A Risk Analysis of Rocket Pool Low Ether Bonded (LEB) Minipools

Ken Smith (@shtimseht) | Nextblock Solutions

May 2022

DRAFT VERSION

NEXTBLOCK SOLUTIONS

Abstract

This is an analysis of low ether bonded (LEB) minipools which are Rocket Pool (RP) Ethereum validators formed with less than 16 ether (ETH) as the node operator (NO) deposit. This analysis attempts to quantify a risk profile by first defining a set of performance specifications as the minimal design criteria of a LEB minipool. To determine the smallest NO deposit that is acceptable we first performed a series of predictions on how a set of minipools might perform based upon known and predicted probabilities.

In order to accomplish this modeling, we estimated the number of beacon chain validators for the purposes of determining average annual staking rewards. We assume that the current ethereum protocol parameter setting for rewards and penalties continues to be constant during the time modeled. We next performed a monte carlo prediction using historical Ethereum proof-of-work mining block rewards to calculate future Ethereum staking (beacon chain) proposer payment value (PPV) rewards. PPV includes both priority fees and inclusion payments that are made payable to the `coinbase` address. Many of the analytical approaches and computer python codes were derived from the earlier work of pintail.xyz. This report also details certain attack strategies that would reduce the APR of the rETH staking derivatives and measures the impact that such attacks would have based on a selected level of success.

Finally, a risk matrix is assembled that compares the predicted risk quantification to the established design specifications to determine the smallest NO deposit of a LEB that will meet the pre-determined risk profile. This minimal NO deposit can be used to lower the financial barriers to forming a minipool while maintaining RP's risk profile.

Tl;dr A NO deposit of 4 ETH and 1.6 Ξ_{RPL} meets the minimum design specification of a LEB minipool and is predicted to return 9.72% APR with only a 10% node commission.

Introduction

The Reasoning for LEB Minipools

For a significant part of the RP protocol's 6-month lifetime, there has been stronger demand for the liquid staking (rETH) option than there has been for minipools. There are numerous theories that opine why this is. The majority of them cite the amount of financial capital (16 ETH) plus 1.6 ether worth of RPL (1.6 Ξ_{RPL}) as a contributing reason.

In order to grow the adoption of RP, there need to be mechanisms to increase the number of active minipools supporting the protocol. The proposal to lower the amount of ETH needed as a NO deposit has been promulgated by members of the RP community. Although such an approach would achieve the objective of reducing the financial barrier to forming a minipool, it also creates a leveraged opportunity for a malicious or dishonest person to obtain an ethereum validator at a greatly reduced capital cost. A quantification of what those risks are is included in this report.

Ethereum Protocol Penalty Scenarios

One action already taken by the RP developers in consideration of LEB minipools was the submittal of a design specification issue as a prelude to an Ethereum Improvement Proposal (EIP) to increase the validator ejection balance[1]

[1] https://github.com/ethereum/consensus-specs/issues/2883

from 16 effective ether (ETH_{eff}). Because of the way the beacon chain rounding mechanism is designed 16 ETH_{eff} corresponds to an actual validator balance of 15.75 ETH.

The following table attempts to delineate the various risk factors that are altered by this proposal was to be accepted and the ejection balance set to 30 ETH_{eff}.

Although such a change to increase the ejection balance would greatly reduce the risks to all validators on the network it is not an essential prerequisite to the creation of a LEB minipool. EIP proposals take a long time to develop and reach acceptance. It could be many years before such a proposal is accepted, if ever, reaches implementation in a forked update of the core protocol. Regardless of whether or not the Ethereum protocol accepts the EIP to raise the ejection balance, RP independently can alter its own protocol settings to enable the formation of a LEB minipool. This analysis continues without incorporating the ejection balance increase as a mitigation factor.

Table 1
Maximum Loss of Deposit to a Validator Under Ethereum Penality Conditions

Scenario	Description	Probability	Max loss to validator		Can validate in exit queue	Exited	Solution
			16 ETH$_{eff}$ Ejection Balance	30 ETH$_{eff}$ Ejection Balance			
Validator Off-line	Validator is off-line for short durations (<1 week)	0.5%[2]	0 ETH	0 ETH	Yes	n/a	Correct off-line condition and rewards will offset penalities
Validator Abandoned	Validator is off-line for years while the chain is finalizing. Loses ¾ of an epoch reward per epoch. (3 * base reward)	Low 116 µmort[3]	16.25 ETH	2.25 ETH	Yes	Once the eff balance reaches below the ejection limit the validator is entered into the exit queue. The validator will continue to leak until exited.	High ejection balance can help resolve this. NO can come alive and start to attest exiting (i.e., in the exit queue). .
Offline During Non-finality	Off-line losses are quadratically resulting in a node reaching the ejection balance in days not years.[4]	Low *Non-finality has not occurred on mainnet yet.*	16.25 - 32 ETH	2.25 ETH - 32 ETH	Yes	The exit queue[5] may be very long if there are a number of other nodes exiting the queue in this window. However, the quadratic leak will stop to all offline nodes, including those ejected and still in the exit queue, once the chain has reached finality. During non-finality, the inactivity penalties are quadratically increased until the chain re-obtains finality.	Any validator that has been ejected or volunteered to exit and thus is in the exit queue can prevent all leakage by attesting. Nodes that are not attesting can perform a software hotfix or switch clients and use snap sync to attest again and help reestablish consensus.
Isolated Slashing	One or a small number of validators are slashed in a 26-day window	Low 473 µmort[6]	1 ETH + 26 days of leaking	1 ETH + 26 days of leaking	No	A validator that is not exited will still be listed in duties; however, attestations from slashed validators are not valid and will be ignored. As such, it may as well stop bothering to attest as soon as it is slashed.	None. Share a postmortem with the staking community as a best practice.
Correlated Slashing	Multiple validators are slashed within a 36-day period. We used 5% of the total validators as the number of slashed validators.[7]	Very-low	1 ETH + 26 days of leaking + 3 ETH special penalty	1 ETH + 26 days of leaking + 3 ETH special penalty	No	A slashed validator incurs leakage penalties for 36 days or more if the exit queue is >36 days.	None.

[2] Calculated as 1 - participation_rate from https://www.rated.network/

[3] Calculated as all validators with a balance less than 32 ETH and stats "Active Red" divided by all total validators. 43 / 367,611 as of 5.5.2022 Data source: https://beaconcha.in/validators

[4] https://github.com/ethereum/consensus-specs/issues/2883#issuecomment-1116700339

[5] https://www.reddit.com/r/ethstaker/comments/jc8asx/inactive_validators_should_be_ejected_much_faster/

[6] 174 validators have been slashed out of 367,611 total validators as of 5.5.2022

[7] THe amount of the special penalty depends on the total number of validators shared in a 36-day window.

LEB Minipool Design Criteria

To start the evaluation of what risk LEB minipools present to RP and its token (rETH and RPL) holders, we first need to establish *a priori* the acceptable risk criteria for an established set of scenarios. RP currently incentivizes NO to operate their node with a high degree of performance and honest behavior by a combination of immutable smart contracts and financial incentives (rewards and penalties).

Currently, a NO needs to deposit the same amount of ether (16 ETH) as they are *borrowing* (16 ETH) from the deposit pool. The ether from the deposit pool was contributed by regular stakers who staked via the liquid staking derivative rETH. Combined these two contributions equal the required 32 ETH to form a beacon chain validator.

In addition, an RP NO also needs to deposit a bond in the amount of 1.6 ether worth of RPL (Ξ_{RPL}) tokens. This helps mitigate the risk of validator abandonment ultimately ending with the abandoned minipool exiting with less than the NO deposit. The inclusion of RPL in the RP tokenomics model provided a way to add value to the token beyond a mere governance token use. The needs to act as a second layer of insurance and the need to continually add demand for RPL are two of the reasons why LEB minipools will continue to integrate the use of RPL tokens. Although the amount of ether for a LEB minipool is reduced from the typical 16 ETH amount there has been no desire to reduce the amount of RPL bonding. The LEB minipools modeled in this report are at the proposed minimum of 1.6 Ξ_{RPL}.

Current aggregate RP performance metrics as tracked by rated.network[8] (see figure 1) continue to show that the RP protocol performs on par (above the 50th percentile) with larger centralized staking providers.

[8] https://www.rated.network

Figure 1
30-Day Aggerate Network Validator (Minipool) Performance of Rocket Pool
source: www.rated.network/e/Rocketpool

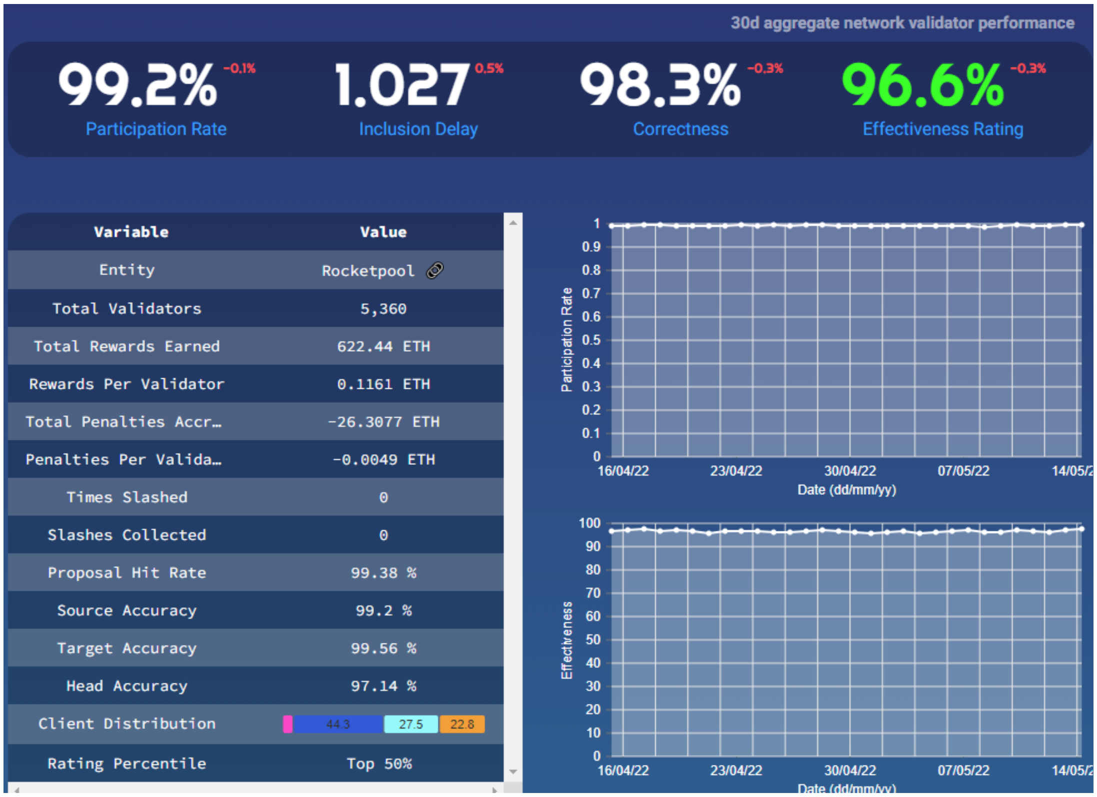

30d aggregate network validator performance

99.2% -0.1%	1.027 0.5%	98.3% -0.3%	96.6% -0.3%
Participation Rate	Inclusion Delay	Correctness	Effectiveness Rating

Variable	Value
Entity	Rocketpool 🔗
Total Validators	5,360
Total Rewards Earned	622.44 ETH
Rewards Per Validator	0.1161 ETH
Total Penalties Accr…	-26.3077 ETH
Penalties Per Valida…	-0.0049 ETH
Times Slashed	0
Slashes Collected	0
Proposal Hit Rate	99.38 %
Source Accuracy	99.2 %
Target Accuracy	99.56 %
Head Accuracy	97.14 %
Client Distribution	44.3 27.5 22.8
Rating Percentile	Top 50%

The *a prior*i design specification of a LEB minipool is proposed as the following performance criteria based on the reasons specified in table 2. Many of these specifications were discussed and refined in the LEB #theory-crafting thread in the RP Discord.

