

Towards Decentralized Searcher Competition in MEV Markets

Roozbeh Sarenche ✉ 

COSIC, KU Leuven, Belgium

Yunwen Liu ✉ 

COSIC, KU Leuven, Belgium

Abstract

Centralization in maximal extractable value (MEV) markets is a significant concern for blockchain systems, as persistent concentration of economic power can weaken competition, reduce openness, and undermine the decentralization goals of permissionless protocols. While much of the existing analysis has focused on builders, validators, and block-building markets, this paper studies centralization from the perspective of searcher competition. We develop a heterogeneous model in which searchers differ in opportunity coverage and execution efficiency, and we analyze how auction design affects fairness, decentralization, and security among searchers competing for the same MEV opportunity.

To evaluate searcher competition, we introduce two metrics: a Shapley-weighted Jain fairness index, which measures whether rewards are proportional to searchers' marginal contributions, and an expected-reward Herfindahl–Hirschman Index (HHI), which measures concentration in long-run searcher rewards. Using these metrics, we first analyze the standard first-price, winner-take-all auction as a benchmark. Our analysis shows that, under searcher heterogeneity, first-price competition can reward rank dominance rather than marginal contribution, leading to concentrated rewards and weaker contribution-adjusted fairness.

Motivated by these limitations, we propose an entry-filtered Shapley-capped auction mechanism that distributes searcher rewards more fairly and broadly among admitted high-quality submissions. Designing such a mechanism in a permissionless blockchain environment is challenging: searchers may create Sybil identities by submitting copied or degraded versions of the same execution strategy, and validators may collude with searchers to increase joint payoff. We address these concerns through Bayesian security constraints for copied-code Sybil deviations and validator–searcher coalition deviations, and show that suitable parameter choices make the Shapley-capped mechanism secure against both.

We complement the theoretical analysis with numerical benchmarks and Ethereum on-chain data. The theoretical benchmarks show that the proposed mechanism is especially effective in centralized opportunity classes, where first-price competition concentrates rewards among a small set of searchers. They further indicate that the proposed mechanism is a potential improvement over first-price competition in empirically motivated opportunity types.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory and mechanism design; Security and privacy → Economics of security and privacy

Keywords and phrases Maximal extractable value, MEV, searcher competition, mechanism design, Sybil resistance, Shapley value, reward decentralization

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.21

Related Version *Full Version:* <https://arxiv.org/abs/2608.05011> [20]

Supplementary Material *Software (Source Code):* https://gitlab.esat.kuleuven.be/mev/searcher_game

Funding *Roozbeh Sarenche:* Supported by the Flemish Government through the Cybersecurity Research Program Flanders with grant number: VOEWICS02.

Yunwen Liu: Partially supported by KU Leuven STG/25/039 and CTVG R&D grant (QJW-410118-LIU-CTVG).



© Roozbeh Sarenche and Yunwen Liu;

licensed under Creative Commons License CC-BY 4.0

8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyropoulou; Article No. 21; pp. 21:1–21:24

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In Ethereum’s Maximal Extractable Value (MEV) supply chain, *searchers* are specialized actors who identify profitable execution opportunities and submit transactions or bundles to capture them [4, 5]. Their submissions enter auction-like competition for block inclusion, where searchers bid part of the generated value to block producers – builders and validators.

Prior work suggests that searcher-level MEV extraction is far from a symmetric market. Searchers differ in opportunity access, latency, information, capital, and execution quality, while specialized bot behavior and heterogeneous competitive pressure further amplify these differences [10, 26]. For atomic MEV strategies, including sandwiching and decentralized-exchange arbitrage, early studies already documented sophisticated bot specialization and intense competition [4]. More recent evidence shows that concentration can be severe: in centralized–decentralized exchange arbitrage, Heimbach et al. [9] find that eleven searchers account for more than 80% of identified arbitrage volume, with the two largest searchers alone accounting for 49.2%.

These observations suggest that searcher markets can become highly skewed in the presence of players with specialized infrastructure, private information, capital, or privileged access. This skewness, together with current searcher-auction design, has consequences beyond searcher profits and can become a bottleneck for blockchain decentralization. Persistent reward concentration weakens open competition and raises effective barriers to entry. This concern is amplified on high-throughput chains where MEV competition often takes the form of large-scale speculative transaction submission [25]. It can also reinforce builder-side concentration when valuable searcher flow is routed privately or integrated with particular builders [9, 12, 15]. Searcher concentration can then propagate through the whole MEV supply chain: dominant searchers become preferred counterparties, builders with access to their flow gain an advantage, and validators face a block-building market shaped by privileged order-flow channels.

Most work on MEV concentration has focused on the block-building layer, analyzing how private order flow, vertical integration, latency advantages, and strategic bidding drive builder concentration under proposer-builder separation [2, 8, 9, 12, 14, 15, 24, 28, 29, 30, 31, 32]. Within this view, searchers are typically modeled merely as sources of transaction bundles or private flow that confer competitive advantages to builders. Consequently, the internal mechanics of the searcher auction market have received considerably less attention. While recent studies have begun to examine searcher mechanism design and bidding strategies [11], the specific implications of these auction formats for fairness and decentralization in the searcher market remain largely unexplored. This motivates our central question: how fair and decentralized is the current first-price, winner-take-all searcher auction, and can an alternative mechanism design improve reward decentralization while preserving the rigorous security guarantees required for permissionless MEV competition?

Decentralizing rewards in a permissionless environment is fundamentally difficult. If a mechanism pays multiple submissions, a searcher may split a single execution strategy into multiple identities, submitting copied or intentionally degraded versions of the same code. Consequently, any mechanism that naively rewards participation inherently invites Sybil attacks. Moreover, validators and searchers have strong incentives to coordinate outside the protocol if identity splitting can increase their joint payoff. Therefore, a decentralized searcher auction must solve two problems simultaneously: it must distribute rewards more broadly and fairly, while strictly preserving the security guarantees required for permissionless MEV execution.

This paper studies that tradeoff. We develop a model of heterogeneous searcher competition, analyze the standard first-price auction, and introduce a Shapley-based mechanism as a more decentralized alternative. The main message is that the current winner-take-all design is secure against copied-code deviations, but can be economically centralizing under realistic searcher heterogeneity. By contrast, a carefully capped and burned Shapley-based mechanism can improve contribution-adjusted fairness and reward decentralization, while controlling Sybil and validator–searcher coalition incentives through explicit Bayesian security constraints.

1.1 Our Contribution

Our contribution includes:

System model and evaluation criteria. We introduce a model of searcher competition in which searchers differ in both opportunity coverage and execution quality. For each opportunity class, a searcher may or may not discover and submit for the opportunity, and conditional on participation its submitted code has a certified normalized net value. This lets us separate two important sources of searcher advantage: seeing more opportunities and executing a given opportunity more efficiently. We then define economic and security criteria for evaluating searcher auctions. The economic criteria are a Shapley-weighted Jain fairness index, which measures whether rewards are aligned with marginal contribution, and an expected-reward HHI, which measures reward concentration. The security criteria are Bayesian copied-code Sybil resistance and Bayesian validator–searcher coalition resistance.

Analysis of the first-price auction. We analyze the first-price auction as the benchmark for current winner-take-all searcher competition. Under monotone searcher-specific bidding, copied or degraded versions of a searcher’s own code cannot improve either the searcher’s payoff or the joint payoff of the searcher and validator. Thus, the first-price auction has strong copied-code security properties. Economically, however, its rewards are governed by rank dominance: searchers are rewarded for being the highest-bidding active participant, not for their marginal contribution. We show that this distinction becomes important under heterogeneity. When execution quality is unequal, first-price rewards can concentrate on a small dominant class; in the limiting case of a single dominant searcher, the effective number of rewarded searchers approaches one.

Shapley-Capped Auction. We propose the Shapley-Capped Auction as an alternative design. The mechanism still executes the highest certified-value submission, but when the number of admitted submissions is below a cap, it distributes retained rewards according to the Shapley value of the validator-dependent max game. When the number of admitted submissions exceeds the cap, it switches to a conservative fallback branch and burns the residual. The cap and burn parameters are used to control the incentives created by reward sharing. We prove budget feasibility and provide a Bayesian security characterization showing that it is sufficient to consider a bounded range of Sybil sizes while optimizing over all feasible copied or degraded submissions. This yields a feasible region for the security parameters under which the mechanism satisfies copied-code Sybil resistance and validator–searcher coalition resistance.

Real-world data and empirical calibration. We analyze Ethereum arbitrage data to connect the model to observed searcher behavior. The goal is to estimate practical patterns of searcher participation, profit concentration, and opportunity coverage, and then use these

patterns to evaluate the mechanism under realistic parameters. This empirical component is not meant to fully identify every real-world searcher entity, since a single entity may control multiple contracts, but it provides a data-driven basis for calibrating the participation and execution-quality primitives used in the theoretical and numerical analysis.

1.2 Related work

Searcher’s MEV activities. To produce competitive bundles, searchers rely on a combination of advantageous resources: efficient algorithms and heuristics to detect MEV opportunities, substantial capital to execute transactions across domains, and information asymmetries to shield lucrative opportunities from competitors. The strategic landscape has grown considerably more complex over recent years. Early MEV extraction was focused on *atomic* strategies, where all transactions constituting an attack execute within a single chain or block, guaranteeing all-or-nothing execution. Prominent examples include cyclic arbitrage across decentralized exchanges (DEX) and sandwich attacks on pending user trades [4, 17]. Later, the proliferation of Layer-2 systems and the frictions between heterogeneous market domains have increasingly pushed searchers toward *non-atomic* strategies, giving rise to arbitrage opportunities spanning multiple layers, chains, and off-chain venues such as centralized exchanges (CEX) [7, 13, 22].

Centralization in the MEV market. Block building on Ethereum has become highly concentrated: the leading builder, Titan, alone accounts for around 50% of blocks in May 2026. A prominent study explaining builder centralization is that of Öz et al. [12], which links builders’ market share to order-flow diversity and their profitability to access to exclusive order-flow providers. Recent measurement studies on block builders also suggest that builder dominance is shaped not only by builder-side strategies, but more importantly by access to valuable searcher-generated bundles [9, 27, 31]. These bundles may derive their value from private information [14], more effective search and bidding strategies [10], or vertical integration between searchers and builders [15].

Compared with the extensive literature on block builders, systematic studies of searchers remain relatively recent. Searchers are heterogeneous in both strategy and market access. At the strategy level, searchers specialize in different forms of MEV extraction, including arbitrage, sandwich, liquidation, and mixed strategies [10]. At the access level, searchers may rely on public mempool transactions, private or exclusive order flow, MEV-Share signals¹, and off-chain information.

Moreover, searcher competition is unevenly distributed across opportunities and actors. Mamageishvili et al. [11] showed that validator revenue increases with searcher competition: when many searchers identify similar opportunities, competition pushes most of the value to the validator, while searchers retain more value when competition is weak. Empirical studies further confirmed that the searcher market is highly skewed toward large or specialized actors. For non-atomic arbitrage, Heimbach et al. [9] and Wu et al. [27] found that a small number of searchers account for most of the identified arbitrage volume and value.

