

To Wait or to Probe: Arbitrage Competition on High-Throughput Blockchains

Fei Wu*   

King's College London, London, UK

Burak Öz  

Flashbots, Berlin, Germany

Abstract

Maximal Extractable Value (MEV) on high-throughput blockchains can be captured through *targeted search*, where bots identify opportunities off-chain and submit route-committed transactions, or through *probabilistic search*, where bots submit repeated attempts that discover opportunities during on-chain execution. This distinction has direct implications for spam, blockspace consumption, and protocol revenue. We model how ordering granularity, fee floors, and opportunity-access shocks shape competition between these architectures. Using cyclic arbitrage data on Base from June 2025 to February 2026, we develop a trace-level classifier for search architectures and show that the resulting labels correspond to distinct execution behavior. In our sample, probabilistic search accounts for only 23% of arbitrage activity but produces 95% of spam and consumes 20% of Base gas. We test the model across three episodes: Flashblocks selects against broad on-chain probabilistic scanners; token-launch opportunity shocks temporarily revive probabilistic search; and higher fee floors select against probabilistic bots whose opportunity flow cannot sustain repeated attempts. After Base's configuration changes, protocol revenue shifts toward successful arbitrages and away from spam, probabilistic bots pay higher priority fees, and spam consumes a smaller share of blockspace.

2012 ACM Subject Classification Applied computing → Electronic commerce

Keywords and phrases Maximal Extractable Value, cyclic arbitrage, spam, Base

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.24

Related Version *Full Version*: <https://arxiv.org/abs/2606.00720> [35]

1 Introduction

A book collector offers \$500 for a first-edition copy of *The Hobbit*. A local bookstore will receive the book at an unknown time and sell it for \$400, with customers allowed to add a gratuity for priority service. The book becomes purchasable as soon as it enters the bookstore's inventory, but the public website updates only every ten minutes. Alice waits for the public update, then submits a purchase offer with a high gratuity. Bob repeatedly queries live inventory and buys if one check reveals the book first, imposing load on the store's system. Who wins the arbitrage? What would the bookstore prefer: targeted purchase attempts with higher gratuities, or repeated checks that might reveal opportunities earlier but consume system resources? Would the answer change if updates became more frequent or repeated checks became more costly?

This toy example captures a central design problem for Maximal Extractable Value (MEV) search on high-throughput blockchains: should searchers *wait* for information and submit targeted execution attempts, or *probe* the system with transactions whose profitability is resolved during execution? On Ethereum, MEV strategies are typically associated with targeted search: searchers identify and simulate an opportunity *off-chain*, then submit a

*Fei Wu performed work in part during an internship at Flashbots in Spring 2026.



© Fei Wu and Burak Öz;

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. 24; pp. 24:1–24:24

Leibniz International Proceedings in Informatics



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

transaction committed to a specific execution route [7, 14, 26, 27, 34, 36]. On high-throughput blockchains, fast block production, low fees, and limited visibility into pending transactions can make a second architecture viable: searchers submit many transactions that read live state and resolve opportunity discovery *on-chain*, proceeding to further execution only when a profitable route is found [6, 11, 16, 28, 32].

We refer to these two architectures as *targeted search* and *probabilistic search*. Because opportunity discovery occurs at different stages, the architectures differ in their sensitivity to ordering rules and fees, and consequently in their attempt efficiency, blockspace usage, and contribution to protocol revenue. Probabilistic searchers submit speculative attempts that resolve opportunity discovery during execution; many fail without producing useful state changes and are therefore treated as *spam*. These attempts still consume blockspace and pay fees, so failed search can be both a congestion externality and a source of protocol revenue. Targeted searchers instead submit route-committed transactions against identified opportunities with potentially higher fees, but can also fail if the opportunity disappears before execution. The design question is therefore which architecture a protocol encourages, and how the resulting bot composition shapes blockspace consumption and protocol revenue.

This question has become increasingly relevant as high-throughput chains experiment with ordering rules and fee designs to reduce spam externalities and improve protocol revenue [4, 5, 21, 22]. However, the equilibrium effect of such changes is not obvious. Existing work characterizes probabilistic search timing [16] and models equilibrium spam volume as a function of block capacity, gas fee floor, and fee-based ordering [32], but typically focuses on a single search architecture. In this paper, we formalize the distinction between targeted and probabilistic search and study how protocol design shapes competition between them.

We develop a reduced-form equilibrium model of architecture competition in which ordering granularity, fee floors, and opportunity access shape the relative viability of targeted and probabilistic search (cf. Section 3). The model highlights two margins through which protocol changes affect spam: which architectures remain active, and how intensively active bots submit attempts.

We test these predictions on Base, which has documented spam activity [11, 28, 32] and underwent major configuration changes during our June 2025–February 2026 sample period [3]. We identify cyclic arbitrages and associated spam, and develop a trace-level classifier that distinguishes search architectures by the extent to which route discovery is observable before swap execution in successful arbitrage traces (cf. Section 4).

Our empirical analysis proceeds in three steps. We first validate that the classifier separates economically meaningful search architectures by showing that they differ in execution behavior, attempt efficiency, and spam production (cf. Section 5.1). We next study how arbitrage competition and the architecture relative viability evolve around three episodes: Flashblocks, the September surge, and the minimum-base-fee escalation (cf. Section 5.2–5.4). Finally, we aggregate bot-level responses to chain-level outcomes, including priority-fee revenue and blockspace consumption (cf. Section 5.5). We summarize our contributions as follows:

1. We formalize the distinction between *targeted* and *probabilistic* search architectures and develop a reduced-form equilibrium model in which ordering granularity, fee floors, and opportunity access shape the relative viability of the two architectures.
2. We develop a trace-level classifier that distinguishes the architectures through observable route discovery within successful arbitrages, and validate that the labels correspond to coherent transaction designs, spam behavior, and efficiency profiles.
3. We test the model’s predictions on Base. Flashblocks selects against scan-intensive probabilistic bots and toward leaner survivors; a temporary opportunity shock revives probabilistic activity through short-lived entrants; and fee floor escalation selects against high-fee-exposure bots whose opportunity flow no longer supports repeated attempts.

4. We quantify the chain-level consequences. After configuration changes, priority-fee revenue shifts away from failed attempts toward realized arbitrages and higher-bidding probabilistic search attempts, while spam consumes a smaller share of blockspace.

1.1 Related work

MEV competition by targeted search has been documented on Ethereum Layer-1 [7, 14, 17, 26, 27, 29, 33, 34, 36], alternate Layer-1s [31, 38], Layer-2 networks [8], and cross-network [12, 20, 25]. Recent work documents probabilistic search on high-throughput blockchains: [6] introduces it as a viable strategy on Solana under faster block times and lower fees, and [11, 28] show that similar strategies consume substantial blockspace on Base and Optimism, though [28] notes a different pattern on Arbitrum without formally distinguishing the two architectures. A related strand studies reverts and mechanisms for reducing failed MEV attempts [13, 18, 37]. Beyond empirical analysis, [32] derives equilibrium spam volume under varying block capacity, minimum gas price, and priority-fee ordering, while [16] analyzes the timing game underlying probabilistic search competition. These models focus on probabilistic search in isolation. Our paper distinguishes targeted and probabilistic search architectures, connects failed attempts to architecture choice, and studies how ordering granularity, fee floors, and opportunity access shift competition between them.

2 Preliminaries: Arbitrage Search Architectures

Section 1 introduced targeted and probabilistic search as two economic architectures. A central setting for our analysis is cyclic arbitrage, where a bot routes a token through two or more decentralized exchange (DEX) pools and returns to the starting token with a positive net profit. On high-throughput chains, bots often submit candidate transactions before knowing with certainty whether execution will be profitable. Failed attempts can therefore arise under both architectures: a route-committed transaction may find that the targeted opportunity has disappeared, while a probabilistic search transaction may fail to find any profitable route during execution. We use *failed attempts* as an umbrella term for arbitrage attempts that do not execute a profitable swap sequence. Within this category, *probes* terminate after state reads without executing swaps, while *reverts* revert during execution. Both consume execution resources and are commonly treated as spam. We next describe how these architectures appear at the *implementation level*, focusing on where route discovery occurs and what evidence is visible in transaction traces.

Off-chain discovery with direct execution. The cleanest implementation of targeted search is a transaction that proceeds directly from entry to swap execution with no pool-state read before the first swap.¹ The route, DEX venues, and input amount are embedded in the transaction calldata, indicating that the opportunity was discovered, simulated, and parameterized off-chain before submission.

Route-confined on-chain evaluation. In the second implementation, the bot also begins with a pre-selected route embedded in the calldata, but reads live state from the relevant pools before committing to swaps.² These reads can be interpreted as a profitability check:

¹ Consider example tx hash: <https://basescan.org/tx/0x1d93f62a009974da3ae575e7c623d88f09693b3e4e64ff0eabbac821f3fa091e>.

