# Blockspace Under Pressure: An Analysis of Spam MEV on High-Throughput Blockchains

Wenhao Wang*
Yale University, IC3

Aditya Saraf*
Cornell University, IC3

Lioba Heimbach
Category Labs

Kushal Babel
Category Labs

Fan Zhang
Yale University, IC3

## Abstract

On high-throughput, low-fee blockchains, a qualitatively new form of maximal extractable value (MEV) has emerged: searchers submit large volumes of speculative transactions, whose profitability is resolved only at execution time. We refer to this as *spam MEV*. On major rollups, it can at times consume more than half of block gas, even though only a small fraction of probes ultimately results in a trade. Despite growing awareness of this phenomenon, there is no principled framework for understanding how blockchain design parameters shape its prevalence and impact.

We develop such a framework, modeling spam transactions competing for on-chain opportunities under a competitive equilibrium that drives their profits to zero, and deriving equilibrium spam volumes as a function of block capacity, minimum gas price, and the transaction fee mechanism. Empirical evidence from Base and Arbitrum supports the model: spam grew sharply as block capacity was scaled up and fell when minimum gas prices were introduced. Our analysis yields three main insights. First, spam is always costly: when block capacity is scarce, it displaces users and drives up gas prices; as block capacity grows, it increasingly consumes execution resources, raising network externality, i.e., the cost of provisioning and processing blocks. We show that spam takes an increasing share of each additional unit of block capacity, so capping it before all users are included creates a favorable trade-off: forgoing a small amount of user welfare eliminates disproportionate spam and externality. Second, we extend the analysis to approximate priority fee ordering and show that this reduces spam, as spammers pay more to reach early block positions. Third, as user demand grows and blockspace is scaled accordingly, spam's share of block capacity plateaus rather than growing indefinitely.

## CCS Concepts

• **Security and privacy → Blockchain**.

## Keywords

Spam MEV, Maximal Extractable Value, Blockchain

## 1 Introduction

Maximal extractable value (MEV) has been a central topic in blockchain research over the past several years [7, 21, 52]. Much of this work has focused on the Ethereum Layer 1, which, for a long time, was the primary home of decentralized finance (DeFi) and thus the main arena for MEV extraction. On the Ethereum Layer 1, the most widely studied MEV strategies, such as sandwich attacks, cyclic arbitrage, liquidations, and non-atomic arbitrage, are precise and targeted: a searcher identifies a specific on-chain opportunity through off-chain computation, constructs a transaction to capture it, and submits it with high confidence of success. These strategies rely on the ability to observe pending transactions in a mempool, i.e., the public waiting area for transactions awaiting inclusion, or to monitor state changes with sufficient time to react. Today, competition among searchers on Ethereum is mediated by block builder auctions [33], in which searchers bid for the right to capture an opportunity. Only the winning transaction is included on-chain; losing bids are largely filtered out before execution, keeping most failed MEV attempts off the blockchain.

A qualitatively different form of MEV has recently emerged on high-throughput blockchains. Rather than targeting a specific opportunity identified off-chain, searchers submit high volumes of speculative transactions whose profitability is resolved only at execution time. We refer to this class of strategies as **spam MEV**. Related work has called this phenomenon *optimistic MEV* [53] or *probabilistic MEV* [43], emphasizing the searcher's uncertainty about whether a transaction will be profitable. Our terminology instead emphasizes that they consume shared block space and infrastructure resources, whether or not each transaction succeeds.

In spam MEV, both the detection and execution of opportunities reside largely in on-chain smart-contract logic [53]: a transaction probes whether a profitable opportunity exists at the moment of execution and, if so, captures it. When no opportunity is found, the transaction consumes gas, i.e., the computational cost of executing transactions on the blockchain, searching without ever executing a trade. This encompasses a range of speculative strategies, though cyclic arbitrage is the most prominently studied among them. To illustrate the contrast: as studied on the Ethereum Layer 1, a cyclic arbitrage bot precomputes a profitable path and executes a single atomic transaction, whereas on many high-throughput chains, the same type of bot speculatively submits transactions without knowing whether an opportunity exists, repeatedly probing liquidity pools in the hope of capturing small gains and accepting a high failure rate in exchange for marginal profits.

Spam MEV creates a different security concern from the targeted MEV studied on Ethereum Layer 1. Classic MEV can harm users through transaction-ordering manipulation and has been linked to consensus-instability concerns [21, 52]. Spam MEV, on the other hand, is closer to the traditional problem of spam in communication systems: each individual message or transaction may be valid, but high-volume unsolicited traffic abuses a shared resource [24]. In

*Part of this work was conducted during the Summer 2025 Research Internship at Category Labs.

blockchains, the shared resource is block space and the infrastructure needed to execute, publish, and index transactions. Therefore, even when spam MEV does not threaten consensus safety directly, it can reduce the effective throughput available to genuine users and impose costs on the broader blockchain infrastructure, creating an availability and resource-abuse problem that the ecosystem is currently trying to understand and address.

First documented on Solana [58], spam MEV has since been identified on Ethereum Layer 2 rollups, including Base, Optimism, and Arbitrum [32, 53]. The scale of the phenomenon is striking: on Base and Optimism, spam MEV transactions consumed over 50% of block gas in Q1 2025 [53], yet only 6–12% of these speculative probes result in an actual trade. On top of that, despite consuming more than half of on-chain gas, spam MEV transactions pay less than a quarter of total fees [53], as they carry low priority fees.

What makes spam MEV viable and prevalent? A combination of architectural and economic features common to modern high-throughput chains likely plays a role. Low transaction fees reduce the cost of failed speculative transactions, making repeated probing profitable even at low success rates. Fast block times, often below one second on rollups, can leave insufficient time for searchers to observe state changes and submit targeted transactions before the next block is produced, favoring continuous speculative submission over deliberate off-chain computation. Similarly, most high-throughput chains lack a public mempool, e.g., rollups use centralized sequencers and Solana forwards transactions directly to the current block producer, limiting the information available for targeted extraction. Designs such as encrypted mempools [4, 8, 27], when deployed, are likely to only increase spam MEV by further reducing the information available for targeted MEV strategies.

While the prevalence of spam MEV has been documented, there is no principled framework to study the interplay between design parameters and spam volumes, and the impact of spam on various stakeholders. As a result, the response from various blockchain designers has been largely reactive and ad-hoc (c.f. Section 1.1).

In this work, we fill this gap by developing a framework to determine the equilibrium volume of spam as a result of block space and fee design choices, and to analyze the impact of resulting spam on user welfare, validator revenue, and network externality, i.e., the cost of provisioning block capacity and processing transactions.We characterize the tradeoffs involved in design choices and provide guidance for blockchain designers to navigate these tradeoffs.

In the remainder of this section, we first highlight the industry's response to spam MEV, describe the design levers and metrics we focus on in this work, and then summarize our contributions.

## 1.1 Ecosystem Response

The prevalence of spam MEV has not gone unnoticed. The blockchain community has become increasingly aware of the scale of the problem, and a debate has emerged about how harmful spam truly is and what should be done about it. Flashbots' thesis that MEV fundamentally limits scaling [44], i.e., that added throughput capacity is absorbed by spam rather than passed on to genuine users, has been particularly influential in framing the discussion.

The most visible case study is Base, the Ethereum Layer 2 rollup with the highest DeFi total value locked (TVL) [23]. Following the Dencun upgrade, which drastically reduced Layer 2 data availability costs, Base increased its block gas limit. However, the additional capacity was largely consumed by spam MEV rather than benefiting genuine users [32, 53]. As Figure 1a shows, spam gas grew disproportionately faster than non-spam gas as Base progressively raised its gas target after Dencun, absorbing the majority of the added capacity. When the community took notice [32, 44, 53], Base reduced the gas target from 70M to 50M. Figure 1b shows that comparing the 30 days before and after the change, spam gas fell by 34% and non-spam gas by 24%, i.e., both declined but spam absorbed a larger share of the reduction. Beyond block space consumption, spam also imposes costs on the broader network infrastructure. Multiple blockchain indexing services, including Etherscan, announced they would discontinue free API access for Base, citing the significant infrastructure costs driven by the high transaction volumes [26, 38, 39]. Base ultimately raised its minimum transaction fee, explicitly citing spam reduction as the rationale [9]. As Figure 2 shows, spam gas on Base trended downward thereafter, coinciding with successive increases in the protocol minimum gas price, though other factors may have also contributed.

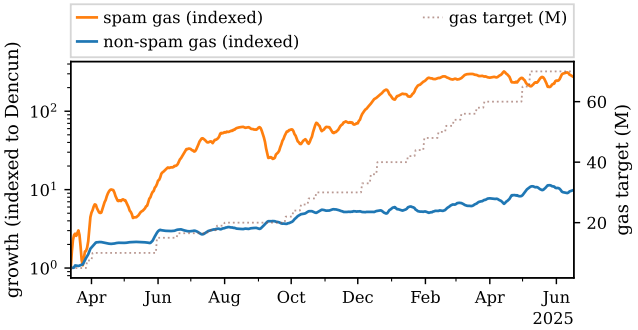


**(a) Spam and non-spam gas on Base indexed to Dencun (March 13, 2024). Spam gas grew 122× from Dencun to its peak (April 2025), while non-spam gas grew 11.2×.**

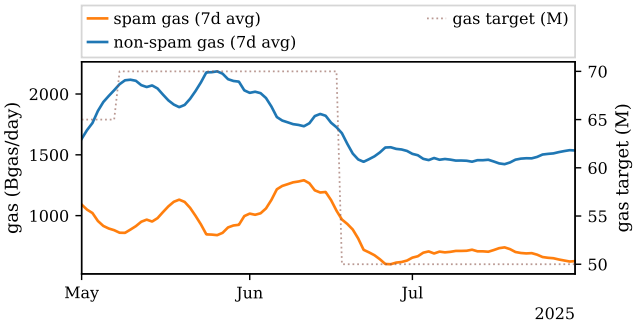


**(b) Spam and non-spam gas on Base around the gas target reduction from 70M to 50M on June 18, 2025. In the 30 days after the change, spam gas fell by 34% and non-spam gas by 24%; spam's share of total gas dropped from 36% to 32%.**

**Figure 1: Spam and non-spam gas on Base over time (7-day moving averages). The right axes show the gas target.**

Other chains have moved in the same direction. Arbitrum, the rollup with the second-highest DeFi TVL [23], had set a minimum

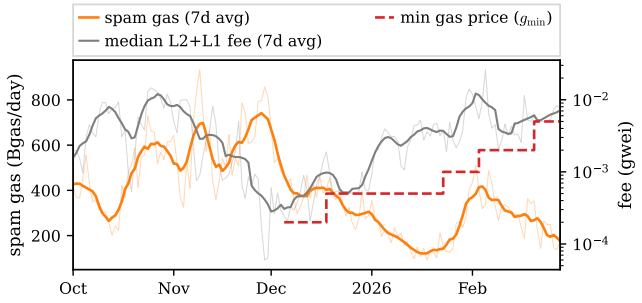


**Figure 2: Daily spam gas on Base from October 2025 onward (7-day moving average), alongside the median fee and protocol minimum gas price (log scale, right axis). Average spam gas fell from 450 Bgas/day in October 2025 to 302 Bgas/day in February 2026 (−33%).**

gas price meaningfully above 1 wei well before spam attracted broader attention, and recently doubled this floor as part of a broader fee mechanism overhaul [6, 62]. Outside the EVM ecosystem, Aptos similarly raised its minimum gas fee, citing spam reduction [5]. Newer chains have launched with spam in mind: Monad, for instance, sets a non-trivial minimum gas price and charges based on the gas limit rather than gas consumed, making spam MEV with large unused gas allocations more costly [16].

These interventions share a common intuition: spam is an unwanted form of resource abuse that reduces useful throughput and imposes costs on the network. Yet the responses have been largely reactive and chain-specific, without formal analysis of how design parameters interact to determine the equilibrium level of spam. A key factor is that, unlike Ethereum L1 where congestion pricing naturally emerges, high-throughput chains can have slack capacity where the market-clearing gas price approaches zero, while the network still bears the cost of processing every transaction. A minimum gas price $g_{\min}$ ensures that each transaction pays at least a meaningful contribution toward its marginal cost and, as our analysis shows, serves as a key lever against spam in this regime.

## 1.2 Design Levers

We now describe the blockchain design choices that influence the presence and characteristics of spam transactions. These are the parameters that our framework analyzes.

**Block Space Limit $B_{\max}$.** The block space limit affects inclusion costs. When the limit $B_{\max}$ is low relative to transaction demand, competition for inclusion in the block intensifies, increasing gas prices $g$ and may therefore reduce the expected revenue of each spam transaction. Conversely, when $B_{\max}$ is large enough to accommodate all transactions, gas prices $g$ decrease, and therefore increase the expected profit of spam transactions.

**Minimum Gas Price $g_{\min}$.** A blockchain can impose a minimum gas price $g_{\min}$ per transaction, independent of block space availability. Higher values of $g_{\min}$ increase the cost when block capacity is slack, influencing the economic outcome for both spam and non-spam transactions. Without a gas price floor, the equilibrium gas price $g$ would approach zero when block capacity exceeds demand, resulting in a disproportionately high volume of spam transactions.

**Transaction Fee Mechanism.** The design of the transaction fee mechanism (TFM) impacts transaction submission behavior. For instance, spam transactions often specify a large gas limit but typically consume only a fraction of it when executed without profitable arbitrage opportunities [53]. Therefore, a TFM that imposes higher costs for specifying large gas limits can change the spam MEV strategy. Additionally, a TFM can dictate the ordering of transactions within a block, such as through priority-based ordering, which affects the resulting spam volumes.

## 1.3 Our Framework and Contributions

We develop a framework for analyzing spam MEV and its interaction with blockchain design choices. We model spam transactions as competing for on-chain opportunities, and study how equilibrium spam volume depends on the block space limit $B_{\max}$, the minimum gas price $g_{\min}$, and the transaction fee mechanism. We evaluate the impact of spam on user welfare, validator revenue, and network externality. Our main results are:

(1) **Closed-form equilibrium spam volume and gas prices** (Section 3.1): We characterize spam equilibrium under random transaction ordering. Depending on block capacity and the minimum gas price, the outcome is either no spam entry, a slack equilibrium pinned down by $g_{\min}$, or a congested equilibrium in which spam raises the gas price.
(2) **User welfare, revenue and externality characterization along with trade-offs** (Section 3.2): We analyze user welfare, validator revenue, and network externality in equilibrium.
(3) **Parameter-setting guidance** (Section 3.3): We characterize the choice of $(B_{\max}, g_{\min})$, and show how to use the proportion of spam at the margins as a practical rule for selecting the design parameters without admitting disproportionate spam.
(4) **Priority fee ordering (PFO)** (Section 4): We extend the analysis to an approximate PFO TFM. We show that PFO can reduce spam by making earlier positions in the block more expensive.
(5) **Demand-scaling analysis** (Section 5 and Section 6): We study whether spam limits blockchain scaling [44]. We estimate that MEV opportunity size grows approximately linearly with non-spam activity using empirical data. We show that as demand grows, spam remains a nontrivial share of included gas, but its share plateaus rather than growing without bound.
(6) **Empirical evidence and case studies** (Section 6): We present case studies of spam on Base and Arbitrum. Our findings support several model predictions: spam falls when gas price floors rise, grows with block capacity, absorbs disproportionate marginal block space, and, on Base, appears later in the block as predicted by the PFO analysis.
(7) **Mitigations** (Section 7): We discuss practical mitigations with the insights from our model and analysis.

# 2 Model

We present our model for analyzing spam MEV. The design levers (block space limit $B_{\max}$, minimum gas price $g_{\min}$, and transaction fee mechanism) were introduced in Section 1.2. We consider blockchains run by validators; we refer to the parties that send transactions as

users. There are two types of users in the model: genuine users and spam bots. Genuine users submit transactions because they receive intrinsic utility from execution. Spam bots, by contrast, are strategic agents that submit transactions in order to search and capture MEV opportunities that are determined at the time of execution. They have no intrinsic utility, and their transaction volume is determined endogenously in equilibrium.