These results suggest that searchers do not compete in a single homogeneous market – some MEV opportunities are highly contested, whereas others are repeatedly captured by searchers with persistent advantages. From a game-theoretic perspective, Mamageishvili et al. [11] provide an important foundation by modeling searcher competition as a cooperative

¹ <https://github.com/flashbots/mev-share>

submodular game, where the value function maps each coalition of searcher bundles to the highest block value achievable from that coalition. What remains less understood is how to model searcher competition in a way that captures heterogeneous opportunity access, variation in execution quality, and measurable competition intensity across the searcher market.

Game-Theoretic Approaches to MEV Auction Design. Recent game-theoretic analyses of MEV auctions increasingly emphasize the tension between allocative efficiency, decentralization, and resilience against strategic manipulation [19, 6]. While traditional winner-take-all procurement auctions inherently drive market concentration, alternative mechanisms that penalize monopoly allocations must strictly account for Sybil attacks in permissionless environments [6]. Pan et al. [16] demonstrate that the second-price auction with symmetric tie-breaking is the only non-wasteful, incentive-compatible, and Sybil-proof direct mechanism in single-parameter settings, severely limiting the design space for secure alternative distributions. To circumvent these constraints and achieve fair surplus allocation without compromising security, recent work applies cooperative game theory directly to order flow auctions (OFAs) [11]. Specifically, Shapley value-based frameworks have been proposed to equitably redistribute matchmaking revenues among transaction creators according to their marginal contributions, utilizing randomized approximation to handle the computational complexity of the MEV supply chain [18]. Complementing these static allocation approaches, Braga et al. [3] propose a dynamic MEV-sharing mechanism that adjusts the protocol-level extraction rate over time to balance participation incentives between users and MEV extractors. Together, this literature highlights a critical shift from purely competitive block building toward robust, Sybil-proof mechanisms and cooperative revenue-sharing models.

2 Model, Metrics, and Mechanisms

This section introduces a heterogeneous model of MEV searcher competition. Searchers differ in both opportunity coverage and execution efficiency. We evaluate mechanisms along two economic dimensions – contribution-adjusted fairness and reward decentralization – and two security dimensions – copied-code Sybil resistance and validator–searcher coalition resistance.

2.1 System Model

We model searcher competition for MEV opportunities in a setting with heterogeneous opportunity discovery, heterogeneous execution quality, and permissionless submissions. We focus on on-chain competition and abstract away from exogenous off-chain or cross-domain factors, such as centralized-exchange price movements and MEV opportunities originating from off-chain systems.

Entities and opportunity classes. There is one validator (proposer) and a finite set of potential searchers $\mathcal{S} = \{1, \dots, n\}$. MEV opportunities are indexed by $o \in \mathcal{O}$. Each opportunity belongs to a class $\tau(o) \in \mathcal{C}^2$.

² A class may represent a broad MEV type, such as arbitrage, liquidation, sandwiching, or a finer subclass, such as arbitrage over a particular family of AMMs or routes.

Opportunity coverage. Searcher $i \in \mathcal{S}$ can participate in opportunity o only if it discovers the opportunity and can submit *execution code* before the deadline. We call such searchers *active searchers* for opportunity o . Let $D_{i,o} \sim \text{Bernoulli}(\alpha_{i,\tau(o)})$ denote this event. The parameter $\alpha_{i,\tau(o)} \in [0, 1]$ is searcher i 's coverage probability for class $\tau(o)$. The realized active set for opportunity o is

$$\mathcal{S}_{\text{act}}(o) = \{i \in \mathcal{S} : D_{i,o} = 1\}, \quad N_o = |\mathcal{S}_{\text{act}}(o)|.$$

How opportunities are clustered into classes, and how $\mathcal{S}_{\text{act}}(o)$ is empirically identified, is outside the formal model.³

Execution code, certified values, and normalization. Any execution code is identified by two certified quantities under the canonical simulation rule: its gross realized value and its execution cost. Let $\tilde{v}_{i,o}$ denote the gross certified value generated by searcher i 's code, and let $\tilde{c}_{i,o}$ denote its certified execution cost. The cost $\tilde{c}_{i,o}$ captures execution losses that are objectively computable from the submitted code, such as base-fee gas costs, realized slippage, and flash-loan fees; it does not include transfers to the validator. We define the certified net value of this code as $\tilde{z}_{i,o} = \tilde{v}_{i,o} - \tilde{c}_{i,o}$.

For each opportunity o , let $V_o > 0$ denote a public opportunity-scale parameter, known before the submission deadline. This scale represents the commonly observable or commonly estimated size of the opportunity and is used to normalize certified values across opportunities. We normalize these quantities by V_o : $v_{i,o} = \frac{\tilde{v}_{i,o}}{V_o}$, $c_{i,o} = \frac{\tilde{c}_{i,o}}{V_o}$, $z_{i,o} = \frac{\tilde{z}_{i,o}}{V_o}$. We assume V_o is chosen so that $v_{i,o} \leq 1$ and $z_{i,o} \leq 1$.

Submissions and copied-code deviations. For opportunity o , only active searchers with $D_{i,o} = 1$ can submit. A submission b consists of execution code and, depending on the auction format, may also include a mechanism-specific message m_b , such as a transfer offer. Each submission has a certified net value $\tilde{z}_b$ and certified normalized net value $z_b = \tilde{z}_b/V_o$, computed under the canonical simulation rule. In a permissionless setting, a searcher may attempt to increase its payoff by representing one execution code through multiple submissions. We model this through *copied-code Sybil deviations* and denote the finite set of submissions made by active searcher i by $\mathcal{B}_i(o)$.

► **Definition 1** (Copied-code Sybil deviation). *Fix an opportunity o , an active searcher $i \in \mathcal{S}_{\text{act}}(o)$, and a submission $b \in \mathcal{B}_i(o)$ with certified normalized net value z_b . A k -submission copied-code Sybil deviation from b , with $k \geq 2$, is a collection of submissions $\sigma_{i,o} = (\hat{b}_1, \dots, \hat{b}_k)$ by the same real searcher i , with certified normalized net values $(z_{\hat{b}_1}, \dots, z_{\hat{b}_k})$, such that $z_{\hat{b}_\ell} \leq z_b$ for every $\ell = 1, \dots, k$. The reference submission b is called the original submission.*

We denote by $\mathcal{B}(o)$ the set of all submissions for opportunity o , including both original submissions and Sybil submissions: $\mathcal{B}(o) = \bigcup_{i \in \mathcal{S}_{\text{act}}(o)} \mathcal{B}_i(o)$, and $t_o = |\mathcal{B}(o)|$. The validator can include at most one submitted execution code for opportunity o .

Bayesian information structure. Before the submission deadline, an active searcher i observes its own original execution code and its certified normalized net value. It does not observe the realized participation decisions, execution codes, certified values, or mechanism-

³ One practical approach is to place submissions in the same active set $\mathcal{S}_{\text{act}}$ when executing the highest-net-value submission would reduce the residual profitability of the remaining submissions below a specified threshold.

specific messages of other searchers⁴. It only knows the corresponding distributional primitives. For each $j \neq i$, $D_{j,o} \sim \text{Bernoulli}(\alpha_{j,\tau(o)})$. Conditional on $D_{j,o} = 1$, the certified normalized net value of searcher j 's original code for opportunity o , denoted by $z_{j,o}^*$, is distributed according to the cumulative distribution function $F_{j,\tau(o)}$: $z_{j,o}^* \sim F_{j,\tau(o)}$, $z_{j,o}^* \in [\underline{z}_{j,\tau(o)}, \bar{z}_{j,\tau(o)}] \subset (0, 1)$. The distribution $F_{j,\tau(o)}$ is searcher- and class-specific: it describes searcher j 's normalized performance across opportunities of class $\tau(o)$. After the submission deadline, the set of submissions is fixed. The mechanism then evaluates the submitted execution codes under the canonical simulation rule and observes their realized certified normalized net values. This interim information structure is the basis for the Bayesian-form security properties discussed later in the paper.

Auction mechanism. For each opportunity o , a searcher auction mechanism $\mathcal{M}$ takes the submitted execution codes, their certified net values, and any mechanism-specific messages as input. It selects one submission for execution and assigns payoffs to the real searchers and the validator. For each submission $b \in \mathcal{B}(o)$, let $\tilde{z}_b = \tilde{v}_b - \tilde{c}_b$ denote its certified net value. Let m_b denote its mechanism-specific message. Write $\tilde{z}_o = (\tilde{z}_b)_{b \in \mathcal{B}(o)}$ and $m_o = (m_b)_{b \in \mathcal{B}(o)}$. The mechanism specifies a winner-selection rule $\text{Win}_o^{\mathcal{M}}(\tilde{z}_o, m_o) \in \mathcal{B}(o)$, which selects a single winning submission. The validator includes the execution code associated with the selected submission. The mechanism also specifies payoff rules $U_i^{\mathcal{M}}(\tilde{z}_o, m_o)$ for each $i \in \mathcal{S}_{\text{act}}(o)$, and $U_{\text{val}}^{\mathcal{M}}(\tilde{z}_o, m_o)$ for the validator. The payoff $U_i^{\mathcal{M}}$ is the total payoff assigned to real searcher i , aggregating over all submissions in $\mathcal{B}_i(o)$.

2.2 Evaluation Criteria

We evaluate a searcher auction mechanism $\mathcal{M}$ along two dimensions: economic performance and security. The economic metrics measure whether rewards are fairly and broadly distributed among real searchers. The security properties require that these rewards cannot be profitably manipulated through copied-code Sybil submissions or validator–searcher coalitions. Let $U_i^{\mathcal{M}}$ denote searcher i 's realized payoff and $U_{\text{val}}^{\mathcal{M}}$ the validator's realized payoff. Define searcher i 's expected reward under $\mathcal{M}$ as $R_i^{\mathcal{M}} = \mathbb{E}[U_i^{\mathcal{M}}]$.

2.2.1 Economic metrics

The first economic metric measures whether expected rewards are aligned with searchers' contributions. We use the Shapley value as the contribution benchmark. The Shapley value is a standard solution concept from cooperative game theory: it assigns to each participant its average marginal contribution over all possible orders in which participants could join a coalition. In our setting, this provides a natural way to measure how much each searcher contributes to the expected surplus generated by competition for an opportunity.

Let $\mathcal{N} = \mathcal{S} \cup \{\text{val}\}$. For $A \subseteq \mathcal{N}$, define the expected value of coalition A in the validator–searcher game as

$$W^{\text{VS}}(A) = \begin{cases} \mathbb{E}[V_o \max_{i \in A \cap \mathcal{S}: D_{i,o}=1} z_{i,o}^*], & \text{val} \in A, \\ 0, & \text{val} \notin A, \end{cases}$$

⁴ This information structure can be implemented, for example, by an encrypted mempool or commit–reveal protocol in which validators and searchers do not observe other submissions before the epoch deadline for decryption or reveal.

with the convention that the maximum over an empty set is zero. Searcher i 's validator–searcher Shapley contribution [21] is

$$\phi_i^{\text{VS}} = \sum_{A \subseteq \mathcal{N} \setminus \{i\}} \frac{|A|!(n - |A|)!}{(n + 1)!} (W^{\text{VS}}(A \cup \{i\}) - W^{\text{VS}}(A)).$$

Assume $\phi_i^{\text{VS}} > 0$ for all searchers included in the metric. Define the reward-to-contribution ratio $r_i^{\mathcal{M}} = R_i^{\mathcal{M}} / \phi_i^{\text{VS}}$. We then measure searcher-side fairness by applying Jain's index to these ratios: a mechanism is fair when searchers receive similar rewards per unit of validator-dependent Shapley contribution.

