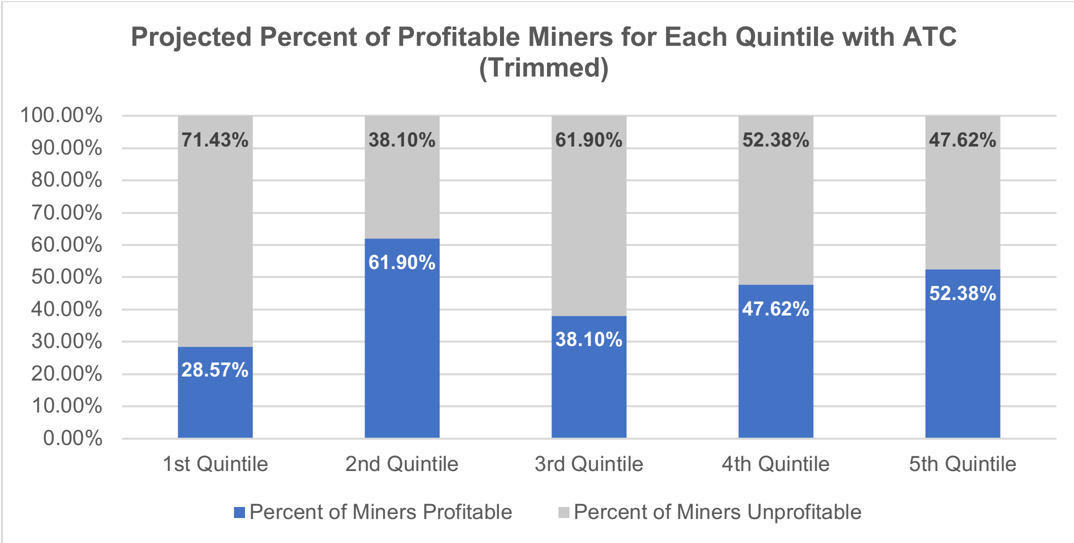

These results were directionally consistent with the untrimmed data as well.

Plainly, ATC in its current form would not be a workable solution for miners, given the sharp decrease in expected profits. With ATC implemented, we would likely expect lower miner commitments per block and likely fewer miners participating in sortition.

At its core, the issue with ATC is that the coinbase from the 6% of missed blocks is never distributed to miners, even though value was committed by miners for that coinbase.

In the next section we'll examine a proposed modified version of ATC that allows for the eventual distribution of the coinbase from a block with a missed sortition.

3.3.3 Assumed Total Commitment with Carryforward

Given that ATC appears unworkable in its standard form, we propose a slightly modified version of ATC whereby a missed sortition results in the coinbase from that block being added to the coinbase in equal increments over some number of following blocks. We're calling this "ATC carryforward."

3.3.3.1 Methodology

We've applied the same methodology here as with ATC, above, with the added assumption that the coinbase for a missed block will be distributed in equal increments over the following 5 blocks. So if a block has a Stacks coinbase of 1000 blocks and is missed, then 200 additional Stacks will be added to the coinbase of each of the following 5 blocks.

3.3.3.2 Results

The results of ATC carryforward end up looking a lot like the results of PCR. Again, there's a slight uptick in overall profitability, likely the result of MEV awards being allocated to long-term miners as a result of implementing ATC carryforward.

The distribution of ATC carryforward profits between quintiles again favors all but the bottom quintile, though the change for the bottom quintile with ATC carryforward is an approximately 82% decrease rather than a 98% decrease with PCR.