² Consider example tx hash: <https://basescan.org/tx/0x47a91fe3b9c31818f6c7f6699f6675297a8b926af48084e0c17cf8504a44df14>.

if the route remains profitable, the transaction proceeds to swap execution; otherwise, it terminates early or reverts. Importantly, the pre-swap reads are confined to the venues that the transaction later executes against. This pattern is economically closer to targeted search than to probabilistic search, but it is intrinsically ambiguous from traces alone: it could also reflect discovery over a confined venue set that happens to coincide with the executed venues.

Broad on-chain discovery. The canonical implementation of probabilistic search is a transaction that scans a wider set of pools than it ultimately executes on.³ The bot reads state from candidate venues to identify which subset offers a profitable route, then executes the selected route, or terminates or reverts if no profitable route is found. This strict superset relation is the trace-level signature of route discovery during execution.

Bots read pool state through several trace-visible mechanisms, including direct pool reads, balance-based reads, revert-based quotes, and helper-contract calls. We use these mechanisms to identify pre-swap state reads in transaction traces (see full version [35]).

3 A reduced-form equilibrium model

This section presents a stylized equilibrium model of how ordering granularity, transaction costs, and opportunity access shape the relative competitiveness of arbitrage bot architectures. Entry adjusts until expected profits are driven to zero.

Ordering regimes and bot architecture. We consider two ordering regimes $r \in \{B, F\}$, where B denotes the baseline regime and F denotes a finer ordering regime. The relevant ordering unit is a regular full block in regime B , but a smaller sub-block slice in regime F .

We next describe two representative bot architectures. Following the implementation-level vocabulary of Section 2, we refer to them as *off-chain discovery*, denoted by O , and *on-chain discovery*, denoted by C , corresponding to targeted and probabilistic search, respectively. Bots compete over candidate arbitrage opportunities. We assume that under the chain's ordering regime, an opportunity is exhausted by the first valid claimant in the earliest profitable ordering unit.

Same-slice competitiveness. On-chain discovery transactions perform state reads to evaluate candidate routes, so they consume more blockspace than route-committed off-chain transactions targeting the same opportunity. Under finer ordering, early sub-block slices have smaller effective capacity, so heavier on-chain transactions are more likely to be displaced, allowing a lean off-chain transaction to capture the opportunity in the earlier slice.⁴

We summarize this disadvantage by a reduced-form *same-slice competitiveness* function

$$\eta_r \in [0, 1), \quad r \in \{B, F\},$$

where η_r is the probability that an on-chain discovery transaction remains competitive in the same earliest profitable slice as an off-chain discovery transaction in regime r . The key restriction is $\eta_F < \eta_B$: on-chain discovery is less likely to remain competitive in the same profitable slice under a finer ordering regime than under the baseline regime.

³ Consider example tx hash: <https://basescan.org/tx/0xc020e534c39e203e9f3a42ea5f89ca023e486c38e15ca55e2a183ead1f4c885d>.

⁴ Under Flashblocks on Base [2], transactions exceeding 14M gas cannot fit in the first Flashblock.

Opportunity access and free-entry payoffs. The two architectures differ in opportunity access. An off-chain discovery transaction may encode several pre-identified routes, but its candidate route set is fixed before submission. An on-chain discovery transaction instead reads the current state and can select among a state-contingent set of candidate routes during execution. We summarize this potential breadth advantage through architecture-specific opportunity-access rates. Let λ_i denote the expected number of candidate opportunities per model period accessible to architecture $i \in \{O, C\}$. With average gross value V conditional on successful capture, $\lambda_i V$ is the gross opportunity value flow accessible to architecture i , before competition and transaction costs.⁵ We allow $\lambda_C \geq \lambda_O$ to represent the potentially broader access of on-chain discovery.

Let m_i denote the effective mass of independent competing strategies using architecture i . Here, “mass” is the standard continuum-entry measure of participation rather than a literal count of unique addresses or contracts observed on-chain. In a fraction η_r of competitive states, an on-chain discovery transaction remains in the earliest profitable ordering slices and competes with off-chain discovery transactions. Conditional on this case, the opportunity is allocated uniformly across the total active mass $m_O + m_C$. In the remaining states, the on-chain discovery transaction is displaced, and only off-chain discovery bots compete, uniformly across m_O . Each bot incurs an expected cost $K_i(b) > 0$ per modeled period, where $b > 0$ is the minimum base fee. Applying this reduced form of the sharing rule to each architecture’s accessible opportunity flow yields the payoffs below.

The expected payoff of an off-chain discovery bot in regime r is

$$\Pi_O^r(m_O, m_C; b) = \eta_r \frac{\lambda_O V}{m_O + m_C} + (1 - \eta_r) \frac{\lambda_O V}{m_O} - K_O(b).$$

The first term is the off-chain bot’s share when both architectures remain competitive in the earliest profitable slice; the second is its share when on-chain discovery is displaced; the final term is its expected cost. An on-chain-discovery bot earns opportunity value only in the fraction η_r of states in which its heavier transaction remains competitive in the earliest profitable slice:

$$\Pi_C^r(m_O, m_C; b) = \eta_r \frac{\lambda_C V}{m_O + m_C} - K_C(b).$$

A competitive equilibrium in regime r is a pair $(m_O^{*,r}, m_C^{*,r})$ such that no bot architecture earns positive expected profit, and any active bot architecture earns zero expected profit:

$$\Pi_i^r(m_O^{*,r}, m_C^{*,r}; b) \leq 0 \quad \text{for } i \in \{O, C\}, \quad m_i^{*,r} > 0 \Rightarrow \Pi_i^r(m_O^{*,r}, m_C^{*,r}; b) = 0.$$

Thus, active architectures earn zero profit, and inactive architectures cannot enter profitably.

► **Remark 1 (Blockspace competitiveness versus opportunity creation).** The model isolates the blockspace competitiveness channel rather than the opportunity-creation channel. Finer ordering may also change the value or arrival rate of arbitrage opportunities by reducing the order flow aggregated into each interval. In our notation, such effects would enter through V or λ_i . We hold V and λ_i fixed to focus on architecture composition.

For simplicity, we next focus on equilibria with active off-chain discovery architecture. Other equilibrium configurations can be derived analogously.

⁵ We assume V to be common across architectures. In reality, off-chain discovery may extract more value through optimized off-chain simulation; this asymmetry would only reinforce the qualitative results.

Define the effective value-cost ratios

$$\kappa_O(b) := \frac{\lambda_O V}{K_O(b)}, \quad \kappa_C(b) := \frac{\lambda_C V}{K_C(b)}.$$

We assume that, over the parameter region of interest, $\kappa_C(b) > \kappa_O(b) > 0$. The condition is imposed to make coexistence possible before applying the ordering competitiveness penalty η_r . If the condition fails, on-chain discovery is already dominated on the branch considered here. The observed coexistence of both architectures in our sample motivates studying this parameter region, but we do not estimate the inequality structurally.

► **Proposition 2** (Within-regime equilibrium). *Fix $b \geq 0$, $r \in \{B, F\}$, and $\kappa_C(b) > \kappa_O(b) > 0$. On the equilibrium branch with active off-chain discovery, the unique equilibrium is*

$$(m_O^{*,r}, m_C^{*,r}) = \begin{cases} \left(\frac{(1 - \eta_r)\kappa_O(b)}{1 - \kappa_O(b)/\kappa_C(b)}, \frac{\eta_r\kappa_C(b) - \kappa_O(b)}{1 - \kappa_O(b)/\kappa_C(b)} \right), & \eta_r\kappa_C(b) > \kappa_O(b), \\ (\kappa_O(b), 0), & \eta_r\kappa_C(b) \leq \kappa_O(b). \end{cases}$$

Thus, on-chain discovery survives iff $\eta_r\kappa_C(b) > \kappa_O(b)$, i.e., its access-adjusted value-cost advantage is large enough to overcome its ordering disadvantage.

Next, we show how a decline in same-slice competitiveness, from η_B to η_F , shifts equilibrium participation away from on-chain discovery and toward off-chain discovery.

► **Theorem 3** (Finer ordering reduces the competitiveness of on-chain discovery). *Fix $b \geq 0$, and suppose $\eta_F < \eta_B$ and both regimes are on the coexistence branch. Then*

$$m_C^{*,F}(b) < m_C^{*,B}(b), \quad m_O^{*,F}(b) > m_O^{*,B}(b).$$

Switching to the finer ordering regime reduces on-chain participation and increases off-chain participation on the coexistence branch.

To capture selection within on-chain discovery, we allow subtypes to differ in their *per-attempt scan footprint*, which measures how much venue-search work a single attempt performs. It is distinct from *attempt intensity*, i.e., how many attempts a bot sends. A larger scan footprint may broaden the candidate set evaluated within an attempt, but also makes the attempt heavier, reducing same-slice competitiveness under finer ordering and raising gas cost. This trade-off creates a selection margin within the on-chain discovery category.