► **Definition 2** (Shapley-weighted Jain fairness). *The Shapley-weighted Jain fairness index of mechanism $\mathcal{M}$ is*

$$\mathcal{J}_{\text{Sh}}^{\mathcal{M}} = \frac{(\sum_{i=1}^n r_i^{\mathcal{M}})^2}{n \sum_{i=1}^n (r_i^{\mathcal{M}})^2} = \frac{(\sum_{i=1}^n R_i^{\mathcal{M}} / \phi_i^{\text{VS}})^2}{n \sum_{i=1}^n (R_i^{\mathcal{M}} / \phi_i^{\text{VS}})^2}.$$

The index satisfies $1/n \leq \mathcal{J}_{\text{Sh}}^{\mathcal{M}} \leq 1$, provided that at least one searcher receives positive expected reward. It equals 1 in a fully contribution-fair allocation, where all searchers receive the same reward per unit of Shapley contribution; equivalently, $R_i^{\mathcal{M}} = \lambda \phi_i^{\text{VS}}$ for all i and some $\lambda > 0$. At the other extreme, it equals $1/n$ when only one searcher has a positive reward-to-contribution ratio, meaning that rewards are maximally concentrated relative to contribution. Note Jain's original index measures equality among the quantities to which it is applied. Since we apply it to $R_i^{\mathcal{M}} / \phi_i^{\text{VS}}$, this metric measures equality of contribution-adjusted rewards.

The second economic metric measures how concentrated long-run expected rewards are among searchers. Whenever $\sum_{j=1}^n R_j^{\mathcal{M}} > 0$, let $s_i^{\mathcal{M}} = \frac{R_i^{\mathcal{M}}}{\sum_{j=1}^n R_j^{\mathcal{M}}}$ be searcher i 's expected reward share.

► **Definition 3** (Effective reward decentralization). *The expected-reward Herfindahl–Hirschman Index (HHI) and effective number of rewarded searchers are*

$$\text{ER-HHI}^{\mathcal{M}} = \sum_{i=1}^n (s_i^{\mathcal{M}})^2, \quad N_{\text{eff}}^{\mathcal{M}} = \frac{1}{\text{ER-HHI}^{\mathcal{M}}}.$$

The Herfindahl–Hirschman Index (HHI) is a standard concentration measure. Here we apply it to expected reward shares rather than market shares, so $\text{ER-HHI}^{\mathcal{M}}$ measures concentration in searcher rewards. It ranges from $1/n$, when all searchers receive equal expected reward shares, to 1, when one searcher receives all expected rewards. The reciprocal $N_{\text{eff}}^{\mathcal{M}}$ gives the effective number of rewarded searchers: it equals n under equal reward shares and equals 1 when rewards are fully concentrated on a single searcher. Unlike $\mathcal{J}_{\text{Sh}}^{\mathcal{M}}$, this metric does not adjust for contribution; it measures reward decentralization directly.

2.2.2 Security properties

The first security property concerns identity splitting by a searcher that can only copy or degrade its own execution code. For the security definitions, write $\mathcal{B}_{-i}(o) = \bigcup_{j \neq i} \mathcal{B}_j(o)$. We use $U_i^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o))$ to denote the payoff of real searcher i when it submits only the original submission b , and $U_i^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o))$ to denote its payoff when it replaces b by a copied-code Sybil deviation σ_i . The payoff $U_i^{\mathcal{M}}$ aggregates over all submissions controlled by real searcher i .

A strong benchmark is ex-post copied-code Sybil resistance [16]. This property requires that copied-code identity splitting is never profitable for any realized profile of rival submissions.

► **Definition 4** (Ex-post copied-code Sybil resistance). *A mechanism $\mathcal{M}$ is ex-post copied-code Sybil-resistant if, for every opportunity o , every real searcher i , every original submission $b \in \mathcal{B}_i(o)$, every realized rival submission set $\mathcal{B}_{-i}(o)$, and every copied-code Sybil deviation σ_i from b , $U_i^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) \geq U_i^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o))$.*

This security requirement states that, once the rival submissions $\mathcal{B}_{-i}(o)$ are fixed, a real searcher cannot increase its total payoff by replacing one original submission b with any feasible collection of copied-code submissions σ_i .

We next introduce a weaker security requirement that evaluates copied-code deviations in expectation rather than state by state [16]. Under this requirement, a copied-code Sybil deviation is ruled out if it does not increase the deviating searcher's expected payoff conditional on its own participation and certified normalized net value.

► **Definition 5** (Bayesian copied-code Sybil resistance). *A mechanism $\mathcal{M}$ is Bayesian copied-code Sybil-resistant if, for every real searcher i , every $z \in [\underline{z}_{i,\tau(o)}, \bar{z}_{i,\tau(o)}]$, every original submission b with $z_b = z$, and every copied-code Sybil deviation σ_i from b ,*

$$\mathbb{E}[U_i^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) \mid D_{i,o} = 1, z_{i,o}^* = z] \geq \mathbb{E}[U_i^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o)) \mid D_{i,o} = 1, z_{i,o}^* = z].$$

This weaker Bayesian definition is appropriate for our introduced deadline-based system model (see Sect. 2.1). Before the submission deadline, a deviating searcher observes its own execution code and certified normalized net value, but cannot condition its copied-code Sybil strategy on the realized submissions of its competitors.

The second security property extends copied-code Sybil resistance to validator–searcher coalitions. It requires that replacing an original submission with copied-code submissions cannot increase the combined payoff of the deviating searcher and the validator. This property is important because, if a copied-code deviation can increase joint surplus, the searcher and validator may have a direct incentive to coordinate the deviation and share the gain. As above, we define both an ex-post version and a Bayesian version of this requirement.

► **Definition 6** (Ex-post validator–searcher coalition resistance). *A mechanism $\mathcal{M}$ is ex-post validator–searcher coalition-resistant under copied-code deviations if, for every opportunity o , every real searcher i , every original submission $b \in \mathcal{B}_i(o)$, every realized rival submission set $\mathcal{B}_{-i}(o)$, and every feasible copied-code Sybil deviation σ_i from b ,*

$$U_i^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) + U_{\text{val}}^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) \geq U_i^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o)) + U_{\text{val}}^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o)).$$

This security requirement rules out copied-code deviations that increase the joint payoff of the deviating searcher and the validator, evaluated state by state for each realized rival submission set.

► **Definition 7** (Bayesian validator–searcher coalition resistance). *A mechanism $\mathcal{M}$ is Bayesian validator–searcher coalition-resistant under copied-code deviations if, for every real searcher i , every $z \in [\underline{z}_{i,\tau(o)}, \bar{z}_{i,\tau(o)}]$, every original submission b with $z_b = z$, and every feasible copied-code Sybil deviation σ_i from b ,*

$$\begin{aligned} & \mathbb{E}[U_i^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) + U_{\text{val}}^{\mathcal{M}}(\{b\} \cup \mathcal{B}_{-i}(o)) \mid D_{i,o} = 1, z_{i,o}^* = z] \\ & \geq \mathbb{E}[U_i^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o)) + U_{\text{val}}^{\mathcal{M}}(\sigma_i \cup \mathcal{B}_{-i}(o)) \mid D_{i,o} = 1, z_{i,o}^* = z]. \end{aligned}$$

The Bayesian version applies the same joint-payoff comparison from the interim perspective of the deviating searcher. It rules out copied-code deviations that increase the expected combined payoff of the searcher–validator coalition under the deadline-based information structure.

3 First-Price Auction

We first analyze the sealed first-price auction (FPA), which is the natural benchmark for current private MEV bundle markets. In this mechanism, each active searcher submits a transfer offer to the validator; the highest offer wins, and the winner pays its own offer.

3.1 FPA Mechanism

For the FPA, the mechanism-specific message is the unnormalized transfer offer to the validator, $m_b = \widetilde{\text{bid}}_b$. We write $\text{bid}_o = (\text{bid}_b)_{b \in \mathcal{B}(o)}$ for the vector of submitted transfer offers. The first-price auction selects the submission with the highest transfer offer, breaking ties by a fixed priority order, such as encrypted-mempool arrival order, fixed before reveal⁵:

$$\text{Win}_o^{\text{FPA}}(\tilde{z}_o, \widetilde{\text{bid}}_o) \in \arg \max_{b \in \mathcal{B}(o)} \widetilde{\text{bid}}_b.$$

When the argmax contains multiple submissions, $\text{Win}_o^{\text{FPA}}$ is the highest-priority submission among them. If submission $b^* = \text{Win}_o^{\text{FPA}}(\tilde{z}_o, \widetilde{\text{bid}}_o)$ wins, the validator includes the execution code associated with b^* . The owner $s(b^*)$ receives payoff $U_{s(b^*)}^{\text{FPA}}(\tilde{z}_o, \widetilde{\text{bid}}_o) = \tilde{z}_{b^*} - \widetilde{\text{bid}}_{b^*}$, and the validator receives $U_{\text{val}}^{\text{FPA}}(\tilde{z}_o, \widetilde{\text{bid}}_o) = \widetilde{\text{bid}}_{b^*}$. All losing searchers receive zero.

For the FPA, each submission $b \in \mathcal{B}_i(o)$ made by searcher i has normalized certified net value z_b and submits the normalized transfer offer $\text{bid}_b = \beta_{i,\tau(o)}(z_b)$, where $\beta_{i,\tau(o)}$ is searcher i 's bidding function for opportunities of class $\tau(o)$. The corresponding unnormalized transfer offer is $\widetilde{\text{bid}}_b = V_o \text{bid}_b = V_o \beta_{i,\tau(o)}(z_b)$.

The formal security analysis of the FPA is deferred to Appendix A of the full version [20]. There we prove both copied-code Sybil resistance and validator–searcher coalition resistance in the strong ex-post sense; the corresponding Bayesian guarantees follow immediately.

3.2 Expected Rewards and Economic Properties under FPA

This section characterizes the economic allocation induced by the FPA. The key feature is that rewards are assigned only to the submission with the highest transfer offer. Hence, under monotone bidding, the FPA rewards the event of being the highest-bidding active searcher rather than marginal contribution to the opportunity.

Let $F_i = F_{i,\tau(o)}$ and $\alpha_i = \alpha_{i,\tau(o)}$. Suppose first that the value distributions are continuous and the bidding functions $\beta_{i,\tau(o)}$ are strictly increasing, so that, for any transfer offer bid in the range of $\beta_{i,\tau(o)}$, the inverse bidding function $\beta_{i,\tau(o)}^{-1}(\text{bid})$ is well defined. Conditional on searcher i being active and having normalized certified net value z , its interim win probability is

$$\Pi_i^{\text{win}}(z) = \prod_{j \neq i} \left(1 - \alpha_j + \alpha_j F_j \left(\beta_{j,\tau(o)}^{-1}(\beta_{i,\tau(o)}(z)) \right) \right).$$

⁵ For copied-code deviations, copied submissions cannot obtain higher tie-breaking priority than the original submission they copy.

