On the Impossibility of Transparent and Decentralized Trading on Permissionless Blockchains

— PRELIMINARY — COMMENTS WELCOME —

Hanna Halaburda*

December 1, 2025

Abstract

Blockchain technology promises a transparent, decentralized, and fair architecture for peer-to-peer interaction, including finance. This paper explores the inherent tension between this decentralized architecture and the economic equilibrium that emerges on top of it. We argue that this equilibrium is centralized and extractive. We develop a generalized, protocol-agnostic model of strategic interaction between payoff-maximizing agents (e.g., block builders, miners) who compete to create blocks. We show that free entry, transparency, and latency-bounded settlement create a natural and powerful tendency towards centralization and rent extraction. This tendency is driven by three main forces: (1) transparency and latency-bounded settlement create predictable opportunities for MEV extraction, so even small differences in agents' extraction capabilities generate persistent heterogeneity in effective rewards; (2) block production across PoW, PoS, and PBS is a winner-take-all contest that strictly sorts agents by reward and magnifies small advantages, producing increasing concentration as contest decisiveness rises or reward dispersion widens; and (3) rational traders avoid extraction in the public mempool by routing privately to the most reliable execution agent, introducing strong network complementarities that make monopoly private order flow the only stable outcome under any positive multihoming cost.

Together, these forces generate an impossibility result: transparent and decentralized trading cannot coexist in equilibrium. All robust equilibria are centralized, while decentralized allocations require knife-edge conditions or fail to exist. The architecture intended to eliminate intermediaries instead recreates and entrenches them; the very features meant to guarantee transparency and decentralization ultimately erode both.

Keywords: blockchain, DeFi, market design, MEV, contest theory, platform economics, centralization
JEL: D47, L13, G23, O30, D83

*NYU, email: hhalaburda@gmail.com

1 Introduction

A core promise of blockchain-based systems is that they can support market activity without centralized intermediaries. This promise rests on three design principles: all market participants observe the same information, the underlying protocol treats all potential service providers identically, and anyone can enter and compete. These features are meant to deliver an execution environment where no single actor can dominate, and where market power cannot arise.

The empirical record tells a different story. Across all major blockchain systems, the activity of assembling and confirming transactions has become highly concentrated. In early proof-of-work systems, a small number of mining pools quickly came to dominate the creation of new blocks. When blockchains transitioned to proof-of-stake, large staking pools captured a disproportionate share of validation rights. More recently, new designs that separate the role of proposing blocks from the role of building them (implemented e.g. on Ethereum) have intensified this pattern: a handful of specialized firms now regularly win the auctions that determine which transactions appear first, with the top five entities securing more than 90% of these opportunities.

Why does a market design that is transparent, open, and decentralized at the protocol level generate concentrated outcomes in equilibrium? This paper argues that these outcomes are not accidents of implementation or the result of flawed engineering. Instead, they arise from structural features inherent to permissionless blockchains — public transparency, time-bounded settlement, and open contests for block production — which in turn activate three economic forces — rent seeking from publicly revealed opportunities, convexity in competitive success, and network complementarities in order routing — that together steer the system away from decentralization and toward dominance by a small set of actors.

These architectural features are not ad hoc engineering choices; they are what allow blockchains to pursue credible decentralization in the first place. First, transactions are publicly visible before they are finalized because the system is designed to let anyone verify correctness and detect manipulation. Public visibility allows an open set of participants to agree on a common state without trusting any particular intermediary. Second, blocks are added to the blockchain at regular intervals. This time-bounded settlement is needed to guarantee liveness: transactions must clear even when participants are anonymous, geographically dispersed, and not coordinated through any central authority. Without predictable settlement, the system could stall or be strategically delayed. Third, the right to assemble the next batch of transactions must be allocated through an open, competitive process. Whether based on computational effort (as in early proof-of-work), financial stake (in proof-of-stake),

or auctions for block construction (as in Ethereum's more recent designs), all permissionless blockchains rely on contest-like mechanisms to prevent any single actor from unilaterally controlling execution. These contests are essential for resisting sybil attacks, maintaining decentralization of participation, and securing the system against adversarial actors.

These architectural commitments are fundamental to permissionless operation and cannot be easily removed or softened. Yet the same features that support decentralization at the protocol level create economic frictions that shape competition among block creators.

First, transparency of pending transactions and fixed settlement intervals activate classic rent-seeking behavior: transparency reveals profitable opportunities—such as arbitrage or front-running—that can be exploited by whoever controls the next batch of executed transactions. Even though all agents observe the same public information, they differ in their ability to convert that information into profits through speed, computational sophistication, or superior engineering. Transparency therefore generates publicly visible opportunities whose value is unequally extractable, allowing for heterogeneity in effective rewards despite symmetry in rules and information. Time-bounded settlement amplifies this effect by providing a predictable window in which such opportunities can be exploited.

Second, the open selection of block creators—whether through proof-of-work, proof-of-stake, or builder auctions—relies on contest-like mechanisms. These contests exhibit convexity of competitive success: small differences in effective reward translate into disproportionately large differences in winning probability. As contest decisiveness increases, e.g. from proof-of-work to auctions in Ethereum's most recent PBS design, symmetric protocols magnify technological heterogeneity into persistent concentration. Rather than neutralizing differences, the competitive structure inherent in permissionless blockchains magnifies them.

Third, traders respond strategically to the risk of having their transactions exploited. Because the public queue exposes their intentions in advance, traders increasingly deliver their transactions privately to the entity most likely to execute them quickly and with minimal interference. This behavior introduces a coordination externality: the more traders route through a given entity, the larger is its effective reward, which increases that entity's likelihood of assembling the next block and, in turn, its attractiveness relative to smaller competitors. Even modest frictions in sending orders to multiple entities lead to a unique stable outcome in which nearly all valuable order flow consolidates around a single intermediary. Attempts to escape extraction simply reinforce the dominance of the leading block creator.

Together, these mechanisms create a centralizing dynamic. Transparency produces public information whose value can be harvested with unequal efficiency. Contest dynamics amplify these differences into concentrated block production. Trader behavior and routing

complementarities eliminate decentralized equilibria altogether. We show formally that no equilibrium that is stable with respect to small perturbations can remain both transparent and decentralized. Decentralized outcomes exist only under knife-edge conditions and collapse when subjected to even slight deviations. Centralization in blockchain-based markets is therefore not a by-product of flawed engineering or immature technology. It is a structural consequence of the very features required for permissionless operation. The architecture intended to eliminate intermediaries instead recreates and entrenches them. The very features meant to guarantee transparency and decentralization ultimately erode both: in equilibrium, the system becomes neither transparent nor decentralized.

The remainder of the paper develops this logic formally. Section 3 explains how transparency and latency-bounded settlement give rise to MEV, why block creators are able to capture it, and how they compete for the next block. Section 4 then introduces the general contest model that unifies PoW, PoS, and PBS environments. Section 5 derives the sorting theorem showing that higher-reward agents endogenously dominate. Section 6 characterizes how the contest's competitiveness characteristics and reward dispersion increase centralization of block creation. Section 7 connects these theoretical mechanisms to empirical patterns of MEV and private order flow observed in Ethereum and related blockchain systems. Section 8 models traders' routing decisions and shows that the only stable equilibrium of private order flow allocation is monopolistic. Together, these results yield an *impossibility result* in Section 10: under transparency and free entry, decentralized trading cannot remain decentralized in equilibrium. Section 11 concludes by discussing the broader implications for the design of transparent markets and the limits of credible decentralization.