We analyze TFMs that charge a transaction for its gas limit rather than gas used, and block size is counted in terms of transaction gas limits, similar to high-throughput blockchains such as Solana and Monad [16]. Our results can be straightforwardly adapted to a mechanism that charges only for the used gas. For our initial analysis in Section 3, the TFM orders transactions in a random order, mimicking the operations of a low-latency blockchain. We extend our analysis to priority-fee based ordering in Section 4.

Now we formalize block space and spammers' utility model and define the metrics we use to evaluate equilibrium outcomes.

## 2.1 Block Demand and Spam Model

We analyze transaction dynamics within a single block. Let $B \leq B_{\max}$ denote the amount of block space actually occupied in the block. Let $Q_u \leq B$ represent the amount of included user gas, i.e., gas consumed by genuine (non-spam) transactions. The demand from genuine users follows a demand function $D(g)$, which specifies how much genuine-user gas is demanded at any given gas price $g$. This demand curve does not include spam demand. Each spam transaction carries a gas limit of $s$, so if $S$ spam transactions are included, occupied block space decomposes as $B = Q_u + sS$. Equivalently, $S = (B - Q_u)/s$ in terms of $B$ and $Q_u$.

In the TFM model used in Section 3, the TFM gives a clearing price $g^* \geq g_{\min}$ that is paid by all included transactions. In the PFO model in Section 4, the same logic applies sub-block by sub-block, with different clearing prices for different execution positions. A single arbitrage opportunity worth $r$ is created by some random genuine user transaction and can be successfully claimed by a spam transaction if it is the first spam transaction to be sequenced after the particular user transaction. We model the opportunity size $r$ to be increasing in the amount of included user gas in Section 5.

We assume a free-entry environment for spam bots: any number of agents can submit spam transactions. There is no fixed cost of becoming a spammer; the spammer only pays for the transaction fees of submitted spam transactions. Strategic spam bots enter as long as their expected net payoff is non-negative. At equilibrium, included spam transactions earn zero expected net profit (full rent dissipation [37, 51]), and an additional spam transaction cannot profitably enter. Given a TFM and an ordering rule, the clearing price and the amount of spam are endogenously determined by genuine user demand function and this zero-profit entry condition.

The utility of a spam transaction is the revenue it captures ($r$ if successful and 0 otherwise) minus the fees $s \cdot g^*$ it pays. The utility of a genuine user transaction is defined as $u - f$ if included in the block, and 0 otherwise, where $u$ is its valuation of execution and $f$ is the fee it pays. For included genuine users, $f \leq u$. Looking ahead, user valuations are represented by the demand function, while fees are determined by the TFM.

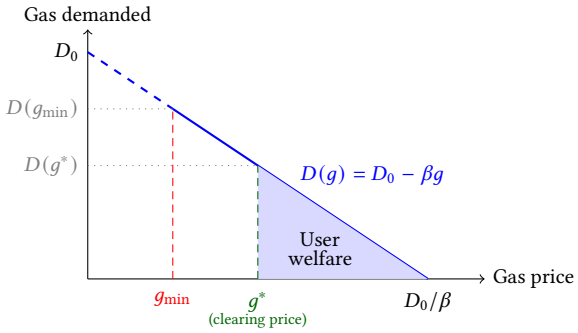


**Figure 3: Linear demand curve $D(g) = D_0 - \beta g$ for genuine users. At clearing price $g^*$, the shaded triangle is the user welfare (aggregate surplus of included users).**

## 2.2 Metrics

We now describe the metrics used to quantify the impact of spam transactions on the blockchain. In addition to the welfare and cost metrics below, we analyze the *spam share* of included gas, defined as the fraction of included gas consumed by spam transactions. A higher spam share means that a larger fraction of the block is unavailable to genuine users. It is therefore a direct measure of the resource-abuse aspect of spam MEV.

**User welfare.** The user welfare is defined as the aggregate utility of all genuine users. As shown in Fig. 3 for a linear demand curve, the user welfare is the area under the genuine-user demand curve above the clearing price $g^*$. A higher user welfare indicates a greater overall benefit to users of the blockchain. Spam can reduce user welfare by reducing block space for genuine users, or by increasing the inclusion price and pricing out users with lower valuations.

**Validator revenue.** Validator revenue is the total gas fees collected from all included transactions (both users and spammers), written as $R = g^* \cdot B$, where $g^*$ is the equilibrium clearing gas price and $B$ is total gas used. This quantity represents the direct income to the network from transaction processing. For the purposes of this work, we do not consider burning of fees or block rewards, as they are merely a way to change allocation of revenue between token holders and validators, and do not affect the net economic incentives for spamming or the overall welfare analysis.

**Network externality.** The network externality is the cost borne by the network for provisioning block capacity and processing transactions, written as $E(B_{\max}) = c_1 B_{\max} + c_2 B$, where $c_1$ is the per-unit cost of providing block capacity (e.g., bandwidth provisioning, data availability, node hardware), and $c_2$ is the per-unit cost of computation (e.g., execution, state access). The capacity component $c_1 B_{\max}$ depends on the provisioned block capacity, not actual usage, while the compute component $c_2 B$ scales with actual gas consumed. Network externality captures costs such as increased barriers to decentralization and additional pressure on infrastructure operators (e.g., full nodes), which are not easily quantifiable but are important considerations in blockchain design, especially as it relates to spam transactions. In this sense, externality captures the security and

reliability cost of allowing high-volume traffic that consumes shared blockchain resources. [1]

## 3 Analysis with Random Ordering

In this section, we study a simple random-ordering TFM, which is intentionally minimal. There is a block of capacity $B_{\max}$, a single opportunity of value $r$, and a volume of spam transactions that try to claim that opportunity. Each spam transaction reserves $s$ gas and is charged based on its gas *limit*, not on ex post gas used. Unless specified otherwise, we use the simple linear demand function ($D(g) = D_0 - \beta g$) for non-spam users' gas consumption to expose the equilibrium characteristics. Recall that $Q_u$ denotes the amount of user gas included in the block. In this section, we assume that the opportunity size $r$ is linear in the included user gas, given by $r = \frac{Q_u}{D_0} \cdot r_0$. Finally, we impose a floor $g_{\min}$, which lets us study the design choice of a minimum inclusion price.

### 3.1 Spam Volume at Equilibrium

The first quantity we need is the probability that spam transactions claim the opportunity. Suppose there are $S$ spam transactions in the block. Only the relative position of the opportunity among these $S$ spam transactions matters. For example, when $S = 3$, we can visualize the situation by placing the three spam transactions as triangles and considering the four possible slots of the opportunity:


The circles represent the possible positions where the opportunity can appear relative to the spam transactions. In the first three positions, there is at least one spam transaction after the opportunity, so a spam transaction can claim it. In the last position, the opportunity appears after all spam transactions, so it cannot be claimed. More generally, there are $S + 1$ equally likely relative positions for the opportunity, and exactly one of them fails. Therefore,

$$\Pr[\text{opportunity claimed by spam}] = \frac{S}{S+1}.$$

We next compute the clearing price at a fixed spam volume. Recall that the user demand curve is characterized by $D = D_0 - \beta \cdot g$ where $g$ is the gas price. In the absence of spam, the market-clearing price implied by the demand curve and block size $B_{\max}$ is $g_1 := \frac{D_0 - B_{\max}}{\beta}$. However, since the protocol enforces a price floor $g_{\min}$, the actual clearing price without spam is $\max\{g_{\min}, g_1\}$. If $S$ spam transactions are included in the block and each reserves $s$ gas, then the remaining space available to non-spam users is $B_{\max} - Ss$. The resulting clearing price is

$$g(S) = \max\left\{g_{\min}, g_1 + \frac{s}{\beta}S\right\}.$$

The term $\frac{s}{\beta}S$ captures the price increase caused by spam taking up block space, and the outer maximum applies the gas price floor.

[1] Buterin [13],drawing analogy to pollution, puts forth a similar approach to modeling the externality of block sizes on the ecosystem, where the externality is linear in block size for modern blockchains.

Given this price, the expected utility of each spam transaction at spam volume $S$ is

$$u(S) = \frac{1}{S+1}r - sg(S).$$

The first term is the expected MEV revenue of each spam transaction. The second term is the inclusion cost. Recall that in a competitive equilibrium, spammers enter until the expected utility of each spam transaction is 0 (see Section 2.1). Solving the zero-profit condition ($u(S) = 0$) gives

$$S^* = \begin{cases} 0, & r_0 \frac{Q_u^0}{D_0} \le sg^0, \\ \frac{r_0 D(g_{\min})}{D_0 s g_{\min}} - 1, & r_0 \frac{Q_u^0}{D_0} > sg^0 \text{ and } g^0 = g_{\min}, \\ \frac{\sqrt{(s - \Delta + \beta r_0/D_0)^2 + 4\beta r_0} - (s + \Delta + \beta r_0/D_0)}{2s}, & \text{otherwise,} \end{cases}$$

where $\Delta = D_0 - B_{\max}$, $g^0 = \max\{g_{\min}, g_1\}$ is the clearing price in the absence of spam, and $Q_u^0 = \min\{B_{\max},\ D(g_{\min})\}$ is the amount of user gas included in the absence of spam.

The first case says that spam does not enter when the effective opportunity value in the spam-free world is below the per-transaction inclusion cost. The second case is the slack-at-the-floor regime, where the gas price remains $g_{\min}$, and the block is not congested. In this regime, spam depends on $r_0$, $g_{\min}$, and the amount of non-spam demand that remains at the floor. The third case is the congested regime, where spam raises the clearing price above the floor. In this regime, spam depends on the baseline opportunity parameter $r_0$, the block size, and user demand.

*Remark* 3.1 (Connection to Mazorra et al.). In the slack case, $S^* = r/(sg_{\min}) - 1 = 1/c - 1$ where $c = (D_0 s g_{\min})/(r_0 \cdot D(g_{\min}))$ is the cost-to-reward ratio. This matches the lower bound on total spam from the timing game model of Mazorra et al. [43], who show that equilibrium spam in an $n$-player timing game lies in $[1/c - 1, 1/c]$ and converges to $1/c - 1$ as $n \to \infty$. Our competitive equilibrium corresponds to this large-$n$ limit derived using a completely different approach from their work.

For our welfare analysis later, we now derive the minimum blockspace, denoted by $B_{\text{plat}}$, needed for the clearing gas price to remain at the floor $g_{\min}$. Looking ahead, the benefit of provisioning blockspace plateaus out beyond $B_{\text{plat}}$. $B_{\text{plat}}$ can be written as

$$B_{\text{plat}} = D(g_{\min}) + \left(\frac{r_0 D(g_{\min})}{D_0 g_{\min}} - s\right)_+.$$

In this equation, the first term is the amount of non-spam demand that remains at the floor price. The second term is the equilibrium spam volume at the floor, converted into gas units. Therefore, $B_{\text{plat}}$ is the amount of block space needed to fit both user demand at $g_{\min}$ and the spam that is still profitable at that price. For $B_{\max} < B_{\text{plat}}$, the block is still effectively scarce, so adding capacity lowers the clearing price and increases the spam volume. Once $B_{\max}$ reaches $B_{\text{plat}}$, the gas price settles at the $g_{\min}$. At that point, spam equilibrium volume reaches its peak level and no longer changes with further increases in block capacity $B_{\max}$. Figure 4 shows this pattern.

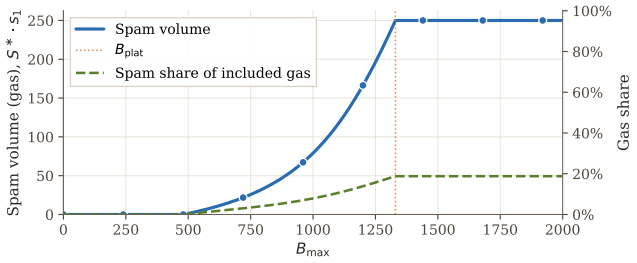


**Figure 4: Equilibrium spam volume as a function of block size $B_{\max}$, and the gas share of spam transactions out of total included gas. Spam is 0 when the block is small and the clearing price is high. As $B_{\max}$ grows, entry becomes more profitable and spam begins to rise. Once $B_{\max}$ reaches $B_{\text{plat}}$, the gas price is pinned at $g_{\min}$ and spam plateaus. In this figure, $D_0 = 1200$, $\beta = 6$, $s = 20$, $r_0 = 6000$, and $g_{\min} = 20$.**

## 3.2 Welfare, Revenue, and Externality Analysis

We next study how the metrics in Section 2.2 vary with the design parameters. To assess the impact of spam, we compare two worlds evaluated at the same parameter values: (1) the realized world with spam, and (2) a counterfactual world without spam. Holding parameters fixed isolates the effect of spam on users, validators, and the cost of the blockchain to process blocks. Our analysis here, especially on welfare and externality, provides the quantities needed later for choosing blockchain parameters wisely.

Figure 5 summarizes the equilibrium levels of user welfare, validator revenue, and externality in the spam world and in the spam-free counterfactual world. For compactness, the maths expressions of the user welfare, validator revenue, and externality, and the explicit counterfactual deltas and the characterization of where the welfare gap is largest are deferred to Section A.

## 3.3 Parameter-Setting Guidance

The analysis above gives the system designer two main parameters against spam, namely the block size $B_{\max}$ and the price floor $g_{\min}$. These two parameters affect user inclusion, validator revenue, and the system cost of supporting larger blocks. A good choice should therefore increase user welfare, but should not keep increasing $B_{\max}$ once the extra capacity mostly creates spam or idle space. Figures 5 and 12 depict the tradeoff between welfare on one side, and revenue and externality on the other side. We address each parameter in turn: here we first discuss choosing $B_{\max}$ for a fixed $g_{\min}$, and we discuss choosing $g_{\min}$ for a fixed $B_{\max}$ in Section B.

We present two rules for choosing $B_{\max}$ given a fixed gas price floor $g_{\min}$: a simple baseline that scales to the point where user welfare plateaus, and a refined rule that stops earlier by requiring that marginal capacity primarily serves users.

**A baseline rule: stop scaling at the plateau $B_{\text{plat}}$.** Choose a gas price floor $g_{\min}$. For this fixed floor, user welfare increases with $B_{\max}$ until the system reaches the slack threshold $B_{\text{plat}}(g_{\min})$. At this point, the block has enough space to fit both the non-spam demand that remains at the floor and the spam that is still profitable at that floor. Once $B_{\max} \geq B_{\text{plat}}(g_{\min})$, further increase in block size does not increase user welfare, because the user side is already capped by $D(g_{\min})$. The revenue also does not increase, and only the externality increases. Therefore, if the designer fixes $g_{\min}$ and wants to maximize user welfare without creating unnecessary externality, the natural choice is $B^{\dagger}_{\max}(g_{\min}) = B_{\text{plat}}(g_{\min})$. This gives a relatively simple benchmark rule: for a fixed gas price floor $g_{\min}$, increase block size until the user welfare plateaus.[2]

**A refined rule: require that marginal capacity does not admit disproportionate spam.** The benchmark rule $B^{\dagger}_{\max}(g_{\min}) = B_{\text{plat}}(g_{\min})$ is useful, but it can be too permissive: near $B_{\text{plat}}$, much of the newly provisioned capacity may be absorbed by spam, as shown in Fig. 6. Therefore, in practice, the designer may want a stricter rule: if the block size is increased by a small amount, a sufficient fraction of that extra capacity should go to useful user gas rather than to spam.

To make this precise, we define the marginal user share as

$$m_{\text{user}} := \frac{\partial Q^*_u(B_{\max}, g_{\min})}{\partial B_{\max}},$$

which measures the fraction of an increase in block capacity that becomes useful user gas. Inside the entry-and-congested region, this derivative can be written as

$$m_{\text{user}} = \frac{1}{2}\left(1 - \frac{B_{\max} - D_0 + s + \beta r_0/D_0}{\sqrt{\left(B_{\max} - D_0 + s + \beta r_0/D_0\right)^2 + 4\beta r_0}}\right).$$

This expression is strictly decreasing in $B_{\max}$ as shown in Theorem 3.2; equivalently, spam takes an increasing share of marginal capacity. Once $B_{\max} \geq B_{\text{plat}}$, $m_{\text{user}}$ becomes zero, because extra capacity is no longer used by either users or spam.