Table 2
Design Specifications for LEB Minipools

Design Criteria	Justification
A. The NO deposit is greater than the offline penalty for short durations during chain consensus.	This will allow the NO to go off-line periodically without capital exposure to the rETH holders.
B. The NO deposit is greater than the penalty for neglect for 2 years or more during chain consensus.	Although without an increase in the ejection balance to the core protocol the validator will eventually lose all funds that exceed the NO deposit. This is true even in 16 ETH minipools. The overall risk to the rETH holders is small if this continues to be a rare occurrence.
C. The NO deposit is greater than the beacon chain penalty (initial, special, and leakage during an empty exit queue) for an isolated (uncorrelated) slashing violation during chain consensus.	This fully backs any slashing lost by the NO deposit. rETH holders will incur no loss of deposited funds.
D. The NO deposit is greater than a 1% correlated slashing event during finality.	Multiple validators are slashed within a 36-day period. We chose 1% of the total Ethereum network as the number of slashed validators in the lookback window.
E. The NO deposit is greater than the penalty for non-responsiveness for 7 days during nonfinality.	This is adequate time for a NO to be made aware of the blockchain non-finality event and to switch clients to one that is responsive. If the cause is not related to a specific client this will give all devs and client teams time to correct the issue.
F. The NO deposit plus their earned share of the beacon ETH rewards are larger than the expected average PPV for 5 years of validating. (625k Validators)	This is to prevent a long-con attack in which a NO immediately decides to be dishonest and intends to steal all PPV starting on day one. The period of time (5 years) was chosen to add an opportunity cost of locking the ETH deposit.
G. The NO deposit is greater than a 2-sigma (97.7 percentile) probability of total PPV for 1 year of validating. (400k Validators)	This is to prove a low probability of a short-term lucky validator attack.
H. The modified NO deposit is greater than a 95% probability of maximum PPV in a block for 1 year of validating. (400k Validators)	This is to establish an upper-risk threshold against a very lucky high PPV block (i.e. a lottery block) and validator defection only when such a block is obtained. Although there is no protection against these lottery blocks the high probability threshold was chosen so that the overall effect of lost PPV blocks on rETH holders is minimized.

The minimally acceptable amount of ETH for use as the NO deposit of a LEB would be evaluated against these design specifications.

Time Until Penalized Balance Exceeds NO Deposit

Table 3 displays the time until the amount of penalty exceeds the initial NO deposit based on a number of adverse ethereum staking scenarios. The length of time can be thought of as a risk quantification, the shorter the duration the greater the risk. The duration is the length of time that the NO would need to be able to react to the inactivity error condition or the amount of time that the NO deposit will have underwrote the loss of funds to rETH holders in the event of one of the following adverse scenarios.

Table 3
Time Unitil Penalized Balance Exceeds NO Deposit Under Ethereum Penality Conditions

			Node Operator Deposit				
Chain	**Scenario**	**Description**	**16 ETH + 1.6 Ξ_{RPL}**	**8 ETH + 1.6 Ξ_{RPL}**	**6 ETH + 1.6 Ξ_{RPL}**	**4 ETH + 1.6 Ξ_{RPL}**	**2 ETH + 1.6 Ξ_{RPL}**
Reaching Consensus	**Validator Off-line**	Validator is off-line for short durations (<1 week)	Infinity	Infinity	Infinity	Infinity	Infinity
	Validator Abandoned[9]	Validator is off-line for a long duration. Loses ¾ of an epoch reward per epoch. (3 * base reward)	1905642 epochs 23.2 years	848890 epochs 10.3 years	645077 epochs 7.9 years	457497 epochs 5.6 years	283754 epochs 3.5 years
	Isolated Slashing[10]	One or a small number of validators are slashed in a 26-day window. Assume an initial penalty of 1 ETH and a special penalty of 0 ETH.	1744579 epochs 21.2 years	744781 epochs 9.1 years	549421 epochs 6.7 years	369025 epochs 4.5 years	201462 epochs 2.5 years
	Correlated Slashing	Multiple validators are slashed within a 36-day period. Assumes 1% of the network is slashed with the window.[11]	1593827 epochs 19.4 years	645077 epochs 7.9 years	457497 epochs 5.6 years	283754 epochs 3.5 years	121946 epochs 1.5 years
In Non-finality	**Validator Off-line**[12]	Validator is off-line. Off-line losses are quadratically increasing when the chain is in non-finality. This results in a node reaching the ejection balance in days not years.[13] Ejectional balance is 16 ETH$_{\text{eff}}$.	5176 epochs 23.0 days	3459 epochs 15.4 days	3016 epochs 13.4 days	2540 epochs 11.3 days	2001 epochs 8.9 days
	Isolated Slashing[14]	One or a small number of validators are slashed in a 36-day window. Assume an initial penalty of 1 ETH and a special penalty of 1 ETH.	4953 epochs 22.0 days	3240 epochs 14.4 days	2783 epochs 12.4 days	2281 epochs 10.1 days	1686 epochs 7.5 days

[9] Calculated with est_InActivity_Leakage.py.
[10] Calculated with est_InActivity_Leakage.py using an initial penalty of 1 ETH.
[11] Time reported is the length of time that the NO deport can support in the exit queue after payment of the initial and special penalties.
[12] Calculated with estLeakage_NonFinality.py.
[13] https://github.com/ethereum/consensus-specs/issues/2883#issuecomment-1116700339
[14] Calculated with estLeakage_NonFinality.py using an initial penalty of 1 ETH and a special penalty of 1 ETH.

Modeling Assumptions

To determine both the potential risks due to a dishonest NO and to calculate the potential APR of a LEB minipool we need to model the potential rewards that a minipool can receive. The value of potential rewards that any Ethereum validator, including RP minipools, will receive post-merge is dependent on two factors. First is the number of block proposals that a validator is randomly selected to produce. Second is the value of *priority fees* and *inclusion payments* made to the proposer's `coinbase` address. Collectively we will refer to the sum of priority fees and inclusion payment rewards as proposer payment value or PPV.

Beacon Chain (Consensus Layer) Validating Rewards

All validators who actively participate in consensus duties on the beacon chain are paid a per-formula share of Ethereum's proof-of-stake (POS) rewards. These rewards are calculated using a straight probability whose odds are determined by the total number of validators (*n*) on the Ethereum network. The current *n* of Ethereum validators is approximately 400,000 validators. It is anticipated that the number of validators on the Ethereum network will increase as we approach the Merge and continue to grow as the post-merge staking APR increases due to the inclusion of PPV for validating duties. The second *n* value of 625,000 was chosen as a reasonable increase due to the continued growth of the beacon chain.

To model this, we will use a standard binomial distribution based on the length of time that the minipool is operating. We have modeled the likely number of block proposals that a minipool will receive over both a 1-year and 3-year time span. For each of these time periods, we have two probability scenarios. The first is a 1 in 400,000 chance and the second is a 1 in 625,000 chance probability of the validator being selected as a proposer per slot. Choosing the number of validators at the current number of validators is a conservative assumption as both the amount of beacon chain rewards and block proposer rewards per validator will decrease as the total validator count (*n*) increases. Combined, we have four total scenario combinations. The reader can model longer and shorter duration using the python code provided.

With a total number of validators at `n = 400,000` the median validator will be selected six times a year to propose a block. With `n = 620,000` the median minipool will be selected four times a year to propose. The histogram of block proposing chances for both sets of total validators is displayed in figure 2.

Figure 2
Probability Mass Function of Block Proposer Payments in ETH

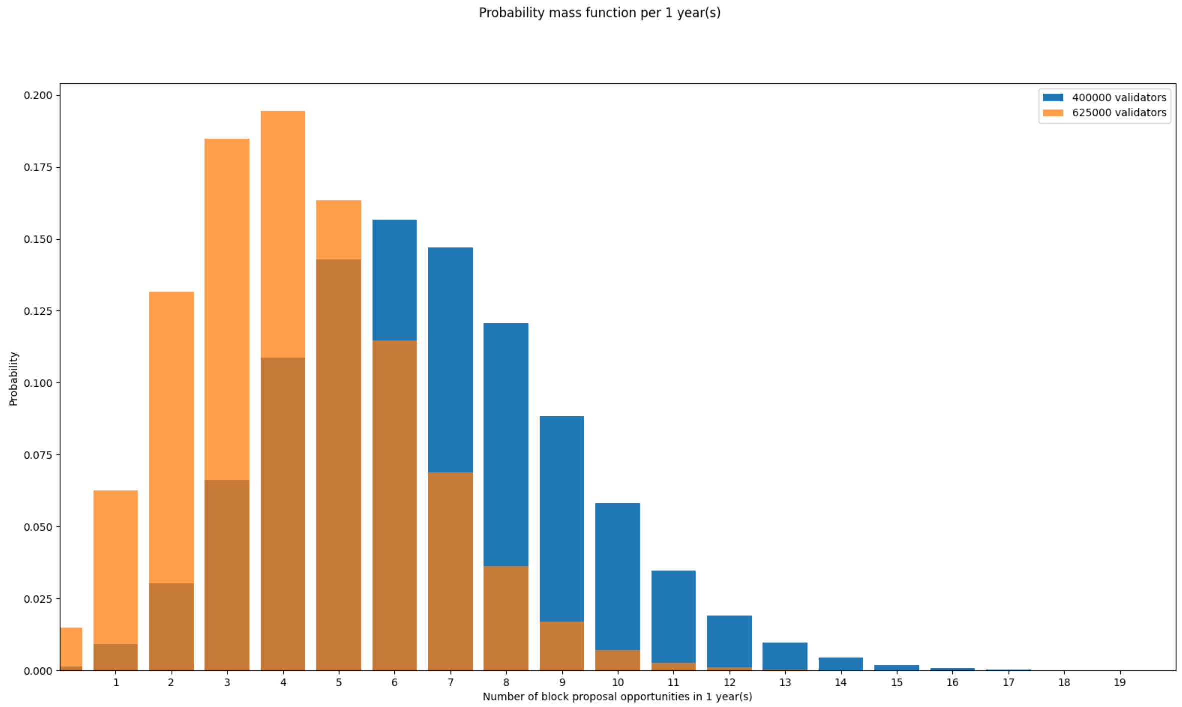

Based on the above assumptions of the total number of validators, we can calculate the average amount of ether that each active minipool will receive from its beacon chain consensus duties (*b*). Table 4 lists the values of *b* and the APR rate for a standard 32 ETH solo validator. For 400,000 and 625,000 validators, the values of ***b*** **are 1.39 and 1.11 ETH/yr** respectively.

Table 4
Annual Beacon Chain Validating Rewards (*b*) and APR.

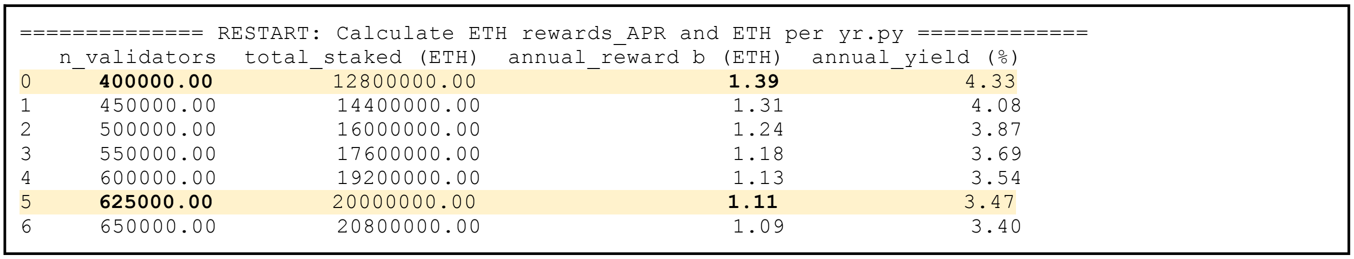

```
============== RESTART: Calculate ETH rewards_APR and ETH per yr.py =============
   n_validators  total_staked (ETH)  annual_reward b (ETH)  annual_yield (%)
0      400000.00         12800000.00                   1.39              4.33
1      450000.00         14400000.00                   1.31              4.08
2      500000.00         16000000.00                   1.24              3.87
3      550000.00         17600000.00                   1.18              3.69
4      600000.00         19200000.00                   1.13              3.54
5      625000.00         20000000.00                   1.11              3.47
6      650000.00         20800000.00                   1.09              3.40
```

.

Proposer Payments Value (PPV)

The second prediction that we need to calculate is the amount of proposer payment value (PPV) expected per minipool . This value is determined by network conditions and arbitrage opportunities when the minipool is selected to propose a block.

In order to quantify PPV we first need to understand what "maximum extractable value" (MEV) is. MEV refers to the total value that can be extracted from the blockchain by altering the sequence of transactions and injecting a bot-generated transaction to take advantage of the EVM blockchain state. Projects like Flashbots MEV Explore[15] have attempted to quantify and measure a subset[16] of MEV that is occurring in the Ethereum network (see figure 3). Most of the transactions observed by Flashbots use timed arbitrage transactions to front-run an automated market maker (AMM) transaction.