Therefore searcher i 's expected reward under the FPA is

$$R_i^{\text{FPA}} = \mathbb{E}[U_i^{\text{FPA}}] = \alpha_i V_o \int (z - \beta_{i,\tau(o)}(z)) \Pi_i^{\text{win}}(z) dF_i(z),$$

and the validator's expected payoff is

$$R_{\text{val}}^{\text{FPA}} = \mathbb{E}[U_{\text{val}}^{\text{FPA}}] = V_o \sum_{i=1}^n \alpha_i \int \beta_{i,\tau(o)}(z) \Pi_i^{\text{win}}(z) dF_i(z).$$

The economic metrics $\mathcal{J}_{\text{Sh}}^{\text{FPA}}$, $\text{ER-HHI}^{\text{FPA}}$, and $N_{\text{eff}}^{\text{FPA}}$ are obtained by substituting R_i^{FPA} into the definitions in Sect. 2.2. To obtain sharper comparisons, we use the proportional-bidding benchmark $\beta_{i,\tau(o)}(z) = \rho z$ for all searchers and opportunity classes, where $\rho \in (0, 1)$. Then the FPA ranks submissions by certified normalized net value, while the winner keeps fraction $1 - \rho$ of its certified net value. In this benchmark,

$$R_i^{\text{FPA}} = V_o(1 - \rho)\Theta_i, \quad \Theta_i = \alpha_i \int z \prod_{j \neq i} (1 - \alpha_j + \alpha_j F_j(z)) dF_i(z).$$

Thus all searcher reward shares, and hence $\text{ER-HHI}^{\text{FPA}}$ and $N_{\text{eff}}^{\text{FPA}}$, are determined by the rank-dominance scores $(\Theta_i)_{i=1}^n$.

► **Proposition 8** (Symmetric setting). *Suppose $\alpha_i = \alpha$ and $F_i = F$ for all i . Under proportional bidding, $R_i^{\text{FPA}} = R_j^{\text{FPA}}$ for all i, j . Consequently,*

$$\text{ER-HHI}^{\text{FPA}} = \frac{1}{n}, \quad N_{\text{eff}}^{\text{FPA}} = n.$$

Moreover, since the Shapley contributions are also symmetric, $\mathcal{J}_{\text{Sh}}^{\text{FPA}} = 1$.

Proof. Symmetry gives $\Theta_i = \Theta_j$ for all i, j , hence equal expected rewards and equal reward shares. The HHI and effective-number claims follow immediately. The game $W^{\text{VS}}(\cdot)$ is symmetric across searchers, so $\phi_i^{\text{VS}} = \phi_j^{\text{VS}}$ for all i, j , and therefore $R_i^{\text{FPA}}/\phi_i^{\text{VS}}$ is constant. ◀

The symmetric case provides a natural benchmark in which the FPA is decentralized and contribution-fair. Once searchers differ in coverage or execution quality, rewards become concentrated around the searchers most likely to win the rank competition.

► **Proposition 9** (Dominant execution-quality class with heterogeneous coverage). *Suppose $\alpha_i = \alpha_H \in (0, 1)$ for every $i \leq m$ and $\alpha_j = \alpha_L \in (0, 1)$ for every $j > m$, with $\alpha_H \geq \alpha_L$. For a distribution F , let $\text{supp}(F)$ denote the set of normalized net values attainable with positive probability. Under proportional bidding, suppose there exist $m \in \{1, \dots, n\}$ and $\underline{z}_H, \bar{z}_L \in (0, 1)$, with $\underline{z}_H > \bar{z}_L$, such that $F_i = F_H$ and $\text{supp}(F_H) \subseteq [\underline{z}_H, 1]$ for every $i \leq m$, while $\text{supp}(F_j) \subseteq [0, \bar{z}_L]$ for every $j > m$. Then,*

$$N_{\text{eff}}^{\text{FPA}} \leq m \left(\frac{(1 - (1 - \alpha_H)^m) \underline{z}_H + (1 - \alpha_H)^m (n - m) \alpha_L \bar{z}_L}{(1 - (1 - \alpha_H)^m) \underline{z}_H} \right)^2.$$

Consequently, as $\alpha_H \rightarrow 1$, $N_{\text{eff}}^{\text{FPA}} \rightarrow m$ and $\mathcal{J}_{\text{Sh}}^{\text{FPA}} \rightarrow m/n$.

Proof. Let $p_H = 1 - (1 - \alpha_H)^m$ be the probability that at least one high-quality searcher is active. Since $\underline{z}_H > \bar{z}_L$, whenever at least one high-quality searcher is active, the FPA winner is in the high-quality class. Hence the total expected reward of the high-quality class is at least $V_o(1 - \rho)p_H \underline{z}_H$. A low-quality searcher can win only when no high-quality searcher is active, which occurs with probability $(1 - \alpha_H)^m$. Since each low-quality searcher is active

with probability α_L and has normalized net value at most $\bar{z}_L$, the total expected reward of the low-quality class is at most $V_o(1-\rho)(1-\alpha_H)^m(n-m)\alpha_L\bar{z}_L$. Therefore, if s_H denotes the total reward share of the high-quality class, then

$$s_H \geq \frac{p_H \bar{z}_H}{p_H \bar{z}_H + (1-\alpha_H)^m(n-m)\alpha_L \bar{z}_L}.$$

The m high-quality searchers are symmetric, since they have the same coverage probability and the same distribution F_H . Thus they receive equal expected rewards, so each has reward share s_H/m . Hence $\text{ER-HHI}^{\text{FPA}} = \sum_i (s_i^{\text{FPA}})^2 \geq m(s_H/m)^2 = s_H^2/m$. Since $N_{\text{eff}}^{\text{FPA}} = 1/\text{ER-HHI}^{\text{FPA}}$, we obtain

$$N_{\text{eff}}^{\text{FPA}} \leq m \left(\frac{p_H \bar{z}_H + (1-\alpha_H)^m(n-m)\alpha_L \bar{z}_L}{p_H \bar{z}_H} \right)^2.$$

As $\alpha_H \rightarrow 1$, we have $p_H \rightarrow 1$ and $(1-\alpha_H)^m \rightarrow 0$, so the upper bound converges to m . Low-quality reward shares converge to zero, while symmetry implies that the m high-quality searchers have equal limiting reward shares $1/m$. Hence $N_{\text{eff}}^{\text{FPA}} \rightarrow m$.

It remains to prove the Jain-index limit. For any low-quality searcher $j > m$, as $\alpha_H \rightarrow 1$,

$$R_j^{\text{FPA}} \leq (1-\rho)V_o\alpha_L(1-\alpha_H)^m\bar{z}_L \rightarrow 0.$$

On the other hand, using only the Shapley marginal contribution from the predecessor coalition $\{\text{val}\}$,

$$\phi_j^{\text{VS}} \geq \frac{1}{n(n+1)} \left(W^{\text{VS}}(\{\text{val}, j\}) - W^{\text{VS}}(\{\text{val}\}) \right) = \frac{V_o\alpha_L \mathbb{E}_{F_j}[z]}{n(n+1)} > 0.$$


Hence $R_j^{\text{FPA}}/\phi_j^{\text{VS}} \rightarrow 0$ for every low-quality searcher $j > m$. For the m high-quality searchers, symmetry implies equal expected rewards and equal validator-searcher Shapley contributions; their common reward-to-contribution ratio has a positive limit. Therefore the Jain vector converges to one with m equal positive components and $n-m$ zero components. Its Jain index is $\frac{(mc)^2}{nmc^2} = \frac{m}{n}$, for any $c > 0$. Thus $\mathcal{J}_{\text{Sh}}^{\text{FPA}} \rightarrow m/n$. ◀

The proposition shows that FPA reward decentralization is governed by the size of the dominant class, not by the total number of searchers. When α_H is high, at least one dominant searcher is almost always active, so lower-quality searchers almost never win. Thus, the FPA concentrates rewards on the top class: on m searchers in general, and on a single effective monopolist when $m = 1$.

► **Corollary 10** (Single dominant searcher). *Under the assumptions of Proposition 9 with $m = 1$, if searcher 1 is the unique dominant searcher, then, as $\alpha_H \rightarrow 1$, $N_{\text{eff}}^{\text{FPA}} \rightarrow 1$, $\mathcal{J}_{\text{Sh}}^{\text{FPA}} \rightarrow \frac{1}{n}$.*

Figure 1 illustrates how increasing dominance affects reward concentration under the FPA. The horizontal axis varies the participation probability α_H of the high-quality class, while different curves correspond to different numbers m_H of elite searchers. As α_H increases, FPA rewards become increasingly concentrated within the dominant class: contribution-adjusted fairness declines toward m_H/n , and the effective number of rewarded searchers declines toward m_H .

The results show that strong asymmetry can substantially reduce the economic decentralization of the FPA. With coverage heterogeneity, FPA rewards concentrate around searchers who access opportunities more often. With execution-quality heterogeneity, concentration is sharper: rewards are captured by the dominant execution class, and when this class contains only a few searchers, the FPA converges toward monopolistic concentration. Thus, the winner-take-all structure of the FPA can amplify centralization in realistic MEV environments where both opportunity access and execution quality are unequal.



■ **Figure 1** Effect of dominant-class participation on FPA fairness and reward decentralization.

4 The Entry-Filtered Shapley-Capped Auction

This section introduces a Shapley-based searcher auction mechanism designed to distribute MEV rewards more fairly and broadly among competing searchers. The goal is to align rewards with searchers' marginal contributions, reduce winner-take-all concentration, and maintain security against copied-code Sybil deviations and validator–searcher coalition deviations in a permissionless MEV environment.

4.1 Design Rationale

Our primary goal is to design a mechanism that, unlike a winner-take-all auction, distributes the searcher-side MEV reward among searchers that contribute to identifying or executing the same opportunity, rather than assigning all surplus to the top-ranked submission. The immediate difficulty is Sybil manipulation. Once rewards are distributed across multiple submissions, a searcher may submit copied or degraded versions of the same execution strategy to inflate its apparent contribution and capture a larger share. A natural solution is a Tullock-style contest [23], where proportional allocation with costly bids can be Sybil-proof because multiple bids can be aggregated into a single equivalent bid [23, 6]. However, this security relies on an all-pay structure: participants may obtain negative utility and must be willing to lose their bids. In an on-chain searcher auction, this would require deposits or enforceable losses, making the mechanism difficult to implement under current MEV infrastructure.

The broader obstruction is formal. Pan et al. show that any non-wasteful, symmetric, incentive-compatible, and ex-post Sybil-proof direct mechanism must reduce to a second-price auction with symmetric tie-breaking; equivalently, strong Sybil-proofness rules out nontrivial distributional objectives under those axioms [16]. Their Bayesian relaxation shows that this impossibility can be softened: an increasing-ticket-price mechanism can be Bayesian Sybil-proof by allocating the object through a lottery among threshold-clearing bidders whenever at least two such bidders exist, and otherwise falling back to a second-price auction [16]. Nevertheless, this construction is not well matched to MEV searcher competition. Its Bayesian Sybil deterrence relies on an ex-ante ticket threshold dictated by an upper-tail cutoff: the point beyond which the bid distribution's tail becomes no heavier than a uniform benchmark. While this condition is easily satisfied by distributions with flat or increasing right tails, it fails under the concentrated, declining upper tails typical of MEV signal models, such as the truncated log-normal distribution [1, 28]. In these cases, the required cutoff is

pushed to the maximum feasible valuation. Consequently, the ticket price also reaches this maximum, the lottery branch occurs with probability zero under continuous bids, and the mechanism entirely collapses back into a standard, winner-take-all second-price auction [16]. Appendix B of the full version [20] discusses Pan et al.'s characterization and the reduction of their mechanism to the second-price auction under our truncated-lognormal specification.