We now introduce a design parameter $\eta \in (0, 1]$, which we call the *minimum marginal user share* (MMUS). It requires that at least an $\eta$ fraction of the next unit of capacity go to users. Given a fixed $g_{\min}$, the refined rule chooses

$$B^*_{\max}(g_{\min}, \eta) \in \arg\max_{B_{\max}} W_{\text{user}}(B_{\max}, g_{\min}) \text{ subject to } m_{\text{user}} \geq \eta.$$

Because $W_{\text{user}}(B_{\max}, g_{\min})$ is nondecreasing in $B_{\max}$, this rule picks the largest block size for which marginal capacity still primarily benefits users. It is a stricter version of the plateau rule: $B^*_{\max}(g_{\min}, \eta) \leq B_{\text{plat}}(g_{\min})$ whenever $\eta > 0$. Lower $\eta$ allows more aggressive scaling, while higher $\eta$ stops scaling earlier.

Proposition 3.2. *Fix $g_{min}$ and suppose the equilibrium is in the entry-and-congested region. Let $g^*$ denote the gas price at equilibrium. Then, for any twice continuously differentiable and strictly decreasing demand curve $D(g)$, if $g^*D(g^*)D''(g^*) + 2D(g^*)D'(g^*) - 2g^*\left(D'(g^*)\right)^2 < 0$ ($D'$ is $\frac{dD}{dg}$), (which holds for the linear demand function), then $m_{user}$ is decreasing in $B_{max}$.*

The proof can be found in Section E. We note that the condition in Theorem 3.2 holds for the linear demand function, and for the exponential demand function (i.e., $D(g) = D_0 \cdot e^{-\lambda g}$) as well.

[2] In practice, the designer may want to keep some slack in blockspace, to accommodate bursts in user demand, i.e., demand function steepening in our model. Some blockchains do this by choosing a block capacity larger than the demand at $g_{\min}$, so that demand spikes during volatile periods can be absorbed without congestion.

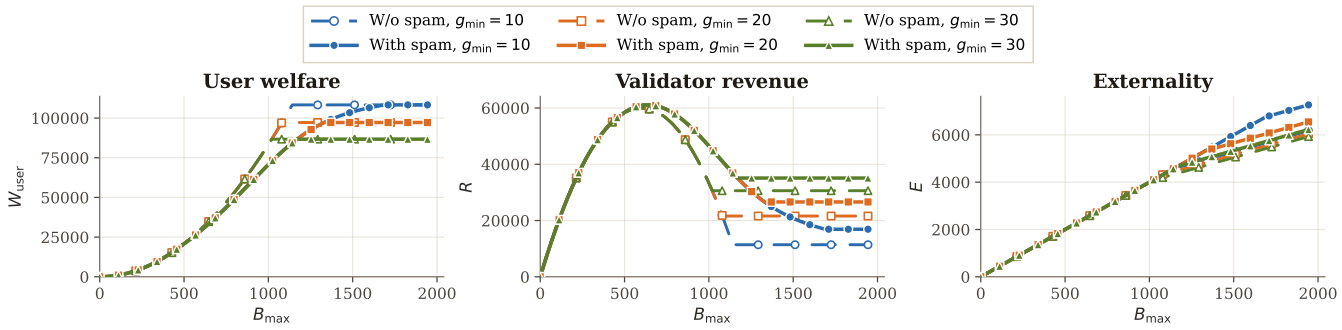


**Figure 5: Levels of user welfare, validator revenue, and externality, with and without spam, as functions of block size. Each panel compares the spam world to the spam-free counterfactual at the same $(B_{\max}, g_{\min})$.**

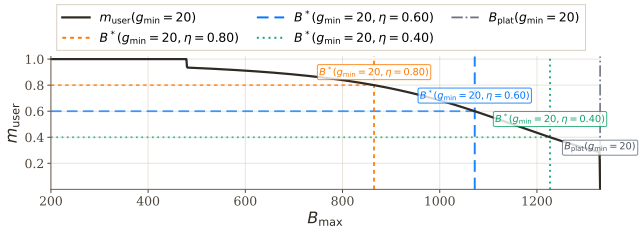


**Figure 6: Marginal allocation of additional block capacity. The x-axis represents the maximum block size, and the y-axis represents the marginal user share. Here we fix $g_{\min} = 20$.**

Figures 6 visualize this rule. It plots $B^*_{\max}$ as a function of $\eta$ when $g_{\min}$ is fixed, and shows the corresponding $B^*_{\max}$ values for representative $\eta$ values.

**Takeaway (Marginal spam share).** Spam takes an increasing share of each additional unit of block capacity. Setting $B_{\max}$ below $B_{\text{plat}}$ or raising $g_{\min}$ creates a favorable trade-off: forgoing a small amount of user welfare eliminates a disproportionate amount of spam, substantially reducing network externality.

## 4 Analysis with (Approximate) Priority Fee Ordering

We now turn to a TFM with priority fee ordering (PFO), where transactions are executed in decreasing order of their bids rather than randomly. PFO closely approximates the ordering observed on many deployed chains, where transactions paying higher priority fees consistently land earlier in the block. It may also help reduce spam: under random ordering, spam transactions can occupy early positions in the block without paying more than others, whereas PFO makes such placement costly.

To capture this effect tractably, we use a block-position-specific demand model. The idea is that users differ not only in their valuation, but also in which part of the block they are willing to occupy. In practice, the top of the block generally serves high-value, time-sensitive flow, e.g., professional DeFi traders, while later positions are filled by users who only need inclusion. We model this by partitioning the block into ordered regions and associating each region with its own demand curve. High-value users only demand early positions, since by the time later positions execute their opportunity has already passed. Lower-value users have no such preference and would happily occupy any block position, but they are priced out of the early region because it clears at a higher gas price.

In the main body, we focus on the two-sub-block case. We split the block into an early sub-block and a late sub-block, with the first executing before the second. Transactions sort into sub-blocks by bid: higher bids land in the first sub-block, lower bids in the second. Within each sub-block, ordering is random. We defer the general $n$-sub-block formulation, where sub-blocks have equal capacity and the demand curve is evenly partitioned across them, to Section C.1.

We use a single parameter $v \in [0, 1]$ to specify the two-sub-block approximation. The first sub-block receives fraction $v$ of block capacity and is associated with the upper $v$ fraction of the original valuation distribution; the second sub-block receives fraction $(1-v)$ of block capacity and is associated with the remaining lower $(1-v)$ fraction. Thus, the sub-block capacities are $C_1 = vB_{\max}$ and $C_2 = (1-v)B_{\max}$. Formally, users in the first sub-block have inverse demand $P_1^{(v)}(Q) = (D_0 - Q)/\beta$ for $0 \le Q \le vD_0$, while users in the second sub-block have inverse demand $P_2^{(v)}(Q) = ((1-v)D_0 - Q)/\beta$ for $0 \le Q \le (1-v)D_0$. Equivalently, the direct demand curves are $D_1^{(v)}(g) = (\min\{vD_0, D_0 - \beta g\})^+$ and $D_2^{(v)}(g) = \big((1-v)D_0 - \beta g\big)^+$. These two demand slices sum to the original linear demand, $D_1^{(v)}(g) + D_2^{(v)}(g) = D_0 - \beta g$.

Let $g_1$ denote the clearing price of the first sub-block and let $\bar{g}$ denote the clearing price of the second sub-block. Whenever both sub-blocks are nonempty, we have $g_1 \ge \bar{g} \ge g_{\min}$. Spam bots may target either sub-block and pay the corresponding price. As in Section 3, we assume that the opportunity value scales linearly with total included user gas, $\bar{r} = r_0 \cdot Q_u/D_0$, where $Q_u$ is the total user gas included in the block. This captures the idea that more user activity creates more MEV opportunities, while spam can reduce those opportunities by displacing users.

### 4.1 Competitive Spam Equilibrium

We now characterize the spam equilibrium in the two-sub-block case. Let $S_1, S_2$ denote the number of spam transactions placed in the first and second sub-blocks, respectively.

Given $S_1$ spam transactions in the first sub-block, the included user gas there is $Q_1(S_1;\bar{g}) = \min\left\{C_1 - S_1 s, D_1^{(v)}(\bar{g})\right\}^+$. The corresponding first-sub-block price is $g_1(S_1;\bar{g}) = \max\left\{\bar{g}, P_1^{(v)}\left(Q_1(S_1;\bar{g})\right)\right\}$. Intuitively, if the top-sub-block demand at the lower price $\bar{g}$ is large enough to fill the post-spam capacity $C_1 - S_1 s$, then the first-sub-block price rises above $\bar{g}$ to clear that capacity. Otherwise, the first sub-block is slack and clears at $\bar{g}$. Similarly, given $S_2$ spam transactions in the second sub-block, the included user gas there is $Q_2(S_2;\bar{g}) = \min\left\{C_2 - S_2 s, D_2^{(v)}(\bar{g})\right\}^+$. The second sub-block clears at the block-level inclusion price, so the equilibrium value $\bar{g}^*$ satisfies the fixed-point condition $\bar{g}^* = \max\left\{g_{\min}, P_2^{(v)}\left(Q_2(S_2^*;\bar{g}^*)\right)\right\}$, with the convention that if $C_2 = 0$, then $\bar{g}^* = g_{\min}$.

Since the opportunity is randomly located among the included user gas, the expected reward attributable to sub-block $i$ is $\frac{r_0}{D_0} Q_i$. Given $S_1$ spam transactions in the first sub-block, the opportunity is captured with probability $S_1/(S_1+1)$, by the same argument as in Section 3.1. The expected utility of spam in the first sub-block is

$$U_1(S_1;\bar{g}) = \frac{r_0}{D_0} Q_1(S_1;\bar{g}) \frac{S_1}{S_1+1} - S_1 s\, g_1(S_1;\bar{g}).$$

The equilibrium spam volume in the first sub-block is the solution $S_1^*$ to $U_1(S_1^*;\bar{g}^*) = 0$, subject to the capacity constraint $S_1^* s \le C_1$.

For the second sub-block, the spam utility has two components. First, the opportunity may originate in the second sub-block and be captured there, contributing $\frac{r_0}{D_0} Q_2(S_2;\bar{g}) \frac{S_2}{S_2+1}$. Second, the opportunity may originate in the first sub-block but fail to be captured there. That happens with probability $1/(S_1^*+1)$. In that case, the opportunity survives into the second sub-block and is captured there by any positive spam. Thus, the expected utility of spam in the second sub-block is

$$U_2(S_2;\bar{g}) = \frac{r_0}{D_0}\left(Q_2(S_2;\bar{g}) \frac{S_2}{S_2+1} + \frac{Q_1^*}{S_1^*+1}\right) - S_2 s\, \bar{g}.$$

The equilibrium spam volume $S_2^*$ is derived by solving $U_2(S_2^*;\bar{g}^*) = 0$, subject to the capacity constraint $S_2^* s \le C_2$.

Together, the two zero-profit conditions and the fixed-point condition for $\bar{g}^*$ determine the equilibrium tuple $(S_1^*, S_2^*, \bar{g}^*)$. The total spam volume is accordingly $S^* = S_1^* + S_2^*$.

Figure 7 illustrates how the resulting spam volume changes with $B_{\max}$ and with the capacity-split parameter $v$. For large $B_{\max}$, once spam reaches its plateau, the $v = 0.5$ curve has less spam than the random-ordering case ($v = 1$). The reason is that spam that is included in the first sub-block pays the higher price $g_1$, while the cheaper lower region has only half of the block capacity. For smaller $B_{\max}$, however, $v = 0.5$ can generate more spam than $v = 1$, since splitting the block creates a lower-priced second sub-block even when the overall block is still relatively scarce. Spam can enter this cheaper tail region, pay only $\bar{g}$, and still capture opportunities that survive from the first sub-block. We note that this is an artifact of our model.

**Takeaway (Priority fee ordering).** Priority fee ordering reduces spam when enough scarce capacity is allocated to a high-value early block region. When the early region is small or has little demand, the lower sub-block can remain cheap, giving spammers an attractive tail region.

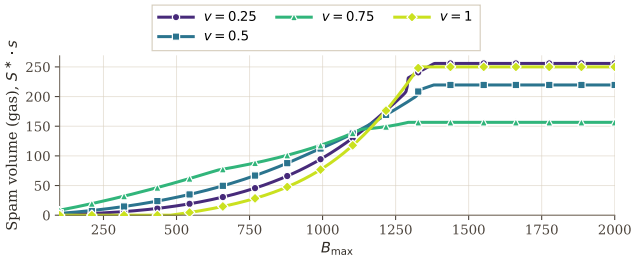


**Figure 7: Equilibrium spam volume under the two-sub-block approximation to PFO for $v \in \{0.25, 0.5, 0.75, 1\}$. We use the same calibration as in Fig. 4. The case $v = 1$ collapses to the random-ordering benchmark.**

## 4.2 Metrics under Approximate PFO

The discussion above focuses on spam volume and where spam is placed inside the block. For completeness, we also evaluate user welfare, validator revenue, and network externality under the two-sub-block PFO model. For the interest of space, we defer the formal expressions and numerical plots to Section C.2.

The main qualitative takeaway is that for any fixed $v$, the pattern in the metrics mirrors the random-ordering model: as $B_{\max}$ rises, user welfare rises and eventually plateaus, validator revenue first increases but then decrease and finally plateaus, and network externality increases. We avoid reading the separate $W_{\text{user}}$, $R$, and $E$ curves as direct comparisons across different values of $v$, since changing $v$ changes both the capacity split and the distribution of user valuations across sub-blocks. The combined quantity $W_{\text{user}} + R$ varies less than user welfare and validator revenue separately as $v$ differs, because changing $v$ partly reallocates surplus between users and the validator.

# 5 Impact of Spam on Blockchain Scaling

One of the central concerns with spam (e.g., the one of Flashbots [44]) is that it may *limit scaling*, as additional block capacity could simply be consumed by spammers instead of benefiting users. Namely, the concern is that the chain may provision more block space, but the effective throughput available to genuine users may not increase proportionally if spam absorbs the added capacity. Recall that spam share (Section 2.2) is a metric measuring the fraction of included gas consumed by spam MEV. We now study whether and to what extent spam may limit the benefit of scaling to be passed down to a rise in user demand.

According to empirical data (see Section G), on both Base and Arbitrum, the total size of MEV profit grows approximately linearly in non-MEV gas. Accordingly, in this section, we assume that the opportunity value scales linearly with included user gas.

To formalize the problem further, let $\lambda \ge 1$ denote a demand-scaling parameter. We scale non-spam demand as $D_\lambda(g) = \lambda(D_0 - \beta g)$, so that the number of users at every price level grows proportionally with $\lambda$. The realized opportunity value is then assumed to be $\bar{r}_\lambda = r_0 \cdot \frac{Q_u}{D_0}$, where $Q_u$ is the equilibrium amount of included user gas under the scaled demand curve. Thus, as more user gas is included, the MEV opportunity grows proportionally. Additionally,

let $\rho_{\text{spam}}$ denote the *spam share* of included gas, i.e., spam gas divided by total included gas. A higher value of $\rho_{\text{spam}}$ means that a larger fraction of the scaled block is unavailable to genuine users.

Here, we focus on the random-ordering benchmark and the two block-size rules from Section 3.3; the demand-scaling analysis under approximate PFO is deferred to Section D. Under the baseline parameter choice, the designer sets $B_{\max}$ equal to the plateau threshold under the scaled demand curve. Since all users willing to pay the floor are then included, the opportunity value at the plateau is $\bar{r}^{\dagger}_{\lambda} = r_0 \cdot D_\lambda(g_{\min})/D_0$, and the corresponding plateau block size is $B_{\text{plat},\lambda} = D_\lambda(g_{\min}) + \left(\frac{r_0 D_\lambda(g_{\min})}{D_0 g_{\min}} - s\right)_+$. Let $S^{\dagger}_{\lambda}$ denote the equilibrium spam volume at $B_{\max} = B_{\text{plat},\lambda}$. The spam share of included gas under the baseline rule is then $\rho^{\dagger}_{\text{spam}}(\lambda) = \frac{s\, S^{\dagger}_{\lambda}}{B_{\text{plat},\lambda}}$.