Figure 3
MEV-Explore Drill Down for Extracted MEV
source: https://explore.flashbots.net/

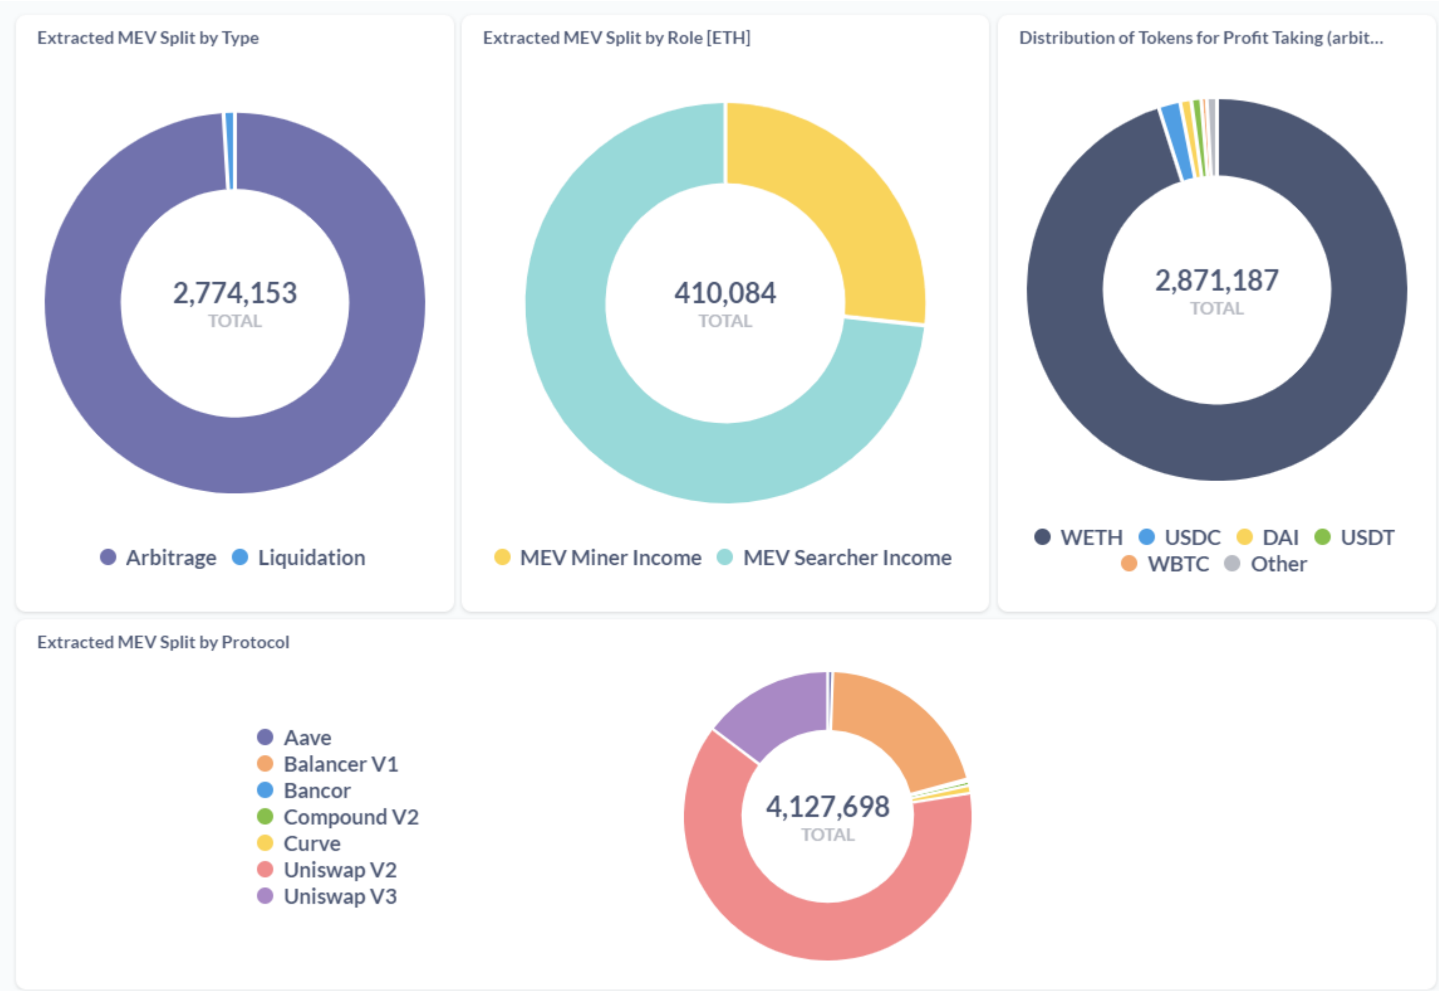

Because the extraction of MEV involves the ordering of transactions within a block a payment has to be made to the block proposer in place the transaction in the correct sequence relative to the other transactions. This payment is made on-chain to a specific address referred to as the `coinbase`. This address will be renamed the `feeReceipent` address post-merge.

[15] https://explore.flashbots.net/

[16] Flashbots inspects only 10 DeFi protocols and tracks profit for single-tx MEV opportunities. It should be considered a lower bound of the total estimate on MEV.

The `feeReceipent` address is configured in the consensus layer client and is provided by the NO at the time of the block proposal. Because the eth1 address that is provided for the `feeReceipent` is under the complete control of the NO, there is an opportunity for a NO to steal all the PPV for themselves. We will expand upon this attack scenario later in the report.

The cumulative sum of measured MEV-related block proposer payments in the Ethereum network over time is shown in figure 4. This graphic was obtained from Flahsbots MEV Explore. As visualized by the flattening of the curve, there has been an overall decline in the number of proposer payments to miners over time. This reduction correlates with a similar reduction in the amount of MEV per time. This reduction might be due to a number of factors including greater adoption of layer-2 (L2) protocols in which layer-1 MEV extraction is not possible and the use of MEV protection mempool relays such as those offered by flashbots[17] and ethermine[18]. Data from MEV-Explorer[19] shows a generally declining value of MEV and MEV-related block proposer payments over time, so we have chosen the most recent 432,000 block timespan to represent what a future MEV opportunity will look like.

Although it is not possible to foresee the actual amount of PPV in the future, we can make a prediction by assuming that the MEV proposer payments obtained in recent transactions are probably representative of proposer payments in the future.

Figure 4
Cumulative Sum of Block Proposer Payments for MEV (ETH)
source: https://explore.flashbots.net/

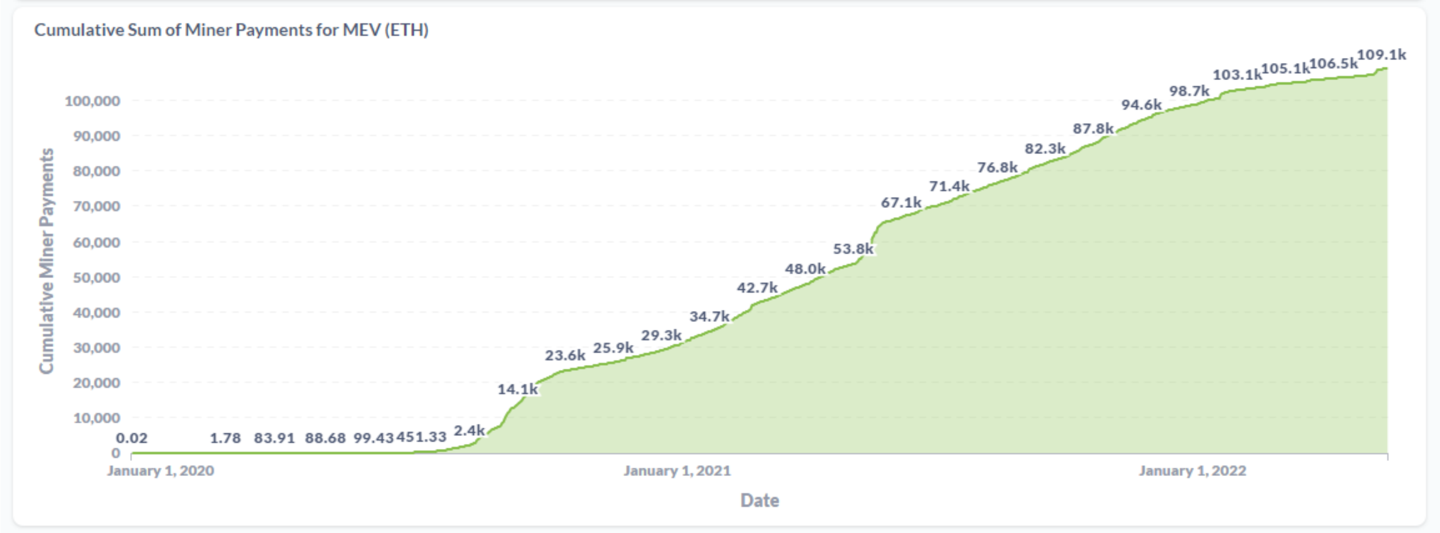

Block proposer payments are only one part of the PPV. The other component is the *priority fee*, sometimes referred to as a proposer tip. During network congestion, Ethereum gas prices increase. If a user is using the public mempool then they will have to pay a higher priority fee to have their transaction included in a near-upcoming block. As shown in Figure 5 the portion of miner tips is highly variable and depends on network usage. Higher network congestion and demand correlate with higher priority fees.

[17] https://github.com/flashbots/pm
[18] https://ethermine.org/mev-relay
[19] https://explore.flashbots.net/

Figure 5
Rewards and Tips per Day (ETH)
source: https://watchtheburn.com/insights

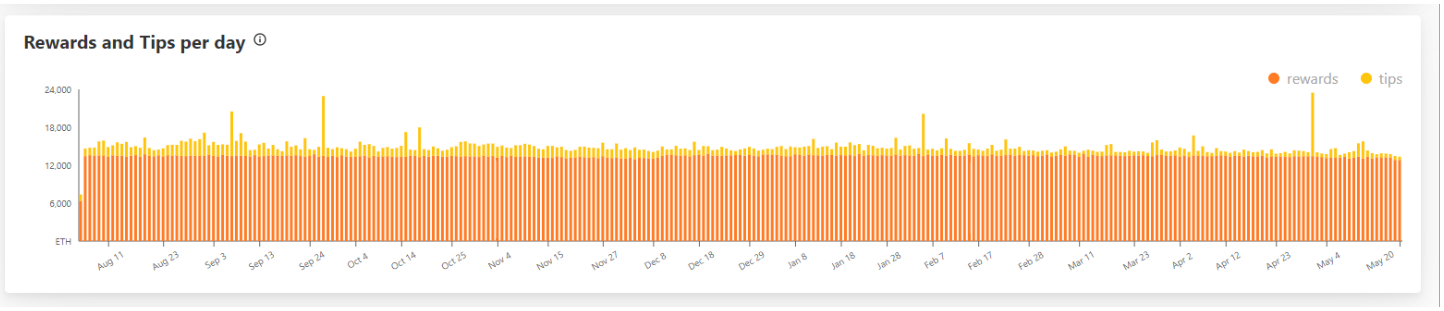

To estimate the amount of PPV in a possible future block, we first need an accurate measurement of priority fees and other inclusion tips made to the block proposer from recent block transactions for use as a surrogate dataset for future MEV possibilities. We downloaded via the etherscan.io API[20] a record of the last n = 432,000 blocks worth of data that corresponded from block 14,771,399[21] to block 14,339,400[22]. These block times varied due to the proof-of-work mining protocol and spanned a total time frame of 67 days, 16 hours, 23 minutes, and 32 seconds. However, we will use this data as a representation covering exactly 60 days as each post-merge as each block will be precisely 12 seconds in duration.

Two ETH (2E18 wei) was subtracted from the `blockReward` field to model only inclusion fees and `coinbase` payments. Even after that correction, it is important to note that this dataset only contains the on-chain payments made to the miner. It does not include any off-chain payments made for block ordering. As such these payments are not included in the forecasts made by the analytical techniques performed as part of this risk assessment.

The modified ehterscan.io dataset will be the basis for monte carlo models. We will use statistical representations of the monte carlo output to also determine the amount of proposal payments (in ETH) that the median (50th percentile) validator is expected to receive in a year. We will refer to this number as m and it has the units of ETH/yr.

We also performed a number of percentile breakdowns about the distribution of the proposer payments. The downloaded etherscan dataset was plotted on a semi-log chart in figure 6.

[20] https://api.etherscan.io/api
[21] Block mined May-14-2022 03:45:52 AM +UTC
[22] Block mined Mar-07-2022 11:22:21 AM +UTC

Figure 6
Probability Mass Function Histogram of Block Proposer Payments in ETH

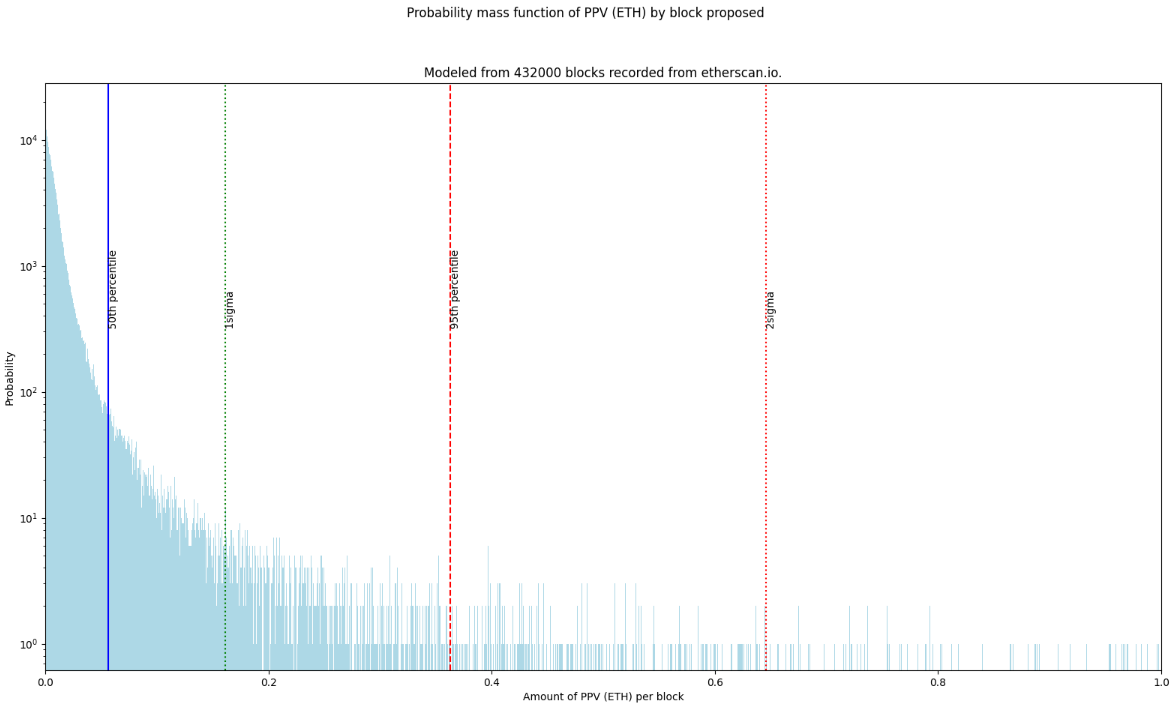

It is important to note that there is a tremendously wide distribution of PPV per block. Figure 6 shows an extremely long tail distribution that is skewed to the right. This is most illustrated by these three statements:

- The average PPV was 0.1711 ETH while the median was only 0.0561 ETH.
- The standard deviation was a whopping 2.2 ETH, 40 times the median value.
- The top 1% of blocks (ordered from smallest PPV to largest) contain 45% of the PPV.