We therefore use the Shapley value as the economic allocation target, since it rewards searchers according to their marginal contribution to the opportunity. Shapley sharing alone, however, prevents neither copied-code Sybil manipulation nor validator–searcher collusion in a permissionless submission environment. The mechanism therefore combines Shapley-based rewards with residual caps and burn functions. The searcher-side burn is used to eliminate the expected gain from copied-code Sybil deviations, while the validator-side burn is calibrated to rule out profitable validator–searcher coalition deviations.

4.2 Entry-Filtered Shapley-Capped Auction Mechanism

The entry-filtered Shapley-capped auction mechanism is denoted by SCA_θ , where θ is the entry-filter threshold. Unlike the FPA, SCA_θ does not use transfer offers; its input is the set of valid submitted execution codes and their certified normalized net values. The mechanism first applies an entry filter: a submitted code is admitted only if its certified normalized net value is at least a protocol threshold $\theta \in [0, 1]$. The threshold is chosen ex ante to control the expected number of admitted submissions. For opportunity class τ , one may choose $\theta = \theta_\tau$ so that $\mathbb{E}[N_o^\theta] = \sum_i \alpha_{i,\tau} \Pr_{F_{i,\tau}}(z_{i,o}^* \geq \theta) \leq \bar{N}$, where $\bar{N}$ is the target upper bound on the expected admitted population.

Fix an opportunity o . Submissions with $z_b < \theta$ are rejected and receive zero. Let $\mathcal{B}^\theta(o) = \{b \in \mathcal{B}(o) : z_b \geq \theta\}$ be the set of admitted submissions, and let $t_\theta := |\mathcal{B}^\theta(o)|$. Their normalized and unnormalized certified net-value profiles are denoted by $z_o^\theta := (z_b)_{b \in \mathcal{B}^\theta(o)}$ and $\tilde{z}_o^\theta := (\tilde{z}_b)_{b \in \mathcal{B}^\theta(o)}$, respectively. If $t_\theta = 0$, no submission is executed; set $z_{(1)}^\theta := 0$, all searcher and validator payoffs to zero, and $B_o^{\text{SCA}_\theta} := 0$. Thus, θ serves as a minimum certified-value requirement for execution.

Otherwise, order admitted submissions by $z_{(1)}^\theta \geq \dots \geq z_{(t_\theta)}^\theta$, where $b_{(r)}^\theta$ is the rank- r admitted submission. Set $z_{(t_\theta+1)}^\theta = 0$ and $\Delta_k^\theta = z_{(k)}^\theta - z_{(k+1)}^\theta$. The executed submission is $\text{Win}_o^{\text{SCA}_\theta}(z_o^\theta) = b_{(1)}^\theta$, with ties broken by mempool arrival order.⁶

For any coalition $A \subseteq \mathcal{B}^\theta(o) \cup \{\text{val}\}$, its realized value under SCA_θ is

$$W_o^\theta(A) = \begin{cases} V_o \max_{b \in A \setminus \{\text{val}\}} z_b, & \text{val} \in A, A \setminus \{\text{val}\} \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

Note that the validator is necessary for a coalition to generate value. For the admitted submissions $\mathcal{B}^\theta(o)$, the Shapley values of the validator and the rank- r admitted submission are

$$\Phi_{\text{val}}^\theta(z_o^\theta) = V_o \sum_{k=1}^{t_\theta} \frac{k \Delta_k^\theta}{k+1}, \quad \Phi_{(r)}^\theta(z_o^\theta) = V_o \sum_{k=r}^{t_\theta} \frac{\Delta_k^\theta}{k(k+1)}.$$

Hence $\Phi_{\text{val}}^\theta(z_o^\theta) + \sum_{r=1}^{t_\theta} \Phi_{(r)}^\theta(z_o^\theta) = V_o z_{(1)}^\theta$.

⁶ If two submissions from different identities have the same certified value, the earlier submission in the encrypted mempool is ranked first. The ordering is fixed before decryption or reveal, so copied submissions cannot improve tie-breaking after observing rivals.

The mechanism includes two branches: the decentralized Shapley-sharing branch and the capped fallback branch. The mechanism places a cap $H \geq 2$ on the number of admitted submissions that can be handled in the decentralized Shapley-sharing branch. This cap is chosen relative to the admitted population, so H should be moderately larger than $\mathbb{E}[N_o^\theta]$. Given retention parameters $\kappa_S, \kappa_V \geq 0$, define $\eta_{t_\theta}^V = e^{-\kappa_V(\min\{t_\theta, H\}-1)}$ and $\eta_{t_\theta}^S = e^{-\kappa_S(\min\{t_\theta, H\}-1)}$.

- If $t_\theta \leq H$, all admitted submissions are rewarded according to retained Shapley values:

$$U_{\text{val}}^{\text{SCA}_\theta}(z_o^\theta) = \eta_{t_\theta}^V \Phi_{\text{val}}^\theta(z_o^\theta), \quad P_{(r)}^{\text{SCA}_\theta}(z_o^\theta) = \eta_{t_\theta}^S \Phi_{(r)}^\theta(z_o^\theta).$$

- If $t_\theta > H$, the number of admitted submissions exceeds the cap, and the mechanism switches to the capped fallback branch:

$$U_{\text{val}}^{\text{SCA}_\theta}(z_o^\theta) = \frac{H\eta_H^V}{H+1} V_o z_{(1)}^\theta, \quad P_{(1)}^{\text{SCA}_\theta}(z_o^\theta) = \frac{\eta_H^S}{H(t_\theta+1)} V_o z_{(1)}^\theta,$$

and $P_{(r)}^{\text{SCA}_\theta}(z_o^\theta) = 0$ for all $r \geq 2$.

Thus, when the number of admitted submissions is at most the cap, the mechanism uses Shapley sharing across admitted submissions. When the number of admitted submissions exceeds the cap, it switches to a fallback rule that pays only the top admitted submission on the searcher side while retaining a cap-dependent burn. For each real searcher i , its total payoff is the sum of the payments assigned to its admitted submissions, $U_i^{\text{SCA}_\theta}(z_o^\theta) = \sum_{b \in \mathcal{B}_i(o) \cap \mathcal{B}^\theta(o)} P_b^{\text{SCA}_\theta}(z_o^\theta)$. Any surplus remaining after paying the validator and all searchers is burned, $B_o^{\text{SCA}_\theta}(z_o^\theta) = V_o z_{(1)}^\theta - U_{\text{val}}^{\text{SCA}_\theta}(z_o^\theta) - \sum_i U_i^{\text{SCA}_\theta}(z_o^\theta)$.

- **Proposition 11** (Budget feasibility). *For every opportunity and every valid submitted profile,*

$$U_{\text{val}}^{\text{SCA}_\theta}(z_o^\theta) + \sum_{i \in \mathcal{S}_{\text{act}}(o)} U_i^{\text{SCA}_\theta}(z_o^\theta) + B_o^{\text{SCA}_\theta}(z_o^\theta) = V_o z_{(1)}^\theta \leq V_o, \quad B_o^{\text{SCA}_\theta}(z_o^\theta) \geq 0.$$

Proof. If $t_\theta = 0$, the claim is immediate. If $1 \leq t_\theta \leq H$, Shapley budget balance gives $\Phi_{\text{val}}^\theta + \sum_{r=1}^{t_\theta} \Phi_{(r)}^\theta = V_o z_{(1)}^\theta$, and $\eta_{t_\theta}^S, \eta_{t_\theta}^V \leq 1$ imply total payments at most $V_o z_{(1)}^\theta$. If $t_\theta > H$, total payments are at most

$$\left(\frac{H}{H+1} + \frac{1}{H(t_\theta+1)} \right) V_o z_{(1)}^\theta \leq \left(\frac{H}{H+1} + \frac{1}{H(H+2)} \right) V_o z_{(1)}^\theta < V_o z_{(1)}^\theta.$$

Thus the residual defining $B_o^{\text{SCA}_\theta}$ is nonnegative, and the displayed identity follows by definition of the burn. ◀

The formal Bayesian security analysis of the entry-filtered SCA is given in Appendix C of the full version [20]. It shows that it is sufficient to consider copy sizes $q = 2, \dots, H+1$, while optimizing over all feasible copied-value profiles. The parameters (κ_S, κ_V) are calibrated so that no such deviation increases the Bayesian interim payoff of either the searcher alone or the searcher-validator coalition. The calibration creates a security-efficiency tradeoff. Increasing κ_S strengthens searcher-side Sybil deterrence but also increases burn. Coalition resistance, by contrast, generally requires κ_V to lie within a feasible interval: identity-expansion deviations may impose a lower bound, while withdrawal deviations may impose an upper bound. When this security region is nonempty, (κ_S^*, κ_V^*) can be chosen to minimize expected burn subject to both Bayesian security requirements. The calibration reported in Appendix D of the full version [20] gives priority to minimizing searcher-side burn.

4.3 Economic Properties of the Entry-Filtered SCA

We analyze the honest one-submission benchmark under the contribution-adjusted metrics defined in Sect. 2.2. The entry filter affects rewards by excluding submissions below the threshold from the reward-sharing game. Thus, for a searcher i , only states in which i is active and satisfies $z_{i,o}^* \geq \theta$ can generate a positive SCA_θ payment.

Let $E_{i,o}^\theta := D_{i,o} \mathbf{1}\{z_{i,o}^* \geq \theta\}$ indicate that searcher i is active and admitted, and let $N_o^\theta := \sum_i E_{i,o}^\theta$ be the admitted population. The cap H is chosen relative to N_o^θ , not the raw active count N_o . In particular, the threshold θ and the cap H should be calibrated jointly so that $\mathbb{E}[N_o^\theta]$ is controlled and H is moderately larger than the typical admitted population.

For a realized profile, let $\Phi_i^\theta(z_o^\theta)$ denote searcher i 's realized Shapley value in the admitted max game, with $\Phi_i^\theta(z_o^\theta) = 0$ when i is not admitted. Let ϕ_i^{VS} denote searcher i 's ex-ante Shapley contribution in the unfiltered validator–searcher game defined in Section 2.2.1. Define the admitted contribution coverage ratio $a_{i,\theta} := \frac{\mathbb{E}[\mathbf{1}\{E_{i,o}^\theta=1, N_o^\theta \leq H\} \Phi_i^\theta(z_o^\theta)]}{\phi_i^{\text{VS}}}$, and the admitted contribution exposure ratio $b_{i,\theta} := \frac{\mathbb{E}[\mathbf{1}\{E_{i,o}^\theta=1\} \Phi_i^\theta(z_o^\theta)]}{\phi_i^{\text{VS}}}$. The ratio $a_{i,\theta}$ is the share of searcher i 's contribution denominator that is admitted and paid in the decentralized branch. The ratio $b_{i,\theta}$ is an upper exposure bound for the contribution available to i after thresholding. These ratios are defined for all searchers and are normalized by ϕ_i^{VS} . Unlike in the unfiltered case, $b_{i,\theta}$ need not be at most one, because removing ineligible rivals can increase an admitted searcher's Shapley value relative to its full-profile Shapley value.