We also study how the refined block-size rule in Section 3.3 behaves under the same scaling model. Fix a floor $g_{\min}$ and a target $\eta \in (0, 1]$ for the minimum marginal user share. Let $B^{*}_{\max}(g_{\min}, \eta)$ denote the resulting MMUS block size under the scaled demand curve, and let $S^{*}_{\lambda}$ be the corresponding equilibrium spam volume. The spam share of included gas is then $\rho^{*}_{\text{spam}}(\lambda) = \frac{s\, S^{*}_{\lambda}}{B^{*}_{\max}(g_{\min},\eta)}$, where, under the refined rule, the block is still in the congested region and so total included gas equals $B^{*}_{\max}(g_{\min}, \eta)$.

Figure 8 plots $\rho^{\dagger}_{\text{spam}}(\lambda)$ and $\rho^{*}_{\text{spam}}(\lambda)$ for $\lambda \in [1, 50]$, with $\eta = 0.6$ and a fixed price floor $g_{\min} = 20$. The main pattern is that, under linear opportunity scaling, the spam share rises at first and then plateaus at a positive level. The intuition is simple: both user demand at the floor and the induced MEV opportunity grow linearly with $\lambda$, so the amount of profitable spam also scales linearly. As a result, scaling does not drive the spam share to zero. Instead, the system approaches a regime in which spam remains a stable fraction of included gas.

We also observe that, for the same $\lambda$, the refined block-size choice $B^{*}_{\max}$ yields a lower spam share than the baseline choice $B_{\text{plat},\lambda}$. The reason is that the MMUS rule stops scaling earlier, namely at the point where the marginal benefit of extra capacity becomes too spam-heavy. Therefore, the refined rule keeps the spam share systematically below the $B_{\text{plat}}$ benchmark.

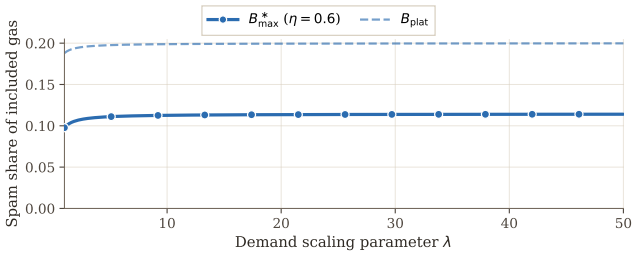


**Figure 8: Impact of demand scaling on the spam share of included gas under random ordering. The x-axis and y-axis are the demand-scaling parameter $\lambda$ and the fraction of included gas consumed by spam, respectively. The dashed curve shows the outcome under the baseline rule that scales block size to $B_{\text{plat},\lambda}$. The solid curve shows the outcome under the refined MMUS rule. We fix $g_{\min} = 20$ and $\eta = 0.6$.**

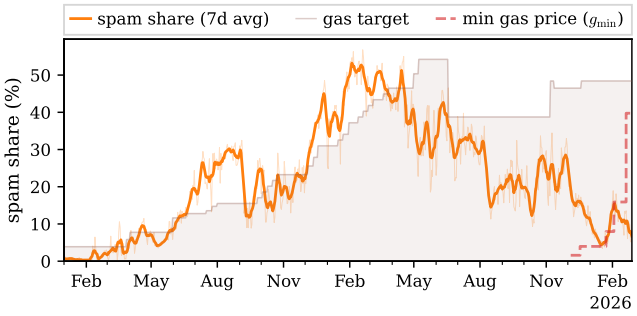


**Figure 9: Spam's share of total gas on Base (7-day moving average, Jan 2024 to Feb 2026). The background shows the gas target and the stepwise protocol minimum gas price. Spam share tracks capacity expansion closely, peaking near 50% in early 2025, then falls sharply after the gas price floor is introduced in Dec 2025.**

**Takeaway (Spam under demand scaling).** When MEV opportunities scale linearly with included user gas, spam's share of included gas does not vanish as demand grows. Instead, under both the plateau rule and the refined MMUS rule, the spam share approaches a sizable positive level.

## 6 Empirical Analysis

Our previous analysis predicts that spam absorbs a disproportionate share of marginal block capacity ($B_{\max}$) and that the gas price floor ($g_{\min}$) directly gates spam profitability. We now examine how these two design levers affect spam in practice by analyzing daily spam gas on two major Ethereum Layer 2 rollups, Base and Arbitrum, which have undergone significant changes to both parameters during our observation period. For the interest of space, we defer how we collect data to Section F.

### 6.1 Spam on Base and Arbitrum

Figure 9 shows spam's share of total gas on Base alongside the gas target and minimum gas price. Three phases are visible. First, after the Dencun upgrade in Mar 2024, Base progressively raised its gas target and spam's share rose sharply, peaking near 50% in early 2025 (cf. Figure 1a for an indexed view of absolute gas). Second, when Base reduced the gas target from 70M to 50M in Jun 2025, spam's share declined (cf. Figure 1b). Third, starting in Dec 2025, Base introduced and progressively raised a protocol minimum gas price; spam's share fell sharply and remained below 15% through Feb 2026 (cf. Figure 2).

Figure 10 shows the corresponding spam share for Arbitrum. Both rollups charge an L2 execution fee plus an L1 data posting surcharge; before Dencun, the L1 component dominated, but blobs made it negligible. The gas price floors discussed here differ in scope: Arbitrum's floor of 0.01 gwei (set at Dencun, down from an effective minimum of approximately 0.1 gwei when L1 costs dominated) covers the total fee, i.e., both execution and data posting, whereas Base's floor (introduced in Dec 2025) applies only to the L2 execution fee. With its higher floor, Arbitrum experienced lower spam than Base throughout the post-Dencun period (mostly 5–25%).

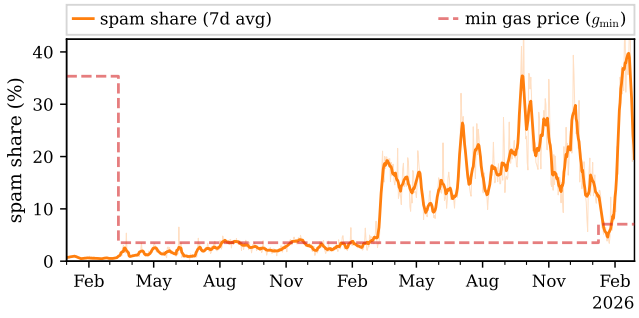


**Figure 10: Spam's share of total gas on Arbitrum (7-day moving average, Jan 2024 to Feb 2026). The background shows the protocol minimum gas price, which was doubled from 0.01 to 0.02 gwei on Jan 9, 2026. Spam share dips briefly after the fee change but rebounds to around 20% by mid-February.**

In Jan 2026, Arbitrum doubled the L2 gas price floor to 0.02 gwei as part of a broader fee mechanism overhaul [6, 62]. As noted in Section 1.1, spam's share dipped briefly after the change but rebounded to around 20% by mid-February.

*Remark* 6.1 (Cross-chain comparisons and ordering.). A raw comparison between Base and Arbitrum should not be read as a test of priority fee ordering. Our approximate PFO analysis is comparative-static: it changes the ordering rule while holding fixed block capacity, gas price floors, user demand, MEV opportunities, and the distribution of demand across block positions. Base and Arbitrum differ along all of these dimensions, so differences in their spam shares do not isolate the effect of ordering alone. We therefore use the empirical evidence mainly in two ways: within-chain changes identify how spam responds to capacity and fee floors, while Base's within-block spam location provides qualitative evidence consistent with the PFO mechanism.

**Spam location within blocks on Base.** We also study where spam transactions appear within Base blocks before the adoption of Flash Blocks in Jul 2025. We collect all transactions sent to the identified spam contracts in each month, rank them by transaction index within the block after excluding the system transaction, and for each $k \in \{0, 1, \ldots, 100\}$ compute the cumulative share of spam gas contained in the first $k$% of block positions. Then we aggregate the monthly curves using spam gas as weights and report the interquartile range (IQR) across months.

Figure 11 shows that the cumulative spam gas curve lies below the uniform benchmark for most of the block. Spam is underrepresented in early block positions and concentrated in later positions. This qualitative pattern is consistent with the analysis in Section 4, as spam tends to appear in later, cheaper parts of the block rather than being spread uniformly. The monthly curves exhibit a consistent qualitative pattern, indicating that the observed effect is persistent throughout the sample period.

## 6.2 Spam and Block Capacity

We regress spam gas on the gas target, i.e., the empirical counterpart of $B_{\max}$, using the full sample of 790 daily observations from Base.

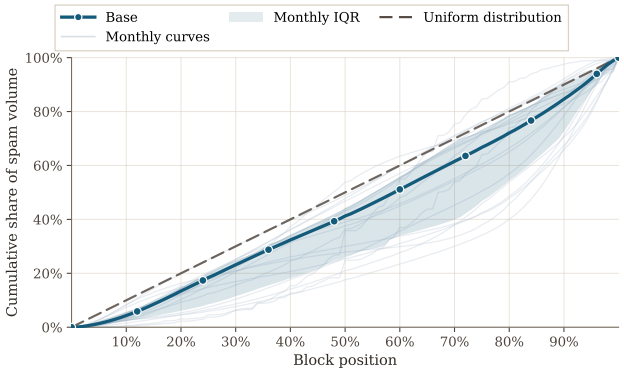


**Figure 11: Distribution of spam gas by normalized position within Base blocks, Jan 2024 to Jun 2025. The x-axis is block position and the y-axis is the cumulative share of spam gas contained at or below that position. The solid line is the gas-weighted aggregate across months, and the shaded band shows the monthly interquartile range (IQR), i.e., the band between the 25th and 75th percentiles of the monthly curves. The dashed 45-degree line is the uniform benchmark. A curve below the dashed line indicates that spam is concentrated toward later positions in the block.**

**Table 1: OLS regression results for spam on Base (790 daily observations, Jan 2024 to Feb 2026). Model 1 regresses** log(**spam gas**) **on** log(**gas target**)**. Model 2 regresses spam's share of total gas on the same regressor. Newey-West standard errors (6 lags) in parentheses.**

| | **Model 1:** log(spam gas) | **Model 2:** spam share |
|---|---|---|
| log(gas target) | $2.27^{***}$ | $0.11^{***}$ |
| | (0.133) | (0.009) |
| Constant | $-13.11^{***}$ | $-1.65^{***}$ |
| | (2.319) | (0.154) |
| $R^2$ | 0.82 | 0.39 |

$^{***}$ $p < 0.001$

We exclude the market-determined clearing fee because it is endogenous: for most of the sample period there was no protocol gas price floor, so the observed minimum fee reflects demand-driven congestion rather than an exogenous policy variable ($g_{\min}$ in our model). Table 1 reports the results.

Model 1 shows that spam gas scales super-linearly with the gas target: a 1% increase in the gas target is associated with a 2.27% increase in spam gas ($R^2 = 0.82$). Model 2 shows that higher gas targets also increase spam's share of total gas (+0.11 per log-unit, $R^2 = 0.39$), consistent with spam disproportionately filling marginal capacity (see Fig. 6 and associated analysis in Section 3.3). The gas target reduction on Jun 18, 2025 (70M to 50M) provides a concrete illustration: spam gas fell by 34% compared to 24% for non-spam gas in the 30-day window (cf. Figure 1b), i.e., both declined, but spam absorbed a larger share of the reduction.

Since daily observations exhibit serial dependence, we report Newey-West standard errors (6 lags) to avoid understating the true uncertainty. Both coefficients remain highly significant ($p < 0.001$) after this correction. These regressions capture only Base; Arbitrum did not change its gas target during the observation period.

## 6.3 Spam and The Minimum Gas Price

The capacity regressions above deliberately exclude the minimum fee because the observed fee is endogenous for most of the sample. To isolate the effect of the protocol gas price floor ($g_{\min}$ in our model), we focus on periods where an exogenous floor was introduced or raised and the gas target was largely constant.

**Base.** Base introduced a protocol minimum gas price in Dec 2025 and raised it in a series of steps (0.0002, 0.0005, 0.001, 0.002, and 0.005 gwei). Since the steps are unevenly spaced (13 to 36 days apart), fixed-length before/after windows would overlap. We therefore use non-overlapping windows: the "after" period runs from one step to the next, and the "before" period is a window of equal length immediately preceding the step. Results are reported in Table 2.

**Table 2: Non-overlapping before-and-after comparisons around each gas price floor step on Base. The "after" window runs from the step to the next step (or end of sample); the "before" window is an equal-length period immediately preceding the step. Δ is the percentage change in daily averages.**

| Gas price floor step | Window | Spam gas Δ | Non-spam gas Δ | Spam share |
|---|---|---|---|---|
| 0 → 0.0002 gwei (Dec 5) | 13d | −41.6% | +9.1% | 26.3% → 16.1% |
| → 0.0005 gwei (Dec 18) | 36d | −55.0% | +17.6% | 20.8% → 9.3% |
| → 0.001 gwei (Jan 23) | 11d | +100.7% | −5.2% | 4.9% → 9.9% |
| → 0.002 gwei (Feb 3) | 17d | +49.0% | −5.4% | 8.2% → 12.3% |
| → 0.005 gwei (Feb 20) | 9d | −27.6% | −1.9% | 11.1% → 8.5% |

Over the full gas price floor period, spam's share of total gas fell from 26% to below 9%, though the decline was not monotonic across individual steps. The introduction of the gas price floor and the first increase to 0.0005 gwei coincided with the largest reductions (−42% and −55%), while non-spam gas grew in both cases (+9% and +18%). The 0.001 and 0.002 gwei steps saw spam rebound, possibly due to confounding factors such as increased price volatility creating more arbitrage opportunities during that period. The subsequent increase to 0.005 gwei again coincided with a 28% spam reduction with negligible non-spam impact (−2%). The data is limited to five steps over a short period, but the overall trend is clear: spam's share declined substantially, and the individual steps where spam fell did so while non-spam gas was unaffected or grew, consistent with the gas price floor raising costs for spam bots while remaining negligible for organic users.

**Arbitrum.** Arbitrum doubled its gas price floor from 0.01 to 0.02 gwei on Jan 9, 2026. Table 3 compares symmetric 50-day windows before and after the change.

Over the 50-day windows, doubling the gas price floor did not reduce spam: spam gas increased by 9%, and its share rose from 19% to 20%. While the first weeks after the change showed a temporary dip (cf. Section 1.1), spam rebounded above pre-change levels by mid-February. Spam did initially decrease after the fee change. The subsequent rebound is largely attributable to three new contracts that went live between Feb 7 and 8, which accounted for 51% of all Arbitrum spam gas in February, and ceased activity around Feb 25. The three contracts went online and offline at the same time, suggesting they may be operated by the same entity. Given their short-lived activity, the rebound likely reflects a transient event rather than a sustained trend.[3]

**Table 3: 50-day before-after comparison around the Arbitrum gas price floor doubling (0.01 to 0.02 gwei) on Jan 9, 2026. Before: Nov 20 to Jan 8; after: Jan 9 to Feb 27.**

| | Before (50d) | After (50d) | Δ |
|---|---|---|---|
| Spam gas (Bgas/day) | 83.9 | 91.1 | +8.6% |
| Non-spam gas (Bgas/day) | 355.9 | 354.1 | −0.5% |
| Spam share | 18.8% | 20.0% | +6.4% |

# 7 Mitigations to Spam

In this section, we discuss mitigation directions suggested by our analysis. Since spam MEV consumes shared block space and infrastructure resources without producing corresponding useful output, the goal is to reduce this resource abuse. This perspective is consistent with ecosystem responses: several chains have made low-value transactions more expensive through gas price floors or minimum gas budgets [5, 6, 9, 62].

We organize mitigations into two broad classes. The first class consists of system-level mechanisms, which operate at the blockchain or block-producer layer. Within this class, *incentive-based mechanisms* make spam more expensive, while *cheap-path mechanisms* make failed spam cheaper to process. The second class consists of *application-specific mechanisms*, where the protocol that creates the opportunity changes its own design to reduce the MEV left for speculative probing. These approaches can be combined in practice.

## 7.1 Incentive-Based Mechanisms

Incentive-based mechanisms reduce spam by increasing its expected cost, so that fewer spam attempts enter in equilibrium. They do not require perfectly identifying spam and can be deployed as coarse policies targeting resource consumption.

**Minimum gas prices.** A simple mechanism is to set a nontrivial floor gas price (i.e., increasing $g_{\min}$). As discussed in Section 1.1, several chains have adopted nontrivial fee floors, either from launch or in response to spam [5, 6, 9, 55, 62]. This is consistent with our analysis in Section 3: a sufficiently high floor gas price caps spam volume and preserves useful block space for users. Empirical evidence for this effect is shown in Section 6.3.