► **Proposition 4** (Finer ordering selects against high-footprint on-chain discovery). *Under the regularity conditions defined in Appendix A.2 in the full version [35], let on-chain discovery subtypes be indexed by a per-attempt scan-footprint parameter x , and let $x_r^*(b)$ denote the largest viable footprint in regime r . Then $x_F^*(b) < x_B^*(b)$, so finer ordering selects against more scan-intensive on-chain-discovery strategies.*

This selection result implies that, among surviving on-chain discovery bots, observable scan-footprint measures, such as the number of pre-swap state reads, should weakly fall under finer ordering. We formalize this prediction in Appendix A.2 in the full version [35].

► **Remark 5** (Off-chain discovery reliability under finer ordering). Finer ordering also shortens the interval between opportunity observation and execution, reducing the stale-state risk faced by off-chain discovery. In the model, this channel would raise the effective value-cost ratio of off-chain discovery by reducing failed attempts. This stale-state channel and the

blockspace-competitiveness channel both predict a shift away from broad on-chain discovery toward route-committed execution, so the composition test in Section 5.2 cannot separately identify them. However, the stale-state channel more directly predicts improved off-chain attempt efficiency. The rise in off-chain discovery success rates in Table 4 is therefore suggestive evidence for this margin, although not separately identified from changes in blockspace competition.

The theorems above isolate the blockspace-competitiveness channel under a finer ordering regime. We then turn to the cost channel. A higher minimum base fee raises the cost of every attempt, so it selects against architectures with higher fee exposure.

► **Proposition 6** (A higher minimum base fee can disproportionately hurt on-chain discovery). *Define architecture i 's fee exposure and on-chain discovery effective advantage as*

$$\chi_i(b) := \frac{\partial \log K_i(b)}{\partial b}, \quad \rho(b) := \frac{\kappa_C(b)}{\kappa_O(b)}.$$

If $\chi_C(b) > \chi_O(b)$, then $\rho(b)$ is strictly decreasing in b . Equivalently, the coexistence condition $\eta_r \kappa_C(b) > \kappa_O(b)$ becomes harder to satisfy as the fee floor rises.

The compositional prediction requires asymmetric fee exposure. If the fee floor scaled every architecture's cost by the same proportional factor, then $\chi_C(b) = \chi_O(b)$, and a higher fee floor would reduce total entry without changing architecture composition. Empirically, this condition is consistent with on-chain discovery consuming more gas per attempt and sending more attempts per successful arbitrage (cf. Section 5.1).

The same logic applies within on-chain discovery: holding opportunity access and same-slice competitiveness fixed, higher fee floors reduce the advantage of on-chain subtypes with more fee-sensitive cost functions. We formalize this in the full version [35].

► **Remark 7** (What fee exposure means empirically). Fee exposure depends not only on per-attempt scan footprint, but also on attempt intensity. Thus, higher fee floors should be interpreted as raising the hurdle rate for repeated probing, not as mechanically eliminating the most scan-intensive bots. High-exposure bots may remain viable if their opportunity flow is valuable enough, while lower-exposure bots may exit if their opportunity flow is weak.

The blockspace and cost channels above hold the opportunity environment fixed. Token launches and fresh DEX deployments can temporarily raise λ_C by creating routes that off-chain discovery architectures have not yet indexed or parameterized. Such shocks can revive on-chain discovery participation even under otherwise unfavorable policy conditions.

► **Proposition 8** (Opportunity-access shocks can revive on-chain discovery). *On the coexistence branch, $m_C^{*,r}$ is strictly increasing and $m_O^{*,r}$ strictly decreasing in λ_C . There exists a viability threshold $\bar{\lambda}_C(r, b)$ such that on-chain discovery is active iff $\lambda_C > \bar{\lambda}_C(r, b)$, and this threshold is strictly increasing in policy frictions. Hence, a temporary increase in λ_C can revive on-chain discovery even when the policy conditions render it inactive.*

► **Remark 9** (Participation and observed spam). Free entry links the model's equilibrium masses to empirical architecture prominence: when a policy change lowers an architecture's expected payoff, the equilibrium supports a smaller mass of participants using that architecture. Empirically, lower architecture mass can be reflected in fewer active bots or a lower architecture share of arbitrage count or volume. Aggregate spam, however, also depends on attempt intensity per bot. Thus, spam may not fall proportionally if surviving or entering bots submit more attempts.

4 Data and Methodology

We use Base as a case study to measure arbitrage competition between search architectures and their response to protocol changes. Base is a natural setting because it has substantial DEX activity and spam volume [11, 28, 32], and because it implemented major configuration changes, most notably Flashblocks and subsequent increases in the minimum base fee [3]. This section describes the data construction and bot-architecture classification methodology.

4.1 Data collection

We collect cyclic arbitrages and associated spam activity on Base from June 1, 2025, through February 28, 2026, covering the periods of major configuration changes on Base [3].

Successful cyclic arbitrages. We extend the cyclic arbitrage detection method of [28]. Using a custom Dune query, we detect 21,374,434 cyclic arbitrages by 4,365 unique arbitrage bot addresses over the observed period. For each transaction, we collect gas usage, fees, USD notional volume, and the realized arbitrage route, including traded tokens and DEX pools.

Failed arbitrage attempts (spam). We separately query reverted and probe transactions that are attributed to the same arbitrage bot addresses. Importantly, the bot does not need to interact with the DEX pool contract to read its state. To capture this broader class of probing behavior, we classify as probes successful transactions that emit no logs from external contracts other than the bot itself and carry ETH value ≤ 1 WEI.⁶

This yields 985,559,360 spam transactions from these cyclic arbitrage bots over the observed period, including 35,653,523 reverts and 949,905,837 probes. Because transaction-level analysis of this full spam set is infeasible, we aggregate it at the bot-week level: for each bot in each week, we collect the count of probes and reverts, corresponding total gas usage and fees, and share of transaction count and gas usage out of the chain’s throughput.

4.2 Bot architecture identification

Transaction-level labeling. Because failed-attempt outcomes can arise from any of the three implementation patterns, distinguishing architectures requires trace-level evidence rather than transaction outcomes alone. We identify architecture from successful arbitrage transactions, where the executed route is observable. The objective is to distinguish whether a transaction arrives with a route already committed, evaluates only the route it eventually executes, or scans multiple candidate venues before choosing an execution route.

Since tracing all 21.37M arbitrages is infeasible, for each bot in each week during the observed period, we randomly sample up to 50 arbitrage transactions, yielding 359,752 sampled transactions across 4,365 bots and 39 weeks. We label these transactions using a trace-level classifier applied to `debug_traceTransaction` call traces with an Alchemy RPC node. The classifier distinguishes three transaction-level labels:

OFFCHAIN_DIRECT transactions show no meaningful pre-swap state read, or use only helper contract or revert-based quoting on the eventual execution route.

ONCHAIN_EVAL transactions perform on-chain evaluation before execution, but that evaluation is confined to the route that is eventually executed.

ONCHAIN_BROAD_SCAN transactions evaluate venues beyond the eventual execution set before committing to the winning route.

⁶ Simple ETH transfer also emits no external logs. Certain arbitrage bots typically attach 1 WEI to their transactions, such as `0xA4B9eDB99806bFB8BbEE89D13F4c296DAF225dEE`. Thus, we set 1 WEI as a buffer.

The classifier’s core is to identify if state reads exist and compare the venues read before swap with the venues used in the successful route. Reads of executed venues are treated as route-confined evaluation; reads of non-executed venues are treated as route discovery.

From transaction labels to bot architecture. We aggregate transaction-level labels into weekly bot architectures using the implementation patterns introduced in Section 2. A bot-week is labeled *on-chain discovery* if its sample contains at least one manually validated `ONCHAIN_BROAD_SCAN` transaction. It is labeled *off-chain discovery* if it contains `OFFCHAIN_DIRECT` behavior and no credible broad-scan evidence after manual validation. It is labeled *on-chain evaluation* only if all sampled transactions are `ONCHAIN_EVAL`. This asymmetry reflects the fact that direct execution and broad venue expansion are clear endpoint behaviors, while confined on-chain evaluation can arise under either off-chain discovery with on-chain re-validation or narrow on-chain discovery.

This rule assumes architecture is persistent within a bot-week, rather than switching transaction by transaction. Broad on-chain scanning requires different execution logic from route-committed execution, so genuine changes to off-chain architecture are more plausibly associated with redeployment or a new bot identity. Because labels are assigned each week, the same bot can still receive different labels across weeks if its observed behavior changes over longer horizons. We then project weekly labels onto the full arbitrage dataset and weekly spam datasets, producing a unified bot-week panel with architecture, successful arbitrage activity, and aggregate spam behavior.

Mixed `ONCHAIN_BROAD_SCAN` and `OFFCHAIN_DIRECT` cases within the same bot-week are manually reviewed to exclude false `ONCHAIN_BROAD_SCAN` classifications. Full classifier details, edge cases, and manual-validation rules are specified in the full version [35].