2 Related Literature

A large and rapidly growing literature studies the economics of blockchain execution, transparency, and maximal extractable value (MEV). Early work by Daian et al. (2019) documented that transparent mempools expose traders' intentions before settlement, enabling frontrunning, sandwich attacks, and copy-trading. Their empirical analysis showed that latency-bounded settlement creates a predictable window in which fast or well-connected actors can extract value at the expense of ordinary traders. This work provided the first clear evidence that blockchain transparency produces an informational externality: public visibility of pending transactions enables predictable extraction. However, while Daian et al. show the existence and magnitude of MEV, they do not explore how these frictions shape market structure or whether decentralized trading can survive in equilibrium under transparency.

A second strand of literature studies the design and consequences of Proposer–Builder

Separation (PBS) in Ethereum blockchain execution. Heimbach et al. (2023) provide an empirical analysis of Ethereum's auction-based PBS implementation and document that block-building has quickly become concentrated among a handful of specialized builders. Wahrstätter et al. (2023) similarly show that builder markets exhibit significant bid shading, reliance on exclusive transactions, and persistent dominance by a small set of actors. Oz, Sui, Thiery, and Matthes (2024) offer complementary evidence that the strongest predictor of success in MEV-Boost auctions is access to exclusive orderflow rather than access to the public mempool, and that such exclusivity systematically drives builder concentration. These papers document the symptoms—centralization of block-building and dominance of exclusive flow—but do not provide a general model explaining why concentration arises or whether such outcomes are avoidable.

A related literature studies private orderflow, orderflow auctions (OFAs), and the proliferation of private mempools. Recent work by Pahari and Canidio (2025) and by Canidio and Pahari (2025) highlights the central role of exclusive orderflow in contemporary block-building. Pahari and Canidio show that between 77% and 84% of builder revenue comes from transactions exclusive to a single builder, using both winning and non-winning proposed blocks submitted to relays. They demonstrate that exclusive transactions, rather than public mempool competition, are the primary determinant of MEV-Boost auction outcomes. Complementary evidence in Canidio and Pahari (2025) shows persistent builder-specific differences in block content, transaction ordering, and MEV execution, indicating that heterogeneous effective rewards are large, systematic, and intrinsic to current block-production dynamics. These empirical findings directly support the central premise of our model: transparency induces economically meaningful heterogeneity in block rewards, which is magnified by contest dynamics and trader routing behavior. Gupta, Pai, and Resnick (2023) and Capponi, Jia and Olafsson (2024) show theoretically that private orderflow can cause builder centralization under PBS and that exclusive flow exacerbates winner-take-all dynamics, due to its auction-based nature. Capponi, Jia, and Wang (2023) demonstrate that private transaction pools degrade price discovery and attract informed traders, reducing the informativeness of the public mempool. Thiery (2023) documents empirically that a large share of block value stems from transactions never visible in the public mempool, and that exclusive routing relationships between searchers and builders now dominate MEV-Boost auctions. Industry analyses such as Kilbourn (2022) and Gosselin and Chiplunkar (2023) argue that OFAs naturally favor dominant builders and centralize routing. Collectively, these papers show that private transaction routing has become central to modern blockchain markets. None, however, endogenize private routing as part of a strategic equilibrium, and none analyze whether decentralized routing can persist when traders seek to protect themselves

from extraction in a transparent environment.

A broader set of empirical studies examines the microstructure of decentralized exchanges (DEXes) and the economics of arbitrage. Milionis, Moallemi, Roughgarden, and Zhang (2024) quantify losses suffered by AMMs due to stale quotes and measure arbitrageur profits arising from latency and slippage. Their results demonstrate predictable arbitrage profits consistent with transparency-driven extraction. Heimbach, Pahari, and Schertenleib (2024) perform the first detailed study of non-atomic arbitrage and find that more than a quarter of all DEX volume is attributable to cross-venue arbitrage, with over eighty percent of this volume executed by fewer than a dozen searchers. Wu et al. (2025) document more than $230 million in CEX–DEX extracted value and show that three searchers capture three-quarters of the total. These studies establish that value extraction is highly concentrated, that actors with superior speed or information dominate, and that orderflow has migrated to private channels where builders and searchers often operate as integrated entities. While these empirical findings strongly reflect the forces emphasized in this paper, existing empirical work does not provide a unifying theoretical explanation for why such concentration arises or whether transparency is fundamentally incompatible with decentralization.

At a higher level, our paper connects to classical contest theory, including the work of Tullock (1980), Nti (1998), and Konrad (2009), which shows that heterogeneity in valuations induces sorting and concentration in rent-seeking tournaments. Arnosti and Weinberg (2022) apply contest-theoretic reasoning to proof-of-work mining and demonstrate that even small asymmetries in cost or technology can generate highly concentrated mining equilibria. However, these contest models treat heterogeneity as exogenous and do not incorporate the transparency and latency frictions that generate MEV, nor do they include a routing layer in which traders strategically respond to extraction risk. Our model generalizes these insights by showing that transparency itself creates endogenous heterogeneity in rewards, and that contest structure interacts with private orderflow to magnify small differences into structural concentration.

Our analysis also connects to the foundational economics of blockchain transparency and verification. Catalini and Gans (2016) emphasize that transparency is a defining feature of blockchain systems, lowering verification costs but reshaping where market power resides. Their framework highlights how the disclosure of transaction information changes the allocation of rents and the incentives of intermediaries. While their analysis is conceptual, our results formalize a structural downside of transparency that their framework foreshadows: when settlement is latency-bounded, transparency generates predictable extraction opportunities that endogenously create heterogeneity among execution agents. Our results build on this intuition by showing that, in blockchain-based finance, transparency itself generates

the conditions under which intermediation reemerges and becomes entrenched.

Finally, our routing analysis is grounded in classic platform economics and the theory of network effects (Katz and Shapiro 1985; Farrell and Saloner 1986; Rochet and Tirole 2003). These models illustrate how coordination externalities and vertical differentiation can tip markets toward dominance by a single intermediary. In our setting, however, these features are not assumed but arise endogenously. Execution agents differ in inclusion probability due to the contest layer, and traders exhibit coordinated routing not because of exogenous platform quality, but because routing affects future block-production probabilities. Thus, the platform structure in our model is not an independent object but emerges jointly from transparency, contest dynamics, and strategic trader behavior. No prior work links the contest layer and the routing layer into a unified equilibrium system.