**Charging for reserved execution gas.** Another lever is to charge based on *gas limit* rather than ex post gas used. This targets strategies that reserve capacity but revert early or do little work when they fail. In our model, this makes spam pay for the capacity it forces the network to reserve, which reduces equilibrium spam volume; Monad, for instance, adopted this approach [16].

[3] The three contracts are `0x7c99...24d9`, `0xb4a1...bf47`, and `0xc487...28ac`.

### 7.2 Cheap-Path Mechanisms

Cheap-path mechanisms reduce the cost to the system of failed spam transactions.

**Filtering when execution is not the bottleneck.** When the block producer can simulate transactions, it can filter out transactions that produce no state changes before including them in a block. This is aligned with Ethereum's block-building pipeline, where builders and relays already simulate candidate blocks [30], and with future proposer-builder separation designs [22]. On rollups such as Unichain, the centralized sequencer can similarly simulate transactions and drop no-effect transactions before posting data to Layer 1 [59]. A practical concern is that block producers themselves can become targets for spam, so rate limits, reputation systems, or out-of-band fees may still be needed [30].

**Cancellation when execution is the bottleneck.** When execution resources are scarce, one possible approach is to allow a failed spam transaction to be *cancelled* without full execution, provided the cancellation is justified in a verifiable way. For example, a sender could provide evidence that the transaction would lead to no state changes, while paying a smaller but nonzero fee.

### 7.3 Application-Specific Mechanisms

Application-specific mechanisms reduce spam by changing the protocol that generates the MEV opportunity, so that the opportunity size is reduced. The application can internalize, auction, or redistribute the value that would otherwise be left to spam transactions. Examples include Chainlink's Smart Value Recapture feeds for oracle-update backrunning [17], and protocol redesigns that redistribute MEV at the application layer [2, 65].

## 8 Related Work

**Empirics of spam MEV.** Spam MEV was first documented on Solana by Umbra Research [58], who observed high failure rates among MEV transactions and provided early evidence that high-throughput, low-fee chains incentivize redundant submissions over priority-fee competition. Solmaz et al. [53] study optimistic cyclic arbitrage patterns on Ethereum Layer 2 rollups and find that cyclic arbitrage bots consumed over 50% of gas on Base and Optimism by early 2025, with only 6–12% of probes resulting in a trade. Gogol et al. [32] study reverted transactions on five rollups, finding that revert rates rose sharply after Dencun; by restricting to reverts, they capture one subtype of spam implementation and strategy, as not all unsuccessful probes revert. These works are primarily empirical. Our work develops a framework that characterizes how block capacity, gas price floors, and the transaction fee mechanism jointly determine spam volume, and analyzes the resulting impact on user welfare, validator revenue, and network externality. We complement the analysis with empirical evidence.

**Modeling spam MEV.** Concurrent with our work, Mazorra et al. [43] model spam as a timing game among a finite number of searchers competing for an opportunity, showing that timing competition fully dissipates expected profits. Their equilibrium spam volume in the large-$n$ limit corresponds exactly to ours, despite the models being derived independently. The two models are complementary: Mazorra et al. characterize when and how often each searcher submits in the timing game, while we focus on how protocol design parameters, i.e., block capacity, gas price floors, and transaction ordering, jointly shape spam volume and its welfare and externality implications.

In related work, Capponi and Zhu [15] model costly duplicate submissions as a latency race among traders competing for time-sensitive opportunities on blockchains, framing the problem as an observable analogue of HFT latency investment. They show that a time-priority auction (Timeboost) reduces redundant submissions and reallocates the waste into platform revenue, and validate this prediction on Arbitrum using a difference-in-differences design. Their focus is on a specific auction mechanism as mitigation, whereas we study how general design parameters jointly shape spam volume and welfare.

**Targeted MEV.** Eskandari et al. [25] taxonomized front-running attacks on Ethereum, and Daian et al. [21] broadened the scope to MEV more generally, documenting priority gas auctions among searchers. Qin et al. [52] and Torres et al. [56] subsequently quantified sandwich attacks, arbitrage, and liquidations at scale. A large body of further work has measured specific strategies and infrastructure on the Ethereum Layer 1 [1, 33, 34, 50, 60, 61, 63, 66, 68, 69], on alternative Layer 1s [48], and across chains and rollups [31, 47, 49, 57]. These works study targeted MEV, i.e., strategies where searchers identify a specific opportunity off-chain and submit a transaction to capture it. Our work, in contrast, focuses on spam MEV, where searchers submit speculative transactions whose profitability is resolved only at execution time.

**MEV prevention.** Proposed mitigations for targeted MEV include protocol-specific parameter tuning [35, 67], trusted third-party ordering via relays or private mempools [20, 28, 29, 46], fair ordering protocols [14, 40–42], commit–reveal and privacy-preserving schemes [12, 45, 64], trusted execution environments [10, 54], and encrypted mempools via threshold encryption [3, 4, 8, 11, 18, 19]; see Heimbach and Wattenhofer [36] for a comprehensive systematization. Spam MEV, however, is largely unaddressed by these mechanisms and may even be worsened: encrypted mempools, for instance, remove the information advantage that enables targeted extraction, potentially shifting searchers toward speculative, high-volume strategies instead. As part of this work, we also propose mitigations specifically targeting spam MEV.

## 9 Conclusions and Future Directions

We develop a framework for analyzing spam MEV under a competitive equilibrium, deriving equilibrium spam volumes and characterizing the impact on user welfare, validator revenue, and network externality. Empirical evidence from Base and Arbitrum validates our predictions. We find that spam crowds out users and inflates network externality, claiming a disproportionate share of marginal block capacity. A lower $B_{\max}$, higher $g_{\min}$, or priority fee ordering in the TFM each curb it, and remain effective as demand scales. Together, $B_{\max}$, $g_{\min}$, and the TFM defend against the availability and resource-abuse problem at the heart of spam MEV.

An important direction for future work is the concrete design of system-level mitigation mechanisms, both making spam less attractive to send and failed probes cheaper to process; how these approaches interact across blockchain architectures remains open.

## References


[1] Austin Adams, Benjamin Y Chan, Sarit Markovich, and Xin Wan. 2024. Don't Let MEV Slip: The Costs of Swapping on the Uniswap Protocol. In *Financial Cryptography and Data Security (FC), Willemstad, Curaçao*. 172–191. doi:10.1007/978-3-031-78676-1_10

[2] Austin Adams, Ciamac C Moallemi, Sara Reynolds, and Dan Robinson. 2025. amamm: An auction-managed automated market maker. In *International Conference on Financial Cryptography and Data Security*. Springer, 93–108.

[3] Amit Agarwal, Kushal Babel, Sourav Das, Babak Poorebrahim Gilkalaye, Arup Mondal, Benny Pinkas, Peter Rindal, and Aayush Yadav. 2025. Weighted Batched Threshold Encryption with Applications to Mempool Privacy. Cryptology ePrint Archive, Paper 2025/2115.

[4] Amit Agarwal, Rex Fernando, and Benny Pinkas. 2025. Efficiently-Thresholdizable Batched Identity Based Encryption, with Applications. In *Advances in Cryptology – CRYPTO 2025, Santa Barbara, CA, USA*.

[5] Aptos. 2026. Tweet announcing minimum gas fee increase. https://x.com/Aptos/status/2024216329826230392.

[6] Arbitrum. 2026. ArbOS Dia: Smoother Fees, Higher Throughput. https://blog.arbitrum.io/arbos-dia/.

[7] Kushal Babel, Philip Daian, Mahimna Kelkar, and Ari Juels. 2023. Clockwork finance: Automated analysis of economic security in smart contracts. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2499–2516.

[8] Kushal Babel, Nerla Jean-Louis, Yan Ji, Ujval Misra, Mahimna Kelkar, Kosala Yapa Mudiyanselage, Andrew Miller, and Ari Juels. 2024. PROF: Protected Order Flow in a Profit-Seeking World. Cryptology ePrint Archive, Paper 2024/1241.

[9] Base. 2026. Tweet announcing Base minimum fee increase. https://x.com/buildonbase/status/2024306165711327610.

[10] Iddo Bentov, Yan Ji, Fan Zhang, Lorenz Breidenbach, Philip Daian, and Ari Juels. 2019. Tesseract: Real-Time Cryptocurrency Exchange Using Trusted Hardware. In *ACM SIGSAC Conference on Computer and Communications Security (CCS), London, UK*.

[11] Jan Bormet, Sebastian Faust, Hussien Othman, and Ziyan Qu. 2025. BEAT-MEV: Epochless Approach to Batched Threshold Encryption for MEV Prevention. In *34th USENIX Security Symposium (USENIX Security 25), Seattle, WA, USA*.

[12] Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. 2018. Enter the Hydra: Towards Principled Bug Bounties and Exploit-Resistant Smart Contracts. In *27th USENIX Security Symposium, Baltimore, MD, USA*.

[13] Vitalik Buterin. 2018. Blockchain Resource Pricing. https://ethresear.ch Accessed: 2026-03-13.

[14] Christian Cachin, Jovana Mićić, Nathalie Steinhauer, and Luca Zanolini. 2022. Quick Order Fairness. In *Financial Cryptography and Data Security (FC), Grenada*.

[15] Agostino Capponi and Brian Zhu. 2025. Auctioning Time to Mitigate Latency Races: Theory and Evidence from Blockchains. arXiv preprint arXiv:2512.10094.

[16] Category Labs. 2025. Monad Initial Specification Proposal. https://category-labs.github.io/category-research/monad-initial-spec-proposal.pdf. Version 2.0.1.

[17] Chainlink. 2026. Smart Value Recapture (SVR) Feeds. https://docs.chain.link/data-feeds/svr-feeds.

[18] Arka Rai Choudhuri, Sanjam Garg, Julien Piet, and Guru-Vamsi Policharla. 2024. Mempool Privacy via Batched Threshold Encryption: Attacks and Defenses. In *33rd USENIX Security Symposium (USENIX Security 24), Philadelphia, PA, USA*.

[19] Arka Rai Choudhuri, Sanjam Garg, Guru-Vamsi Policharla, and Mingyuan Wang. 2025. Practical Mempool Privacy via One-Time Setup Batched Threshold Encryption. In *34th USENIX Security Symposium (USENIX Security 25), Seattle, WA, USA*.

[20] CoW Protocol. 2022. CoW Protocol Documentation. https://docs.cow.fi.

[21] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA*.

[22] Francesco D'Amato, Nico Flaig, Barnabé Monnot, Michael Neuder, Potuz, Justin Traglia, and Terence Tsao. 2024. EIP-7732: Enshrined Proposer-Builder Separation. https://eips.ethereum.org/EIPS/eip-7732.

[23] DefiLlama. [n. d.]. Rollup Chains by TVL. https://defillama.com/chains/rollup. Accessed March 2026.

[24] Cynthia Dwork and Moni Naor. 1992. Pricing via processing or combatting junk mail. In *Annual international cryptology conference*. Springer, 139–147.

[25] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. 2019. SoK: Transparent Dishonesty: Front-Running Attacks on Blockchain. In *Financial Cryptography and Data Security Workshops (FC), Kota Kinabalu, Sabah, Malaysia*.

[26] Etherscan. 2025. Announcement of paid APIs for Base, BNB Chain, Avalanche C-Chain, and OP Mainnet. https://info.etherscan.com/whats-changing-in-the-free-api-tier-coverage-and-why/.

[27] Rex Fernando, Guru-Vamsi Policharla, Andrei Tonkikh, and Zhuolun Xiang. 2025. TrX: Encrypted Mempools in High Performance BFT Protocols. Cryptology ePrint Archive, Paper 2025/2032. https://eprint.iacr.org/2025/2032

[28] Flashbots. 2022. MEV-Boost: Proposer-Builder Separation for Ethereum. https://boost.flashbots.net.

[29] Flashbots. 2023. Flashbots Protect RPC. https://docs.flashbots.net/flashbots-protect/overview.

[30] Flashbots. 2025. MEV-Boost Relay Documentation. https://docs.flashbots.net/flashbots-mev-boost/relay.

[31] Krzysztof Gogol, Johnnatan Messias, Deborah Miori, Claudio Tessone, and Benjamin Livshits. 2024. Cross-Rollup MEV: Non-Atomic Arbitrage Across L2 Blockchains. arXiv preprint arXiv:2406.02172.

[32] Krzysztof Gogol, Manvir Schneider, and Claudio Tessone. 2025. When Priority Fails: Revert-Based MEV on Fast-Finality Rollups. *arXiv preprint arXiv:2506.01462* (2025).

[33] Lioba Heimbach, Lucianna Kiffer, Christof Ferreira Torres, and Roger Wattenhofer. 2023. Ethereum's Proposer-Builder Separation: Promises and Realities. In *ACM Internet Measurement Conference (IMC), Montreal, QC, Canada*.

[34] Lioba Heimbach, Vabuk Pahari, and Eric Schertenleib. 2024. Non-Atomic Arbitrage in Decentralized Finance. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA*.

[35] Lioba Heimbach and Roger Wattenhofer. 2022. Eliminating Sandwich Attacks with the Help of Game Theory. In *ACM ASIA Conference on Computer and Communications Security (AsiaCCS), Nagasaki, Japan*.

[36] Lioba Heimbach and Roger Wattenhofer. 2022. SoK: Preventing Transaction Reordering Manipulations in Decentralized Finance. In *4th ACM Conference on Advances in Financial Technologies (AFT), Cambridge, MA, USA*.

[37] Arye L. Hillman and Dov Samet. 1987. Dissipation of contestable rents by small numbers of contenders. *Public Choice* 54, 1 (1987), 63–82. doi:10.1007/BF00123805

[38] ijaack94. 2026. Tweet on indexers discussing payment terms with Base. https://x.com/ijaack94/status/2025499558889795735.

[39] Lefteris Karapetsas. 2026. Tweet on difficulty of indexing Base due to spam. https://x.com/lefterisjp/status/2025312901905408305.

[40] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. 2023. Themis: Fast, Strong Order-Fairness in Byzantine Consensus. In *ACM SIGSAC Conference on Computer and Communications Security (CCS), Copenhagen, Denmark*.

[41] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. 2020. Order-Fairness for Byzantine Consensus. In *Advances in Cryptology – CRYPTO, Santa Barbara, CA, USA*.

[42] Klaus Kursawe. 2020. Wendy, the Good Little Fairness Widget: Achieving Order Fairness for Blockchains. In *2nd ACM Conference on Advances in Financial Technologies (AFT), Virtual Event*.

[43] Bruno Mazorra, Christoph Schlegel, and Akaki Mamageishvili. 2026. Timing Games: Probabilistic backrunning and Spam. arXiv preprint arXiv:2602.22032.

[44] Robert Miller. 2025. MEV and the Limits of Scaling. https://writings.flashbots.net/mev-and-the-limits-of-scaling. Flashbots Writings.

[45] Peyman Momeni, Sergey Gorbunov, and Bohan Zhang. 2022. FairBlock: Preventing Blockchain Front-Running with Minimal Overheads. In *18th EAI International Conference on Security and Privacy in Communication Networks (SecureComm), Virtual Event*.

[46] Alex Obadia. 2020. Frontrunning the MEV Crisis. https://writings.flashbots.net/frontrunning-mev-crisis.

[47] Alexandre Obadia, Alejo Salles, Lakshman Sankar, Tarun Chitra, Vaibhav Chellani, and Philip Daian. 2021. Unity Is Strength: A Formalization of Cross-Domain Maximal Extractable Value. arXiv preprint arXiv:2112.01472.

[48] Burak Öz, Jonas Gebele, Parshant Singh, Filip Rezabek, and Florian Matthes. 2024. Playing the MEV Game on a First-Come-First-Served Blockchain. arXiv preprint arXiv:2404.07169.

[49] Burak Öz, Christof Ferreira Torres, Jonas Gebele, Filip Rezabek, Bruno Mazorra, and Florian Matthes. 2025. Pandora's Box: Cross-Chain Arbitrages in the Realm of Blockchain Interoperability. arXiv preprint arXiv:2501.17335.