Scope and limitations. We infer architecture from successful transactions because failed attempts often do not reveal a realized execution route, making route-confined evaluation and broad scanning indistinguishable. Projecting the weekly label onto spam therefore assumes that successful and failed attempts within a bot-week are generated by the same underlying architecture. While this cannot be easily validated transaction by transaction, the distinct architecture behavior we identified in the later sections provides indirect support for the projection. The classifier is conservative in two further ways: singleton `PoolManager` designs can hide venue expansion, and sampling up to 50 transactions may miss rare direct-execution or broad-scan transactions. We quantify these effects in Appendix B.2 in the full version [35].

5 Empirical Results

We organize the empirical results around the model’s predictions and Base configuration changes during the sample period [3]. We first establish that the three architecture labels correspond to distinct successful arbitrage and spam profiles. We then study how the prominence of each architecture changes around three empirical episodes: the introduction of Flashblocks, the September on-chain-discovery surge, and the escalation of the minimum-base-fee. Finally, we quantify how each architecture’s response to these episodes translates into chain-level outcomes, including priority-fee payments and blockspace consumption.

■ **Table 1** Bot activity by weekly architecture label, June 1, 2025–February 28, 2026. Bot counts are unique addresses observed with each label at least once; categories are not mutually exclusive because labels can vary across weeks.

Bot label	Bots	Arbs	Spam	Reverts	Probes	Success rate	Median / P95 avg. spam tx gas	Avg. cost per success
Off-chain discovery	1,634	7.67M	23.63M	16.86M	6.78M	42.51%	352K / 1.02M	0.000095 ETH
On-chain evaluation	798	8.83M	27.63M	12.51M	15.12M	18.49%	64.8K / 371K	0.000045 ETH
On-chain discovery	2,049	4.87M	934.30M	6.29M	928.01M	4.58%	277K / 1.39M	0.000253 ETH

5.1 Bot architecture characteristics


We begin by documenting how the three bot architectures differ in their observable behavior, both in successful arbitrages and in spam during our sample period. These differences validate that the classifier separates economically distinct architectures rather than spurious trace clusters, and they support the model’s key primitives: on-chain discovery bots operate over broader venue sets with lower success rates, and have higher fee exposure.

Spam character. The three architectures produce strikingly different spam activities. Table 1 summarizes full-sample activity by weekly bot label. The table reports the number of bots, successful arbitrages, spam transactions, reverts, probes, average bot-level success rate, spam transaction gas usage, and average cost per arbitrage opportunity. Success rate is computed at the bot level as successful arbitrages divided by total attempts, where total attempts include successful arbitrages and spam transactions. Spam gas is computed as each bot-week’s average gas per spam transaction and then summarized across bot-weeks. Average cost per successful arbitrage is total gas cost across successful and failed attempts divided by successful arbitrages.

Off-chain discovery bots produce the least spam and have the highest average success rate at 42.51%. The spam they produce consists predominantly of reverts rather than probes. Because routes are pre-committed, their transactions typically proceed directly to swap execution and revert only when profitability conditions fail close to the end of execution. The resulting reverts traverse part or all of the swap logic, which explains their highest median average spam gas usage (352K) of the three architectures. Their average cost per successful arbitrage is 0.000095 ETH – moderate despite heavy per-spam gas usage, because relatively few failed attempts accompany each successful arbitrage.

On-chain evaluation bots produce spam split roughly evenly between reverts and probes, which we interpret as an implementation choice rather than a meaningful architectural distinction: both outcomes terminate a failed attempt after verifying a confined route set. Their evaluation is lightweight and confined to a narrow candidate venue set, so their average spam gas usage is the lowest of the three architectures (median 64.8K, P95 371K). Their average cost per captured opportunity is correspondingly the lowest at 0.000045 ETH, reflecting both cheap spam and a modest success rate.

On-chain discovery bots generate the largest spam activity with 934.30M total spam transactions, almost all of which are probes. This matches their architecture: the bot reads state across many venues to discover whether an opportunity exists at all, and most such reads find none, leaving a probe transaction on-chain. The median average spam gas usage (277K) is higher than that of on-chain evaluation, consistent with broader venue scanning, and the P95 (1.39M) reflects a tail of very wide route searches. Their average cost per opportunity is 0.000253 ETH, much higher than other architectures – the compound effect of heavy spam transactions and a large number of them per successful arbitrage. The average success rate is correspondingly low at 4.58%.



■ **Figure 1** Arbitrage gas usage CDF.

■ **Table 2** Sampled arbitrage trace-level metrics.

Bot label	Bot-level architecture		Transaction-level footprint	
	Median calldata length	Average / Median venue commitment	Median pre-reads	oracle usage
Off-chain disc.	516	0.607 / 0.971	0	22.61%
On-chain eval.	150	0.535 / 0.886	3	4.61%
On-chain disc.	91	0.184 / 0.000	40	0.08%


Successful arbitrage fingerprints. The architectural distinction also shows up in how successful arbitrage transactions are structured. Figure 1 plots the cumulative distribution function of successful arbitrage gas usage by bot label. Successful arbitrages by on-chain discovery bots are more gas-intensive, consistent with broad venue scanning making these transactions heavier. We also observe that off-chain discovery transactions use more gas than on-chain evaluation transactions. This does not contradict their route-committed classification: off-chain discovery bots may execute more complex routes, use flash loans,⁷ or route through venues such as Balancer, where execution may involve a vault and oracle calls.

Table 2 reports complementary architecture and footprint measures from the sampled successful arbitrage transactions. Calldata length and venue commitment are summarized at the bot level to avoid overweighting high-activity bots; pre-execution reads and oracle usage are summarized at the transaction level because they measure per-transaction footprint. Venue commitment is a conservative address-level proxy for route pre-commitment, defined as the fraction of executed venue addresses embedded in calldata.


Off-chain discovery bots have the longest median calldata and the highest venue commitment, and their transactions have a median of zero pre-execution reads, consistent with pre-committed routes carried in the transaction itself and direct execution behavior. On-chain discovery bots show the opposite pattern with the shortest median calldata, zero median venue commitment, and 40 median pre-execution reads in their transactions, indicating that the route is resolved during execution through dense state reads across many venues. On-chain evaluation bots sit between the two, with modest pre-reads against the confined set of venues they subsequently execute on.

Time-series overview. We next document the time-series evolution of the prominence of each architecture during the observed period. Figure 2 presents the weekly successful arbitrage, USD volume, and spam activity by architecture. Across the full sample, on-chain discovery bots account for 22.8% of successful arbitrages and 28% of volume, but produce nearly 95% of the total spam activity. Figure 3A plots the weekly count of active bots by architecture. In most weeks, 30–50 active on-chain discovery bots generate roughly 20M spam transactions per week. By contrast, off-chain discovery and on-chain evaluation have larger active populations, but generate substantially less spam while accounting for a larger share of successful arbitrages. Figure 3B–C tracks active bots and volume shares by architecture. The series highlights three episodes we study below: Flashblocks on July 7, 2025; a temporary on-chain discovery surge in early September 2025; and the minimum-base-fee escalation beginning December 4, 2025. These episodes correspond to the model’s three comparative statics: finer ordering, opportunity-access shocks, and fee-floor increases.

⁷ Consider example tx hash: <https://basescan.org/tx/0x05d5aeccc0ecc11013b62481ea92f47213c200e9ec6bad0c72598d306be40526>.



■ **Figure 2** Weekly successful arbitrage transaction count, USD volume, and spam transaction count by bot architecture label.



■ **Figure 3** Weekly number of active bots per label, weekly active bot label share, and weekly arbitrage volume share per bot label. Vertical lines mark configuration changes of Flashblocks and fee floors.

5.2 Flashblocks

The Flashblocks activation on July 7, 2025, introduced 200ms sub-block pre-confirmations within Base’s 2-second block cycle, and each Flashblock has an incrementally increasing space capacity [2], substantially finer than the full-block ordering. Our model predicts two patterns around this transition: a decline in on-chain discovery participation (cf. Theorem 3); selection against more complex on-chain discovery strategies (cf. Proposition 4) and towards simpler on-chain discovery survivors. We empirically validate these predictions using a four-week symmetric event window around the rollout: pre weeks June 8 – July 6, 2025, post weeks July 13 – August 10, 2025, with the event week itself excluded.

Population and volume shift. The composition of active bots changes sharply after the Flashblocks introduction. The number of total distinct active bots declines by 53%, falling from 983 in the pre-window to 462 in the post-window. The contraction is concentrated almost entirely in on-chain discovery architecture: its share of active bots falls from 63.1% to 25.1%, while off-chain discovery rises from 23.4% to 46.7% and on-chain evaluation rises from 13.4% to 28.2%. As illustrated in Figure 3A, the on-chain discovery population peaks above 300 weekly active in mid-June, declines through the pre-window, and stabilizes around 30 weekly active bots in the post-window. The off-chain discovery population, in contrast, expands steadily from roughly 85 to over 150 weekly active bots by early August.