Taken together, this literature paints a fragmented picture: transparency produces informational leakage; MEV extraction is substantial; private orderflow undermines the public mempool; builders compete in decisive contests; and private routing increasingly centralizes execution. Yet the key question remains unanswered: *Are these patterns contingent on specific designs, or are they structural and unavoidable?*

Contribution. The contribution of this paper is to provide a unified, protocol-agnostic answer. First, we show that transparency and latency-bounded settlement *endogenously* create heterogeneity in block rewards, even when all agents are symmetric at the protocol level. Second, we prove that any such heterogeneity induces strict sorting in effort and persistent concentration in winner-take-all contests. Third, we endogenize private orderflow routing and show that once traders attempt to avoid extraction, routing equilibria either collapse to monopoly or fail to exist altogether whenever routing duplication (i.e., multihoming) costs are positive. Finally, we combine these mechanisms into an impossibility theorem demonstrating that no equilibrium that is transparent, latency-bounded, and stable with respect to small perturbations can remain decentralized. Existing work has documented centralization, measured MEV, or studied individual protocol components. This paper shows that these phenomena are not design flaws or temporary imbalances but the inevitable outcome of transparency in decentralized settlement. In doing so, it shifts the discussion from "fixing" MEV or "mitigating" centralization to recognizing the structural limits of transparent and decentralized trading itself.

3 From Transparency to MEV Extraction

Decentralized blockchains are built on the principle of transparent and time-bounded settlement. In their canonical design, transactions are broadcast publictly, visible to all partici-

pants before inclusion in a block, and blocks are finalized at fixed intervals — e.g., roughly every 12 seconds on Ethereum. This architecture aims to ensure openness and verifiability, allowing anyone to observe pending activity and validate the chain's integrity. In practice, however, this transparency introduces latency between submission and settlement, creating a window in which strategic agents can react to others' transactions before they are finalized.

The emergence of MEV. Transparency and latency together enable strategic manipulation of transaction ordering. When pending transactions are visible in the public *mempool*, any participant can act on that information before the next block is finalized. By reordering transactions, inserting their own, or censoring others, they can capture predictable value created by other users' trades.

A canonical example is the *sandwich attack*. Suppose a trader submits a large buy order for an asset. A strategic agent observing this order first buys the asset (front-running), allowing the victim's purchase to occur at a higher price, and then sells the asset (back-running). The round-trip profit arises purely from manipulating the order of inclusion—without adding information or liquidity.

Sandwich attacks are only one among many forms of extraction: others include arbitraging cross-exchange price differences, liquidating undercollateralized positions, and capturing on-chain liquidation bonuses or oracle updates. The total value that can be extracted through such transaction ordering, insertion, or censorship is known as *Maximal Extractable Value* (MEV).

MEV generalizes the notion of bid–ask spread to a transparent, time-bounded environment: whoever controls the order of inclusion in the next block captures the surplus generated by others' predictably valuable actions. Unlike traditional intermediaries, block creators are not paid for intermediation—they compete to seize it.

Although MEV remains challenging to measure precisely, several recent studies provide a sense of magnitude. For example, MEV-Explore reports at least ∼$314 million of extracted value on Ethereum since January 2020. More recent work estimates ∼$675 million prior to September 2022. Other research shows that non-atomic arbitrage alone may account for more than 10% of block value in Ethereum, and that a very small number of searchers capture the overwhelming share. These figures confirm that MEV is not a marginal nuisance but a commercially significant phenomenon.[1]

[1] cf. https://writings.flashbots.net/quantifying-mev, https://arxiv.org/abs/2405.17944, https://www.mdpi.com/2813-2203/4/3/23, https://arxiv.org/abs/2401.01622

Why block creators capture MEV. By the nature of blockchains, in each round, only one block is added to the chain. The agent who wins the right to create that block—through mining, staking, or auction bidding—determines the precise ordering of transactions within it. Because all pending transactions are public, the winner can choose the ordering that maximizes their own payoff. This makes MEV an endogenous component of the block reward. Agents differ in their ability to identify and implement profitable reorderings, for example through faster search algorithms or private information about off-chain trades. These differences translate into heterogeneous *effective rewards* from winning the block.

Competition for the next block. Block creation is a recurring contest for control over transaction ordering. In Proof-of-Work (PoW), agents expend computational effort (hashing power) and win the right to create the next block with probability increasing in their share of total hash rate. In Proof-of-Stake (PoS), validators invest capital in staking and are (pseudo-)randomly selected according to protocol rules, with selection probabilities increasing in their staked balance; the selected validator (or its delegate) assembles and publishes the block. In architectures with proposer–builder separation (PBS), a PoS-selected proposer runs a market among specialized builders who submit sealed bids for fully assembled block templates; the proposer accepts the highest bid and publishes that block. In all cases, exactly one block is accepted per round, and the winning *block creator* determines the ordering of included transactions (subject to protocol constraints), thereby capturing any associated MEV. Entry is open, but rewards are highly skewed: a small advantage in skill or speed yields a disproportionately higher probability of success.

Implications. The combination of transparency and latency-bounded settlement therefore transforms an open verification system into a tournament for execution rents. Even without collusion or formal barriers to entry, equilibrium outcomes depend on how small heterogeneities are amplified by the contest's characteristics (defined by the protocol) and by the flow of transactions seeking protection from extraction. These forces motivate the general model developed in the next section.

4 The General Contest Model (Γ)

We model the competition for the right to add a new block (and decide on its content) as a general contest, which we call *environment* Γ.

4.1 Environment Γ

The environment consists of

Players and Actions: There are $N \geq 2$ players (potential block creators). Player i chooses an effort $e_i \geq 0$ and has a valuation (reward) $R_i > 0$ for winning the indivisible prize (the right to build the block).

Contest Success Function (CSF): We use a generalized Tullock CSF. For a fixed $a \geq 1$ (the "decisiveness" parameter), define $A(e) := \sum_{k=1}^{N} e_k^a$. The probability of player i winning is

$$p_i(e) = \begin{cases} e_i^a / A(e), & A(e) > 0, \\ 0, & A(e) = 0. \end{cases}$$

As shown in Lemma 5 (Appendix A), as $a \to \infty$, this CSF approaches a highest-bidder-always-wins auction, where all probability mass is shared equally among the agents with highest efforts, i.e.,

$$p_i^{\infty}(\mathbf{e}) = \begin{cases} 1/|\arg\max_k e_k|, & \text{if } i \in \arg\max_k e_k \text{ and } \max_k e_k > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Costs and Externality: Costs are winner-dependent. The cost for player i is

$$C_i(e_i, e_{-i}) = \alpha(e_i) + \beta(e_i)p_i(e) + \gamma(e)p_i(e)$$

where $\alpha(e_i)$ is an unconditional effort cost (like cost of mining), $\beta(e_i)$ is a win-contingent cost (like own bid), and $\gamma(e) = \tilde{\gamma}(S(e))$ is a win-contingent externality (like a network fee, or other's bid).