[50] Julien Piet, Jaiden Fairoze, and Nicholas Weaver. 2022. Extracting Godl [sic] from the Salt Mines: Ethereum Miners Extracting Value. *arXiv preprint arXiv:2203.15930* (2022).

[51] Richard A. Posner. 1975. The Social Costs of Monopoly and Regulation. *Journal of Political Economy* 83, 4 (1975), 807–827. doi:10.1086/260357

[52] Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying Blockchain Extractable Value: How Dark is the Forest?. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA*.

[53] Ozan Solmaz, Lioba Heimbach, Yann Vonlanthen, and Roger Wattenhofer. 2025. Optimistic MEV in Ethereum Layer 2s: Why Blockspace Is Always in Demand. In *7th Conference on Advances in Financial Technologies (AFT), Pittsburgh, PA, USA (LIPIcs)*. doi:10.4230/LIPIcs.AFT.2025.28

[54] Chrysoula Stathakopoulou, Signe Rüsch, Marcus Brandenburger, and Marko Vukolić. 2021. Adding Fairness to Order: Preventing Front-Running Attacks in BFT Protocols Using TEEs. In *40th International Symposium on Reliable Distributed Systems (SRDS), Chicago, IL, USA*.

[55] Sui. 2025. Gas in Sui. https://docs.sui.io/concepts/tokenomics/gas-in-sui.

[56] Christof Ferreira Torres, Ramiro Camino, and Radu State. 2021. Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain. In *30th USENIX Security Symposium, Virtual Event*.

[57] Christof Ferreira Torres, Albin Mamuti, Ben Weintraub, Cristina Nita-Rotaru, and Shweta Shinde. 2024. Rolling in the Shadows: Analyzing the Extraction of MEV across Layer-2 Rollups. arXiv preprint arXiv:2405.00138.
[58] Umbra Research. 2023. MEV on Solana. https://www.umbraresearch.xyz/writings/mev-on-solana.
[59] Unichain. 2025. Revert Protection on Unichain. https://docs.unichain.org/docs/technical-information/advanced-txn.
[60] Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to Bribe: Measuring Block Construction Market. *arXiv preprint arXiv:2305.16468* (2023).
[61] Ye Wang, Yan Chen, Haotian Wu, Liyi Zhou, Shuiguang Deng, and Roger Wattenhofer. 2022. Cyclic Arbitrage in Decentralized Exchanges. In *The Web Conference Companion (WWW), Virtual Event.*
[62] A.J. Warner. 2025. Tweet on Arbitrum minimum fee increase. https://x.com/ajwarner90/status/2027048105644876069.
[63] Ben Weintraub, Christof Ferreira Torres, Cristina Nita-Rotaru, and Radu State. 2022. A Flash (Bot) in the Pan: Measuring Maximal Extractable Value in Private Pools. In *ACM Internet Measurement Conference (IMC), Nice, France.*
[64] Haoqian Zhang, Louis-Henri Merino, Ziyan Qu, Mahsa Bastankhah, Vero Estrada-Galiñanes, and Bryan Ford. 2023. F3B: A Low-Overhead Blockchain Architecture with Per-Transaction Front-Running Protection. In *5th ACM Conference on Advances in Financial Technologies (AFT), Princeton, NJ, USA.*
[65] Mengqian Zhang, Sen Yang, and Fan Zhang. 2025. RediSwap: MEV Redistribution Mechanism for CFMMs. In *Proceedings of the 2025 Workshop on Decentralized Finance and Security.* 27–36.
[66] Liyi Zhou, Kaihua Qin, Antoine Cully, Benjamin Livshits, and Arthur Gervais. 2021. On the Just-In-Time Discovery of Profit-Generating Transactions in DeFi Protocols. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA.*
[67] Liyi Zhou, Kaihua Qin, and Arthur Gervais. 2021. A2MM: Mitigating Frontrunning, Transaction Reordering and Consensus Instability in Decentralized Exchanges. *arXiv preprint arXiv:2106.07371* (2021).
[68] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc Viet Le, and Arthur Gervais. 2021. High-Frequency Trading on Decentralized On-Chain Exchanges. In *IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA.*
[69] Brian Z Zhu, Xin Wan, Ciamac C Moallemi, Dan Robinson, and Brad Bachu. 2024. Quantifying the Value of Revert Protection. *arXiv preprint arXiv:2410.19106* (2024).

## Ethical Considerations

This work studies spam MEV as a resource-abuse problem for high-throughput blockchains. Our empirical analysis uses publicly available on-chain data and public indexers; we do not collect private user data, interact with users, run experiments on live systems, or submit transactions to any blockchain. All measurements are based on transactions that have already been included on-chain. We focus on aggregate trends, contract-level behavior, and system-level outcomes, and do not attempt to deanonymize users/searchers or attribute activity to real-world identities.

Our spam classification is a heuristic intended for measurement. Contracts identified as spam candidates should not be interpreted as malicious actors in a legal sense; rather, they match the high-volume probing behavior studied in our model.

The main goal of the paper is defensive: to help blockchain designers understand when spam consumes effective throughput, how it affects user welfare and network externality, and which mechanisms can potentially reduce the resulting resource abuse.

## Open Science

The empirical analysis in this paper is based on Dune Analytics queries over public on-chain data. All Dune SQL queries used to collect the data are publicly accessible online on this page.

## A Metrics for Random Ordering Model

This appendix records the level quantities and counterfactual gaps used in Section 3.2. Throughout, we compare two worlds evaluated at the same design parameters $(B_{\max}, g_{\min})$: the equilibrium world with spam, and the counterfactual world without spam. This isolates the effect of spam while holding the block size and gas price floor fixed.

**Spam world and spam-free world.** In the spam-free world, all block space is available to genuine users. The clearing price is

$$g^0(B_{\max}, g_{\min}) = \max\left\{g_{\min}, \frac{D_0 - B_{\max}}{\beta}\right\},$$

and the amount of user gas included in the block is

$$Q_u^0(B_{\max}, g_{\min}) := \min\left\{B_{\max}, D(g_{\min})\right\}.$$

In the spam world, some block space may be taken up by spam. Using the equilibrium spam volume $S^*$ from Section 3.1, the amount of included user gas is

$$Q_u^*(B_{\max}, g_{\min}) := \min\left\{B_{\max} - sS^*(B_{\max}, g_{\min}), D(g_{\min})\right\}.$$

Thus, users get only the block space left after spam. We have $Q_u^* \le Q_u^0$, and the gap $Q_u^0 - Q_u^*$ is the useful block space displaced by spam.

Let $g^*(B_{\max}, g_{\min})$ denote the clearing price in the spam world, and let

$$G^*(B_{\max}, g_{\min}) := Q_u^*(B_{\max}, g_{\min}) + sS^*(B_{\max}, g_{\min})$$

denote the total gas sold in equilibrium. Equivalently, in the congested region $G^* = B_{\max}$, while in the slack-at-the-floor region $G^* = D(g_{\min}) + sS^*$.

**User welfare.** Recall that user welfare is the sum of genuine users' utilities. Under the linear demand model, it depends only on how much user gas is included. Therefore, in the spam-free world,

$$W_{\text{user}}^0(B_{\max}, g_{\min}) = \frac{(Q_u^0(B_{\max}, g_{\min}))^2}{2\beta},$$

while in the spam world,

$$W_{\text{user}}^*(B_{\max}, g_{\min}) = \frac{(Q_u^*(B_{\max}, g_{\min}))^2}{2\beta}.$$

These level expressions show the main mechanism: under random ordering, spam lowers user welfare by reducing useful user gas.

**Validator revenue.** Validator revenue is the amount of gas sold times the clearing price. Without spam, revenue comes only from users:

$$R^0(B_{\max}, g_{\min}) = g^0(B_{\max}, g_{\min})Q_u^0(B_{\max}, g_{\min}).$$

In the spam world, validator revenue is

$$R^*(B_{\max}, g_{\min}) = g^*(B_{\max}, g_{\min})G^*(B_{\max}, g_{\min}).$$

Comparing $R^*$ and $R^0$ captures how spam can increase validator revenue even when it harms users.

**Externality.** We note that when choosing the parameters, we cannot ignore the cost that the parameter choices have on the overall blockchain ecosystem. Along the lines of Buterin [13], we model the ecosystem cost of processing blocks as

$$E(B_{\max}, g_{\min}) = c_1 B_{\max} + c_2 G,$$

where $c_1$ is the cost of provisioning block capacity and $c_2$ is the per-gas execution cost. In the spam-free world,

$$E^0(B_{\max}, g_{\min}) = c_1 B_{\max} + c_2 Q_u^0(B_{\max}, g_{\min}),$$

while in the spam world,

$$E^*(B_{\max}, g_{\min}) = c_1 B_{\max} + c_2 G^*(B_{\max}, g_{\min}).$$

The capacity term is shared by both worlds, while the execution term can be larger in the spam world. Note that this additional burden is not only borne by the validators, but by the overall ecosystem, such as higher hardware and bandwidth requirements for full nodes, which has second-order externalities on decentralization. We use the term *externality* following Buterin [13], who draws a parallel to environmental pollution: although the block producer is compensated for including transactions, the cost of disseminating, executing, and storing them is borne by every full node in the network, much like pollution produced by one factory is suffered by the entire surrounding area. The quantity $E(B_{\max}, g_{\min})$ captures this collective burden.

## A.1 Counterfactual Delta

We now record the gaps between the spam world and the spam-free world. These deltas isolate the effect of spam itself while holding $(B_{\max}, g_{\min})$ fixed.

**User welfare delta.** The user-welfare change caused by spam is

$$\begin{aligned}\Delta W_{\text{user}}(B_{\max}, g_{\min}) &= W^*_{\text{user}}(B_{\max}, g_{\min}) - W^0_{\text{user}}(B_{\max}, g_{\min}) \\ &= \frac{(Q^*_u(B_{\max}, g_{\min}))^2 - (Q^0_u(B_{\max}, g_{\min}))^2}{2\beta}.\end{aligned}$$

This quantity is weakly negative. It is zero when spam does not enter, and also when the block is large enough that spam no longer displaces users.

**Validator revenue delta.** The change in validator revenue caused by spam is

$$\Delta R_{\text{val}}(B_{\max}, g_{\min}) = R^*(B_{\max}, g_{\min}) - R^0(B_{\max}, g_{\min}).$$

In our model, this term is always weakly positive: spam can raise validator revenue by increasing the clearing price, by filling otherwise empty block space, or by doing both.

**Externality delta.** The increase in ecosystem cost caused by spam is

$$\begin{aligned}\Delta E(B_{\max}, g_{\min}) &= E^*(B_{\max}, g_{\min}) - E^0(B_{\max}, g_{\min}) \\ &= c_2\big(G^*(B_{\max}, g_{\min}) - Q^0_u(B_{\max}, g_{\min})\big).\end{aligned}$$

The capacity term $c_1 B_{\max}$ cancels because both worlds use the same block size. Thus, $\Delta E$ measures the additional execution burden created by spam.

Figure 12 summarizes these counterfactual gaps.

**User welfare loss peaks at block capacity of $D(g_{\min})$.** The following result characterizes where spam is most harmful to users.

Proposition A.1 (Welfare Loss Peaks at $D(g_{\text{MIN}})$). *Under the linear demand model $D(g) = D_0 - \beta g$, the user welfare loss $\Delta W_{user}(B_{max})$ is most negative (i.e., users are most harmed) at $B_{max} = Q_{\min}$, where $Q_{\min} := D(g_{min})$.*

The proof is in Section E. For $B_{\max} > Q_{\min}$, the spam-free welfare is constant while the spam-world welfare increases, so the gap shrinks. For $B_{\max} < Q_{\min}$, every marginal unit of capacity serves a user in the spam-free world, but some is absorbed by spam, so the gap widens.

# B Choosing $g_{\min}$ for a fixed $B_{\max}$

The discussion above fixes $g_{\min}$ and chooses $B_{\max}$. In practice, the opposite situation also occurs: the designer may want to keep $B_{\max}$ fixed and choose the gas price floor instead. This is natural when block size is costly to change, or when the system operator wants to adjust spam volume through the gas price floor rather than through throughput.

**A baseline rule.** Fix a block size $B_{\max}$. The simplest rule is to choose the smallest $g_{\min}$ such that the current block size already equals the plateau threshold, that is,

$$g^\dagger_{\min}(B_{\max}) \text{ such that } B_{\text{plat}}(g^\dagger_{\min}) = B_{\max}.$$

Because $B_{\text{plat}}(g_{\min})$ is strictly decreasing in $g_{\min}$, this rule defines a unique gas price floor. If the floor is already high enough that spam does not enter, then the baseline choice is simply

$$g^\dagger_{\min}(B_{\max}) = \frac{D_0 - B_{\max}}{\beta}.$$

When the solution lies in the slack-with-spam region, we have $g^\dagger_{\min}(B_{\max}) = \frac{-(B_{\max} - D_0 + s + \beta r_0/D_0) + \sqrt{(B_{\max} - D_0 + s + \beta r_0/D_0)^2 + 4\beta r_0}}{2\beta}$. If the gas price floor were lower than $g^\dagger_{\min}(B_{\max})$, then the current block size would be too small relative to the amount of spam and non-spam demand that wants to enter. If the gas price floor were higher than $g^\dagger_{\min}(B_{\max})$, then the block would be slack, but the system would be using a stronger fee floor than needed to reach that point.

**Choosing $g_{\min}$ with the refined rule.** The baseline rule makes the current block just sufficient, but it does not control how the newly admitted block usage is split between users and spam as the gas price floor is lowered. Similar to the previous analysis, we define the user share of the newly admitted used capacity when the gas price floor is reduced from $g_{\min}$ to $g_{\min} - \Delta g$ (assume $\Delta g \to 0$):

$$\mu_{\text{user}} := \frac{-\frac{\partial Q^*_u(B_{\max}, g_{\min})}{\partial g_{\min}}}{-\frac{\partial Q^*_u(B_{\max}, g_{\min})}{\partial g_{\min}} - \frac{\partial (s\, S^*(B_{\max}, g_{\min}))}{\partial g_{\min}}}.$$

The minus signs appear because we are studying the effect of lowering the gas price floor.

Proposition B.1. *Fix $B_{max}$. Then the local user share from lowering $g_{min}$ has the following form with the linear demand function:*

*(1) If $g_{min} \geq \max\left\{\frac{D_0 - B_{max}}{\beta}, \frac{r_0 D_0}{D_0 s + \beta r_0}\right\}$, then no spam enters and $\mu_{user} = 1$. In this region, lowering the gas price floor only admits additional users.*

*(2) If $g^\dagger_{min}(B_{max}) \leq g_{min} < \frac{r_0 D_0}{D_0 s + \beta r_0}$, then the block remains slack and spam enters. In this region, lowering $g_{min}$ strictly decreases the fraction of newly admitted used capacity that goes to users.*

*(3) If $\frac{D_0 - B_{max}}{\beta} < g_{min} < g^\dagger_{min}(B_{max})$, then the spam world is congested. In this region, the denominator in the definition of $\mu_{user}$ is zero.*

The proof can be found in Section E. Theorem B.1 gives a simple refined rule for choosing $g_{\min}$. Fix a target $\eta \in (0, 1]$ and require that at least an $\eta$ fraction of newly admitted used capacity go to users. Inside the slack-with-spam region, the condition $\mu_{\text{user}} \geq \eta$ becomes

$$\frac{\beta\, g^2_{\min}}{\beta\, g^2_{\min} + r_0} \geq \eta,$$

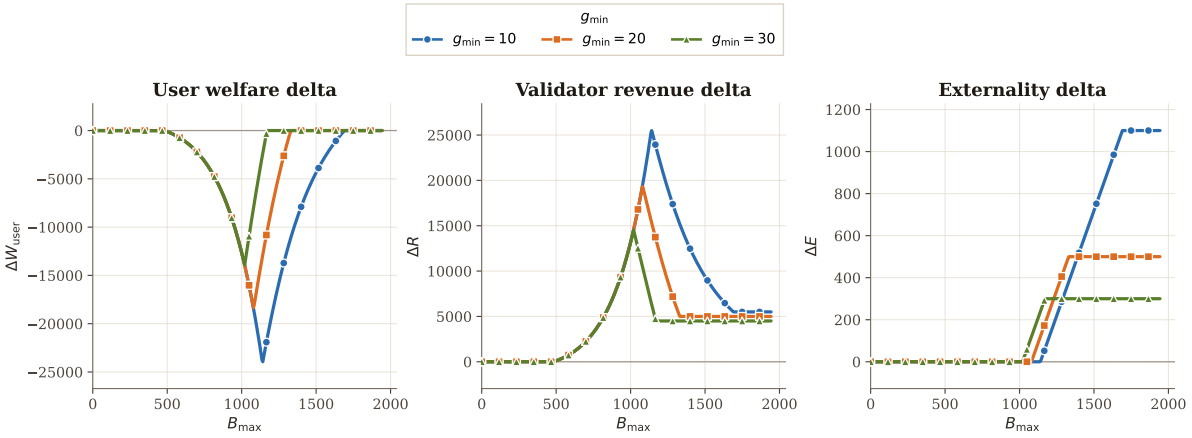


**Figure 12: Counterfactual impact of spam as a function of block size. Negative $\Delta W_{\text{user}}$ means that spam lowers user welfare. Positive $\Delta R_{\text{val}}$ and $\Delta E$ mean that spam raises validator revenue and network externality, respectively.**

which is equivalent to

$$g_{\min} \geq \sqrt{\frac{\eta r_0}{\beta(1-\eta)}}.$$

Therefore, if the block size $B_{\max}$ is fixed, a refined gas price floor choice is

$$g^{\star}_{\min}(B_{\max}, \eta) = \max\left\{ g^{\dagger}_{\min}(B_{\max}),\ \min\left\{ \frac{r_0 D_0}{D_0 s + \beta r_0}, \sqrt{\frac{\eta r_0}{\beta(1-\eta)}} \right\} \right\}.$$

In this equation, the term $g^{\dagger}_{\min}(B_{\max})$ is the baseline floor that makes the current block size just sufficient, and the second term is the user-share threshold. Therefore, the designer lowers $g_{\min}$ only until either the current block size would cease to be sufficient, or the newly added capacity would become too spam-heavy.