Table 5
Statistical Analysis of blockRewards data form etherscan.io

```
============ RESTART: RP_PPV_blockReward_Histogram.py ============
 number of blocks analyzed = 432000
 The timespan of block sampled is 60.0 day(s)
   mean        0.1711 ETH in a block
   median      0.0561 ETH in a block
   std          2.2285 ETH
   sum     73900.80   ETH


Value of PPV reported in ETH
   Cutoff    PPV/blk   lowerSum  lowETH% upperSum      upETH%     sum
-----------------------------------------------------------------------
[[   50.        0.056  5245.174     7.098 68655.627    92.902 73900.801]
 [   84.1       0.161 19382.128    26.227 54518.673    73.773 73900.801]
 [   95.        0.363 30207.79     40.876 43693.01     59.124 73900.801]
 [   97.7       0.646 35710.342    48.322 38190.459    51.678 73900.801]
 [   99.        1.265 40599.572    54.938 33301.229    45.062 73900.801]
 [   99.9      12.644 52172.177    70.598 21728.623    29.402 73900.801]]
```

As visualized in figure 6 for the plot of block values most blocks have very little PPV, while some rare blocks in the long-tail distribution have a significant amount of PPV. The most lucrative blocks where the PPV is greater than or equal to the NO deposit will be referred to as *lottery blocks*.

Methods Used

Monte Carlo Estimation of PPV

We next performed a monte carlo analysis predicting 100,000 tries where each modeled validator was assigned a probability of proposing and then randomly assigned one of the representative 432,000 historic blocks as their estimated PPV reward.

A statistical analysis of the aggregate monte carlo tries was then performed to make general statements about the probable returns that a typical (the median 50th percentile) validator would receive. We also computed the probabilities for more lucky validators; in increasing order, these are 1-sigma, 90th percentile, 2-sigma, 99th percentile, and finally a 3-sigma levels of confidence.

The modeling algorithm is as follows:

1. Determine some future period of time that we wish to model our risks over. We used 1 and 3 years.
2. Perform a try. Guess the number of blocks that a validator would randomly be selected to propose over that future period of time (e.g., one year.) using a binomial distribution and assuming one of two scenarios depending on the number of active validators on the network. (*n* = 400,000 and 625,000 validators).
3. For each predicted block proposal of this try, randomly choose one of the historically mined blocks from the ehterscan.io dataset containing 432,000 of the most recent blocks. Use this selected modified `blockRewards` value as the surrogate to predict PPV.
4. Record information about this try - The minimum PPV in a block, maximum PPV in a block, average PPV/block, and the sum of all PPV for that simulated minipool over that future period. Record this data to an `OUTPUTdata` array.
5. Repeat this monte carlo simulation 100,000 times.
6. Finally, display a plot and summary of the analysis of the results.

The use of the monte carlo method will allow us to determine the value *m* that represents the average amount of PPV (in units of ETH/yr) that a minipool is expended to earn in a year. We will use *m* in our further analysis of potential LEB minipool returns and risks to the RP protocol. The three-year model calculated an ***m* of 1.12 and 0.72 ETH/yr** for 400,000 and 625,000 network validators.

Note that this number *m* that we will use in our predictions is an average value, not the median amount. The average is substantially larger than what the typical (medial) minipool is expected to receive in short time duration and its use in the forthcoming model results in a conservative estimate of the amount of PPV that a likely validator will receive. This is because most of its value contained in *m* is obtained from a very small set of the luckiest mini pools that receive exceptionally large PPV blocks. As the modeling period lengthens, assuming all conditions are equal, all validators have a random chance of being one of the lucky validators on the network and the median converges towards the average.

Modeling Results: 1 Year(s) Validating

Figure 7
Predicted Probability of Total PPV per Minipool for 1 Year(s); 400k Validators

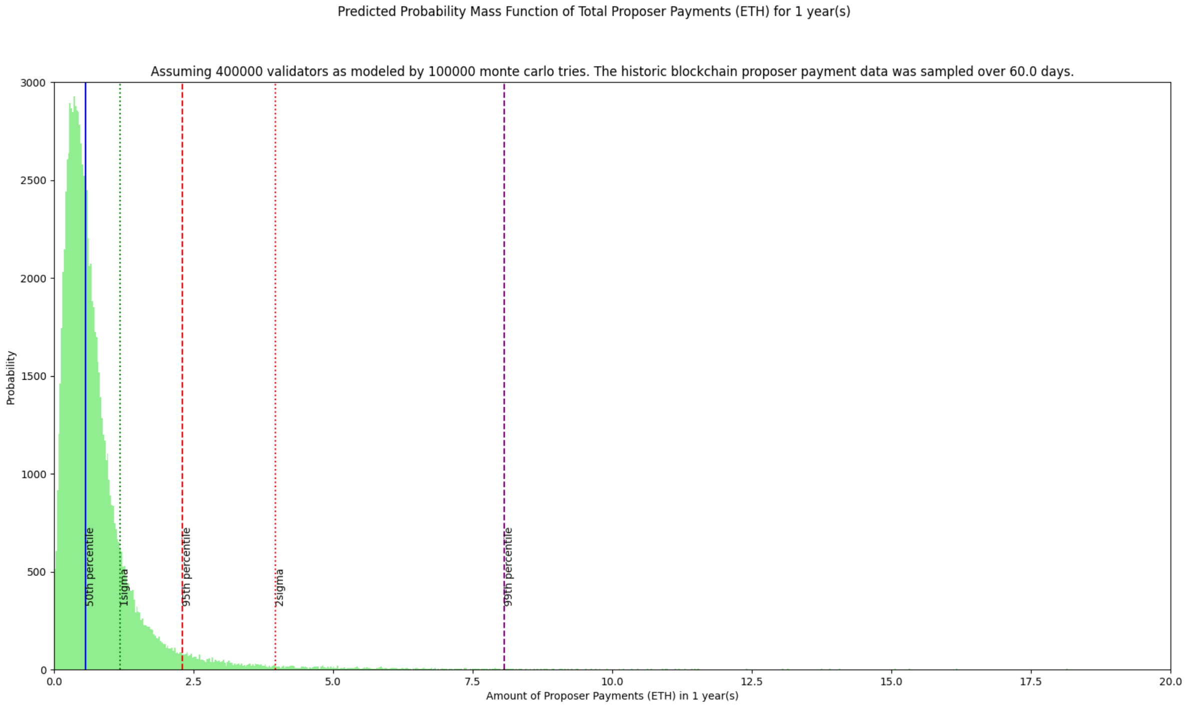

Table 6
Statistical Analysis Total PPV per Minipool for 1 Year(s); 400k Validators

```
================ RESTART: C:\Users\Ken\Dropbox\python\crypto\RP_PPV_monte _carlo.py
================
Enter the number of years you will be node operating: 1
The number of validators assumed was 400000.
The number of monte carlo tries evaluated was 100000.
 number of blocks analyzed = 432000
 The timespan of block rewards sampled is 60.0 day(s)

Proposer Payment Statistical Analysis:

The mean (average) stats of the OUTPUTdata are :
   mean min 0.0165 ETH in a block
   mean max 0.7584 ETH in block
   mean avg 0.1722 ETH/block
   mean sum 1.1236 ETH in 1 year(s)
   mean avg 1.1236 ETH/yr
```

```
        m = 1.12 ETH/yr

The median stats of the OUTPUTdata are :
   median min 0.0084 ETH in a block
   median max 0.2156 ETH in a block
   median avg 0.0853 ETH/block
   median sum 0.5594 ETH in 1 year(s)
   median avg 0.5594 ETH/yr

Based on the mean avg 0.1722 ETH/block we expect 8677 ETH in PPV for 7D of PPV
Based on the mean avg 0.1722 ETH/block we expect 37187 ETH in PPV for 1M of PPV

The         50% minipool will block propose 6.0 times a year.
The         50% min PPV is 0.0084 ETH.
The         50% max PPV is 0.2156 ETH.
The         50% avg PPV is 0.0853 ETH/proposal.
The         50% sum PPV is 0.5594 ETH over 1 year(s).
The         50% avg PPV is 0.5594 ETH/yr.

The 1sigma 84.1% minipool will block propose 9.0 times a year.
The 1sigma 84.1% min PPV is 0.0256 ETH.
The 1sigma 84.1% max PPV is 0.5846 ETH.
The 1sigma 84.1% avg PPV is 0.1647 ETH/proposal.
The 1sigma 84.1% sum PPV is 1.1778 ETH over 1 year(s).
The 1sigma 84.1% avg PPV is 1.1778 ETH/yr.

The        95% minipool will block propose 11.0 times a year.
The        95% min PPV is 0.0486 ETH.
The        95% max PPV is 1.5655 ETH.
The        95% avg PPV is 0.3273 ETH/proposal.
The        95% sum PPV is 2.2924 ETH over 1 year(s).
The        95% avg PPV is 2.2924 ETH/yr.

The 2sigma 97.7% minipool will block propose 12.0 times a year.
The 2sigma 97.7% min PPV is 0.0677 ETH.
The 2sigma 97.7% max PPV is 3.1846 ETH.
The 2sigma 97.7% avg PPV is 0.5849 ETH/proposal.
The 2sigma 97.7% sum PPV is 3.9677 ETH over 1 year(s).
The 2sigma 97.7% avg PPV is 3.9677 ETH/yr.

The        99% minipool will block propose 13.0 times a year.
The        99% min PPV is 0.0924 ETH.
The        99% max PPV is 7.2222 ETH.
The        99% avg PPV is 1.2478 ETH/proposal.
The        99% sum PPV is 8.0627 ETH over 1 year(s).
The        99% avg PPV is 8.0627 ETH/yr.

The 3sigma 99.9% minipool will block propose 16.0 times a year.
The 3sigma 99.9% min PPV is 0.2351 ETH.
The 3sigma 99.9% max PPV is 80.6677 ETH.
The 3sigma 99.9% avg PPV is 12.5882 ETH/proposal.
The 3sigma 99.9% sum PPV is 82.1807 ETH over 1 year(s).
The 3sigma 99.9% avg PPV is 82.1807 ETH/yr.
```

Figure 8
Predicted Probability of Total PPV per Minipool for 1 Year(s); 625k Validators