Define the global lower and upper reward-ratio bounds $\lambda_\theta := e^{-\kappa_S(H-1)} \min_i a_{i,\theta}$ and $\Lambda_\theta := \max_i b_{i,\theta}$. Thus λ_θ and Λ_θ are bounds on the reward-to-contribution ratios $R_i^{\text{SCA}_\theta} / \phi_i^{\text{VS}}$ across all searchers.

► **Proposition 12** (Entry-filtered SCA reward bounds). *Under honest one-submission play, assume the entry threshold is chosen so that the admitted population is bounded by $t_{\max}$, and suppose the fallback searcher retention satisfies $\eta_H^S \leq H/t_{\max}$. Then, for every searcher i ,*

$$\lambda_\theta \leq \frac{R_i^{\text{SCA}_\theta}}{\phi_i^{\text{VS}}} \leq \Lambda_\theta.$$

Proof. On the decentralized branch $N_o^\theta \leq H$, the entry-filtered SCA pays $\eta_{N_o^\theta}^S \Phi_i^\theta(z_o^\theta)$. Since $\eta_{N_o^\theta}^S = e^{-\kappa_S(N_o^\theta-1)} \geq e^{-\kappa_S(H-1)}$, we have

$$R_i^{\text{SCA}_\theta} \geq e^{-\kappa_S(H-1)} \mathbb{E}[\mathbf{1}\{E_{i,o}^\theta = 1, N_o^\theta \leq H\} \Phi_i^\theta(z_o^\theta)] = e^{-\kappa_S(H-1)} a_{i,\theta} \phi_i^{\text{VS}}.$$

Taking the minimum over i gives the lower bound.

For the upper bound, $U_i^{\text{SCA}_\theta} \leq \Phi_i^\theta(z_o^\theta)$ pointwise. This is immediate in the decentralized branch. In fallback, non-winners receive zero, and the winner receives

$$\frac{\eta_H^S}{H(t_\theta + 1)} V_o z_{(1)}^\theta \leq \frac{1}{t_\theta(t_\theta + 1)} V_o z_{(1)}^\theta \leq \Phi_{(1)}^\theta(z_o^\theta),$$

where the first inequality uses $t_\theta \leq t_{\max}$ and $\eta_H^S \leq H/t_{\max}$. Hence $R_i^{\text{SCA}_\theta} \leq \mathbb{E}[\mathbf{1}\{E_{i,o}^\theta = 1\} \Phi_i^\theta(z_o^\theta)] = b_{i,\theta} \phi_i^{\text{VS}}$. Taking the maximum over i gives the upper bound. ◀

Define the contribution-side effective number by $N_{\text{eff}}^{\phi, \text{VS}} := \frac{(\sum_{i=1}^n \phi_i^{\text{VS}})^2}{\sum_{i=1}^n (\phi_i^{\text{VS}})^2}$, which measures how many economically meaningful searchers are represented by the validator–searcher Shapley contribution vector $(\phi_i^{\text{VS}})_{i=1}^n$.

■ **Table 1** Asymmetric searcher populations. Scenarios become more centralized from S1 to S4.

Scenario	n	$\mathbb{E}[N_o]$	Class structure	$m(\alpha, \tilde{c}, [\underline{c}, \bar{c}])$
S1: broad heterogeneity	20	9.60	4(0.60, 0.15, [0.10, 0.25]), 8(0.50, 0.30, [0.20, 0.45]), 8(0.40, 0.52, [0.35, 0.70])	
S2: aligned coverage–quality advantage	20	10.10	5(0.85, 0.14, [0.10, 0.25]), 7(0.55, 0.30, [0.20, 0.45]), 8(0.25, 0.55, [0.40, 0.70])	
S3: dominant high-quality class	20	8.00	3(0.95, 0.10, [0.05, 0.20]), 7(0.45, 0.30, [0.18, 0.45]), 10(0.20, 0.52, [0.35, 0.70])	
S4: single elite searcher	20	9.43	1(0.98, 0.10, [0.05, 0.20]), 9(0.55, 0.30, [0.18, 0.45]), 10(0.35, 0.52, [0.35, 0.70])	

► **Theorem 13** (Economic performance of the entry-filtered SCA). *Under the assumptions of Proposition 12, and assuming that $\sum_{i=1}^n R_i^{\text{SCA}_\theta} > 0$,*

$$\mathcal{J}_{\text{Sh}}^{\text{SCA}_\theta} \geq \frac{4\lambda_\theta \Lambda_\theta}{(\lambda_\theta + \Lambda_\theta)^2}, \quad N_{\text{eff}}^{\text{SCA}_\theta} \geq \left(\frac{\lambda_\theta}{\Lambda_\theta} \right)^2 N_{\text{eff}}^{\phi, \text{VS}}.$$

Proof. Let $q_i := R_i^{\text{SCA}_\theta} / \phi_i^{\text{VS}}$. By Proposition 12, $q_i \in [\lambda_\theta, \Lambda_\theta]$ for every i . Therefore, for any given i , $q_i^2 \leq (\lambda_\theta + \Lambda_\theta)q_i - \lambda_\theta \Lambda_\theta$, because $(q_i - \lambda_\theta)(q_i - \Lambda_\theta) \leq 0$. Summing over i gives $\sum_i q_i^2 \leq (\lambda_\theta + \Lambda_\theta) \sum_i q_i - n\lambda_\theta \Lambda_\theta$. For fixed bounds $[\lambda_\theta, \Lambda_\theta]$, the Jain index is minimized when the ratios are split between the two endpoints. Equivalently, the worst case satisfies $\mathcal{J}(q) = \frac{(\sum_i q_i)^2}{n \sum_i q_i^2} \geq \frac{4\lambda_\theta \Lambda_\theta}{(\lambda_\theta + \Lambda_\theta)^2}$. Hence $\mathcal{J}_{\text{Sh}}^{\text{SCA}_\theta} \geq \frac{4\lambda_\theta \Lambda_\theta}{(\lambda_\theta + \Lambda_\theta)^2}$. For the effective number, using $R_i^{\text{SCA}_\theta} = q_i \phi_i^{\text{VS}}$, $N_{\text{eff}}^{\text{SCA}_\theta} = \frac{(\sum_i q_i \phi_i^{\text{VS}})^2}{\sum_i (q_i \phi_i^{\text{VS}})^2} \geq \frac{\lambda_\theta^2 (\sum_i \phi_i^{\text{VS}})^2}{\Lambda_\theta^2 \sum_i (\phi_i^{\text{VS}})^2} = \left(\frac{\lambda_\theta}{\Lambda_\theta} \right)^2 N_{\text{eff}}^{\phi, \text{VS}}$. ◀

The entry threshold creates a security–coverage tradeoff. Increasing θ reduces the expected admitted population, $\mathbb{E}[N_o^\theta] = \sum_i \alpha_i \Pr(z_{i,o}^* \geq \theta)$, which weakens copied-code Sybil incentives and can reduce the amount of searcher-side burn required for security. However, increasing θ can reduce $a_{i,\theta}$, because some genuine low-value contributions become ineligible for rewards. Thus the threshold changes how much of each searcher’s contribution can receive a positive mechanism payment. For $\theta > 0$, the mechanism remains Shapley-based among admitted submissions, but avoids paying the entire long tail of low-certified-value submissions.

4.4 Asymmetric Entry-Filtered Benchmark

To evaluate the mechanisms under realistic searcher heterogeneity, we model execution costs by a truncated lognormal distribution. Conditional on participating in opportunity o , searcher i ’s normalized execution cost is

$$C_i \mid D_{i,o} = 1 \sim \text{TruncLogNormal}(\mu_i, \sigma^2; [\underline{c}_i, \bar{c}_i]), \quad z_i^* = 1 - C_i.$$

Table 1 introduced the four asymmetric searcher populations. Searchers differ in both participation probability and cost efficiency. Lower median cost corresponds to a higher certified normalized net value. Each benchmark scenario has three searcher classes. A class entry $m(\alpha, \tilde{c}, [\underline{c}, \bar{c}])$ means m searchers with participation probability α , median cost $\tilde{c}$, and truncated-lognormal support $[\underline{c}, \bar{c}]$. The four scenarios are ordered from less concentrated to more centralized competition.

Table 2 compares the FPA with entry-filtered SCA_θ across different admission targets. For the entry-filtered SCA_θ , we report three threshold levels. The threshold θ is selected so that the expected admitted population satisfies $\mathbb{E}[N_o^\theta] \in \{8, 6, 4\}$. The cap is then chosen relative to the admitted population, using $H = \lceil \mathbb{E}[N_o^\theta] \rceil + 3$. For each θ , κ_S^* is the least searcher-side security parameter satisfying the Bayesian copied-code Sybil-resistance constraints over every realized value z , every copy size $q = 2, \dots, H + 1$, and every feasible copied-value profile $\mathbf{z}^c \in \mathcal{Z}_q^c(z)$. Notably, across all twelve benchmark calibrations, the binding deviation is the

■ **Table 2** FPA versus entry-filtered SCA. Within each scenario, the SCA rows show increasingly selective admission rules; the FPA row is reported once as the winner-take-all benchmark.

Scenario	Mechanism	$\mathbb{E}[N_o^\theta]$	θ	H	κ_S^*	$\mathcal{J}_{Sh}$	N_{eff}	Note
S1	SCA $_\theta$	8	0.495	11	0.371	0.948	11.10	broad admission
	SCA $_\theta$	6	0.620	9	0.293	0.706	9.67	intermediate admission
	SCA $_\theta$	4	0.719	7	0.165	0.574	7.19	selective admission
	FPA	—	—	—	—	0.244	4.26	winner-take-all benchmark
S2	SCA $_\theta$	8	0.580	11	0.415	0.661	8.72	broad admission
	SCA $_\theta$	6	0.709	9	0.356	0.536	6.96	intermediate admission
	SCA $_\theta$	4	0.790	7	0.265	0.288	5.21	selective admission
	FPA	—	—	—	—	0.251	5.00	winner-take-all benchmark
S3	SCA $_\theta$	8	0.300	11	0.406	0.996	6.07	broad admission
	SCA $_\theta$	6	0.598	9	0.355	0.693	5.45	intermediate admission
	SCA $_\theta$	4	0.734	7	0.284	0.413	4.01	selective admission
	FPA	—	—	—	—	0.150	3.00	winner-take-all benchmark
S4	SCA $_\theta$	8	0.467	11	0.388	0.970	3.84	broad admission
	SCA $_\theta$	6	0.597	9	0.322	0.726	3.78	intermediate admission
	SCA $_\theta$	4	0.683	7	0.201	0.500	3.23	selective admission
	FPA	—	—	—	—	0.075	1.04	winner-take-all benchmark

two-copy attack at the admission boundary, with $q = 2$, $z = \theta$, and $\mathbf{z}^c = (\theta, \theta)$. Hence, in these benchmark designs, the numerically worst deviation is a two-copy attack, although this need not hold in general.