Assumption 1 (Primitives)**.** *$\alpha, \beta : \mathbb{R}_+ \to \mathbb{R}_+$ are continuous, nondecreasing, and weakly convex, with $\alpha(e_i) + \beta(e_i) + \gamma(e_i, e_{-i}) > 0$ if $e_i > 0$ (no free lunch) and $\alpha(0) = 0$.*

Assumption 2 (Symmetric-statistic externality)**.** *The externality $\gamma(e)$ depends on a continuous, permutation-invariant (i.e. label-neutral) statistic $S : \mathbb{R}_+^N \to \mathbb{R}_+$ (e.g., max, sum, k-th order statistic) via a nondecreasing function $\tilde{\gamma} : \mathbb{R}_+ \to \mathbb{R}_+$ such that*

$$\gamma(\mathbf{e}) = \tilde{\gamma}\big(S(\mathbf{e})\big),$$

and S is almost everywhere differentiable with respect to each coordinate, with symmetric

partials: for any two indices i, j and any $\mathbf{e}$ where the partials exist,

$$\frac{\partial S}{\partial e_i}(\mathbf{e}) = \psi(e_i; \textit{multiset}(e_{-i})),$$

i.e. the functional form of the marginal is the same across labels.

Examples. *Any symmetric statistic qualifies: means, power means, ℓ_p norms, soft-maxes, or order statistics (max, kth largest, median, etc.). Order statistics are nondifferentiable only on tie profiles, which have Lebesgue measure 0. (For $a = \infty$ (hard–max), Lemma 2 below shows that strictly profitable players never tie with positive probability, so subgradients suffice.)*

Payoffs: The player's payoff is therefore

$$u_i(e; R_i) = p_i(e)(R_i - \beta(e_i) - \tilde{\gamma}(S(e))) - \alpha(e_i) \tag{1}$$

Mixed strategies and expected objective: Fix a (possibly mixed) opponents' profile σ_{-i}. For own action $e \geq 0$ define

$$\Phi_i(e; R_i \,|\, \sigma_{-i}) := \mathbb{E}_{e_{-i} \sim \sigma_{-i}}\Big[p_i\big((e, e_{-i})\big)\,\big(R_i - \beta(e) - \tilde{\gamma}(S(e, e_{-i}))\big) - \alpha(e)\Big].$$

Assumption 3 (Regularity for differentiation under the expectation)**.** *There exists $\bar{e} < \infty$ such that equilibrium efforts lie in $[0, \bar{e}]$ almost surely, and $\alpha', \beta', \tilde{\gamma}'$ and $\partial S/\partial e_i$ are bounded on these ranges. Then dominated convergence justifies interchanging $\partial/\partial e$ and the expectation on $(0, \bar{e}]$.*

Assumption 4 (Uniqueness on the profitable region)**.** *For every i, every σ_{-i}, and every $R > 0$, whenever $\sup_{e \geq 0} \Phi_i(e; R \,|\, \sigma_{-i}) > 0$, the maximizer over $e \in (0, \infty)$ is unique.*

4.2 Mapping to Canonical Environments

This general model nests several canonical environments:

- **PoW/PoS (Tullock Contest):** $a = 1$, $\alpha(e_i) = c\, e_i$, $\beta \equiv 0$, $\gamma \equiv 0$.
- **All-pay Auction:** $a \to \infty$, $\alpha(e_i) = e_i$, $\beta \equiv 0$, $\gamma \equiv 0$.
- **First-price Auction (PBS):** $a \to \infty$, $\alpha \equiv 0$, $\beta(e_i) = s_i$, $\gamma \equiv 0$.
- **Second-price Auction:** $a \to \infty$, $\alpha \equiv 0$, $\beta \equiv 0$, $\gamma(e_i, e_{-i}) = \max_{k \neq i} e_k$.

4.3 Idealistic Benchmark: Symmetric Agents

We first analyze the "idealistic" benchmark with free entry and equal rewards, $R_i = R$ for all i.

Proposition 1 (Symmetric PoW)**.** *In the PoW case ($a = 1$, $\alpha(e_i) = c\,e_i$), the payoff is $\Pi_i = e_i(\frac{R}{\sum_k e_k} - c)$. Any effort profile with $\sum_k e_k = R/c$ is a Nash Equilibrium.*

In this equilibrium, all active agents get an expected payoff of 0. Critically, this equilibrium allows for any subset of agents to be active, permitting (almost) arbitrary concentration.

Proposition 2 (Symmetric PBS)**.** *In the PBS (first-price) case ($a \to \infty$, $\beta(e_i) = e_i$), the payoff is $\Pi_i = p_i(e)[R - e_i]$. Any agent's best response is $e_i^* \in \{0, R\}$. Any profile with two or more agents bidding $e_i = R$ is an equilibrium.*

Again, all agents get 0 payoff. While concentration is capped at 50% (as 2 bidders are required), the key insight holds: **even under ideal symmetry, the architecture allows decentralization, but the equilibrium does not guarantee it**.

5 Asymmetric Rewards (R_i) and Concentration

The symmetric equilibrium is fragile. Here we show how different valuations R_i may lead to concentration.

5.1 Monotone Sorting by Reward

Intuitively, an agent with a higher reward R_i has a higher "net-prize" or "effective prize". Since the marginal cost of effort is the same, a higher R_i shifts the marginal benefit curve up, leading to a higher optimal effort e_i^*.

We now formalize this intuition. First, we establish that the objective function $\Phi_i(e; R) = \mathbb{E}[u_i(e, e_{-i}; R)]$ exhibits increasing differences in effort e and reward R.

Lemma 1 (Increasing differences)**.** *Under Assumptions 1–3, for every $e > 0$,*

$$\frac{\partial^2}{\partial e\,\partial R}\,\Phi_i(e; R \,|\, \sigma_{-i}) = \mathbb{E}\left[\frac{a\,e^{a-1}\,A_{-i}}{(e^a + A_{-i})^2}\right] \;\geq\; 0,$$

with strict inequality whenever $\Pr[A_{-i} > 0] > 0$, where $A_{-i} := \sum_{k \neq i} e_k^a$.

Proof. $\partial\Phi_i/\partial R = \mathbb{E}[p_i((e, e_{-i}))]$. Differentiate in e inside the expectation using Assumption 3. □

This property, combined with the uniqueness of the best reply (Assumption 4), ensures that an agent's best-reply effort is increasing in their own valuation.

Proposition 3 (Monotone Best Reply in R)**.** *Under Assumptions 1–4, for any fixed opponent profile σ_{-i}, the unique interior best reply $e_i^*(R|\sigma_{-i})$ is weakly increasing in R, and strictly increasing whenever $\Pr[A_{-i} > 0] > 0$.*

Proof. In the Appendix. □

This allows us to prove a strict sorting theorem for any Nash equilibrium. Because the cost and externality functions $(\alpha, \beta, \tilde{\gamma})$ are type-independent and the externality statistic $S(e)$ is symmetric (Assumption 2), the only thing distinguishing the FOCs of two profitable players is their reward R_i.

Theorem 1 (Strict sorting by R_i among strictly profitable players)**.** *Let σ be any Nash equilibrium in environment Γ satisfying Assumptions 1–4. For any two players i, j with strictly positive expected profits ($\mathbb{E}[u_i(\sigma)] > 0$ and $\mathbb{E}[u_j(\sigma)] > 0$),*

$$R_i > R_j \implies e_i^* > e_j^* \implies \mathbb{E}[p_i] > \mathbb{E}[p_j]$$

Proof. Proof in the Appendix. □

This sorting also holds in the limit as $a \to \infty$.