Figure 3C presents a consistent story in terms of arbitrage volume. On-chain discovery’s share of weekly arbitrage USD volume falls from roughly 50-60% pre-Flashblocks to around 30% post-Flashblocks and remains at that level for the remainder of the observation window. Off-chain discovery expands from roughly 20% to over 50% of volume, and on-chain evaluation remains broadly stable. These movements match Theorem 3: finer ordering reduces participation from on-chain discovery architecture and increases off-chain participation.

■ **Table 3** On-chain discovery cohort composition around Flashblocks. Bot-level metrics first summarize each bot within the relevant window and then summarize across bots. Transaction-level metrics summarize sampled arbitrage transactions and spam bot-week medians within the same four-week pre- and post-event windows.

Cohort	Bots	Bot-level architecture			Transaction-level footprint		
		Median calldata length	Average venue commitment	Median venue commitment	Median pre-reads	Median arbitrage tx gas	Median of avg. spam tx gas
Selected out (pre only)	741	36	0.044	0.000	57	855K	671K
Entrants (post only)	120	244	0.258	0.000	46	630K	272K
Survivors (pre → post)	37	163 → 302	0.345 → 0.340	0.000 → 0.000	39 → 12	662K → 437K	226K → 103K

A natural concern is that on-chain discovery was already declining before Flashblocks, falling from 309 weekly active bots by June 15 to 125 by July 6. This decline likely reflects the unwinding of a May-June opportunity cycle around AI-agent tokens on Base.⁸ However, a purely demand-side explanation is less plausible: on-chain discovery’s volume share remains stable during the pre-window, and around Flashblocks, aggregate arbitrage activity rebounds (cf. Figure 2), yet on-chain discovery bots continue to contract from 179 in the week before Flashblocks to 38 three weeks after.

Selection against on-chain discovery. The composition shift around Flashblocks is driven primarily by selection rather than within-bot label switching: among the 174 bots active in both pre- and post-windows, 167 retain the same label, so the shift mainly reflects differential exit and replacement. The selection margin is sharpest within on-chain discovery. To characterize it, we compare three cohorts within the on-chain discovery architecture: bots active only in the pre-window (*selected out*), only in the post-window (*entrants*), and in both windows under the same label (*survivors*). Table 3 reports the bot-level architecture and transaction trace-level footprint metrics calculated from sampled arbitrages for three cohorts, using only their bot-weeks labeled as on-chain discovery.

The 741 selected-out on-chain discovery bots have the highest scan footprint in the dataset. Their median bot-level calldata length is only 36 bytes, median venue commitment is zero, and their transaction-level footprint is the heaviest, with 57 median pre-reads and the highest gas usage for both arbitrage and spam transactions. In other words, the bots most heavily dependent on short-calldata, low-commitment, high-read scanner logic are precisely the ones that vanish after Flashblocks.

Relative to the selected-out bots, the entrant cohort is markedly less scan-intensive – longer calldata, higher (average) venue commitment, fewer pre-reads, and lower gas usage on successful arbitrages and spam transactions. The survivors are simpler still. Even in the pre-window, surviving on-chain discovery bots exhibit the least search behavior across the three cohorts, suggesting that survival is concentrated among bots whose on-chain search is already relatively lean. Moreover, the survivors become even less scan-intensive after Flashblocks, indicating within-bot adaptation toward leaner scanning.

The same selection and adaptation patterns hold in a two-week window around Flashblocks, suggesting that the result is not driven by the choice of event window or by later post-event weeks. By contrast, the corresponding architecture changes for off-chain discovery and on-chain evaluation population are modest around Flashblocks (see Appendix C.1 in the full version [35]). These results confirm Proposition 4 that finer ordering selects against on-chain discovery with higher per-attempt scan-footprint and towards less scan-intensive ones.

⁸ Examples include KTA, RDAC, VIRTUAL, and AIXBT. These provide context for a contemporaneous opportunity cycle; the empirical results do not depend on this attribution.

■ **Table 4** Attempt efficiency and spam production around Flashblocks. Success rates and gas cost per successful arbitrage are bot-level averages within each window. Spam per active bot-week normalizes spam count by the number of labeled bot-week observations. Bot counts are not mutually exclusive because labels may vary across weeks.

Bot label	Success rate pre → post	Gas cost / success pre → post	Total spam txs pre → post	Spam per active bot-week	Active bots pre → post
Off-chain disc.	35.8% → 43.6%	0.000140 → 0.000113	3.73M → 1.78M (-52.3%)	11.13K → 4.83K	134 → 207
On-chain eval.	26.1% → 23.1%	0.000015 → 0.000027	3.66M → 2.54M (-30.7%)	18.70K → 10.45K	86 → 113
On-chain disc.	5.7% → 3.8%	0.000293 → 0.000350	68.43M → 70.50M (+3.0%)	72.49K → 314.71K	782 → 157

Attempt efficiency and spam concentration. The selection documented above concerns the transaction-level footprint of arbitrage attempts. It does not imply that aggregate spam activity should fall mechanically, because spam also depends on how many failed attempts each bot sends. Table 4 compares spam production and attempt efficiency by architecture over the four-week pre- and post-Flashblocks windows.

The spam response is highly asymmetric. Off-chain discovery and on-chain evaluation both reduce total spam and spam per active bot-week. On-chain discovery, however, remains spam-intensive: despite active bots falling from 782 to 157, total spam rises from 68.43M to 70.50M, and spam per active bot-week rises from 72.49K to 314.71K. Because on-chain discovery dominates spam in both windows, aggregate spam is nearly flat.

Attempt efficiency moves in a similar direction. Flashblocks improves attempt efficiency for off-chain discovery: average success rate rises from 35.8% to 43.6%, and gas cost per successful arbitrage falls from 0.000140 to 0.000113 ETH. This is consistent with reduced stale-state risk between off-chain opportunity discovery and execution (cf. Remark 5). On-chain discovery moves in the opposite direction, with lower success rates and higher gas cost per success. Despite fewer failed attempts, the attempt efficiency of on-chain evaluation declines, as their successful arbitrages fall (1.403M pre-window → 0.855M post-window).

The persistence of on-chain discovery spam reflects a response in attempt intensity among the bots that remain. Flashblocks selects against many broad scanner bots and leaves less scan-intensive survivors, but it does not directly price attempt frequency. By shortening the window in which each opportunity remains available for on-chain discovery, finer ordering can lower per-attempt yield; surviving discovery bots can respond by probing more frequently. In this sense, Flashblocks changes who participates in on-chain discovery and narrows the form of surviving search, but does not by itself eliminate high-frequency probing.

Summary. The Flashblocks event study supports the model’s central mechanism: finer ordering reduces the competitiveness, and therefore the participation, of on-chain discovery, especially among the most scan-intensive bots. The dominant margin is selection: broad scanners exit, same-label survivors are already narrower and simplify further, and the market shifts toward route-committed off-chain discovery. Off-chain discovery becomes more attempt-efficient after Flashblocks, consistent with the stale-state channel: shorter ordering intervals make route-committed execution more reliable. However, this architectural selection does not translate into a proportional reduction in total spam. Aggregate spam remains nearly flat because on-chain discovery spam becomes concentrated among fewer, higher attempt-intensity entrants and survivors. Thus, Flashblocks reduces broad on-chain discovery and improves off-chain execution efficiency, but its effect on aggregate spam is partially offset by higher attempt intensity among the remaining on-chain discovery bots.




Figure 4 AVNT and MIRROR arbitrage activity around the surge window. Panel A shows weekly AVNT/MIRROR volume relative to total volume, with the x-axis showing the week-start date. Panels B and C show AVNT and MIRROR daily volume share, respectively, by bot architecture.

5.3 The September surge

The post-Flashblocks on-chain discovery population stabilizes at roughly 30-50 weekly active bots through August 2025. However, this number spikes sharply in early September: active on-chain discovery bots rise to 220 in the week beginning September 14, before quickly falling back to 62 by September 28 and stabilizing around 35 in early October (cf. Figure 3A). A smaller spike also appears among on-chain evaluation, while off-chain discovery remains comparatively stable. Because the aggregate time series does not show a broad, sustained market-wide expansion (cf. Figure 2), we interpret the episode as a localized entry wave around newly attractive venues rather than a broad expansion in arbitrage opportunities. For the remainder of this section, we refer to the four weekly buckets from September 7–13 through September 28–October 4 as the September surge window.

A short-lived entrant wave. The entrant composition shows that the surge is bimodal: a few serious entrants and a much larger mass of short-lived triers. Across the September surge window, 504 unique on-chain discovery bots are active, of which 464 (92.1%) are both first and last seen within these weeks. Most are active only briefly: 426 bots appear only in one week, and 470 appear in at most two weeks. The same pattern appears in transaction intensity: median transient on-chain discovery bots execute only 29 surge-window arbitrages, compared with 262.5 for non-transient bots (see Appendix C.2 in the full version [35]).