Table 2 supports the main decentralization claim. As competition becomes more concentrated, the FPA reward allocation collapses: N_{eff}^{FPA} falls from 4.26 in S1 to 1.04 in S4, and $\mathcal{J}_{Sh}^{FPA}$ falls from 0.244 to 0.075. By contrast, the entry-filtered SCA $_\theta$ consistently preserves a broader reward base. In the most centralized case, S4, the FPA is effectively controlled by one elite searcher, while SCA $_\theta$ with $\mathbb{E}[N_o^\theta] = 6$ achieves $N_{eff}^{SCA_\theta} = 3.78$ and $\mathcal{J}_{Sh}^{SCA_\theta} = 0.726$.

The other side of the result is the security–coverage tradeoff. Lower thresholds admit more searchers and improve decentralization, but require larger κ_S^* . Higher thresholds reduce κ_S^* , but exclude more low-value submissions from reward sharing. Thus, relative to the FPA, the entry-filtered SCA $_\theta$ offers substantially better reward decentralization, with the threshold θ controlling the cost of Bayesian Sybil resistance.

Additional numerical results on increasing elite execution advantage, together with validator-side security calibration, are reported in Appendix D of the full version [20].

5 Empirical Benchmarks

We evaluate the performance of searchers by analyzing historical MEV transactions on Ethereum⁷. In practice, a searcher entity may deploy multiple contracts to identify and capture MEV opportunities. However, establishing correlations between different contracts controlled by the same searcher entity is outside the scope of this paper. For the purposes of this section, the terms searcher and contract are treated as interchangeable.

Dataset. We use a dataset [10] developed by Luo et al., which contains 577,264 arbitrage transactions from block 16308203 to 18908892. In total, there are 855 contracts generating a total revenue of 37.79 million dollars and a total profit of 8.88 million dollars.

⁷ Our implementation is available at https://gitlab.esat.kuleuven.be/mev/searcher_game.

■ **Table 3** Summary of Searcher Strategies.

Strategy	#Contracts	#Tx	Revenue Share	Profit Share	Avg. Profit / Tx
C1	74	3,323	4.35%	14.02%	\$374.74
C2	88	510,620	76.29%	73.57%	\$12.79
C3	693	63,321	19.36%	12.40%	\$17.39

Searcher strategies. Arbitrage searchers implement one or more trading strategies to capture MEV opportunities. These strategies often involve constructing cyclic transactions across decentralized exchanges, in which two or more tokens are swapped within a single atomic execution. When multiple searchers identify the same opportunity, they compete for inclusion by submitting higher bids.

We identify two commonly implemented arbitrage strategies. One prioritizes selectivity, targeting scarce opportunities leading to high profit margins. The other prioritizes volume, executing a large number of transactions with relatively low profit per transaction.

The contracts in the dataset are therefore classified into three strategy profiles: (C1) margin-oriented, which contains contracts in the top 10% by average profit per transaction; (C2) volume-oriented, which are those in the top 10% by total transaction count; and (C3) mixed-profile contracts, whose observed performance does not show a dominant margin-oriented or volume-oriented profile.


Among the classified contracts, ten satisfy both the margin-oriented and volume-oriented criteria. We assign these overlapping contracts to the volume-oriented group, since high transaction volume more directly captures repeated execution behavior. In addition, two contracts `0x03c6...67cc` and `0xeeaa...3376` are excluded from margin-oriented group after identifying volume-oriented traits in their behavior.

A summary of the classified contracts is shown in Table 3. The profit retained by a searcher varies substantially depending on the type of opportunity covered, with the average profit per transaction in the C1 group being approximately 30 times that of the C2 group.

Execution efficiency across strategies. Within each strategy group, contracts are further classified into three tiers based on their contribution to the group’s total realized profit. Contracts are first ranked in descending order by their total profit within the strategy. We then compute the cumulative share of positive realized profit accounted for by this ranking. Elite contracts are defined as the smallest set of top-ranked contracts whose cumulative positive profit reaches 60% of the strategy’s total positive profit. Middle contracts are the next set of ranked contracts needed to raise the cumulative positive-profit share from 60% to 90%. Tail contracts are all remaining contracts, including contracts with zero or negative realized profit. For each tier, we show the median profit-to-revenue ratio, which measures the share of generated revenue retained by searchers. A higher profit-to-revenue ratio indicates that searchers retain more value after paying the costs required to win inclusion.

The performance of each tier is summarized in Table 4. A strong concentration is observed within each strategy group, with a small fraction of top-tier contracts accounting for over 60% of total realized profit. A stark contrast emerges between the two strategies: the bottom tier of C1 retains a higher median profit-to-revenue ratio than the top tier of C2, 35.66% and 17.67% respectively. This hints that C1 faces considerably lower competition intensity.

Figure 2 illustrates the transaction distribution across profit-to-revenue ratio ranges. The profit-to-revenue ratio measures the fraction of arbitrage revenue that remains as realized profit after fees and bids, and thus provides an observable proxy for the retained margin of included transactions. Within each strategy, elite contracts generally exhibit higher



■ **Figure 2** Transaction distribution over profit-to-revenue ratio bins, grouped by profit tier. The x-axis denotes profit-to-revenue ratio bins, while the y-axis reports the percentage share of transactions.

median profit-to-revenue ratios than middle and tail contracts, indicating that the contracts accounting for the largest share of realized profit also tend to retain a larger fraction of revenue on a per-transaction basis. Across strategies, C1 displays substantially higher retained margins than C2 and C3. This pattern is consistent with the interpretation that C1 searchers target arbitrage opportunities with lower relative execution or bidding costs, whereas C2 searchers operate in a more cost-intensive environment associated with high-frequency opportunity capture. Since the dataset contains only successful transactions and does not observe failed bids, the profit-to-revenue ratio should be interpreted as an ex-post measure of realized margin rather than a direct measure of auction-level bidding power.

Empirical calibration and limitations. Table 5 reports the truncated-lognormal cost distributions extracted from the Ethereum arbitrage data. The opportunity types C_1 and C_2 correspond to different empirical arbitrage environments: C_1 is margin-oriented, while C_2 is volume-oriented. Within each opportunity type, contracts are grouped into elite, middle, and tail tiers. A class entry $m(\alpha, \tilde{c}, [\underline{c}, \bar{c}])$ represents m contracts with relative coverage α , median normalized cost $\tilde{c}$, and truncated-lognormal support $[\underline{c}, \bar{c}]$.

■ **Table 4** Summary of contract tiers within each strategy. (M) denotes median.

Tier	#Contracts	#Tx	Profit/Tx (M)	Profit Share	Profit/Revenue (M)
C1-elite	9	1,950	\$561.76	62.08%	71.15%
C1-middle	17	988	\$491.89	28.15%	40.28%
C1-tail	48	385	\$274.35	9.77%	35.66%
C2-elite	11	53,719	\$128.11	61.65%	17.67%
C2-middle	19	99,565	\$38.32	28.86%	5.41%
C2-tail	58	357,336	\$1.04	9.48%	1.26%
C3-elite	26	14,836	\$45.37	62.12%	27.25%
C3-middle	51	15,782	\$23.47	30.68%	18.74%
C3-tail	616	32,703	\$0.02	7.20%	1.11%

A limitation is that the dataset contains only successful on-chain transactions. We therefore observe included winners, but not the full set of searchers competing for each opportunity, unsuccessful submissions, or builder- and relay-side auction messages. Since the model is defined over the full opportunity-level competition environment, the empirical calibration should be interpreted as a contract-level proxy rather than a fully identified structural estimate. Builder, relay, or private order-flow data would permit a more precise calibration of participation probabilities and value distributions.

Table 6 compares the FPA and the entry-filtered SCA_θ under these empirical settings. The raw expected active populations are $\mathbb{E}[N_o] = 14.565$ for C_1 and $\mathbb{E}[N_o] = 78.735$ for C_2 . For each opportunity type, θ is selected so that $\mathbb{E}[N_o^\theta] = 14$, and the cap is set to $H = 17$. Thus, the target corresponds to near-full admission in C_1 , but selective admission in C_2 .

For each θ , κ_S^* is the smallest searcher-side parameter satisfying the Bayesian copied-code Sybil-resistance inequalities for every participating class, every realized value $z \geq \theta$, every $q = 2, \dots, H + 1$, and every feasible copied-value profile $\mathbf{z}^c \in \mathcal{Z}_q^c(z)$. The copied-value profile is selected before the rival profile is realized. In both empirical calibrations, the numerical supremum is attained by a two-copy exact-copy deviation: for C_1 , by a middle-tier contract at $z \approx 0.143 > \theta$, and for C_2 , by a middle-tier contract at $z = \theta$.

The empirical comparison suggests that SCA_θ is most valuable when FPA rewards are relatively concentrated. This is the case for C_1 : the contribution-adjusted Jain index increases from 0.175 under the FPA to 0.983 under SCA_θ , while N_{eff} increases from 9.10 to 11.11. The increase in N_{eff} is more modest because the calibration treats the nine elite C_1 contracts as distinct searchers with identical class-level primitives. In practice, these contracts may be heterogeneous or share common ownership, implying stronger FPA reward concentration.

For C_2 , the FPA is already highly decentralized at the contract level, with $\mathcal{J}_{\text{Sh}}^{\text{FPA}} = 0.997$ and $N_{\text{eff}}^{\text{FPA}} = 79.54$. The entry-filtered SCA therefore produces little additional contribution-adjusted fairness, but increases the effective reward base to $N_{\text{eff}}^{\text{SCA}_\theta} = 84.05$ while satisfying the reported Bayesian copied-code Sybil-resistance constraints. Overall, the empirical evidence supports the interpretation that the entry-filtered SCA is most consequential in relatively concentrated opportunity classes such as C_1 . In already diffuse classes such as C_2 , its economic effect is smaller, and its primary role is to preserve broad reward allocation under an explicit searcher-side Bayesian security calibration.

6 Conclusion

This paper argues that searcher competition is a distinct and important locus of MEV centralization. While the current first-price, winner-take-all auction is robust against copied-code identity splitting, its economic allocation is inherently rank-based: rewards accrue to the searcher most likely to submit the highest-value execution, rather than to searchers in proportion to their marginal contribution. Under heterogeneous access and execution quality, rank dominance can amplify small advantages into persistent reward concentration.

■ **Table 5** Empirically extracted truncated-lognormal cost distributions.

Opportunity type	Tier	m	α	$\bar{c}$, $[\underline{c}, \bar{c}]$
C_1	Elite	9	0.950	0.2885, $[0.01, 0.95]$
C_1	Middle	17	0.255	0.5972, $[0.05, 1.00]$
C_1	Tail	48	0.035	0.6434, $[0.01, 1.00]$
C_2	Elite	11	0.753	0.8233, $[0.01, 1.00]$
C_2	Middle	19	0.808	0.9459, $[0.10, 1.00]$
C_2	Tail	58	0.950	0.9874, $[0.20, 1.00]$

■ **Table 6** FPA versus entry-filtered SCA for the empirical Ethereum-calibrated settings with $\mathbb{E}[N_o^\theta] = 14$. The SCA parameter κ_S^* is calibrated using the full Bayesian copied-value supremum.