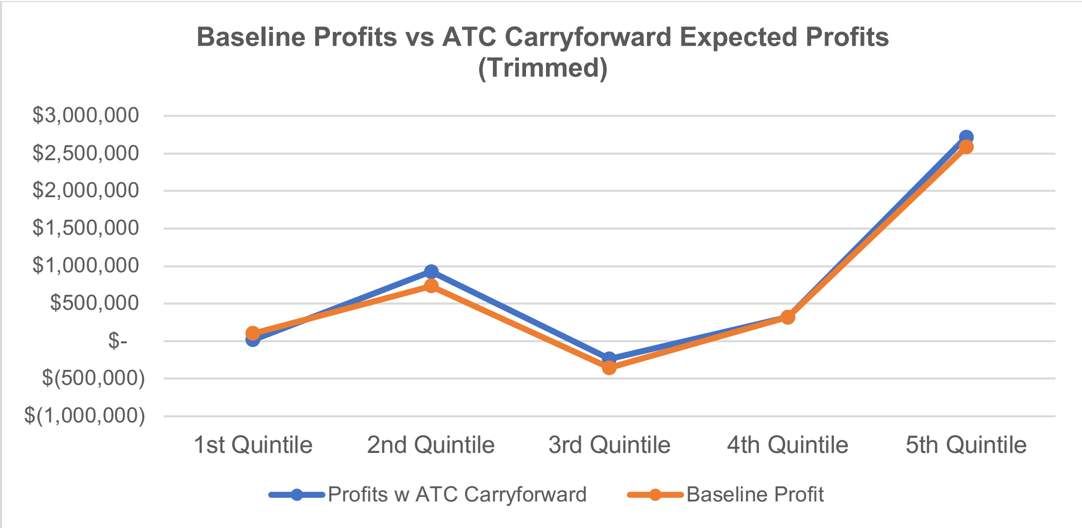

The expected percent of profitable miners for each quintile also looks very similar to PCR. In fact, the numbers are identical except the 3rd quintile with ATC carryforward has a slightly higher percent of profitable miners than PCR.

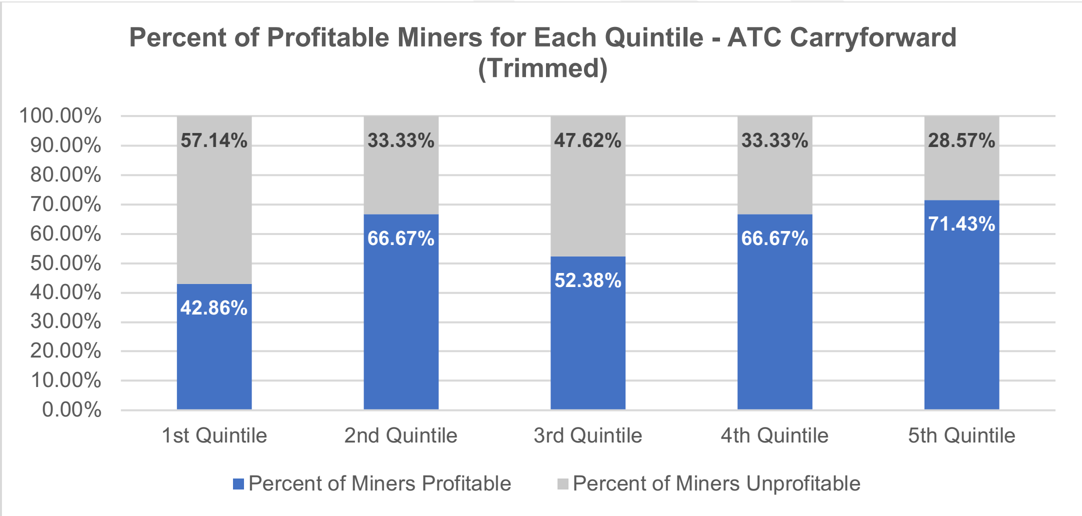

3.3.3.3 <u>Impact on MEV Miner Expected Value</u>

Given that ATC carryforward was not introduced in Part I of this paper, we need to examine the expected value for MEV Miners in implementing such a solution.

If we adopt the same assumptions from Section 2.2.3, we can quickly arrive at an expected value. In large part, the analysis remains the same except we can anticipate that with some frequency an MEV Miner will encounter a block with coinbase rewards greater than the standard coinbase. In other words, the MEV Miner will occasionally be able to mine blocks that are within 5 blocks of a missed sortition, and therefore offer 120% of the typical coinbase.