AVNT/MIRROR opportunity wave. The entrant wave is centered on two newly launched tokens: AVNT, launched on September 9 [1], and MIRROR, launched on September 8 [19]. Figure 4A shows that AVNT and MIRROR form a visible share of arbitrage volume from September 7 to September 28, and account for much of the contemporaneous increase in total volume. After September 28, AVNT/MIRROR arbitrage volume falls quickly, coinciding with the disappearance of the temporary on-chain discovery entrant population.

AVNT and MIRROR account for a large share of arbitrage activity during the surge window. From September 7 to October 4, AVNT or MIRROR appears in 420,541 arbitrages (19.3% of all arbitrages) and \$109.1M in volume (19.0%). The concentration is much stronger within on-chain discovery: 42.8% of arbitrages and 50.2% of volume involve AVNT/MIRROR, spanning 412 of the 504 active on-chain discovery bots. Most of this activity is driven by AVNT (\$79.1M vs. \$0.5M for MIRROR).

The concentration is especially strong among short-lived triers. Of the 464 transient on-chain discovery bots, 84.3% execute at least one AVNT/MIRROR arbitrage, and 54.7% of them have a majority (>50%) of their activity in these tokens during the surge window. By contrast, persistent on-chain discovery bots are more diversified: median AVNT/MIRROR share is only 2.0%, and 35 of 40 have a majority of their activity outside these tokens.

The swarm does not persist. Figure 4B–C show that on-chain discovery’s share of AVNT/MIRROR volume falls quickly post-surge, while the other two architectures absorb the bulk of the flow. The opportunity itself does not disappear – AVNT continues to generate arbitrage activity post-surge (Oct 5, 2025 to Feb 28, 2026), with 120,872 transactions across 308 bots, compared with 396,957 transactions across 619 bots during the surge weeks. What disappears is the surge-window on-chain discovery cohort: only 13 of the 504 September on-chain discovery bots remain active on post-surge AVNT flow, and all on-chain discovery bots account for 29.5% of post-surge AVNT volume, while off-chain discovery bots capture the largest share of post-surge AVNT volume (41.4%). The MIRROR pattern is similar, with only the 7 surge-window cohort bots surviving post-surge. Off-chain discovery bots capture the largest 60.9% post-surge MIRROR volume, while on-chain discovery bots fall to 24%.

We interpret this transition as evidence that on-chain discovery is temporarily valuable when new token routes are not yet routinized. This need not mean that existing on-chain discovery contracts discover new pools faster. Instead, early token-launch markets can generate dispersed retail flow, uncertain venue quality, and many short-lived route candidates, making probing attractive. Bots may pass venue parameters through calldata, use generic scanning logic, or deploy cloned contracts specialized to new route families [28]. As the routes become easier to monitor and parameterize off-chain, activity shifts toward on-chain evaluation and off-chain discovery, while most short-lived on-chain discovery entrants exit.

Spam burden. The entrant wave also generates a temporary spam shock: on-chain discovery bots produce 147M spam transactions during the surge window (cf. Figure 2C). Most of the level comes from short-lived entrants. The 464 transient bots generate 105.7M spam transactions, or 72% of on-chain discovery spam. The remaining 40 non-transient bots generate 41.3M spam transactions, but are far more intensive on a per-bot basis and account for more arbitrage activity per bot (403,464 arbitrages across 40 bots, versus 185,997 across 464 transient bots). Thus, the surge combines a mass-entry spam shock from transient bots with a smaller persistent core of high-intensity scanners.

Summary. The September surge provides an example of the model’s opportunity-shock channel (cf. Proposition 8). On-chain discovery briefly revives around two newly launched token families in an otherwise stable post-Flashblocks regime, then recedes as launch-specific activity fades and the remaining flow is increasingly handled by other architectures. The associated spam surge comes mostly from short-lived triers that exit afterward. Since the episode overlaps the September 17 per-transaction gas-limit change [3], we do not interpret it as a clean event study of that policy. A later November rebound contrasts with September: volatility in already active token families mainly expands off-chain discovery, while a smaller on-chain discovery response still generates a visible spam increase (see full version [35]).

5.4 Minimum base fee escalation

Base introduced a minimum base fee on December 4, 2025 and increased it in five steps: 200K WEI on December 4, 500K WEI on December 18, 1M WEI on January 22, 2M WEI on February 2, and 5M WEI on February 19. In the model, the fee floor operates through




Figure 5 On-chain discovery bots active in the pre-step window for each fee increase, plotted by attempts per successful arbitrage and USD volume per successful arbitrage on log scales. Color indicates whether the bot remains active in the post-step window (orange) or is absent (blue). Marker size scales with total volume. The vertical and horizontal dashed lines mark cohort medians. Leftward indicates lower attempt intensity; upward indicates higher value per success, so the upper-left region corresponds to stronger value-cost margins.

the cost channel: architectures with greater fee exposure per captured opportunity become less viable as the fee floor rises. We expect the fee ramp to disproportionately affect on-chain discovery, where repeated probing and low success rates make the value-cost ratio especially fee-sensitive (cf. Proposition 6 and Remark 7). Because the fee changes are tightly spaced, we analyze the overall weekly population shift and use one-week local comparisons around each step to study selection on value-cost economics.

Population shift. The fee ramp coincides with a gradual reallocation toward off-chain discovery. Figure 3B shows on-chain discovery’s active-bot share falling steadily across the ramp, while off-chain discovery rises, especially from late January 2026 onward. By the final sample week, on-chain discovery accounts for 8.4% of active bots and off-chain discovery for 71.1%, with on-chain evaluation comparatively stable. The shift is mainly selection and replacement rather than label switching: only three surviving bots switch labels during the fee escalation regime. The remaining on-chain discovery bots still capture non-trivial arbitrage flow, indicating that a smaller high-productivity core remains viable.

Selection on value-cost margin. Remark 7 predicts that fee floors raise the hurdle rate for on-chain discovery participation: bots whose opportunity access does not justify repeated attempts at higher cost should exit, regardless of their attempt intensity in isolation. Figure 5 plots pre-step on-chain discovery bots in attempts per success and USD volume per success at the 200K, 500K, and 5M fee steps, with marker size proportional to each bot’s total pre-step arbitrage volume. Volume per success proxies for the value side of the value-cost ratio: large-size successful arbitrages imply more gross opportunity value available to cover failed-attempt costs. We discuss the 1M and 2M steps separately because they overlap a mid-ramp opportunity shock (see Appendix C.3 in the full version [35]).

The selection pattern is two-dimensional: neither attempt intensity nor volume per success alone predicts survival, with both surviving and selected-out bots present across most regions of each panel. Bots with high volume per success and low attempts per success tend to survive: they generate value without incurring high failed-attempt costs. Bots with low volume per success and high attempts per success are disproportionately selected out because

their value-cost ratio is unfavorable, and the fee floor pushes them below viability. Bots with both high attempts and high volume per success survive only when their successful opportunities are valuable enough to offset the cumulative cost of repeated attempts; in the 5M panel, this region is anchored by a small cluster of markers operating at \$10,000–100,000 per successful arbitrage. Thus, the fee floor does not select against attempt intensity alone; it selects against attempt intensity unsupported by sufficient opportunity value.

A cohort-level value-per-attempt comparison before and after each fee step, reported in Appendix C.3 in the full version [35], supports this interpretation. At the clean 200K and 5M steps, surviving bots’ aggregate USD volume per attempt rises from 0.36 to 0.87 and from 0.46 to 1.31, respectively, consistent with pruning low-yield attempts or shifting toward higher-value opportunities. The 500K step moves in the opposite direction, but this window coincides with a broad market contraction in late December.

Unlike Flashblocks, the fee ramp does not appear to select for less scan-intensive on-chain discovery attempts at the trace-level. Therefore, broad scanning can still be profitable when opportunities are valuable enough, but low-yield probing becomes harder to sustain. The endpoint trace-level comparison is reported in Appendix C.3 in the full version [35].

Summary. These results support the model’s cost-channel in Proposition 6. As fee floors raise per-attempt costs, on-chain discovery, the highest-exposure architecture, contracts from 22.0% to 8.4% of active bots. The selection is against repeated attempts unsupported by sufficient opportunity value, while broad scanning remains viable when opportunity value is high. Thus, higher fee floors shift the on-chain discovery viability frontier along the value-cost margin, distinct from Flashblocks’ blockspace channel, which selects against heavy venue scanning in early ordering slices.


5.5 Protocol revenue and blockspace efficiency

The previous sections documented how bot architectures responded to Base’s configuration changes. We now ask whether these responses improved chain-level outcomes, including protocol revenue composition and blockspace consumption.⁹ These outcomes speak to both protocol revenue quality and the blockspace externality caused by spam.