Opportunity type	Mechanism	$\mathbb{E}[N_o^\theta]$	θ	H	κ_S^*	$\mathcal{J}_{Sh}$	N_{eff}	Note
C_1	SCA $_\theta$	14	0.141	17	0.527	0.983	11.11	near-full admission
	FPA	–	–	–	–	0.175	9.10	winner-take-all benchmark
C_2	SCA $_\theta$	14	0.392	17	0.457	0.999	84.05	selective admission
	FPA	–	–	–	–	0.997	79.54	winner-take-all benchmark

Motivated by this limitation, we introduced the entry-filtered Shapley-capped auction as one possible alternative. The mechanism preserves efficient execution by selecting the highest certified-value submission, but reallocates part of the retained surplus according to contribution among admitted high-quality submissions. Its cap, entry filter, and burn parameters expose the central design tradeoff: broader reward sharing improves contribution-adjusted fairness and reward decentralization, but must be disciplined by explicit Bayesian security constraints against copied-code Sybil deviations and validator–searcher coalitions. Overall, the results show that decentralization in MEV markets is not only a matter of infrastructure or order-flow access but also a mechanism-design problem requiring fairness, efficiency, and permissionless security to be addressed jointly.

References

- 1 Aleksei Adadurov, Sergey Barseghyan, Anton Chtepine, Antero Eloranta, Andrei Sebyakin, and Arsenii Valitov. Open vs. sealed: Auction format choice for maximal extractable value. *arXiv preprint*, 2026. [arXiv:2603.16333](#).
- 2 Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. Centralization in block-building and proposer-builder separation. In Jeremy Clark and Elaine Shi, editors, *Financial Cryptography and Data Security - 28th International Conference, FC 2024, Willemstad, Curaçao, March 4-8, 2024, Revised Selected Papers, Part I*, Lecture Notes in Computer Science, pages 331–349. Springer, 2024. doi:10.1007/978-3-031-78676-1_19.
- 3 Pedro Braga, Georgios Chionas, Piotr Krysta, Stefanos Leonardos, Georgios Piliouras, and Carmine Ventre. Who gets the maximal extractable value? a dynamic sharing blockchain mechanism. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, pages 2171–2173. International Foundation for Autonomous Agents and Multiagent Systems, 2024. doi:10.5555/3635637.3663097.
- 4 Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy, SP 2020, San Francisco, CA, USA, May 18-21, 2020*, pages 910–927. IEEE, 2020. doi:10.1109/SP40000.2020.00040.
- 5 Flashbots. Searching Post-Merge, 2022. URL: <https://writings.flashbots.net/searching-post-merge>.
- 6 Pranav Garimidi, Michael Neuder, and Tim Roughgarden. Beyond winner-take-all procurement auctions. *arXiv preprint*, 2026. [arXiv:2603.27779](#), doi:10.48550/arXiv.2603.27779.
- 7 Krzysztof Gogol, Johnnatan Messias, Deborah Miori, Claudio Tessone, and Benjamin Livshits. Cross-rollup MEV: Non-atomic arbitrage across L2 blockchains. *arXiv preprint*, 2024. [arXiv:2406.02172](#).
- 8 Tivas Gupta, Mallesh M. Pai, and Max Resnick. The centralizing effects of private order flow on proposer-builder separation. In Joseph Bonneau and S. Matthew Weinberg, editors, *5th Conference on Advances in Financial Technologies, AFT 2023, Princeton, NJ, USA, October 23-25, 2023*, LIPIcs, pages 20:1–20:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.AFT.2023.20.

- 9 Lioba Heimbach, Vabuk Pahari, and Eric Schertenleib. Non-atomic arbitrage in decentralized finance. In *IEEE Symposium on Security and Privacy, SP 2024, San Francisco, CA, USA, May 19-23, 2024*, pages 3866–3884. IEEE, 2024. doi:10.1109/SP54263.2024.00256.
- 10 Feng Luo, Zihao Li, Wenxuan Luo, Zheyuan He, Xiapu Luo, Zuchao Ma, Shuwei Song, and Ting Chen. Light into darkness: Demystifying profit strategies throughout the MEV bot lifecycle. In *33rd Annual Network and Distributed System Security Symposium, NDSS 2026, San Diego, California, USA, February 23-27, 2026*. The Internet Society, 2026. URL: <https://www.ndss-symposium.org/ndss-paper/light-into-darkness-demystifying-profit-strategies-throughout-the-mev-bot-lifecycle/>.
- 11 Akaki Mamageishvili, Christoph Schlegel, and Benny Sudakov. Searcher competition in block building. In Rainer Böhme and Lucianna Kiffer, editors, *6th Conference on Advances in Financial Technologies, AFT 2024, Vienna, Austria, September 23-25, 2024*, LIPIcs, pages 21:1–21:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.AFT.2024.21.
- 12 Burak Öz, Danning Sui, Thomas Thiery, and Florian Matthes. Who wins Ethereum block building auctions and why? In Rainer Böhme and Lucianna Kiffer, editors, *6th Conference on Advances in Financial Technologies, AFT 2024, Vienna, Austria, September 23-25, 2024*, LIPIcs, pages 22:1–22:25. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.AFT.2024.22.
- 13 Burak Öz, Christof Ferreira Torres, Christoph Schlegel, Bruno Mazorra, Jonas Gebele, Filip Rezabek, and Florian Matthes. Cross-chain arbitrage: The next frontier of MEV in decentralized finance. *Proc. ACM Meas. Anal. Comput. Syst.*, 9(3):51:1–51:33, 2025. doi:10.1145/3771566.
- 14 Vabuk Pahari and Andrea Canidio. How exclusive are Ethereum transactions? Evidence from non-winning blocks. *arXiv preprint*, 2025. doi:10.48550/arXiv.2509.16052.
- 15 Mallesh M. Pai and Max Resnick. Structural advantages for integrated builders in mev-boost. In Jurlind Budurushi, Oksana Kulyk, Sarah Allen, Theo Diamandis, Ariah Klages-Mundt, Andrea Bracciali, Geoffrey Goodell, and Shin'ichiro Matsuo, editors, *Financial Cryptography and Data Security. FC 2024 International Workshops - Voting, DeFI, WTSC, CoDecFin, Willemstad, Curaçao, March 4-8, 2024, Revised Selected Papers*, Lecture Notes in Computer Science, pages 128–132. Springer, 2024. doi:10.1007/978-3-031-69231-4_9.
- 16 Minghao Pan, Bruno Mazorra, Christoph Schlegel, and Akaki Mamageishvili. On sybil-proof mechanisms. *arXiv preprint*, 2026. arXiv:2407.14485.
- 17 Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*, pages 198–214. IEEE, 2022. doi:10.1109/SP46214.2022.9833734.
- 18 Rasheed, Pujan Desai, Yash Chaurasia, and Sujit P. Gujar. Shapley value-based approach for redistributing revenue of matchmaking of private transactions in blockchains, 2025. arXiv:2502.15420.
- 19 Tim Roughgarden. Transaction fee mechanism design. *ACM SIGecom Exchanges*, 19(1):52–55, 2021. doi:10.1145/3476436.3476445.
- 20 Roozbeh Sarenche and Yunwen Liu. Towards decentralized searcher competition in MEV markets, 2026. doi:10.48550/arXiv.2608.05011.
- 21 Lloyd S. Shapley. A value for n-person games. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 307–317. Princeton University Press, 1953.
- 22 Christof Ferreira Torres, Albin Mamuti, Ben Weintraub, Cristina Nita-Rotaru, and Shweta Shinde. Rolling in the shadows: Analyzing the extraction of MEV across layer-2 rollups. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, CCS 2024, Salt Lake City, UT, USA, October 14-18, 2024*, pages 2591–2605. ACM, 2024. doi:10.1145/3658644.3690259.

- 23 Gordon Tullock. Efficient rent seeking. In James M. Buchanan, Robert D. Tollison, and Gordon Tullock, editors, *Toward a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, 1980.
- 24 Shuzheng Wang, Yue Huang, Wenqin Zhang, Yuming Huang, Xuechao Wang, and Jing Tang. Private order flows and builder bidding dynamics: The road to monopoly in Ethereum’s block building market. In Guodong Long, Michale Blumestein, Yi Chang, Liane Lewin-Eytan, Zi Helen Huang, and Elad Yom-Tov, editors, *Proceedings of the ACM on Web Conference 2025, WWW 2025, Sydney, NSW, Australia, 28 April 2025- 2 May 2025*, pages 2144–2157. ACM, 2025. doi:10.1145/3696410.3714754.
- 25 Wenhao Wang, Aditya Saraf, Lioba Heimbach, Kushal Babel, and Fan Zhang. Blockspace under pressure: An analysis of spam MEV on high-throughput blockchains. *arXiv preprint*, 2026. doi:10.48550/arXiv.2604.00234.
- 26 Ben Weintraub, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. A flash(bot) in the pan: measuring maximal extractable value in private pools. In Chadi Barakat, Cristel Pelsser, Theophilus A. Benson, and David R. Choffnes, editors, *Proceedings of the 22nd ACM Internet Measurement Conference, IMC 2022, Nice, France, October 25-27, 2022*, pages 458–471. ACM, 2022. doi:10.1145/3517745.3561448.
- 27 Fei Wu, Danning Sui, Thomas Thiery, and Mallesh M. Pai. Measuring CEX-DEX extracted value and searcher profitability: The darkest of the MEV dark forest. In Zeta Avarikioti and Nicolas Christin, editors, *7th Conference on Advances in Financial Technologies, AFT 2025, Pittsburgh, PA, USA, October 8-10, 2025*, LIPIcs, pages 26:1–26:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.AFT.2025.26.
- 28 Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. Strategic bidding wars in on-chain auctions. In *IEEE International Conference on Blockchain and Cryptocurrency, ICBC 2024, Dublin, Ireland, May 27-31, 2024*, pages 503–511. IEEE, 2024. doi:10.1109/ICBC59979.2024.10634354.
- 29 Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. To compete or collude: Bidding incentives in Ethereum block building auctions. In *Proceedings of the 5th ACM International Conference on AI in Finance, ICAIF 2024, Brooklyn, NY, USA, November 14-17, 2024*, pages 813–821. ACM, 2024. doi:10.1145/3677052.3698629.
- 30 Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. From competition to centralization: The oligopoly in Ethereum block building auctions. In Inês Lynce, Nello Murano, Mauro Vallati, Serena Villata, Federico Chesani, Michela Milano, Andrea Omicini, and Mehdi Dastani, editors, *ECAI 2025 - 28th European Conference on Artificial Intelligence, 25-30 October 2025, Bologna, Italy - Including 14th Conference on Prestigious Applications of Intelligent Systems (PAIS 2025)*, Frontiers in Artificial Intelligence and Applications, pages 3387–3394. IOS Press, 2025. doi:10.3233/FAIA251209.
- 31 Sen Yang, Kartik Nayak, and Fan Zhang. Decentralization of Ethereum’s builder market. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *IEEE Symposium on Security and Privacy, SP 2025, San Francisco, CA, USA, May 12-15, 2025*, pages 1512–1530. IEEE, 2025. doi:10.1109/SP61157.2025.00157.
- 32 Ao Zhang, Yunwen Liu, Ren Zhang, Yingdi Shan, and Yongwei Wu. Order flow exclusivity and value extraction mechanisms: An analysis of Ethereum builder centralization. *arXiv preprint*, 2026. doi:10.48550/arXiv.2605.04471.