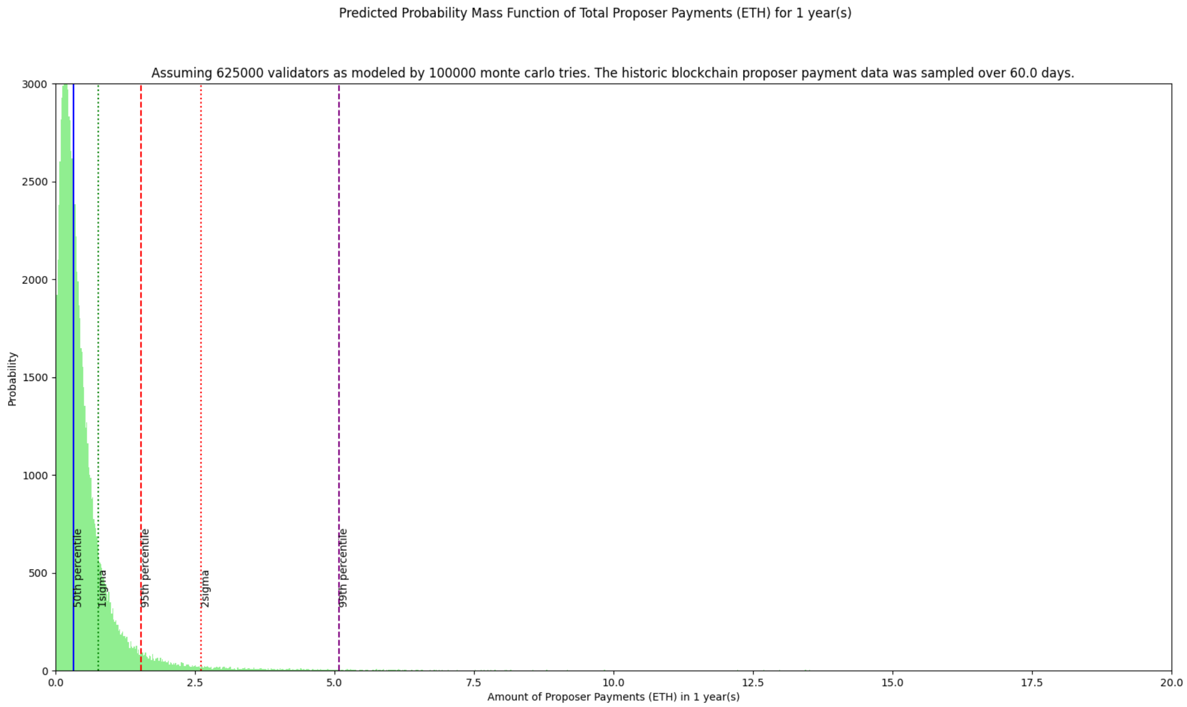

Table 7
Statistical Analysis Total PPV per Minipool for 1 Year(s); 625k Validators

```
================= RESTART: RP_PPV_monte _carlo.py ================
Enter the number of years you will be node operating: 1
The number of validators assumed was 625000.
The number of monte carlo tries evaluated was 100000.
 number of blocks analyzed = 432000
 The timespan of block rewards sampled is 60.0 day(s)

Proposer Payment Statistical Analysis:

The mean (average) stats of the OUTPUTdata are :
   mean min 0.0283 ETH in a block
   mean max 0.5364 ETH in block
   mean avg 0.1673 ETH/block
   mean sum 0.7276 ETH in 1 year(s)
   mean avg 0.7276 ETH/yr

        m = 0.73 ETH/yr

The median stats of the OUTPUTdata are :
   median min 0.0132 ETH in a block
   median max 0.1577 ETH in a block
   median avg 0.0784 ETH/block
   median sum 0.3247 ETH in 1 year(s)
   median avg 0.3247 ETH/yr
```

```
Based on the mean avg 0.1673 ETH/block we expect 8433 ETH in PPV for 7D of PPV
Based on the mean avg 0.1673 ETH/block we expect 36141 ETH in PPV for 1M of PPV

The         50% minipool will block propose 4.0 times a year.
The         50% min PPV is 0.0132 ETH.
The         50% max PPV is 0.1577 ETH.
The         50% avg PPV is 0.0784 ETH/proposal.
The         50% sum PPV is 0.3247 ETH over 1 year(s).
The         50% avg PPV is 0.3247 ETH/yr.

The 1sigma 84.1% minipool will block propose 6.0 times a year.
The 1sigma 84.1% min PPV is 0.0433 ETH.
The 1sigma 84.1% max PPV is 0.4227 ETH.
The 1sigma 84.1% avg PPV is 0.1633 ETH/proposal.
The 1sigma 84.1% sum PPV is 0.7718 ETH over 1 year(s).
The 1sigma 84.1% avg PPV is 0.7718 ETH/yr.

The        95% minipool will block propose 8.0 times a year.
The        95% min PPV is 0.0853 ETH.
The        95% max PPV is 1.0806 ETH.
The        95% avg PPV is 0.3350 ETH/proposal.
The        95% sum PPV is 1.5330 ETH over 1 year(s).
The        95% avg PPV is 1.5330 ETH/yr.

The 2sigma 97.7% minipool will block propose 9.0 times a year.
The 2sigma 97.7% min PPV is 0.1250 ETH.
The 2sigma 97.7% max PPV is 2.1137 ETH.
The 2sigma 97.7% avg PPV is 0.5878 ETH/proposal.
The 2sigma 97.7% sum PPV is 2.6127 ETH over 1 year(s).
The 2sigma 97.7% avg PPV is 2.6127 ETH/yr.

The        99% minipool will block propose 10.0 times a year.
The        99% min PPV is 0.1922 ETH.
The        99% max PPV is 4.6119 ETH.
The        99% avg PPV is 1.1313 ETH/proposal.
The        99% sum PPV is 5.0795 ETH over 1 year(s).
The        99% avg PPV is 5.0795 ETH/yr.

The 3sigma 99.9% minipool will block propose 12.0 times a year.
The 3sigma 99.9% min PPV is 0.7754 ETH.
The 3sigma 99.9% max PPV is 74.5120 ETH.
The 3sigma 99.9% avg PPV is 15.5390 ETH/proposal.
The 3sigma 99.9% sum PPV is 74.7601 ETH over 1 year(s).
The 3sigma 99.9% avg PPV is 74.7601 ETH/yr.
```

Modeling Results: 3 Year(s) Validating

Figure 9
Predicted Probability of Total PPV per Minipool for 3 Year(s); 400k Validators

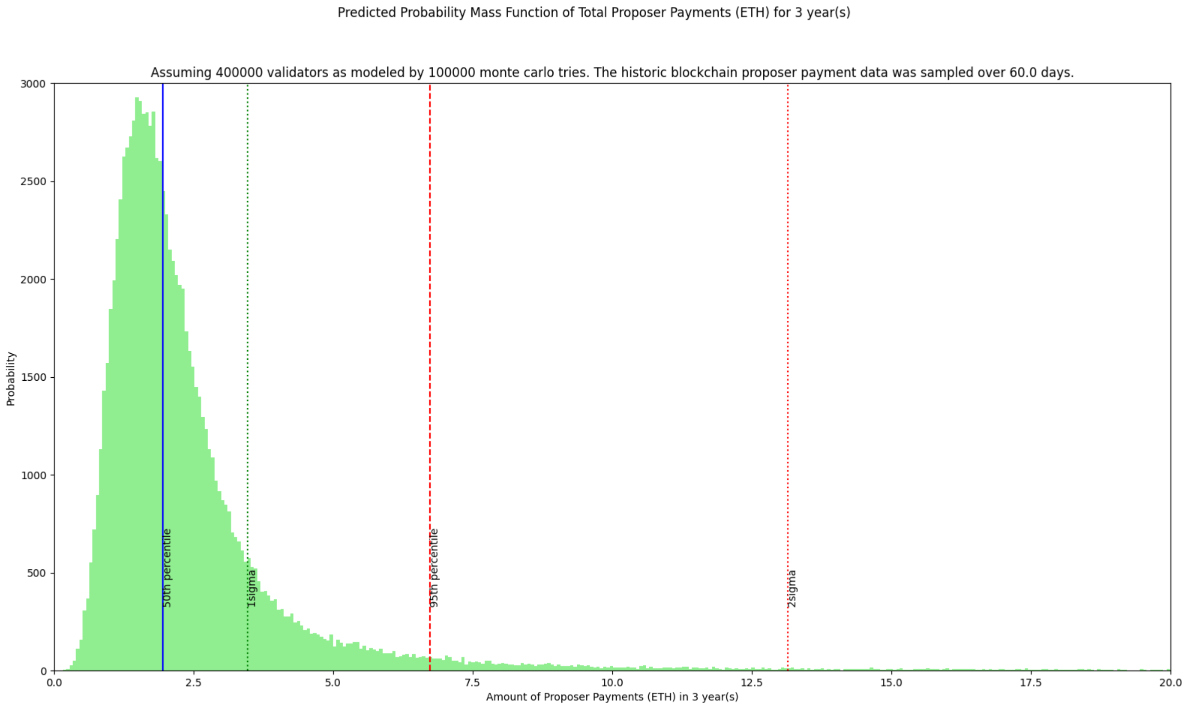

Table 8
Statistical Analysis Total PPV per Minipool for 3 Year(s); 400k Validators

```
================= RESTART: RP_PPV_monte _carlo.py ================
Enter the number of years you will be node operating: 3
The number of validators assumed was 400000.
The number of monte carlo tries evaluated was 100000.
 number of blocks analyzed = 432000
 The timespan of block rewards sampled is 60.0 day(s)

Proposer Payment Statistical Analysis:

The mean (average) stats of the OUTPUTdata are :
   mean min 0.0035 ETH in a block
   mean max 1.8397 ETH in block
   mean avg 0.1701 ETH/block
   mean sum 3.3546 ETH in 3 year(s)
   mean avg 1.1182 ETH/yr

        m = 1.12 ETH/yr

The median stats of the OUTPUTdata are :
   median min 0.0019 ETH in a block
   median max 0.4680 ETH in a block
   median avg 0.0985 ETH/block
   median sum 1.9524 ETH in 3 year(s)
   median avg 0.6508 ETH/yr
```

```
Based on the mean avg 0.1701 ETH/block we expect 8574 ETH in PPV for 7D of PPV
Based on the mean avg 0.1701 ETH/block we expect 36744 ETH in PPV for 1M of PPV

The         50% minipool will block propose 6.7 times a year.
The         50% min PPV is 0.0019 ETH.
The         50% max PPV is 0.4680 ETH.
The         50% avg PPV is 0.0985 ETH/proposal.
The         50% sum PPV is 1.9524 ETH over 3 year(s).
The         50% avg PPV is 0.6508 ETH/yr.

The 1sigma 84.1% minipool will block propose 8.0 times a year.
The 1sigma 84.1% min PPV is 0.0073 ETH.
The 1sigma 84.1% max PPV is 1.4096 ETH.
The 1sigma 84.1% avg PPV is 0.1686 ETH/proposal.
The 1sigma 84.1% sum PPV is 3.4619 ETH over 3 year(s).
The 1sigma 84.1% avg PPV is 1.1540 ETH/yr.

The        95% minipool will block propose 9.0 times a year.
The        95% min PPV is 0.0126 ETH.
The        95% max PPV is 4.4202 ETH.
The        95% avg PPV is 0.3381 ETH/proposal.
The        95% sum PPV is 6.7395 ETH over 3 year(s).
The        95% avg PPV is 2.2465 ETH/yr.

The 2sigma 97.7% minipool will block propose 9.7 times a year.
The 2sigma 97.7% min PPV is 0.0164 ETH.
The 2sigma 97.7% max PPV is 10.6279 ETH.
The 2sigma 97.7% avg PPV is 0.6534 ETH/proposal.
The 2sigma 97.7% sum PPV is 13.1431 ETH over 3 year(s).
The 2sigma 97.7% avg PPV is 4.3810 ETH/yr.

The        99% minipool will block propose 10.3 times a year.
The        99% min PPV is 0.0208 ETH.
The        99% max PPV is 28.3026 ETH.
The        99% avg PPV is 1.6089 ETH/proposal.
The        99% sum PPV is 31.0461 ETH over 3 year(s).
The        99% avg PPV is 10.3487 ETH/yr.

The 3sigma 99.9% minipool will block propose 11.7 times a year.
The 3sigma 99.9% min PPV is 0.0350 ETH.
The 3sigma 99.9% max PPV is 153.2545 ETH.
The 3sigma 99.9% avg PPV is 7.6516 ETH/proposal.
The 3sigma 99.9% sum PPV is 154.9695 ETH over 3 year(s).
The 3sigma 99.9% avg PPV is 51.6565 ETH/yr.
```

Figure 10
Predicted Probability of Total PPV per Minipool for 3 Year(s); 625k Validators

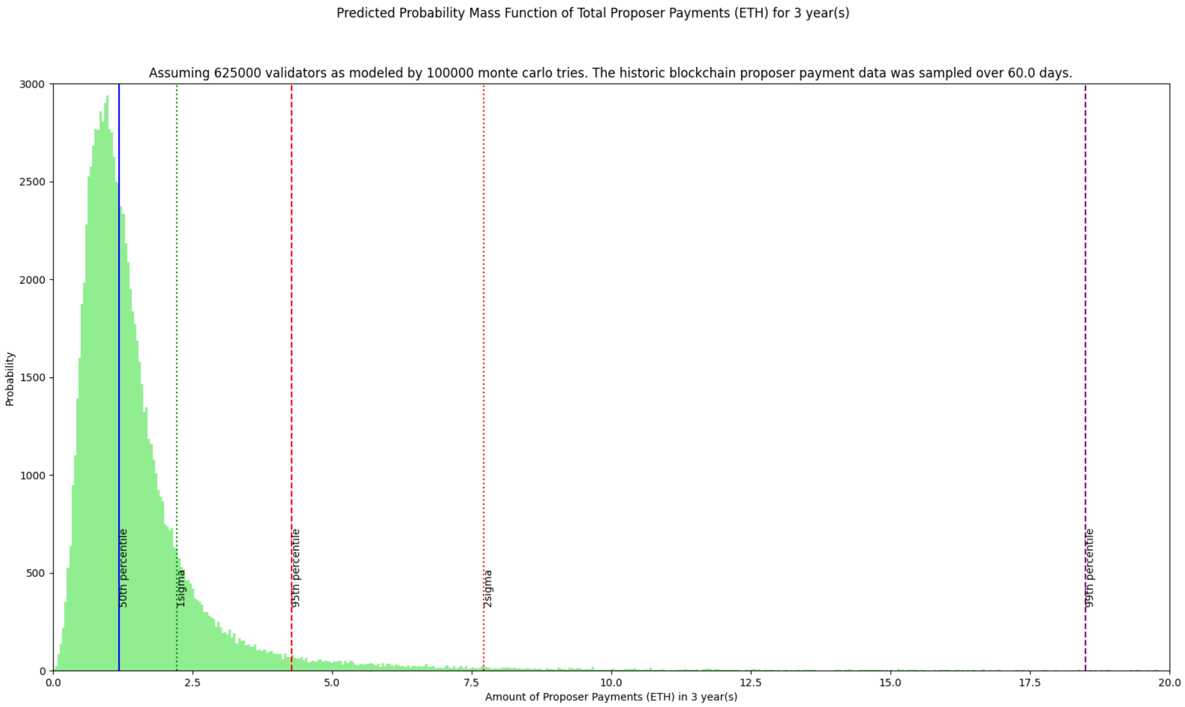

Table 9
Statistical Analysis Total PPV per Minipool for 3 Year(s); 625k Validators