Lemma 2 (Increasing differences and no positive-surplus ties at $a = \infty$)**.** *Let $g(e) := \mathbb{E}[p_i^\infty((e, e_{-i}))]$. Then for $e' \geq e$ and $R' > R$,*

$$\big(\Phi_i^\infty(e'; R') - \Phi_i^\infty(e; R')\big) - \big(\Phi_i^\infty(e'; R) - \Phi_i^\infty(e; R)\big) = (R' - R)\big(g(e') - g(e)\big) \geq 0.$$

Moreover, if ties against i occur with positive probability in equilibrium, then on those tie events necessarily $R_i - \beta(e_i) - \tilde{\gamma}(S(\mathbf{e})) = 0$ and $\mathbb{E}[u_i] = 0$; hence strictly profitable players do not tie with positive probability.

Proof. g is nondecreasing in e statewise and independent of R, which gives increasing differences. The standard ε-outbid argument applies because $\tilde{\gamma}(S(\mathbf{e}))$ is fixed on a tie state when i perturbs e_i slightly upward. □

Theorem 2 (Monotone best reply and strict sorting at $a = \infty$)**.** *Let $a \to \infty$ with success rule $p_i^\infty(\mathbf{e})$. Under assumptions 1–4, for any agent i:*

1. *(*Monotone best reply*) $e_i^*(R)$ is weakly increasing in R.*

2. *(*Strictness*) If* $\Pr_{\sigma_{-i}}[\max_{k\neq i} e_k > 0] > 0$*, then* $e_i^*(R)$ *is strictly increasing in* R *on the profitable region.*

3. *(*Strict sorting in equilibrium*) In any Nash equilibrium* σ *of the hard–max game, if two players* i, j *have strictly positive expected profits, then*

$$R_i > R_j \quad \Longrightarrow \quad e_i^* > e_j^* \quad \Longrightarrow \quad \mathbb{E}_\sigma[p_i^\infty] > \mathbb{E}_\sigma[p_j^\infty].$$

Proof. Proof in the appendix. □

Discussion. The sorting theorem requires only three ingredients:

(i) strictly profitable players are interior and do not mix;

(ii) their first-order condition has the same *label-neutral common symmetric operator* at best replies; and

(iii) the best reply is strictly increasing in own valuation R.

Assumption 2—symmetry of the externality statistic—is exactly what ensures (ii). It applies broadly to symmetric aggregates, including *order statistics* such as the winner's or runner-up's effort. Nondifferentiability at ties poses no difficulty: for finite a ties occur on a null set, while for $a = \infty$ strictly profitable players do not tie in equilibrium.

The sorting result reveals the first step toward centralization. Even when all block creators face identical technology and costs, any heterogeneity in realized rewards — such as a small informational or latency advantage — translates into a persistent rank order in efforts and expected block shares. This mechanism mirrors classic rent-seeking tournaments: openness and transparency do not flatten the field but amplify small asymmetries into large market-share differences. In the blockchain context, this means that protocol-level fairness cannot guarantee equal participation once agents are payoff-maximizing under public information.

6 Drivers of Concentration

We identify two key factors that increase market concentration: the decisiveness of the contest (a) and the dispersion of rewards (R_i).

Decisiveness (a) increases concentration

Moving from PoW ($a = 1$) to PBS (approaching $a = \infty$) makes the contest more decisive. This amplifies existing advantages.

Proposition 4 (Decisiveness Increases Concentration)**.** *Consider environment* Γ *satisfying Assumptions 1–4. Let* $\sigma(a)$ *be a Nash equilibrium with a fixed set of profitable players* P*.*

(i) ***(Odds Magnification)*** *For any* $i, j \in P$*:*

$$\frac{p_i^*(a)}{p_j^*(a)} = \left(\frac{e_i^*(a)}{e_j^*(a)}\right)^a$$

If $e_i^*(a) > e_j^*(a)$*, this odds ratio is strictly increasing in* a*.*

(ii) ***(Concentration Increases)*** *As* a *increases, the sorted probability vector* $p^*(a)$ *becomes more concentrated. Any Schur-convex concentration index (e.g., the Herfindahl index* $H(a) = \sum_i (p_i^*)^2$*) is weakly increasing in* a*. Moreover, this increase is strict if and only if the equilibrium efforts of the profitable players are not all identical (i.e., if* $e_i^*(a) \neq e_j^*(a)$ *for at least one pair* $i, j \in P$*).*

Proof. In the appendix □

Reward Dispersion (R_i) increases concentration

Increasing the advantage of the top player also drives concentration.

Proposition 5 (Concentration and Value Dispersion)**.** *Consider environment* Γ *satisfying Assumptions 1–4. Fix* a*. Assume a Nash equilibrium* $\sigma(R)$ *with a locally fixed set of profitable players* P*. Assume efforts are strategic substitutes (*$\partial e_j^*/\partial R_i \leq 0$ *for* $j \neq i$*) and* $\partial e_i^*/\partial R_i > 0$*. Let player* $1 \in P$ *be the top player (*$R_1 > R_j$ *for* $j \neq 1$*). Then,*

$$\frac{d}{dR_1}\left(\frac{p_1}{p_j}\right) > 0 \quad \text{for all } j \in P \setminus \{1\}$$

Consequently, the within-P Herfindahl index strictly increases with R_1*.*

Proof. In the appendix □

The two comparative-statics results in this section capture distinct, yet complementary, channels through which concentration arises in decentralized block production. The first channel operates through the *decisiveness* of the contest, governed by the parameter a. As a increases, the success function approaches a winner-take-all rule. Small differences in effort or skill then translate into large differences in winning probability. In blockchain terms, moving from a low-decisiveness environment such as Proof-of-Work—where block rewards are allocated probabilistically in proportion to computational power—to a high-decisiveness architecture such as Proposer–Builder Separation effectively sharpens competition into an all-pay auction for block inclusion. Protocol upgrades that aim to increase efficiency or reduce latency thus inadvertently magnify inequality of outcomes: they raise the elasticity of success with respect to effort, so that even minute advantages in speed, capital access, or information lead to outsized gains in block share. The contest becomes more selective, but also more concentrated.

The second channel arises from *reward dispersion* across participants. When some block creators obtain higher effective valuations R_i—because of superior MEV-extraction skill, better search infrastructure, or privileged access to private order flow—the monotone best-reply property ensures that these agents invest more effort, win more frequently, and accumulate larger market shares. Proposition 5 formalizes this intuition: any widening of the reward gap mechanically increases concentration, as measured by standard indices such as the Herfindahl–Hirschman Index. In equilibrium, higher-valued agents expand their dominance until lower-valued rivals are driven to the margin of profitability. The resulting structure resembles that of rent-seeking tournaments or "arms races," where the most capable player endogenously captures a disproportionate share of the prize.

Taken together, these forces reveal why decentralization is inherently fragile. Increasing decisiveness—often motivated by technical improvements in block production—and increasing heterogeneity in extraction capability both tilt the equilibrium toward a smaller number of dominant builders. Importantly, this outcome does not depend on explicit collusion, entry barriers, or strategic manipulation; it is an equilibrium response to transparent competition under free entry. The very mechanisms that make the protocol more efficient and open also amplify small asymmetries into large and persistent concentrations of market power.