Priority fees by architecture. Across the full sample, off-chain discovery successful arbitrages pay the highest median priority fee per transaction, at 8.66×10^{-7} ETH, compared with 4.72×10^{-7} ETH for on-chain evaluation and 5.02×10^{-7} ETH for on-chain discovery. This ordering is consistent with the economic role of off-chain route commitment. Because these bots can simulate routes and optimize input sizes before submitting, they are better positioned to identify higher-value opportunities and bid more aggressively for prioritized inclusion. Thus, the higher priority fees likely reflect both better opportunities and competition for more valuable arbitrages.

Figure 6 plots priority-fee distributions for successful arbitrages before and after Flashblocks by architecture. We note that the pre-Flashblocks window contains 5 weeks, while the post-Flashblocks window contains 33 weeks. The largest distributional change occurs for on-chain discovery. Before Flashblocks, successful on-chain discovery arbitrages exhibit a pronounced low-priority-fee mass, which we interpret as low-willingness-to-pay probabilistic search. After Flashblocks, this low-fee mass shrinks, and the on-chain discovery distribution

⁹ We note that the base fees on Base accrue to the sequencer rather than being burned as on Ethereum.



■ **Figure 6** Priority-fee distributions for successful arbitrages by architecture before and after Flashblocks (3.04M arbitrages pre, 17.40M arbitrages post).



■ **Figure 7** Priority-fee payments by bot architecture and transaction outcome. Panel A reports cumulative priority fees from arbitrage attempts. Panel B reports weekly priority-fee shares. Colors denote bot architecture; solid areas denote successful arbitrages and hatched areas denote spam.

moves closer to the other architectures. This pattern is consistent with Flashblocks selecting out the broadest low-revenue scanners and leaving a smaller population whose attempts are more targeted (cf. Table 3). The same qualitative pattern appears on the spam side (see Appendix C.4 in the full version [35]).


Taken together, the priority-fee evidence suggests that Flashblocks changes not only which bots remain active, but also how surviving on-chain discovery bots bid for inclusion. The low-priority-fee mass of on-chain discovery shrinks after Flashblocks, and the remaining on-chain discovery transactions pay priority fees closer to the other architectures. Thus, the post-Flashblocks on-chain discovery population is smaller but bids more aggressively.

A similar, though less clean, pattern appears across the minimum-base-fee ramp. Median priority fees for successful on-chain discovery arbitrages rise from 4.58×10^{-7} ETH in the pre-ramp endpoint week to 2.69×10^{-6} ETH in the post-ramp endpoint week, consistent with low-value (and therefore low-fee) attempts becoming less viable after repeated attempts become more costly (see Appendix C.4 in the full version [35]). Because we observe only one post-ramp week, we interpret it cautiously.

Protocol revenue composition. We next aggregate priority fees across arbitrage attempts to examine protocol revenue. Figure 7 decomposes cumulative and weekly priority fees by architecture and attempt outcome. The level series in Figure 7A is highly opportunity-driven: protocol revenue accumulates faster during high-volume periods, especially around the October and late-January episodes (cf. Figure 2B). For off-chain discovery and on-chain evaluation bots, cumulative priority fees come predominantly from successful arbitrages

■ **Table 5** Gas usage by architecture.

Architecture	Arbitrage gas share	Spam gas share	Total gas share
Off-chain disc.	0.47%	0.99%	1.45%
On-chain eval.	0.41%	0.34%	0.75%
On-chain disc.	0.44%	19.59%	20.04%
All bots	1.31%	20.92%	22.23%



■ **Figure 8** Arbitrage-related weekly gas usage share.

despite spam transactions being roughly three times more frequent (cf. Table 1). This reflects that their high-value successful captures pay large tips and account for a disproportionate share of cumulative fees. For on-chain discovery, by contrast, spam contributes a large share of cumulative priority fees because the architecture generates so many failed attempts.

Figure 7B shows a post-Flashblocks shift in arbitrage-related priority-fee revenue away from spam and toward successful arbitrages. Before Flashblocks, spam accounts for more than 60% of weekly priority fees; afterward, it falls to roughly 40–45%, even during the September entrant surge when spam transactions nearly double relative to pre-Flashblocks weeks (cf. Figure 2C). The successful-arbitrage share is correspondingly higher after Flashblocks, but remains opportunity-driven: it rises during high-volume periods, especially around the October episode, and declines as aggregate arbitrage volume falls (cf. Figure 2B). Thus, spam continues to contribute to nontrivial protocol revenue, but it becomes a less dominant component after the ordering granularity changes.

The fee-ramp period is harder to interpret because policy changes overlap with low-activity periods and a mid-ramp opportunity shock (cf. Figure 2B). Even so, the successful-arbitrage share rises following the 500K step, despite volume contraction, consistent with selection against low-value, high-attempt-intensity bots. The pattern becomes less clear around the 1M and 2M steps, when the opportunity shock temporarily revives on-chain discovery.

Blockspace efficiency. Despite protocol revenue contributions, failed attempts impose a burden on blockspace. Table 5 documents the gas usage by architecture as a share of total Base gas over the sample period. Although on-chain discovery bots only capture roughly half of successful arbitrages of the other two architectures (cf. Table 1), their arbitrages consume comparable blockspace because they are heavier (cf. Figure 1). Successful arbitrage transactions across all architectures consume only 1.31% of total gas on Base. The blockspace burden is instead driven almost entirely by spam, especially on-chain discovery spam, which accounts for roughly 20% of total Base gas.

Figure 8 decomposes weekly gas usage into arbitrage-related gas and other chain activity. Before Flashblocks, on-chain discovery spam alone consumes roughly one-third to nearly one-half weekly gas. Flashblocks reduces this burden, although not immediately. As shown in Section 5.2, surviving and entering on-chain discovery bots become more attempt-intensive in the four-week post window. The gas share declines further after that window to 15.2%, as a few extremely intensive bots disappear or downscale (see Appendix C.4 in the full version [35]).

Later episodes show the same distinction between architecture selection and opportunity shocks. During the September surge and the November rebound period, spam transaction counts rise substantially, but spam does not return to its early-sample blockspace share.

Through the fee ramp, spam gas share declines further, consistent with higher fees making low-yield probing harder to sustain. We interpret this cautiously because December and early January are periods of low market-wide activity (cf. Figure 2B), and the late-January opportunity shock temporarily raises spam gas even under a higher fee floor.

Taken together, Base’s configuration changes improve blockspace efficiency in a narrow but important sense. Spam continues to occur, while its share of chain gas falls after Flashblocks and declines further through the fee-ramp period, aside from temporary opportunity shocks.

6 Concluding Discussion

Our paper shows that arbitrage competition on high-throughput blockchains involves economically distinct search architectures with different responses to protocol design. A reduced-form model studies how ordering granularity, fee floors, and opportunity-access shocks affect the relative viability of targeted and probabilistic search. Empirical evidence from Base supports these predictions, and traces how protocol changes reshape bot composition, execution behavior, priority-fee payments, and blockspace consumption.

The broader lesson is that protocol design affects not only how much arbitrage MEV is extracted, but also how it is pursued on-chain. Failed attempts are not merely isolated transaction-level inefficiencies; they arise from equilibrium search behavior under incomplete information, low attempt costs, and specific ordering environments. Opportunity shocks further show that protocol changes do not mechanically guarantee permanent reductions in failed attempts. The externality these attempts impose is blockspace congestion that displaces other users. Understanding how search architectures emerge across protocol designs is important for reducing externalities without suppressing useful economic activity such as arbitrage. We conclude by outlining several open questions for future work.

Does the search-architecture framework apply to other forms of MEV? Our analysis focuses on cyclic arbitrage, where the route discovery is visible in transaction traces. Other MEV forms may admit the same dichotomy. Liquidations can be targeted, when searchers wait for oracle updates against known positions, or probabilistic, when they repeatedly probe for liquidatable positions. Cross-domain arbitrage may similarly vary in how information is gathered and how transactions are submitted across domains. Future work could test whether these MEV forms exhibit analogous search architectures, how they respond to protocol changes, and what externalities they impose.

Can the model’s predictions be tested on other blockchains? A natural next step is to compare chains with different execution environments, fees, and information availability. Ethereum likely favors targeted search because failed attempts are costly, and mempool visibility enables off-chain simulation. Arbitrum may also favor targeted search, but for a different reason: its fast block time reduces stale-state risk, consistent with evidence that Arbitrum bots appear less probe-like than bots on Base and Optimism [28]. Optimism is a closer Base comparison because both use the OP Stack [23] and share similar ordering and configuration choices, including Flashblocks. Outside the Ethereum ecosystem, Solana offers a useful contrast as a high-throughput chain with active probabilistic searching [6] and a different information environment, where transactions are forwarded directly to validators who may share pending-transaction information with searchers.