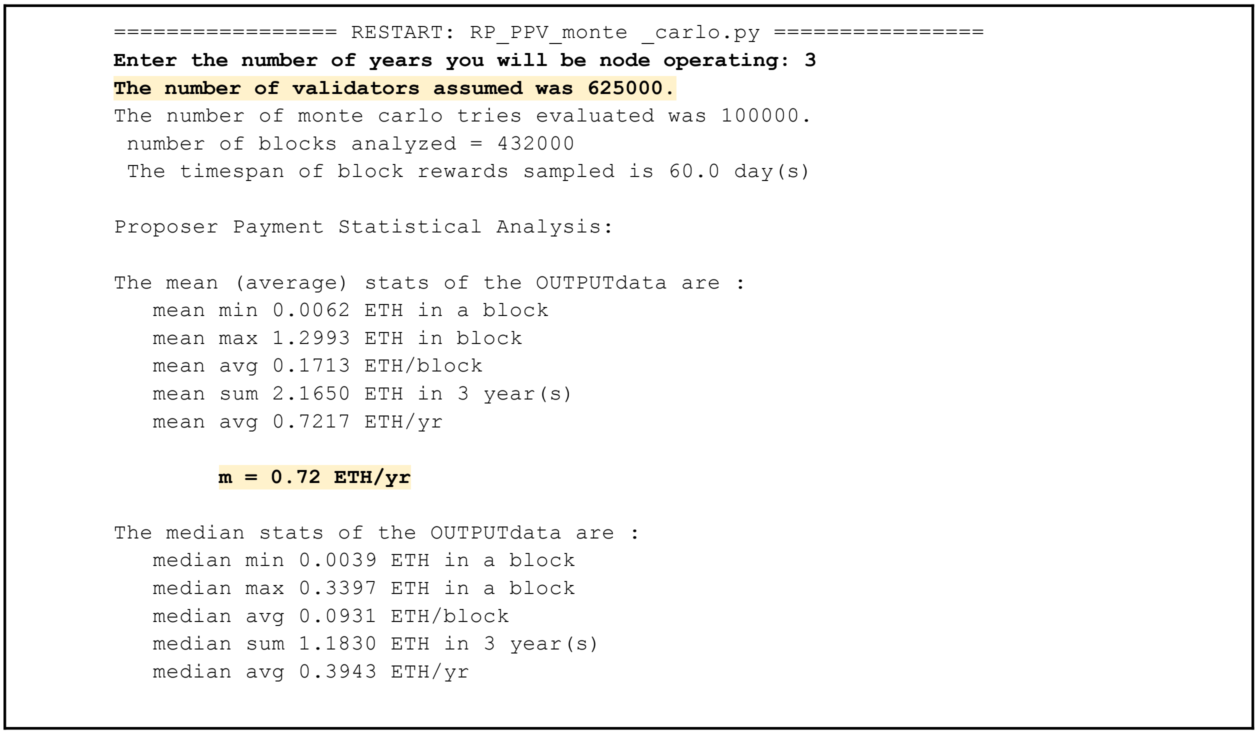

```
================= RESTART: RP_PPV_monte _carlo.py ================
Enter the number of years you will be node operating: 3
The number of validators assumed was 625000.
The number of monte carlo tries evaluated was 100000.
 number of blocks analyzed = 432000
 The timespan of block rewards sampled is 60.0 day(s)

Proposer Payment Statistical Analysis:

The mean (average) stats of the OUTPUTdata are :
   mean min 0.0062 ETH in a block
   mean max 1.2993 ETH in block
   mean avg 0.1713 ETH/block
   mean sum 2.1650 ETH in 3 year(s)
   mean avg 0.7217 ETH/yr

        m = 0.72 ETH/yr

The median stats of the OUTPUTdata are :
   median min 0.0039 ETH in a block
   median max 0.3397 ETH in a block
   median avg 0.0931 ETH/block
   median sum 1.1830 ETH in 3 year(s)
   median avg 0.3943 ETH/yr
```

```
Based on the mean avg 0.1713 ETH/block we expect 8636 ETH in PPV for 7D of PPV
Based on the mean avg 0.1713 ETH/block we expect 37010 ETH in PPV for 1M of PPV

The          50% minipool will block propose 4.0 times a year.
The          50% min PPV is 0.0039 ETH.
The          50% max PPV is 0.3397 ETH.
The          50% avg PPV is 0.0931 ETH/proposal.
The          50% sum PPV is 1.1830 ETH over 3 year(s).
The          50% avg PPV is 0.3943 ETH/yr.

The 1sigma 84.1% minipool will block propose 5.3 times a year.
The 1sigma 84.1% min PPV is 0.0120 ETH.
The 1sigma 84.1% max PPV is 0.9570 ETH.
The 1sigma 84.1% avg PPV is 0.1656 ETH/proposal.
The 1sigma 84.1% sum PPV is 2.2207 ETH over 3 year(s).
The 1sigma 84.1% avg PPV is 0.7402 ETH/yr.

The         95% minipool will block propose 6.3 times a year.
The         95% min PPV is 0.0213 ETH.
The         95% max PPV is 2.7764 ETH.
The         95% avg PPV is 0.3251 ETH/proposal.
The         95% sum PPV is 4.2689 ETH over 3 year(s).
The         95% avg PPV is 1.4230 ETH/yr.

The 2sigma 97.7% minipool will block propose 6.7 times a year.
The 2sigma 97.7% min PPV is 0.0282 ETH.
The 2sigma 97.7% max PPV is 6.0874 ETH.
The 2sigma 97.7% avg PPV is 0.6064 ETH/proposal.
The 2sigma 97.7% sum PPV is 7.7203 ETH over 3 year(s).
The 2sigma 97.7% avg PPV is 2.5734 ETH/yr.

The         99% minipool will block propose 7.3 times a year.
The         99% min PPV is 0.0368 ETH.
The         99% max PPV is 16.4492 ETH.
The         99% avg PPV is 1.4441 ETH/proposal.
The         99% sum PPV is 18.4989 ETH over 3 year(s).
The         99% avg PPV is 6.1663 ETH/yr.

The 3sigma 99.9% minipool will block propose 8.3 times a year.
The 3sigma 99.9% min PPV is 0.0655 ETH.
The 3sigma 99.9% max PPV is 121.2619 ETH.
The 3sigma 99.9% avg PPV is 10.5786 ETH/proposal.
The 3sigma 99.9% sum PPV is 121.9251 ETH over 3 year(s).
The 3sigma 99.9% avg PPV is 40.6417 ETH/yr.
```

Attack Strategies

The introduction of LEB minipools provides a reduced capital cost to execute malicious activities on the beacon chain. This may pose a threat to both the Ethereum network and the RP protocol. In particular, an attacker can adversely affect the total returns of those that staked with RP's liquid staking token, rETH. As part of this risk analysis, it is helpful to think about what vectors of attack LEBs afford a dishonest NO. The strategies evaluated in this report are based on the fact that a lower NO deposit amount reduces the potential size of penalty that the RP protocol can levy against the dishonest NO.

Two such attack strategies have been evaluated in this report, both of which are based on the dishonest NO stealing the PPVand not distributing the fair share to the liquid regular staking rETH holders. The first strategy involves being an honest NO up until the time that they receive a very large PPV block. The second involves a dishonest NO that immediately steals all inclusion and priority fees and relies on a long validating time to recover sufficient PPV rewards to outweigh the assessed penalties. We will provide an analysis of each of those attack vectors and determine the likelihood of its success as measured by the average time needed in order to accomplish such an attack.

Ethereum Network Risks

Although theoretically, LEBs may create an attack vector to the overall Ethereum network, this risk is not present until the percent of RP minipools exceeds a significant (>33%) portion of the overall number of ethereum validators. At present RP minipools (`n = 5443`) represent only 1.4 % of the overall network (`n= 392,8935`). No further analysis of the overall ethereum network is presented in this report. However, RP should re-evaluate this when such time occurs that they start to approach this consensus threshold.

Penalty for Stealing PPV

A short primer on the RP rewards and penalty system would be helpful in understanding these attack scenarios. RP allows a NO to use a contribution of ETH from the deposit to form a minipool. Currently, this is 16 ETH of deposit and 16 ETH of deposit pool contributions. In return, the NO is paid a minipool commission which is currently set at 15%. An honest NO rightly follows the protocol settings and directs the PPV payments to one of three approved addresses that fairly split these rewards to both parties, the NO and the rETH holders, based on their respective shares and the assigned minipool commission.

Currently, the RP protocol is establishing a penalty system in anticipation of the Merge. This mechanism will employ oDAO members to review all of the block proposals submitted by RP minipools and confirm that the `feeRecipient` was set to one of three acceptable RP addresses. These are either the deposit pool address (default), the nodes distributor contract, or the yet-to-be-released opt-in smoothing contract. All three of these addresses ensure that the regular stakers (e.g., the rETH token holders) receive at least their share of the PPV.

Recall that the `feeReceipent` address is configured in the consensus layer client and is provided by the NO at the time of the block proposal. A dishonest NO can alter their copy of the validating software such that they could provide a `feeReceipent` address that is a privately held Externally Owed Account (EOA). This action steals all the PPV for themselves.

To discourage stealing of PPV the RP protocol will assign a penalty that would be up to 100% of the stolen PPV. This is an effective deterrent up to the total amount of funds deposited to the minipool by the NO. This includes the amount of collateral (both ETH and RPL) placed by the NO to form the minipool and all validating ETH and RPL staking rewards offered periodically by the protocols. All of these funds can all be penalized for stealing NO. These funds are only recoverable when the NO claims the rewards (withdrawals from their distributor or the smoothing contract in the case of ETH or claims RPL) or when the NO exits the minipools at the end of its life. It is currently proposed that the maximum penalty for the RP protocol to assess a dishonest NO is 100% of the ETH deposit and beacon rewards.

Lottery Block Attack

In this scenario a NO will initially start validating as an honest NO. They will wait for the opportune time when a block proposal opportunity presents itself such that the PPV in the block exceeds the value of the NO deposit. At that point, the dishonest NO will steal that block for themselves and from that moment on operate as a dishonest NO.

Using the etherscan dataset we can calculate the average number of years it takes a minipool to win a lottery block. The results of that calculation are found in table 10.

Table 10
Average Number of Years for a Minipool to Propose a Lottery Block Assuming 625k Validators.

```
========= RESTART: RP_PPV_blockReward_NOdeposit Values.py ========
 number of blocks analyzed = 432000
 The timespan sampled is 60.0 day(s)
   mean        0.1711 ETH in a block
   median      0.0561 ETH in a block
   std         2.2285 ETH
   sum     73900.80   ETH

Deposit, percentile, year(s) to win lottery block=d, sum of PPV above threshold, upETH%
Assuming 625000 validators.
                    2.0       4.0       6.0       6.4       8.0       16.0
Deposit             2.00      4.00      6.00      6.40      8.00     16.00
Percentile_of_d    99.41     99.71     99.82     99.83     99.86     99.92
Yrs_to_win         40.29     82.73    130.13    138.37    168.31    287.59
UpperSum_of_PPV  30516.75  26953.85  24756.91  24466.44  23519.03  20657.59
upETH%             41.29     36.47     33.50     33.11     31.83     27.95
```

Long-Con Attack

The second scenario involves involves a dishonest NO who decides from the beginning to steal all PPV from the start. We are referring to this as a long-con attack. We can calculate the length of time (Q_t) that is needed for a dishonest NO to operate such that their stolen PPV, is equal to or greater than the earning and returned deposits of an honest validator. We evaluate this by calculating the sum of validator earnings and the deposit refund upon exit for both an honest NO and dishonest NO at some time (t) after the start of validating. The formulas for an honest NO (1) and a dishonest NO (2) are shown below in their expanded form. The yellow highlighted portion calculates the beacon chain rewards; the green highlighted portion calculates the PPV rewards and the blue highlight is the amount of deposit and RPL bonding (in ETH) returned upon the minipool exit from the beacon chain.