# C Priority Fee Ordering

## C.1 Equilibrium Analysis

In this appendix, we generalize the $n = 2$ model from the main body to an arbitrary number $n$ of sub-blocks. Unlike the main body, which uses a free parameter $v$ for the top-block demand share, here we assume an even partition of the original linear demand curve across sub-blocks.

Let $C = B_{\max}/n$ denote the capacity of each sub-block. Let $S_i$ denote the number of spam transactions in sub-block $i$, and let each spam transaction reserve $s$ gas. For $i = 1, \ldots, n$, define the inverse demand of sub-block $i$ by

$$P_i(Q) = \frac{\frac{n-i+1}{n} D_0 - Q}{\beta}, \qquad 0 \leq Q \leq \frac{D_0}{n}.$$

Equivalently, the direct demand is

$$D_i(g) = \left( \min\left\{ \frac{D_0}{n},\ \frac{n-i+1}{n} D_0 - \beta g \right\} \right)^{+}.$$

These demand slices sum to the original linear demand:

$$\sum_{i=1}^{n} D_i(g) = D_0 - \beta g.$$

Thus, each sub-block is associated with a distinct slice of the original valuation distribution, and users from earlier slices do not spill into later sub-blocks.

For a candidate spam profile $\mathbf{S} = (S_1, \ldots, S_n)$ and a candidate block-clearing price $\bar{g}$, define the included user gas in sub-block $i$ by $Q_i(\mathbf{S}; \bar{g}) = \min\{C - S_i s,\ D_i(\bar{g})\}$. The clearing price of sub-block $i$ is then $g_i(\mathbf{S}; \bar{g}) = \max\{\bar{g},\ P_i(Q_i(\mathbf{S}; \bar{g}))\}$. Because the demand slices are nested, these prices are automatically weakly decreasing down the block.

The total included user gas is $Q_u(\mathbf{S}; \bar{g}) = \sum_{i=1}^{n} Q_i(\mathbf{S}; \bar{g})$, and, as in the main body, the opportunity value is $\bar{r} = r_0 \cdot \frac{Q_u(\mathbf{S}; \bar{g})}{D_0}$. The expected reward attributable to sub-block $i$ is therefore $\frac{r_0}{D_0} Q_i(\mathbf{S}; \bar{g})$.

**Top and second top sub-blocks.** Given $S_1$ spam transactions in the top sub-block, the opportunity is captured there with probability $S_1/(S_1 + 1)$. Thus,

$$U_1(S_1) = \frac{r_0}{D_0} Q_1(S_1; \bar{g}) \frac{S_1}{S_1 + 1} - S_1 s\, g_1(S_1; \bar{g}).$$

Setting $U_1(S_1^*) = 0$ gives the equilibrium spam volume in the top sub-block.

For sub-block 2, spam can either capture an opportunity that originates in sub-block 2, or capture an opportunity that originates in sub-block 1 but is not captured there. Hence

$$U_2(S_2) = \frac{r_0}{D_0} \left( Q_2(S_2; \bar{g}) \frac{S_2}{S_2 + 1} + \frac{Q_1^*}{S_1^* + 1} \right) - S_2 s\, g_2(S_2; \bar{g}).$$

Solving $U_2(S_2^*) = 0$ gives the second-sub-block equilibrium.

**General sub-block.** For sub-block $i \geq 3$, suppose that $S_1^*, \ldots, S_{i-1}^*$ have already been computed. The success probability again has two components: the opportunity may originate in sub-block $i$ and be captured there, or it may spill over from an earlier sub-block.

Let $h_i$ be the index of the last earlier sub-block with positive spam, that is, $h_i = \max\{j < i : S_j^* > 0\}$, with the convention $h_i = 0$ if no earlier sub-block has positive spam. Let $k_i = i - 1 - h_i$ be the number of contiguous zero-spam sub-blocks immediately before sub-block $i$. If the opportunity originates in any of those $k_i$ zero-spam sub-blocks, then it reaches sub-block $i$ alive and is captured there by

any positive spam. The contribution of those blocks is therefore $\frac{r_0}{D_0}\sum_{j=h_i+1}^{i-1} Q_j^*$. If $h_i \geq 1$, then there is one additional spillover source: the opportunity may originate in sub-block $h_i$ and fail to be captured there, which occurs with probability $1/(S_{h_i}^*+1)$. Since all sub-blocks between $h_i$ and $i$ have zero spam, that opportunity then reaches sub-block $i$ alive. This contributes $\frac{r_0}{D_0}\frac{Q_{h_i}^*}{S_{h_i}^*+1}$. Putting these terms together, the expected utility of spam in sub-block $i$ is

$$U_i(S_i) = \frac{r_0}{D_0}\left(Q_i(S_i;\bar{g})\frac{S_i}{S_i+1} + \sum_{j=h_i+1}^{i-1} Q_j^* + \mathbf{1}[h_i \geq 1]\frac{Q_{h_i}^*}{S_{h_i}^*+1}\right) - S_i s\, g_i(S_i;\bar{g}).$$

We solve $U_i(S_i^*) = 0$ iteratively for $i = 1,\ldots,n$. If $U_i(S_i)$ remains nonnegative all the way up to the capacity boundary $S_i = C/s$, then the equilibrium for that sub-block is the corner solution $S_i^* = C/s$.

This gives a spam profile $S_1^*,\ldots,S_n^*$. The equilibrium is then closed by the fixed-point condition

$$\bar{g}^* = \max\left\{g_{\min},\ P_n\big(Q_n(S_n^*;\bar{g}^*)\big)\right\}.$$

The total equilibrium spam volume is

$$S^* = \sum_{i=1}^{n} S_i^*,$$

and the total included user gas is

$$Q_u^* = \sum_{i=1}^{n} Q_i^*.$$

## C.2 Derivation of Welfare Metrics

We now record the outcome metrics for approximate PFO. We first give the expressions for the two-sub-block model used in the main body, where the first sub-block has capacity $C_1 = vB_{\max}$ and the second sub-block has capacity $C_2 = (1-v)B_{\max}$. We then state the corresponding expressions for the general $n$-sub-block model from Section C.1.

**Two-sub-block model.** Recall that the first sub-block is associated with the upper slice of the demand curve, with inverse demand $P_1^{(v)}$, while the second sub-block is associated with the lower slice, with inverse demand $P_2^{(v)}$. Changing $v$ therefore changes both the capacity allocated to the early block region and the distribution of user valuations across sub-blocks.

**User welfare.** Because the two sub-blocks correspond to different segments of the original demand curve, welfare is computed separately in each segment. In the spam world, users in the first sub-block pay $g_1^*$, while users in the second sub-block pay $\bar{g}^*$. Therefore,

$$W_{\text{user}}(B_{\max}) = \int_0^{Q_1^*}\left(P_1^{(v)}(q) - g_1^*\right)dq + \int_0^{Q_2^*}\left(P_2^{(v)}(q) - \bar{g}^*\right)dq.$$

The no-spam benchmark is obtained by setting $S_1 = S_2 = 0$ and recomputing the two clearing prices and user quantities.

**Validator revenue.** Each unit of gas in sub-block 1 pays $g_1^*$, while each unit of gas in sub-block 2 pays $\bar{g}^*$. The total gas sold in each sub-block is the sum of user gas and spam gas. Therefore,

$$R(B_{\max}) = g_1^*\cdot\left(Q_1^* + S_1^* s\right) + \bar{g}^*\cdot\left(Q_2^* + S_2^* s\right).$$

**Externality.** As in Section 2.2, we model the cost of supporting larger blocks as a capacity cost plus a per-gas execution cost. In the spam world,

$$E(B_{\max}) = c_1 B_{\max} + c_2\left(Q_1^* + Q_2^* + sS_1^* + sS_2^*\right).$$

The no-spam counterfactual is obtained by setting $S_1 = S_2 = 0$.

Figure 13 compares the spam and no-spam outcomes under approximate PFO.

**General $n$-sub-block model.** For completeness, we also record the corresponding metrics for the general $n$-sub-block model from Section C.1. In that model, each sub-block corresponds to its own slice of the original demand curve. Users in sub-block $i$ have inverse demand $P_i$ and pay price $g_i^* = g_i(S_i^*;\bar{g}^*)$. Their consumer surplus is therefore

$$\int_0^{Q_i^*}\left(P_i(q) - g_i^*\right)dq.$$

Summing over sub-blocks gives total user welfare:

$$W_{\text{user}}(B_{\max}) = \sum_{i=1}^{n}\int_0^{Q_i^*}\left(P_i(q) - g_i^*\right)dq.$$

Each unit of gas in sub-block $i$ pays $g_i^*$. The total gas sold in sub-block $i$ is $Q_i^* + S_i^* s$. Therefore, validator revenue is

$$R(B_{\max}) = \sum_{i=1}^{n} g_i^*\cdot\left(Q_i^* + S_i^* s\right).$$

Using the same system-cost model as above, the externality is

$$E(B_{\max}) = c_1 B_{\max} + c_2\sum_{i=1}^{n}\left(Q_i^* + S_i^* s\right).$$

The no-spam benchmark is obtained by setting $S_i = 0$ for every sub-block and recomputing the induced prices and user quantities.

# D Spam Volume Share under Demand Scaling with Approximate Priority Fee Ordering

We ask whether priority fee ordering changes the scaling picture. To keep the comparison aligned with the random-ordering analysis in the main body, we use the same block-size benchmark. For each value of $\lambda$, we set $B_{\max} = B_{\text{plat},\lambda}$, where $B_{\text{plat},\lambda}$ is the random-ordering plateau size defined in Section 5. Thus, total provisioned capacity is held fixed at the same benchmark level, while the allocation of block capacity across execution positions changes according to the two-sub-block PFO model.

For a given value of $v$, let $S_\lambda^{(v)}$ denote the equilibrium spam volume in the two-sub-block PFO model, and let $Q_{u,\lambda}^{(v)}$ denote the corresponding included user gas. Recall that $v$ is the fraction of both block capacity and user valuation mass assigned to the first, higher-priority sub-block. We define the spam share of included gas as

$$\rho_{\text{spam}}^{(v)}(\lambda) = \frac{s\,S_\lambda^{(v)}}{s\,S_\lambda^{(v)} + Q_{u,\lambda}^{(v)}}.$$

The case $v = 1$ collapses to a single sub-block and is therefore equivalent to the random-ordering benchmark.

Figure 14 compares the two-sub-block PFO outcomes for $v \in \{0.25, 0.5, 0.75, 1.0\}$. We fix $g_{\min} = 20$ and size the block using $B_{\text{plat},\lambda}$. The figure shows that PFO can reduce the spam share relative to the $v = 1$ benchmark.

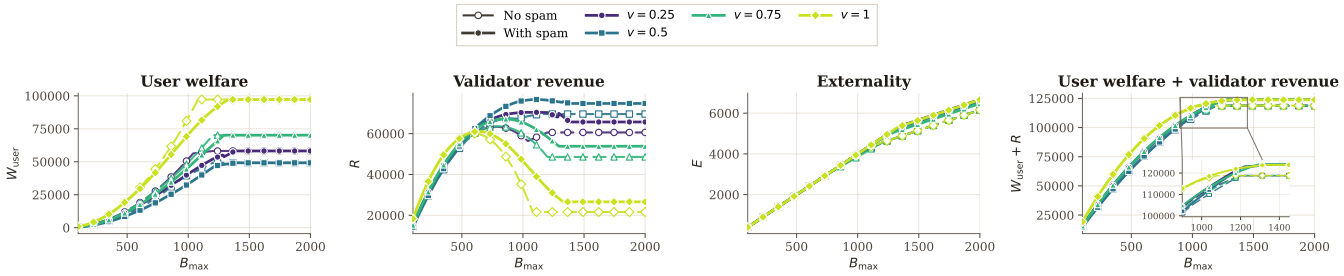


**Figure 13: User welfare, validator revenue, externality, and $W_{\text{user}}+R$ under approximate priority fee ordering in the two-sub-block model. The solid curves show the spam world and the dashed curves show the no-spam counterfactual. Each color corresponds to a different value of $v$.**

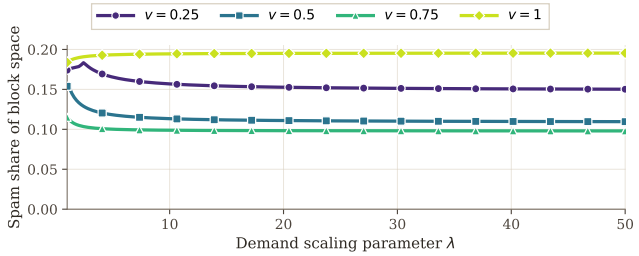


**Figure 14: Spam volume share under demand scaling with approximate priority fee ordering. The x-axis is the demand-scaling parameter $\lambda$. The y-axis is the fraction of included gas consumed by spam. All curves are evaluated at the same block-size benchmark $B_{\max} = B_{\text{plat},\lambda}$. The curves show the two-sub-block PFO outcome for $v \in \{0.25, 0.5, 0.75, 1.0\}$, where $v$ is the fraction of both block capacity and user valuation mass assigned to the first sub-block. The case $v = 1$ collapses to the random-ordering benchmark.**

The qualitative scaling pattern remains similar to the random-ordering case. Across the curves, the spam share approaches a positive plateau as $\lambda$ grows. Thus, approximate PFO changes the level of spam pressure, but does not change the basic conclusion that linear opportunity scaling preserves a nontrivial spam share. The role of PFO is therefore to reduce the share of spam in favorable parameter regimes.

# E Proofs

## E.1 Proof of Theorem A.1

We work with the linear demand $D(g) = D_0 - \beta g$, inverse demand $P(Q) = (D_0 - Q)/\beta$, and $Q_{\min} := D_0 - \beta g_{\min}$. Consumer surplus is $W(Q) = Q^2/(2\beta)$. Recall that the opportunity value is now endogenous and equal to $r = r_0 Q_u / D_0$. If $r_0 Q_{\min}/D_0 \le s g_{\min}$, then spam does not enter for any $B_{\max}$, so $\Delta W_{\text{user}} \equiv 0$ and the claim is trivial. We therefore assume that entry occurs.