7 The Source of Asymmetric Rewards: MEV and POF

We have shown that asymmetries in R_i lead to concentration. Now we ask: why do different R_i emerge in DeFi? The answer is MEV.

7.1 Transparent Market (Public Mempool)

In a classic transparent market, all transactions are in a public pool. As explained in Section 3, block creator's reward depends on a baseline protocol reward R plus their ability to extract value from this public pool:

$$R_i = R + r_i B \,.$$

Here, B is the total value of a block created from the public pool, and r_i is the agent's intrinsic technical skill in extracting that value. Since all block creators have access to the same transaction pool, their rewards are differentiated only by their skill r_i.

7.2 Private Order Flow (POF)

To avoid being exploited in the public pool, traders route transactions directly to block creators. A block creator's reward now depends on the private flow it attracts, T_i, as well as any public market flow, T_m:

$$R_i = R + r_i(T_i v + T_m v)$$

(where v is the value per transaction). The crucial difference is that **an agent's reward R_i is now endogenous** to the private flow T_i it attracts. This creates a direct link between a block creator's popularity and its fundamental reward, setting the stage for a coordination game.

8 The Centralizing Force of POF

We now model the trader's choice, which reveals a powerful network effect that leads to a stable monopoly.

8.1 The Trader's Choice

A rational trader with a transaction of value v chooses between the public mempool and a private flow to a single block creator.

- **Public mempool:** The trader is subject to extraction by whoever wins the block. The payoff is $\Pi_{public} = v(1 - \tau_M)$, where $\tau_M = \sum_i p_i r_i$ is the market-wide average extraction rate.

- **Private to block creator j:** The trader is only subject to extraction if j wins and exercises it. The payoff is $\Pi_{POF}(j) = p_j v(1 - r_j)$.

A rational trader routes directly to the block creator j that offers the maximal "surplus share" $S_j = p_j(1 - r_j)$, provided this is better than the public option ($S_j > 1 - \tau_M$).

8.2 Network Effects and Stability

This choice creates a powerful network effect.

Lemma 3 (Incentive for Trader Coordination under POF)**.** *Fix $v > 0$ and extraction abilities $\{r_i\}_{i=1}^N$. Let a fixed number of traders traders choose destinations among block creators or the public option. For any block creator i, a trader's payoff from routing to i,*

$$\Pi_{POF}(i) \;=\; p_i^*\, v\,(1 - r_i),$$

is strictly increasing in the number of traders T_i who route to i.

Proof. In the appendix. □

This creates a positive externality: more traders joining a POF pool increases the payoff for all members, incentivizing coordination.

8.3 Equilibrium: A Stable Monopoly

Accounting for strategic choice of block creators' r_i, creates a coordination game with two types of candidate equilibria: a monopolistic one where one block creator captures all POF, and a symmetric sharing one where $k \geq 2$ block creators share the flow equally. We find that only the monopoly is stable.

Proposition 6 (Stability of POF Equilibria)**.** *In the POF coordination game (if for some i, $S_i > 1 - \tau_M$):*

*(a) A symmetric sharing allocation with multiple block creators sharing the private order flow is **unstable**.*

*(b) A monopolistic allocation with all POF routed to a single block creator is a **stable** Nash Equilibrium.*

(c) There are no other Nash Equilibria.

Proof. In the appendix. □

Lemma 3 shows that private-order-flow routing introduces a strong network complementarity: a trader's payoff from joining a given builder rises with the number of traders already there. This coordination externality is familiar from platform economics but arises here endogenously from the contest environment—because larger private pools increase a builder's reward and equilibrium inclusion probability. Proposition 6 then establishes that this feedback has a unique stable outcome: monopoly. Any symmetric sharing equilibrium unravels as traders migrate toward the slightly larger or less-extractive builder, and no offsetting force restores balance. The implication is stark: once traders flee the public mempool to avoid MEV, they inevitably centralize on a single trusted intermediary. The market's attempt to escape rent extraction reproduces it.

A natural question is why the outcome in Proposition 6 is not disciplined by Bertrand-style price competition. If traders prefer lower extraction rates, why do rival builders not simply undercut the dominant one? The key reason is that builders do not offer homogeneous goods. Even with full commitment and zero marginal cost of lowering r_i, each builder provides a different effective service quality determined by its inclusion probability p_i^*. A trader's expected payoff $\Pi_{POF}(i) = p_i^*(1 - r_i)v$ depends multiplicatively on both terms, so a small builder with low p_i^* cannot match the dominant builder's value to traders even by setting $r_i = 0$. This is a case of *vertical differentiation*: traders value not only low extraction, but also the reliability that their transaction will be included.

Furthermore, increasing returns from network effects make demand discontinuous. As more traders join a given builder, its p_i^* rises, improving its relative appeal to all others. A tiny improvement in expected payoff thus tips all traders toward the leading builder, leaving no smooth demand curve that would support a standard Bertrand equilibrium. The equilibrium is a corner solution—a single builder capturing nearly all private order flow—rather than a competitive price equilibrium.

Finally, the public mempool provides an outside option that pins down, but does not eliminate, the dominant builder's rent: the monopolist need only offer a surplus share slightly above $(1 - \tau_M)$ to retain all traders. Hence, even with full commitment and zero pricing cost, Bertrand competition cannot restore decentralization once network externalities and vertical differentiation are taken into account.

These results together map a consistent logic: transparency creates observable opportunities for extraction; contest design amplifies small advantages into dominance; and traders' self-protection responses further entrench that dominance. The combination of Sections 5–8 thus delivers an impossibility result: open and transparent trading cannot sustain decentralization in equilibrium.

8.4 Multi-Homing Traders

In the analysis above, traders *single-home*: each privately routes her transaction to one Execution Agent (EA, or block creator) who can include it in the next block. We now relax this and allow traders to send their order to multiple EAs at once, at the cost of some value loss. This extension shows that the centralizing forces identified earlier remain robust—and in fact become stronger—once duplication costs are positive.

Multi-homing Environment. Let a trader choose a subset $S \subseteq \{1, \ldots, N\}$ of EAs to whom she privately routes her transaction. If at least one $i \in S$ wins the right to produce the block, the transaction executes there and is exposed to extraction r_i. Sending to multiple EAs increases redundancy, latency, and privacy leakage; we model these frictions as a per-transaction loss $\ell(|S|)$, with $\ell(1) = 0$, $\ell'(m) \geq 0$, and $\ell''(m) \geq 0$. Let $v > 0$ denote the gross value of a successfully executed trade.