(1) Honest NO Return (y):

```
y = (s*b*t)+((1-s)*c*b*t)+(s*m*t)+((1-s)*c*m*t)+(d)
```

(2) Dishonest NO Return (z):

```
z = (f*s*b*t)+((f*(1-s)*c*b*t)+(m*t)+(f*d)
```

Where:

`b` = Average validator beacon chain rewards (ETH/yr)
`c` = Minipool commission
`d` = NO deposit (ETH)

f = Fraction returned ($f = 1-p$; where p is the % penalty assessed by the oDAO for stealing)
m = Average validator PPV rewards (ETH/yr)
n = Number of validators
p = Stealing penalty
r = rETH deposit
s = NO share of the minipool ($s = d / 32$)
t = Time validating in years

The dishonest NO formula (2) is very similar but it accounts for a penalty that is assessed against both the deposited ETH and any staking rewards that are deposited back to the main pool contract. When a dishonest NO exits the deposed ETH from the regular stakers back to the deposit pool. In such a circumstance no loss of funds occurred to the regular stakers but those funds did not earn the full APR potential had the NO been honest.

We can calculate the time (Q_t) for the accumulated stolen PPVto equal the amount of return that an honest NO would have earned. This would be the estimated length of time needed for the average dishonest NO to make more stealing the PPV than to have operated honestly. The formula for Qt, when y = z, is shown in equation (3).

(3) Time for dishonesty to payoff:

$$Q_t = (d*f-d)/(b*s+b*c-b*c*s+s*m+c*m-c*s*m-b*s*f-b*c*f+b*c*s*f-m)$$

It is expected that the number of validators on the Beacon chain will continue to grow. Modeling of the Qt for the larger number of expected Ethereum validators (n = 625,000) provides us with the results in table 11.

A negative value for Q_t implies that there the average honest rewards will outgrow the dishonest rewards. This is visible in the divergent lines of a NO deposit value of 16 ETH and the dishonest NO shown in figure 11. For all LEB deposit values, the dishonest and honest lines are converging which means that there is some finite (Q_t) time in which it will be more profitable for a NO to act dishonestly.

Table 11
Q_t times by NO Deposit Assuming 625k Validators

```
====== RESTART: RP_REV_Qt_plot2lines.py =====
INPUTS:
Average validator Beacon Chain rewards (ETH/yr)  b = 1.11
Minipool commission                               c = 0.15
Fraction returned (f = 1-p)                       f = 0
Average validator PPV rewards (ETH/yr)           m = 0.72
Number of validators                              n = 625000
Stealing penalty                                  p = 1


The Qt for a  2   ETH NO deposit is -0.00 year(s)
The Qt for a  4   ETH NO deposit is -0.00 year(s)
The Qt for a  6   ETH NO deposit is -0.00 year(s)
The Qt for a  6.4 ETH NO deposit is 0.00 year(s)
The Qt for a  8   ETH NO deposit is -0.00 year(s)
The Qt for a 16   ETH NO deposit is 0.00 year(s)
```

Figure 11
Honest and Dishonest Earnings and Exited Values Assuming 625k Validators

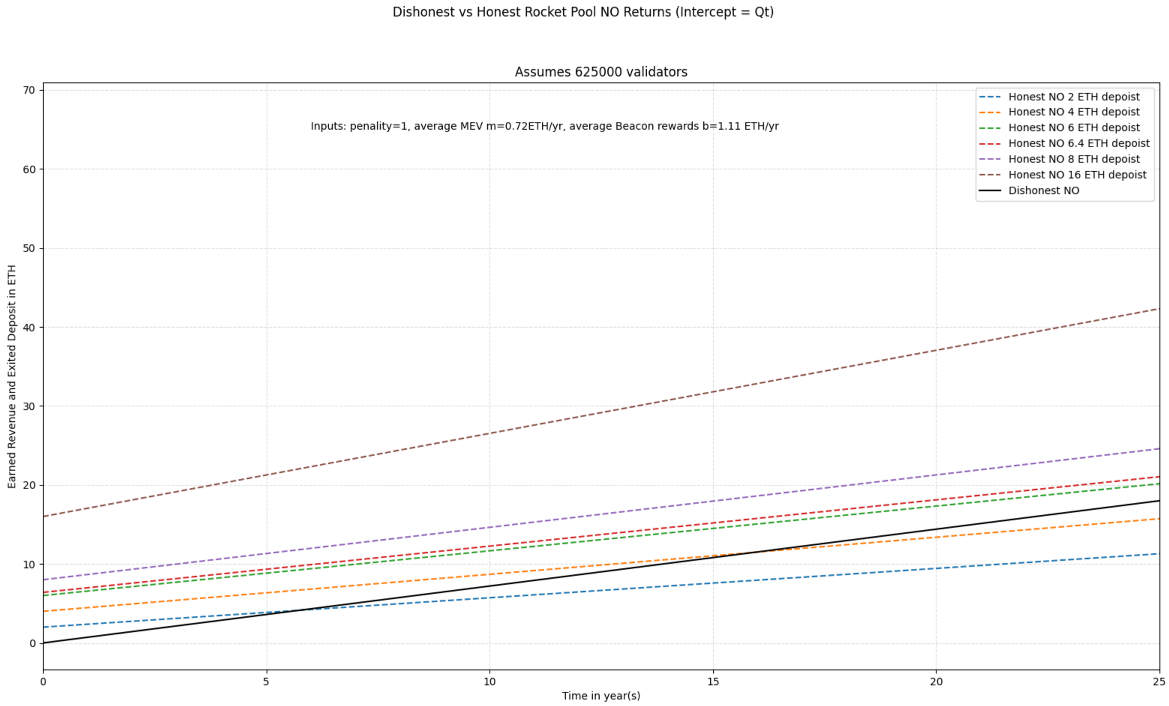

There are a number of risk factors that a dishonest NO would be gambling as part of the long-con theft strategery. These risks include

1. They assume that PPV generation continues to remain at rates comparable to the current rate and does not continue the downward trend of declining PPV over time.

2. A change to the core Ethereum protocol that would prevent MEV (e.g., proposer-builder block separation) does not happen before Q_t.
3. A change that allows the withdrawal credential to execute a beacon chain exit command without the need of the beacon chain signing key does not occur before Q_t.
4. Their dishonest minipool performs in the top 50% of the minipools.

As part of the LEB design criteria a sufficiently long period of Q_t, 10 years, was chosen to minimize the amount of prospective dishonest NO from gambling on this stealing strategy.

Assessing the Impact of Successful Attacks

Although both attack strategies have a low probability of success there is some non-zero chance of them succeeding. Both of those chances rely on the validator's luck in finding more block proposals faster and PPV higher than the average validator.

As a first-order approximation, we can perform a simple calculation that determines the reduction to the rETH APR if all PPV was stolen from blocks over a specified threshold. This simple APR reduction calculation does not account for any recovery of stolen rewards through the RP penalty mechanism. It just simply calculates what would be the overall effect to the rETH APR if hypothetically all of the PPV in blocks that exceed the percentile threshold value were to be stolen

Table 12
rETH APR Reduction Estimates at 15% Minipool Commission

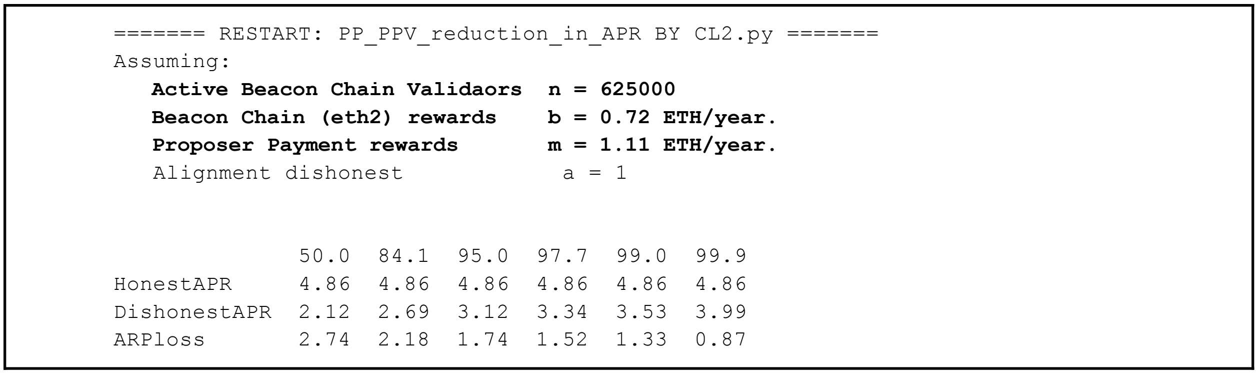

```
======= RESTART: PP_PPV_reduction_in_APR BY CL2.py =======
Assuming:
   Active Beacon Chain Validaors  n = 625000
   Beacon Chain (eth2) rewards    b = 0.72 ETH/year.
   Proposer Payment rewards       m = 1.11 ETH/year.
   Alignment dishonest             a = 1


              50.0  84.1  95.0  97.7  99.0  99.9
HonestAPR     4.86  4.86  4.86  4.86  4.86  4.86
DishonestAPR  2.12  2.69  3.12  3.34  3.53  3.99
ARPloss       2.74  2.18  1.74  1.52  1.33  0.87
```

We can expand on refining the impact of a series of successful attacks on the reduction of the rETH APR by choosing some dishonesty factor, or alignment factor (*a*) that represents the share of NOs, that given the opportunity to steal choose, to do so. We will define *a* such that *a* = 0 is lawful good and *a* = 1 is chaotic evil. Following the Prato principle, we have proposed selecting a = 0.2 in the design specification.

The formulas for calculating the reduction in rETH APR are:

(4) Reduction in rETH APR:

$$APR_{honest} - APR_{dishonest}$$

(5) Honest rETH APR:

$$\text{APR}_{\text{dishonest}} = \frac{((1-s)*(1-c)*b*t)+((1-s)*(1-c)*m*t)}{(32-d)*100}$$

(6) Dishonest rETH APR:

$$\text{APR}_{\text{dishonest}} = \frac{((1-s)*(1-c)*b*t)+((1-s)*(1-c)*m*t*(1-stolen))}{(32-d)*100}$$

(7) Portion of PPV stolen:

$$\text{stolen} = \frac{(\sum_{d}^{\infty} PPV)*a}{\sum_{0}^{\infty} PPV}$$

A two-dimensional plot of the reduction in regular staking APR vs the percentile of PPV that temps a NO to steal the block vs their alignment factor (*a*) is shown in figure 12. The contour lines mark the APR reduction levels of 25 basis points (bp), 50 bp, 75bp, and 100 bp.

Table 13
rETH APR Reduction Estimates at 10% Minipool Commission and Alignment a = 20%

```
======== RESTART: RP_PPV_reduction_in_APR BY CL2.py ========
Assuming:
   Active Beacon Chain Validaors  n = 625000
   Beacon Chain (eth2) rewards    b = 0.72 ETH/year.
   Proposer Payment rewards       m = 1.11 ETH/year.
   Alignment dishonest            a = 0.2


              50.0  84.1  95.0  97.7  99.0  99.9
HonestAPR     4.86  4.86  4.86  4.86  4.86  4.86
DishonestAPR  4.31  4.43  4.51  4.56  4.60  4.69
ARPloss       0.55  0.44  0.35  0.30  0.27  0.17
```

Figure 12
Reduction in rETH APY Plotted Against the Stealing Threshold and the Alignment of the NO

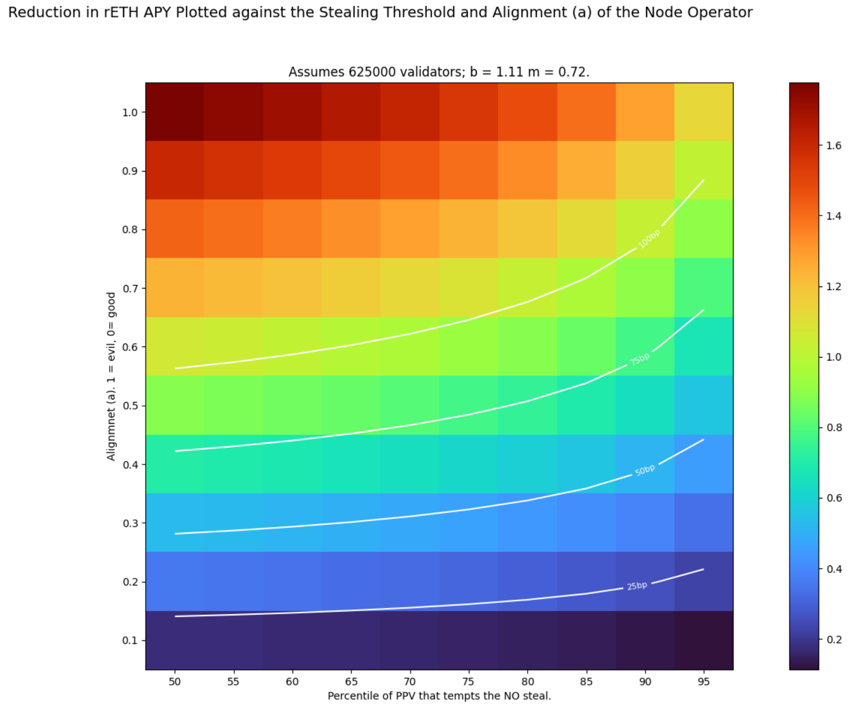

APR Estimates for LEB Minipools

Although the adoption of LEB minipool will help expand the RP protocol by increasing the amount of ETH that can be staked as rETH, there needs to be an incentive for a NO to form a LEB minipool versus a typical 16 ETH minipool. The major incentive for a NO to run multiple LEB minipools versus a typical 16 ETH minipool is the increased APR earned. The increase in APR is due to the increased leveraging of their NO deposit investment. By borrowing a larger portion of the validator's needed 32 ETH stake the NO will a substantially larger return on their minipool commission payment. Figure 13 and table 14 indicate the APR that would be expected based on our modeling and the current minipool commission of 15%.