In the congested regime, let $Q := Q_u^*$. Free entry, block fullness, and the linear inverse demand imply $\frac{r_0 Q}{D_0(S^*+1)} = s g^*$, $B_{\max} = Q + s S^*$, and $g^* = \frac{D_0 - Q}{\beta}$. Eliminating $S^*$ and $g^*$ gives

$$\beta r_0 Q = D_0(D_0 - Q)(B_{\max} - Q + s).$$

Differentiating implicitly with respect to $B_{\max}$ yields

$$Q'(B_{\max}) = \frac{D_0(D_0 - Q)}{\beta r_0 + D_0(B_{\max} + s + D_0 - 2Q)}.$$

Hence $0 < Q'(B_{\max}) < 1$, since $Q < B_{\max}$ in the spam world and the denominator exceeds the numerator by

$$\beta r_0 + D_0(B_{\max} + s - Q) > 0.$$

**Right of $Q_{\min}$ (loss shrinks).** The spam-free world is slack, so $W^0 = Q_{\min}^2/(2\beta)$ is constant. For $Q_{\min} < B_{\max} < B_{\text{plat}}$, the spam world is congested and

$$W^* = \frac{Q(B_{\max})^2}{2\beta}, \qquad \frac{dW^*}{dB_{\max}} = \frac{Q(B_{\max})Q'(B_{\max})}{\beta} > 0.$$

Thus $\Delta W_{\text{user}} = W^* - W^0$ increases toward 0 as $B_{\max}$ grows. Once $B_{\max} \ge B_{\text{plat}}$, the spam world is also slack and $\Delta W_{\text{user}} = 0$.

**Left of $Q_{\min}$ (loss grows).** Below the spam-entry threshold, $\Delta W_{\text{user}} = 0$. Once spam enters and $B_{\max} < Q_{\min}$, both worlds are congested:

$$W^0 = \frac{B_{\max}^2}{2\beta}, \qquad W^* = \frac{Q(B_{\max})^2}{2\beta}.$$

Therefore

$$\Delta W_{\text{user}} = W^* - W^0 = \frac{Q(B_{\max})^2 - B_{\max}^2}{2\beta},$$

and

$$\frac{d}{dB_{\max}} \Delta W_{\text{user}} = \frac{Q(B_{\max})Q'(B_{\max}) - B_{\max}}{\beta} < 0,$$

because $Q(B_{\max}) < B_{\max}$ and $0 < Q'(B_{\max}) < 1$. So the welfare loss becomes more negative as $B_{\max}$ increases up to $Q_{\min}$.

Combining the two regions, $\Delta W_{\text{user}}(B_{\max})$ decreases on the left of $Q_{\min}$ and increases on the right of $Q_{\min}$. Therefore it is most negative at $B_{\max} = Q_{\min}$. □

## E.2 Proof of Theorem 3.2

Proof. In the entry-and-congested region, free entry implies

$$\frac{r_0 D(g^*)}{D_0(S^* + 1)} = s\, g^*,$$

so

$$S^* = \frac{r_0 D(g^*)}{D_0 s\, g^*} - 1.$$

Because the block is congested, it is full, so

$$D(g^*) + s\,S^* = B_{\max}.$$

Substituting the expression for $S^*$ gives

$$r_0 D(g^*) = D_0 g^* \big(B_{\max} - D(g^*) + s\big).$$

Differentiating both sides with respect to $B_{\max}$ yields

$$r_0 D'(g^*)\tfrac{\partial g^*}{\partial B_{\max}} = D_0\left[\tfrac{\partial g^*}{\partial B_{\max}}\big(B_{\max} - D(g^*) + s - g^* D'(g^*)\big) + g^*\right].$$

Rearranging and using

$$B_{\max} - D(g^*) + s = \frac{r_0 D(g^*)}{D_0 g^*}$$

from the equilibrium condition above gives

$$\frac{\partial g^*}{\partial B_{\max}} = -\frac{D_0 (g^*)^2}{r_0 D(g^*) - g^*(r_0 + D_0 g^*) D'(g^*)}.$$

Now use $Q_u^* = D(g^*)$, so

$$m_{\text{user}} = \frac{\partial Q_u^*}{\partial B_{\max}} = -\frac{D_0 (g^*)^2 D'(g^*)}{r_0 D(g^*) - g^*(r_0 + D_0 g^*) D'(g^*)}.$$

Differentiating $m_{\text{user}}$ with respect to $g^*$ gives

$$\frac{\partial m_{\text{user}}}{\partial g^*} = -\frac{D_0 g^* r_0\left(g^* D(g^*) D''(g^*) + 2D(g^*)D'(g^*) - 2g^*(D'(g^*))^2\right)}{\left(r_0 D(g^*) - g^*(r_0 + D_0 g^*) D'(g^*)\right)^2}.$$

Multiplying by $\partial g^*/\partial B_{\max}$ and using the formula above yields

$$\frac{\partial m_{\text{user}}}{\partial B_{\max}} = \frac{D_0^2 (g^*)^3 r_0\left(g^* D(g^*) D''(g^*) + 2D(g^*)D'(g^*) - 2g^*(D'(g^*))^2\right)}{\left(r_0 D(g^*) - g^*(r_0 + D_0 g^*) D'(g^*)\right)^3}.$$

The denominator is positive, so the sign is determined by

$$g^* D(g^*) D''(g^*) + 2D(g^*)D'(g^*) - 2g^*(D'(g^*))^2.$$

Thus, whenever

$$g^* D(g^*) D''(g^*) + 2D(g^*)D'(g^*) - 2g^*(D'(g^*))^2 < 0,$$

we have

$$\frac{\partial m_{\text{user}}}{\partial B_{\max}} < 0.$$

For the linear demand curve $D(g) = D_0 - \beta g$, we have $D'(g) = -\beta$ and $D''(g) = 0$, so

$$g D(g) D''(g) + 2D(g)D'(g) - 2g(D'(g))^2 = -2\beta D_0 < 0.$$

Therefore, the sufficient condition above holds for the linear demand function. □

*Remark* E.1. The condition in Theorem 3.2 does not hold for every decreasing demand curve. For example, let

$$D(g) = A \exp\left(-\frac{1 - (1+g)^{-2}}{2}\right).$$

Then $D'(g) = -D(g)/(1+g)^3 < 0$, so the demand curve is strictly decreasing. A direct calculation gives

$$g D(g) D''(g) + 2D(g)D'(g) - 2g(D'(g))^2 = A^2 e^{-1+(1+g)^{-2}} \tfrac{g^3 - 4g - 2}{(1+g)^6}.$$

This expression is positive for all sufficiently large $g$. Therefore, the sufficient condition in Theorem 3.2 can fail even for a smooth strictly decreasing demand curve.

*Remark* E.2 (The condition that the slope of $m_{\text{user}}$ decreasing in $B_{\max}$). Let

$$A(g) = r_0 D(g) - g(r_0 + D_0 g) D'(g),$$

and

$$F(g) = g D(g) D''(g) + 2D(g)D'(g) - 2g(D'(g))^2.$$

We have

$$\frac{\partial m_{\text{user}}}{\partial B_{\max}} = \frac{D_0^2 (g^*)^3 r_0\, F(g^*)}{A(g^*)^3}.$$

Since

$$\frac{\partial g^*}{\partial B_{\max}} < 0,$$

the condition for $\partial m_{\text{user}}/\partial B_{\max}$ to be non-increasing is

$$\frac{\mathrm{d}}{\mathrm{d}g}\left(\frac{g^3 F(g)}{A(g)^3}\right) \geq 0.$$

It holds for the linear demand function.

## E.3 Proof of Theorem B.1

Proof. In the no-entry region, $S^* = 0$, so lowering the gas price floor changes only user inclusion. Therefore the full marginal increase in used capacity goes to users, and $\mu_{\text{user}} = 1$. When the floor binds, no entry means

$$\frac{r_0 D(g_{\min})}{D_0} \leq s\, g_{\min}.$$

Under the linear demand function $D(g) = D_0 - \beta g$, this is equivalent to

$$g_{\min} \geq \frac{r_0 D_0}{D_0 s + \beta r_0}.$$

Together with the condition that the floor binds, $g_{\min} \geq \frac{D_0 - B_{\max}}{\beta}$, this gives the first case.

In the slack-with-spam region, we have

$$Q_u^*(B_{\max}, g_{\min}) = D(g_{\min}) = D_0 - \beta g_{\min}$$

and

$$s\,S^*(B_{\max}, g_{\min}) = s\left(\frac{r_0 D(g_{\min})}{D_0 s\, g_{\min}} - 1\right) = \frac{r_0 D(g_{\min})}{D_0 g_{\min}} - s.$$

Under the linear demand function, this simplifies to

$$s\,S^*(B_{\max}, g_{\min}) = \frac{r_0(D_0 - \beta g_{\min})}{D_0 g_{\min}} - s = \frac{r_0}{g_{\min}} - \frac{\beta r_0}{D_0} - s.$$

Differentiating with respect to $g_{\min}$ gives

$$\frac{\partial Q_u^*}{\partial g_{\min}} = -\beta \qquad \text{and} \qquad \frac{\partial (s\,S^*)}{\partial g_{\min}} = -\frac{r_0}{g_{\min}^2}.$$

Substituting these derivatives into the definition of $\mu_{\text{user}}$ yields

$$\mu_{\text{user}}(B_{\max}, g_{\min}) = \frac{\beta}{\beta + r_0/g_{\min}^2} = \frac{\beta\, g_{\min}^2}{\beta\, g_{\min}^2 + r_0}.$$

Differentiating this expression gives

$$\frac{\partial \mu_{\text{user}}(B_{\max}, g_{\min})}{\partial g_{\min}} = \frac{2\beta r_0\, g_{\min}}{(\beta g_{\min}^2 + r_0)^2} > 0.$$

Therefore $\mu_{\text{user}}$ increases with $g_{\min}$, or equivalently decreases as the gas price floor is lowered.

Finally, in the congested-with-spam region, the equilibrium price is already above the gas price floor. A small change in $g_{\min}$ does not change the local equilibrium allocation, so both the user-gas

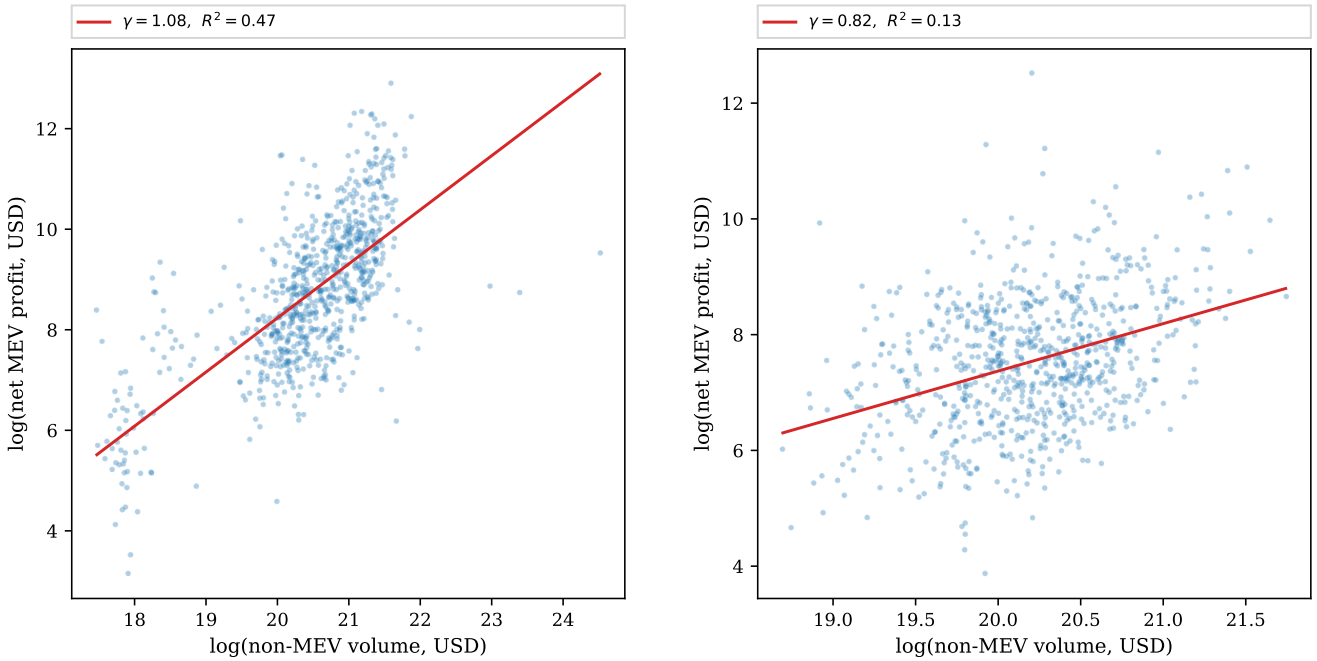


**(a) Base. The estimated elasticity is $\hat{\gamma} = 1.08$ with $R^2 = 0.47$.**

**(b) Arbitrum. The estimated elasticity is $\hat{\gamma} = 0.82$ with $R^2 = 0.13$.**

**Figure 15: Daily net cyclic arbitrage profit versus daily non-MEV trading volume on log-log scales.**

increment and the spam-gas increment are zero. Thus the denominator in the definition of $\mu_{\text{user}}$ vanishes, and the local share is not informative in that region. □

## F Data and methodology

We collect 790 daily observations per chain from January 2024 through February 2026 using Dune Analytics. To classify spam, we use an approach similar to the methodology employed by Flashbots [44] and consistent with the optimistic probing behavior described by Solmaz et al. [53]: we identify contracts that interact with decentralized exchanges (DEXes) but frequently fail to execute trades, i.e., they probe for arbitrage opportunities without producing ERC-20 token transfers. Using on-chain traces, we find all transactions where a contract makes an internal call to a known DEX pool or router address. For each such transaction, we check whether it produced at least one ERC-20 transfer event. Contracts that interact with DEXes but produce no token transfer are probing but not trading. For each contract and month, we compute the fraction of DEX-interacting transactions that produced a transfer. Contracts with a transfer rate below 50% and at least 10 interactions are classified as spam candidates. We rank candidates by total gas consumed and conservatively retain only the top 100 per month. We then attribute all daily transactions from these contracts to spam transactions. Non-spam gas is the total gas minus spam gas. The scripts that we use to gather the data can be found on this page.

## G MEV Revenue Scaling with User Volume

In this appendix, we estimate how the opportunity value $r$ scales with user demand in deployed systems. This matters for the scaling analysis in the main body. If MEV opportunity size grows with user activity, then the amount of spam that can be sustained at equilibrium may also grow with adoption.

*Methodology.* We study cyclic arbitrage on Ethereum Layer 2 rollups. Our main case study is Base, and we also report corresponding estimates for Arbitrum. We use daily observations and identify cyclic arbitrage transactions from decoded DEX swap data on Dune Analytics. A transaction is classified as cyclic arbitrage if it contains at least three swaps that form a cycle, meaning that the first token sold is also the last token bought, and the transaction yields a positive net balance in exactly that token. We exclude transactions routed through known DEX aggregator contracts, since these are more likely to reflect user-initiated routing rather than MEV. For each arbitrage transaction, we compute gross profit from the net token balance and subtract the gas fee to obtain net profit. We then aggregate these net profits by day and use the resulting daily total as a proxy for the available opportunity value $r$. This proxy is imperfect, but it is natural in a competitive setting, since realized arbitrage profits should track the amount of extractable value available to searchers.

To measure user demand, we use daily non-MEV trading volume in USD and estimate the log-log regression

$$\log(P_t) = \alpha + \gamma \log(V_t),$$

where $P_t$ is daily net cyclic arbitrage profit and $V_t$ is daily non-MEV trading volume. The coefficient $\gamma$ is the elasticity of realized cyclic arbitrage profit with respect to user trading volume.

*Base.* We first analyze Base using 790 daily observations from January 2024 through February 2026. The estimate is $\hat{\gamma} = 1.08$ with standard error 0.041, $R^2 = 0.47$, and $p \approx 0$. Thus, on Base, a 1% increase in user trading volume is associated with roughly a 1.08% increase in cyclic arbitrage profit, which is close to linear.

*Arbitrum.* We next apply the same methodology to Arbitrum. For Arbitrum, we estimate $\hat{\gamma} = 0.82$ with $R^2 = 0.13$. The low $R^2$ indicates that the relationship between user trading volume and cyclic arbitrage profit on Arbitrum is considerably weaker and less reliable than on Base, with most of the variation in arbitrage profit unexplained by trading volume alone.