Without getting too precise, we can see that if 6% of blocks have a missed sortition, then about 30% of blocks would have an increased coinbase (that is 20% greater). Therefore an MEV Miner might expect a 70% chance of a standard coinbase and a 30% of an increased coinbase, resulting in an overall expected coinbase of:

$$0.3 * 1.2 * coinbase + .7 * coinbase = 1.06 * coinbase$$

With this, the expected value calculation for ATC carryforward would simplify to:

$$EV = \frac{1.06 * C}{.95} - C$$

This function graphs to:

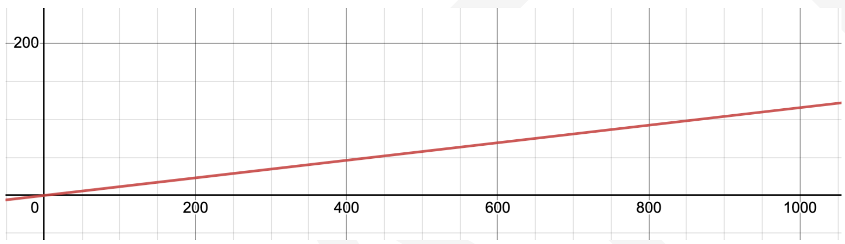

We see that the expected return for an MEV Miner increases from about 5% with standard ATC, to about 11.5% with ATC carryforward. We discuss the implications of this, and in comparison to the expected rate of return of PCR, in Section 3.5 (Recommendation), below.

3.4 Ecosystem Considerations

Given the above analysis, it seems PCR and ATC carryforward are both viable options, both with little expected impact on long-term miner profitability.

As we discussed in Section 2.1, though, we also want to consider any solution in light of the following questions:

- How fundamental or disruptive is the change to the current mining process?
- Is the change expected to alter miner behavior in significant ways?
- Does the change disproportionately affect smaller miners versus large miners?
- Is the change expected to alter stacking payout amounts or behavior?
- Will the change cause mining to further centralize?

We'll discuss each of these, in turn, for both PCR and ATC carryforward.

3.4.1 Mining Process

We want to consider how fundamental or disruptive each change is to the current mining process.

PCR creates a fairy significant change to the allocation of rewards, with additional complexity. While our instinct is that this change will not significantly alter miner behavior, it's difficult to predict.

ATC carryforward, on the other hand, keeps the mining process relatively similar as it exists today. One way to think about ATC is that it's exactly like current mining, except there's always a "null miner" with some baseline probability of winning (and resulting in no sortition if the null miner wins). The main change with ATC carryforward from current mining is that some distribution of rewards would be pushed to later blocks in the event of no sortition.

If the aim is to keep have the solution alter the mining process as minimally as possible, we believe that ATC carryforward would likely be the more appropriate solution.

3.4.2 Miner Behavior

We'll now consider whether each change is expected to alter miner behavior in significant ways.

While we believe both solutions are unlikely to alter miner behavior significantly, we have a bit more confidence in that position for ATC carryforward than we do for PCR. Given that we see ATC carryforward as a more straightforward change, we think it's possible miners may act in some unpredictable ways in the event of implementation of PCR.

3.4.3 Small vs. Large Miners

It appears that both ATC carryforward and PCR have similar impact on miners, including across miner size. In particular, both solutions do not appear to negatively affect miner profitability overall, but both do seem to have some negative impact on the lowest quintile miners.

As discussed above, this is not surprising, given that it's likely the case that the bottom quintile miners would be most likely to earn rewards during blocks with the smallest cumulative commitments. Both ATC carryforward and PCR, either directly or probabilistically, allocate rewards from blocks with small cumulative commitments to other blocks.

ATC carryforward appeared to have a slightly less negative affect on profits for the bottom quintile miners, but we do not believe the difference was significant enough to serve as a basis for making a decision between the two.

We'll further note that the bottom quintile miners account for only ~13% of commitments overall, and so it may not be appropriate to base a decision only on the impact to the bottom quintile miners.

3.4.4 Stacking

Given that we would not expect significant differences in miner behavior, we believe it follows that Stacking behavior would be similarly unaffected. Again, however, we have higher confidence this is the case with ATC carryforward than with PCR.