It is important that this increased APR does not come at the expense of the APR of the regular staking (assuming that dishonestly does not occur). In the honest scenarios, the rate of return for an rETH holder is the same if the minipools staking the funds are typical 16 ETH minipools or if they are the proposed LEB minipools. This is what makes the proposition of LEBs attractive to both NO and regular stakers.

Figure 13
Minipool Node Operator APR Estimates at 15% Minipool Commission

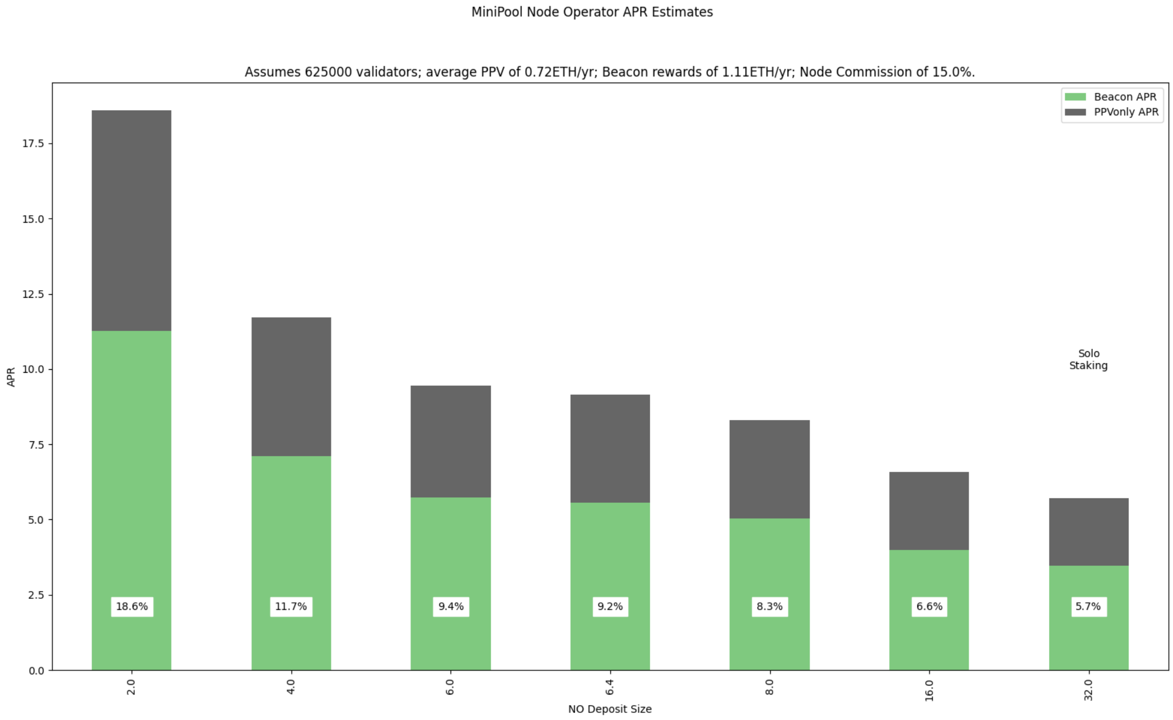

Table 14
Minipool Node Operator APR Estimates at 15% Minipool Commission

```
============= RESTART: C:\Users\Ken\Dropbox\python\crypto\RP_LEB Profit returns.py
============
Assuming :
n = 625000 Validators
c = 15.0% Node Commission
m = 0.72 ETH/yr average PPV per minipool
b = 1.11 ETH /yr average beacon chain (Consensys Layer) rewards per minipool


      NO Deposit  minipool APR  Beacon APR  PPVonly APR
2.0          2.0         18.59       11.27         7.31
4.0          4.0         11.72        7.11         4.61
6.0          6.0          9.44        5.72         3.71
6.4          6.4          9.15        5.55         3.60
8.0          8.0          8.29        5.03         3.26
16.0        16.0          6.58        3.99         2.59
32.0        32.0          5.72        3.47         2.25
```

Mitigation Strategies

The introduction of LEB minipools is not without risk. Each attack scenario has a non-zero chance of succeeding. There could be some incentive or assurance afforded to regular stakers to continue their trust in the rETH investment. Although it can be argued that the likelihood of dishonesty among NO is small given the small probability of success over a short duration we can not rule it out entirely by the mathematical modeling, especially over long time periods. Any incurred losses due to theft of PPV rewards or any loss due to abandonment by the NO would be socialized across the entire rETH holder. As such, no individual rETH investor would be directly at risk but collectively, as a group, they would experience a lower APR than what they could have obtained through other centralized staking services.

One of the original targets proposed in the first tokenomics model for RP was that the minipool commission should target 10% to be competitive with other staking services. Currently, in order to attract more NOs, RP offers a fixed 15% minipool commission. Because of the amount of leverage that a LEB minipool offers it would be possible for RP to reduce the minipool commission to 10% yet still be very attractive to potential NOs. Figure 14 and table 15 display the estimated APR for various NO deposit values. At the time of the modeling, these are some of the highest if not the highest, projected APRs on Ethereum staking that has been advertised.

Figure 14
Minipool Node Operator APR Estimates at 10% Minipool Commission

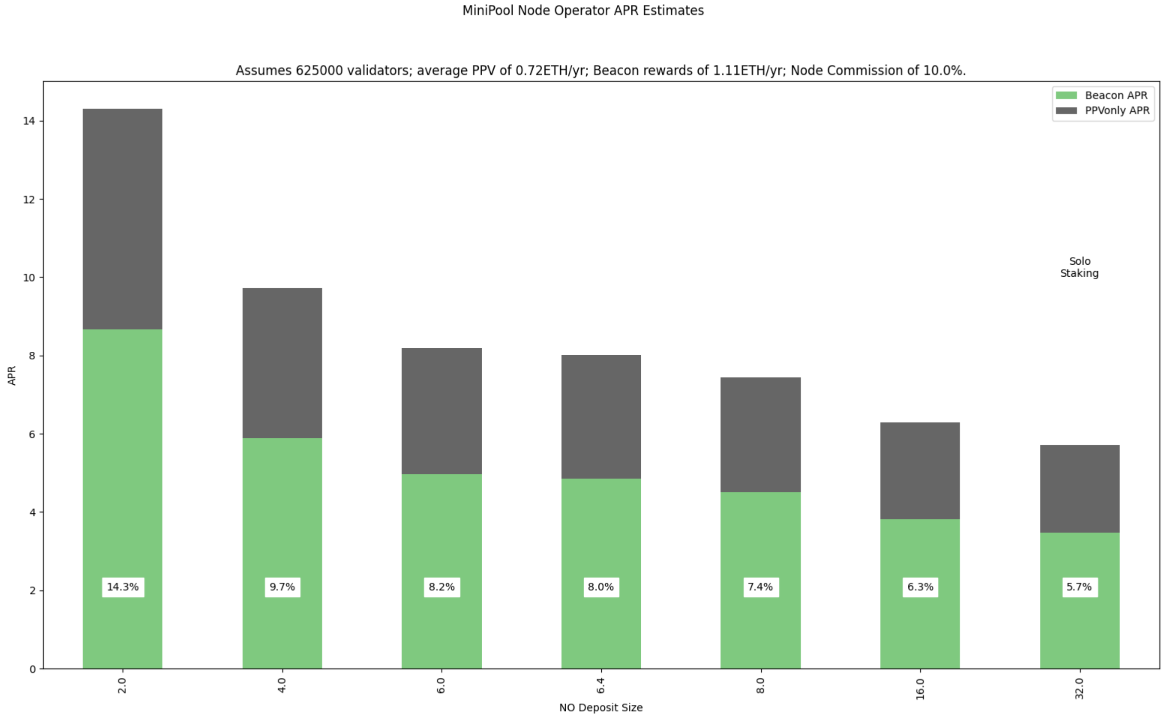

Table 15
Minipool Node Operator APR Estimates at 10% Minipool Commission

```
============ RESTART: C:\Users\Ken\Dropbox\python\crypto\RP_LEB Profit returns.py
============
Assuming :
n = 625000 Validators
c = 10.0% Node Commission
m = 0.72 ETH/yr average PPV per minipool
b = 1.11 ETH /yr average beacon chain (Consensys Layer) rewards per minipool


      NO Deposit  minipool APR  Beacon APR  PPVonly APR
2.0          2.0         14.30        8.67         5.62
4.0          4.0          9.72        5.90         3.82
6.0          6.0          8.20        4.97         3.23
6.4          6.4          8.01        4.86         3.15
8.0          8.0          7.43        4.51         2.92
16.0        16.0          6.29        3.82         2.47
32.0        32.0          5.72        3.47         2.25
```

Future Protections

Additional protocol improvements that would help mitigate the risks of both LEB mini pool and distributed staking protocols in general are:

1. The ability to eject a node operator that acts dishonestly. This "forced ejections" ability envisioned would be possible via an execution layer transaction. If an EIP designing such a mechanism were to be adopted this would ensure that nearly all of the risk and attack vectors evaluated in the report would be essentially prevented.

2. Banning of the node operator that steals from participating in flashbots. The basis of PPV is predicated on the block proposer executing the transaction correctly and not stealing the MEV for themselves. The majority of PPV is paid through third-party block builders like flashbots. RP could partner with flashbots to blacklist any dishonest node operator that steals PPV. This essentially renders both attack strategies useless as without a relay service they would have to perform MEV searcher/builder duties and then build their own MEV extraction routines. This would prove to be difficult and likely not cost-productive for the few times a year that their minipool is selected to block propose.

3. RP can continue to support the core Ethereum proposals of block proposer-builder separations.

4. Since the majority of the impact of dishonestly occurs from stealing the larger value PPV blocks. Tiered node commissions could be explored that increase the higher the PPV award is. This increasing bounty for honest behavior may act as a stronger incentive than the fixed commission rate.

5. Implementation of distributed validator technology (DVT) which splits a BLS beacon chain signing key into multiple fractions and requires a threshold of signatures (e.g., 3 out of 5) to be valid. This would ensure that the `fee_recipient` is an agreed-upon address by the fractional keyholders.

Minimum LEB Minipool Design

In this final matrix (table 16) we compare the original LEB minipool design specifications against a variety of NO deposit values. If the design criteria were met we indicated that with a Boolean response of yes or no. The minimally viable NO deposit that would be acceptable for a LEB minipool would be identified as the smallest deposit amount for which all design specifications are met with a Yes.

Table 16
Design Specification Meet per Node Operator Deposit

			Node Operator (NO) Deposit				
Chain	**Scenario**	**Design Criteria**	**16 ETH+ 1.6 Ξ_{RPL}**	**8 ETH+ 1.6 Ξ_{RPL}**	**6 ETH+ 1.6 Ξ_{RPL}**	**4 ETH+ 1.6 Ξ_{RPL}**	**2 ETH+ 1.6 Ξ_{RPL}**
Consensus	**Validator Off-line**	A. The NO deposit is greater than the offline penalty for short durations during chain consensus.	Yes	Yes	Yes	Yes	Yes
	Validator Abandoned	B. The NO deposit is greater than the penalty for neglect for 2 years or more during chain consensus.	Yes	Yes	Yes	Yes	Yes*
	Isolated Slashing	C. The NO deposit is greater than the beacon chain penalty for an isolated (uncorrelated) slashing violation during chain consensus.	Yes	Yes	Yes	Yes	Yes
	Correlated Slashing	D. The NO deposit is greater than a 1% correlated slashing event during finality.	Yes	Yes	Yes	Yes	Yes
Non-finality	**Offline**	E. The NO deposit is greater than the penalty for non-responsiveness for 7 days during nonfinality.	Yes	Yes	Yes	Yes	Yes
	PPV Theft	F. The NO deposit plus their earned share of the beacon ETH rewards are larger than the expected average PPV for 5 years of validating. (625k Validators)	Yes	Yes	Yes	Yes	Yes
		G. The NO deposit is greater than a 2-sigma (97.7 percentile) probability of total PPV for 1 year of validating. (400k Validators)	Yes	Yes	Yes	Yes	No
		H. The NO deposit is greater than a 95% probability of maximum PPV in a block for 1 year of validating. (400k Validators)	Yes	Yes	Yes	Yes	No

General References

1. Understanding the Validator lifecycle. Jim McDonald Jan 23, 2020.
2. Upgrading Ethereum Edition 0.1: Altair [WIP] by Ben Edgington
3. Github: Ethereum / consensus-specs / slash_validator
4. Github: Ethereum / consensus-specs / slashings
5. Github: Ethereum / consensus-specs / Rewards and penalties
6. Investopedia: Empirical Rule
7. Rewards and Penalties on Ethereum 2.0 [Phase 0] by James BeckMarch 2, 2020