Trader's routing problem. Given EAs' expected inclusion probabilities $\mathbf{p} = (p_1, \ldots, p_N)$ and extraction rates $\mathbf{r} = (r_1, \ldots, r_N)$, the expected payoff to a trader choosing S is

$$\Pi(S) \;=\; v \sum_{i \in S} p_i(1 - r_i) \;-\; v\,\ell(|S|). \tag{2}$$

Define each EA's *surplus share* as

$$S_i \;:=\; p_i(1 - r_i).$$

Sorting them in decreasing order, $S_{(1)} \geq S_{(2)} \geq \cdots \geq S_{(N)}$, the trader's optimal routing set has a cutoff form:

$$S_{(m)} \;\geq\; \Delta\ell(m) \quad \text{for all } m \leq m^*, \quad S_{(m^*+1)} \;<\; \Delta\ell(m^* + 1), \tag{3}$$

where $\Delta\ell(m) := \ell(m) - \ell(m-1)$ is the marginal loss from adding the m-th EA. Hence traders include the top m^* EAs whose marginal contribution exceeds the marginal duplication cost.

Equilibrium structure. EAs' inclusion probabilities p_i depend on the aggregate POF mass they attract. Let q_i denote the amount of trades routed to EA i. We consider a SPNE, characterized by a pair $(\mathbf{q}, \mathbf{r})$, of the sequential *fee–routing* game, where:

1. Each EA i chooses r_i optimally given its expected inflow q_i.
2. Each trader's routing choice S maximizes (2) given $\mathbf{p}(\mathbf{q})$ and $\mathbf{r}$.

Equilibrium Pricing-to-the-Margin. In any equilibrium, active EAs will not leave rents on the table. They will raise their extraction rate r_i until the trader's participation constraint binds.

Lemma 4 (Pricing-to-the-Margin)**.** *In any equilibrium where exactly $m \geq 1$ EAs are actively included by traders, each active EA i sets its fee r_i^* such that its surplus share equals the binding threshold:*

$$\Phi_i \;=\; p_i(q_i)\left(1 - r_i^*\right) \;=\; T_m, \qquad \textit{hence} \qquad r_i^* \;=\; 1 - \frac{T_m}{p_i(q_i)}.$$

Any inactive EA j must satisfy $\Phi_j \leq \Delta\ell(m+1)$.

Proof. For $i \in A_m$, suppose $\Phi_i > T_m$. Increasing r_i raises profit and keeps i included because T_m is the binding threshold for the last included route: adding/removing the m-th EA changes the trader's payoff by $\Delta\ell(m)$, and the mempool imposes $\Phi_i \geq \Phi_M$; hence the relevant constraint is $\Phi_i \geq T_m$. Thus $\Phi_i > T_m$ leaves slack and is strictly suboptimal. Conversely, $\Phi_i < T_m$ would drop i from traders' optimal set. Therefore the optimum is $\Phi_i = T_m$, yielding (4). For $j \in I_m$, if $\Phi_j > \Delta\ell(m+1)$, traders would strictly benefit from adding a $(m{+}1)$-st EA, contradicting that exactly m are included. □

Pass-through of duplication costs. When the duplication margin binds ($T_m = \Delta\ell(m) \geq \Phi_M$), the fee formula reduces to

$$r_i^* \;=\; 1 - \frac{\Delta\ell(m)}{p_i(q_i)}.$$

Thus any increase in the marginal duplication loss $\Delta\ell(m)$ is passed through one-for-one into higher equilibrium fees, scaled by $1/p_i(q_i)$. Intuitively, multi-homing costs become "rents" that active EAs can extract because they set fees up to the traders' routing cutoff.

Equilibrium with atomic traders under costly mutli-homing. We first consider a setting where traders are *atomic*: each trader $t \in \mathcal{T}$ has mass $\varepsilon_t > 0$. Traders may multi-home by routing the same order to multiple execution agents (EAs), but duplication incurs a convex loss $\ell(m)$ with $\ell(1) = 0$ and strictly positive marginal cost at the second route:

$$\Delta\ell(2) \;:=\; \ell(2) - \ell(1) \;>\; 0.$$

Let q_i denote the total private order flow routed to EA i and EA i's expected inclusion probability be $p_i = p_i(q)$.

Corollary 1 (Properties of inclusion probabilities under Tullock contest)**.** *Under the Tullock contest success function* $p_i(e_i, e_{-i}) = e_i^a / \sum_k e_k^a$ *with* $a \geq 1$*, the induced mapping* $p_i(q)$ *obtained from equilibrium efforts* $e_i^*(q_i)$ *is*

(i) continuous in q*,*

(ii) strictly increasing in own q_i*,*

(iii) weakly decreasing in rivals' q_{-i}*, and*

(iv) satisfies $p_i(0, q_{-i}) = 0$ *and* $0 \leq p_i \leq 1$*.*

Proof. Immediate from differentiating p_i w.r.t. e_i and e_j and using $\partial e_i^*/\partial q_i > 0$ from Proposition 3. □

These properties imply all results below.

Proposition 7 (No Market-Sharing NE with Atomic Traders)**.** *Assume costly multi-homing* ($\Delta\ell(2) > 0$) *and atomic traders. There is* ***no*** *Nash equilibrium in which at least two EAs receive positive private order flow.*

Proof. Suppose, for contradiction, an equilibrium with $m \geq 2$ active EAs exists. By Lemma 4, every active EA i must have $\Phi_i(q) = T_m$. Pick any trader t who routes to at least two EAs, i and j. Consider t's unilateral deviation to *drop* EA j. The new aggregate flow to j is $q'_j = q_j - \varepsilon_t$. Since p_j is strictly increasing in its own flow (by Corollary 1), $p_j(q') < p_j(q)$. This means the new surplus share for j is $\Phi_j(q') = p_j(q')(1 - r_j) < p_j(q)(1 - r_j) = T_m$. The trader t's marginal gain from re-adding j would be $v[\Phi_j(q') - \Delta\ell(2)]$. Since $T_m \geq \Delta\ell(2)$, we have $\Phi_j(q') < T_m \leq \Delta\ell(2)$ (or $\Phi_j(q') < \Delta\ell(2)$), making this deviation strictly profitable. Thus, any trader routing to $m \geq 2$ EAs has a strict incentive to deviate and drop one. No such equilibrium can exist. □

Assumption 5 (No-solo-entry condition)**.** *For any block creator* k*,* $\varepsilon_k < \bar{\varepsilon}$ *where* $\bar{\varepsilon}$ *is defined as*

$$p(\bar{\epsilon}) = \Delta\ell(2)\,.$$

Assumption 5 says: a single atomic trader of size cannot, by themselves, make it privately beneficial to open a brand-new second route. It assumption holds, for example, if EA with no private flow has negligible inclusion chance.[2]

[2]Our model shows forces that lead to concentration even when AEs have no fixed cost of operating (so they enter at any positive margin). Fixed cost limits the number of potential competitors on the fringe and makes concentration more prevalent.