3.4.5 Miner Centralization

Given that we believe overall miner profitability, and the distribution of miner profitability, remains relatively stable with the implementation of both ATC carryforward and PCR, we do not expect an increase of miner centralization as a result of the implementation of either of these solutions.

3.5 Recommendation

Although we believe that both PCR and ATC carryforward are reasonable solutions, we think that ATC carryforward is slightly preferable.

We see little apparent difference in impact on expected profits for long-term miners, with both having fairly minimal impact.

For both PCR and ATC carryforward, however, there remains some positive expected value for MEV Miners to exclude others from their own blocks. The rate of return for each, under the assumptions laid out in Part I, can be compared in the chart below. The blue line is the rate of return for PCR and the red line is the rate of return for ATC carryforward.

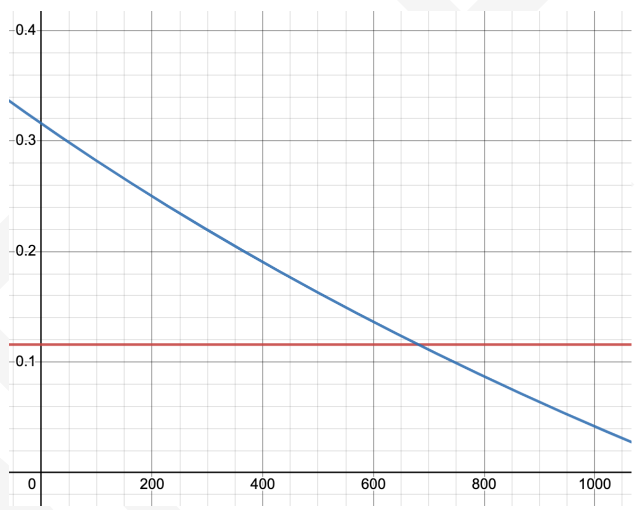

We see that, with PCR, as the MEV Miner's contribution increases, their rate of return decreases, whereas the rate of return for ATC carryforward remains constant. If a miner was contributing more than around $680 of a $1000 Stacks coinbase block, the expected rate of return for PCR would fall below that of ATC carryforward.

We think it's unlikely that an MEV Miner would consistently mine more than 68% of the value of the coinbase, and therefore the rate of return to an MEV Miner would likely be less, potentially significantly, in the case of ATC carryforward.

Lastly, we think ATC carryforward is a slightly less disruptive and more predictable change to the mining process.

3.6 Additional Considerations

As we've examined this MEV Strategy in depth, it's become clear that it would be very difficult, consistent with the current network design, to create a solution that resulted in zero or negative expected value to MEV Miners, that also didn't make it considerably more expensive for Honest Miners to mine.

Ultimately, if a Bitcoin miner wishes to mine the Stacks blockchain, they will likely be able gain some added value by engaging in exclusion of other miners, even with ATC carryforward, PCR or other solutions implemented. However, by implementing one of the solutions discussed in this paper, their returns would be dramatically decreased.

If, even after the implementation of ATC carryforward, PCR or some other solution the Stacks community decides to implement, we see Bitcoin miners engaging in some adaptation of the MEV Strategy in a problematic way, we see three potential ways of addressing the issue, in order of our confidence of them addressing the problem:

- The Stacks community may wish to revisit the MBR solution, implementing a threshold that tracks closely to the value of the coinbase, and consider whether there is a tolerance for making mining more expensive overall.
- A solution that requires miners put up some stake in order to mine and that can be slashed by certain parties, perhaps Stackers whose interests are aligned with the proper operating of the Stacks mining process.
- Removing the value transfer of PoX mining from the Bitcoin blockchain entirely, removing the power of Bitcoin miners to censor Stacks mining transactions.

While we understand that neither ATC carryforward nor PCR are perfect solutions. We believe they are appropriate incremental solutions that will remove the ability of Bitcoin miners to earn significant block rewards from insignificant payments. These solutions may not fully disincentivize all MEV-like behavior, but both appear to be good candidates for initial solutions for this less-than-ideal, though not-catastrophic, issue.