How would priority-access mechanisms affect architecture competition? Our framework can also be extended to ordering mechanisms that allocate priority explicitly. Such mechanisms change the margin of competition from simply reaching the earliest ordering unit to deciding whether priority access is worth paying for and how aggressively to use it. Arbitrum’s Timeboost [21], which sells priority access, tests one such margin: priority access may favor targeted searchers that can forecast and monetize latency advantages, while leaving probabilistic search viable only for operators with enough opportunity flow to justify the cost. Existing evidence that express-lane access is concentrated and that roughly 22% of time-boosted transactions revert suggests that priority access alone does not eliminate spam [18]. Optimism’s recent experiment with stake-based priority ordering [22] provides another useful setting for studying how stake requirements, first-come-first-served ordering, and priority-fee multipliers affect the relative viability of targeted and probabilistic search.

How should spam mitigation differ across architectures? Failed attempts from different architectures call for different mitigation tools. Revert protection can reduce failed route-committed attempts [9, 30], but it does not eliminate early-terminating targeted profitability checks or probabilistic probe spam, where transactions themselves are used to discover opportunities. Reducing probe-like spam requires improving information availability or reducing the incentive to use transactions as information-gathering devices, for example, through richer real-time information channels [15], protocol-run backrunning [24] and privacy-preserving backrunning auctions [10], and charging based on reserved execution gas [32]. On high-throughput chains, these mechanisms face a tight tradeoff between information freshness, auction efficiency, latency, and user protection.

References

- 1 AVANTIS. Avantis launch. URL: <https://x.com/avantisfi/status/1965415885243556092>.
- 2 Base. Base flashblocks configurations, 2025. URL: <https://docs.base.org/base-chain/network-information/block-building#flashblocks>.
- 3 Base. Configuration changelog, 2026. URL: <https://docs.base.org/base-chain/network-information/configuration-changelog>.
- 4 Base. Flashblocks Overview. <https://docs.base.org/base-chain/flashblocks/overview>, 2026. Base Documentation. Accessed: 2026-05-12.
- 5 Base. Tweet announcing base minimum fee increase. <https://x.com/buildonbase/status/2024306165711327610>, 2026.
- 6 Eugene Chen and @0xEDgar. MEV on Solana. <https://www.umbraresearch.xyz/writings/mev-on-solana>, 2023. Umbra Research. Accessed: 2026-05-10.
- 7 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)*, pages 910–927, 2020. doi:10.1109/SP40000.2020.00040.
- 8 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 *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, CCS ’24*, pages 2591–2605, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3658644.3690259.
- 9 Flashbots. MEV Protection Overview. <https://docs.flashbots.net/flashbots-protect/overview>, 2025. Flashbots Documentation. Last updated: 2025-11-04. Accessed: 2026-05-13.
- 10 Flashbots. MEV-Share Introduction. <https://docs.flashbots.net/flashbots-mev-share/introduction>, 2025. Flashbots Documentation. Last updated: 2025-11-04. Accessed: 2026-05-13.

- 11 Flashbots and Robert Miller. Mev and the limits of scaling. URL: <https://writings.flashbots.net/mev-and-the-limits-of-scaling>.
- 12 Krzysztof Gogol, Johnnatan Messias, Deborah Miori, Claudio Tessone, and Benjamin Livshits. Cross-rollup mev: Non-atomic arbitrage across l2 blockchains, 2024. [arXiv:2406.02172](https://arxiv.org/abs/2406.02172).
- 13 Krzysztof Gogol, Manvir Schneider, and Claudio Tessone. When priority fails: Revert-based mev on fast-finality rollups, 2025. [arXiv:2506.01462](https://arxiv.org/abs/2506.01462).
- 14 Lioba Heimbach, Vabuk Pahari, and Eric Schertenleib. Non-Atomic Arbitrage in Decentralized Finance . In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 3866–3884, Los Alamitos, CA, USA, May 2024. IEEE Computer Society. doi:10.1109/SP54263.2024.00256.
- 15 Quintus Kilbourn. Dealing with spam caused by on-chain searching. <https://collective.flashbots.net/t/dealing-with-spam-caused-by-on-chain-searching/3381>, 2024. The Flashbots Collective, Research forum post. Accessed: 2026-05-13.
- 16 Bruno Mazorra, Christoph Schlegel, and Akaki Mamageishvili. Timing games: Probabilistic backrunning and spam, 2026. doi:10.48550/arXiv.2602.22032.
- 17 Robert McLaughlin, Christopher Kruegel, and Giovanni Vigna. A large scale study of the ethereum arbitrage ecosystem. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 3295–3312, Anaheim, CA, August 2023. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/mclaughlin>.
- 18 Johnnatan Messias and Christof Ferreira Torres. The express lane to spam and centralization: An empirical analysis of arbitrum’s timeboost, 2025. doi:10.48550/arXiv.2509.22143.
- 19 MIRROR. Mirror tokenomics & distribution. URL: https://x.com/blackmirror_xp/status/1964605903824703813.
- 20 Alexandre Obadia, Alejo Salles, Lakshman Sankar, Tarun Chitra, Vaibhav Chellani, and Philip Daian. Unity is strength: A formalization of cross-domain maximal extractable value, 2021. [arXiv:2112.01472](https://arxiv.org/abs/2112.01472).
- 21 Offchain Labs. A Gentle Introduction: Timeboost. <https://docs.arbitrum.io/how-arbitrum-works/timeboost/gentle-introduction>, 2026. Last updated: 2026-02-17; Accessed: 2026-05-08.
- 22 Optimism. Stake-based priority ordering on OP Mainnet. <https://docs.optimism.io/notes/stake-based-priority-ordering>, 2026. Accessed: 2026-05-08.
- 23 Optimism. The OP Stack. <https://docs.optimism.io/op-stack/introduction/op-stack>, 2026. Optimism Documentation. Accessed: 2026-05-21.
- 24 Osmosis. Protorev. <https://docs.osmosis.zone/overview/features/protorev/>, 2026. Osmosis Docs. Accessed: 2026-05-13.
- 25 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), December 2025. doi:10.1145/3771566.
- 26 Kaihua Qin, Liyi Zhou, Pablo Gamito, Philipp Jovanovic, and Arthur Gervais. An empirical study of defi liquidations: incentives, risks, and instabilities. In *Proceedings of the 21st ACM Internet Measurement Conference, IMC ’21*, pages 336–350, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3487552.3487811.
- 27 Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 198–214, 2022. doi:10.1109/SP46214.2022.9833734.
- 28 Ozan Solmaz, Lioba Heimbach, Yann Vonlanthen, and Roger Wattenhofer. Optimistic MEV in Ethereum Layer 2s: Why Blockspace Is Always in Demand. In Zeta Avarikioti and Nicolas Christin, editors, *7th Conference on Advances in Financial Technologies (AFT 2025)*, volume 354 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:24, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2025.28.

- 29 Christof Ferreira Torres, Ramiro Camino, and Radu State. Frontrunner jones and the raiders of the dark forest: An empirical study of frontrunning on the ethereum blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1343–1359. USENIX Association, August 2021. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/torres>.
- 30 Uniswap Labs. Live on Unichain: Fair Transaction Ordering and MEV Protection. <https://blog.uniswap.org/rollup-boost-is-live-on-unichain>, 2025. Accessed: 2026-05-08.
- 31 Qin Wang, Ruiqiang Li, Guangsheng Yu, Vincent Gramoli, and Shiping Chen. Mev in binance builder, 2026. doi:10.48550/arXiv.2602.15395.
- 32 Wenhao Wang, Aditya Saraf, Lioba Heimbach, Kushal Babel, and Fan Zhang. Blockspace under pressure: An analysis of spam mev on high-throughput blockchains, 2026. doi:10.48550/arXiv.2604.00234.
- 33 Ye Wang, Yan Chen, Haotian Wu, Liyi Zhou, Shuiguang Deng, and Roger Wattenhofer. Cyclic arbitrage in decentralized exchanges. In *Companion Proceedings of the Web Conference 2022, WWW '22*, pages 12–19, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3487553.3524201.
- 34 Fei Wu, Danning Sui, Thomas Thiery, and Mallesh 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)*, volume 354 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:23, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2025.26.
- 35 Fei Wu and Burak Öz. To wait or to probe: Arbitrage competition on high-throughput blockchains, 2026. doi:10.48550/arXiv.2606.00720.
- 36 Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 428–445, 2021. doi:10.1109/SP40001.2021.00027.
- 37 Brian Z. Zhu, Xin Wan, Ciamac C. Moallemi, Dan Robinson, and Brad Bachu. Quantifying the value of revert protection. In Christina Garman and Pedro Moreno-Sanchez, editors, *Financial Cryptography and Data Security*, pages 72–88, Cham, 2026. Springer Nature Switzerland.
- 38 Burak Öz, Jonas Gebele, Parshant Singh, Filip Rezabek, and Florian Matthes. Playing the mev game on a first-come-first-served blockchain. In *2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 220–224, 2024. doi:10.1109/ICBC59979.2024.10634397.