Theorem 3 (Monopoly POF equilibrium with atomic traders)**.** *Under* $\Delta\ell(2) > 0$ *and Assumption 5, the only SPNE in the fee–routing game is a monopoly with a single active EA. In any such equilibrium, all traders single-home to the same EA* $i^\star$ *and the active EA prices to the mempool margin:*

$$\Phi_{i^\star}(q) \;=\; \max\{\Delta\ell(1),\, \Phi_M\} \;=\; \Phi_M, \qquad r^*_{i^\star} \;=\; 1 - \frac{\Phi_M}{p_{i^\star}(q)}. \tag{4}$$

Proof. (*Existence.*) Fix an EA $i^\star$ and let all traders route exclusively to $i^\star$; thus $q_{i^\star} = \sum_t \varepsilon_t$ and $q_k = 0$ for $k \neq i^\star$. Let $i^\star$ set $r_{i^\star}$ so that $\Phi_{i^\star}(q) = \Phi_M$ (pricing to the outside option); inactive EAs choose arbitrary $r_k \in [0, 1)$. Consider any trader t of mass $\varepsilon_t \leq \bar{\varepsilon}$ and any inactive EA $k \neq i^\star$. If t deviates by adding k as a second route, the net gain is

$$\Pi_t(\{i^\star, k\}) - \Pi_t(\{i^\star\}) \;=\; v \cdot \sup_{r_k \in [0,1)} p_k\big(\varepsilon_t\big)\,(1 - r_k) \;-\; v\,\Delta\ell(2) \;<\; 0,$$

by Assumption 5. Thus no trader gains by adding any second route. Dropping $i^\star$ in favor of the mempool never helps since $\Phi_{i^\star} = \Phi_M$. Hence the profile is a Nash equilibrium.

(*Uniqueness of the number of actives.*) Proposition 7 rules out any equilibrium with $m \geq 2$ actives. Therefore, any equilibrium must be monopoly POF. □

Comparative statics (pricing-to-multi-homing). Holding $p_i(q_i)$ fixed, (4) implies:

$$\frac{\partial r^*_i}{\partial T_m} = \frac{1}{p_i(q_i)} \;>\; 0 \qquad \text{(fees rise one-for-one with the binding threshold)},$$

$$\frac{\partial r^*_i}{\partial p_i} = \frac{T_m\, p_i'(q_i)}{p_i(q_i)^2} \;>\; 0 \qquad \text{(higher inclusion probability permits higher fees)},$$

$$\frac{\partial r^*_i}{\partial \Phi_M} = \begin{cases} \frac{1}{p_i(q_i)} & \text{if } \Phi_M > \Delta\ell(m), \\ 0 & \text{if } \Phi_M < \Delta\ell(m). \end{cases} \qquad \text{(fees increase when the mempool outside option binds)}.$$

Because $\Delta\ell(\cdot)$ is increasing, T_m is weakly increasing in m. Hence, *for a given* p_i, greater multi-homing (larger m) pushes up the pricing target T_m and thus raises r^*_i.[3]

Discussion. Proposition 7 shows that *because* an atomic trader's deviation moves q and hence lowers the dropped EA's attractiveness $\Phi_j(q)$ below the duplication threshold, multi-homing cannot survive in equilibrium when $\Delta\ell(2) > 0$. Theorem 3 complements this with

[3] Endogenously, $p_i(q_i)$ typically falls when m grows (each EA receives a smaller share), partially offsetting this direct effect; the net effect on r^*_i is ambiguous ex ante and depends on how quickly p_i declines relative to the increase in T_m.

a simple existence condition: if a lone atomic trader cannot profitably "seed" a second EA above the duplication margin, the unique equilibrium market structure features a single active EA who prices to the mempool margin.

Profitable Seeding: What if a Second EA Can Be Jump-Started? The no–solo–entry condition ensures that a single trader of arbitrarily small mass cannot profitably "seed" a new execution agent (EA). It acts as a minimal entry barrier: without it, any inactive EA can find a fee low enough to attract a small ε–mass trader who is willing to multihome, thereby opening a second private route. Once this condition is violated, the two–stage routing game admits no subgame–perfect Nash equilibrium (SPNE) with positive private orderflow (POF).

Proposition 8 (No SPNE with positive POF under profitable seeding)**.** *Suppose* $\Delta\ell(2) > 0$ *and Assumption 5 is violated. Then in the two–stage fee–routing game* Γ *there is no subgame–perfect Nash equilibrium—pure or mixed—in which private orderflow is positive in some subgame. In any SPNE of* Γ *(if one exists), POF is never used in equilibrium.*

Proof. In the appendix. □

Economic Interpretation and Implications for Decentralization The economic logic behind Proposition 8 is direct. If a new EA can attract *even one* trader's orderflow profitably, then a monopolist EA can never maintain a positive margin: an entrant can always slightly undercut the incumbent and capture an ε–sized foothold. Transparency ensures that entrants can identify exactly when such seeding is profitable, and the two–stage timing guarantees that traders will immediately reroute in response. The consequence is not heightened competition but *nonexistence of equilibrium*: any fee profile that leaves an exploitable margin is instantly undermined by entry.

Thus transparent decentralized routing becomes inherently unstable once small–scale entry is profitable. Every marginal trader with mass $\varepsilon_k \geq \bar{\varepsilon}$ becomes a vehicle for destabilizing the incumbent, and every EA faces perpetual undercutting pressure. Without some minimal entry barrier—here, the no–solo–entry condition—the routing layer cannot sustain a stable fee vector or a stable market structure. The only SPNE (if one exists at all) is an equilibrium in which POF is never used.

This instability has direct implications for decentralized finance. First, *transparency* makes contestability extreme: any EA can see precisely how low a fee must be to trigger profitable seeding. Second, *atomic traders* "shatter" stability: a single small trader is enough to overturn an incumbent's position and break any attempt at sustaining markups. Third,

because profitable seeding is triggered by arbitrarily small deviations, decentralized routing cannot support either competitive fees or a diversified market structure when entry is frictionless.

Removing the no–solo barrier therefore does not produce "more decentralization." It produces *nonexistence*: a breakdown of the private–routing layer in equilibrium. In practice, this pushes the system toward dynamic dominance, path dependence, and re–centralization, as the only sustainable long–run outcomes either eliminate private orderflow altogether or rely on some form of persistent entry–deterring power for a dominant EA.

Relation to the classical IO literature The structure uncovered in Proposition 8 is closely related to the classical theory of contestable markets. In the canonical models of Baumol, Panzar and Willig (1982), an incumbent firm cannot sustain a price above average cost if entrants can engage in costless "hit–and–run" entry: an entrant that can profitably serve even a small portion of demand has a profitable deviation, and pure or mixed price equilibria may fail to exist unless the incumbent exactly sets the limit price. Similar logics appear in Spence (1977) and Dixit (1979), where potential entrants that can profitably enter at arbitrarily small scale make it impossible to sustain monopoly markups without some form of entry deterrence. As in these models, violation of the no–solo entry condition eliminates any stationary price equilibrium and forces the system toward nonexistence.

This logic is also deeply connected to canonical IO models in which *dynamic instability* replaces static equilibrium. In Edgeworth's classical analysis of price competition with capacity constraints, pure and mixed equilibria generically fail to exist; instead, firms engage in *perpetual undercutting and restorative price increases*, yielding the well-known *Edgeworth price cycles*. More formally, Maskin and Tirole (1988) show that in repeated price competition with small frictions or capacity constraints, the entire dynamic system may cycle: prices rise until undercut, collapse, and then rebuild, generating no stationary symmetric equilibrium. These models mirror the structure here. Small-scale profitable deviations (our "seedind") produce discontinuities in best replies and destroy fixed points of the dynamic operator, producing cycles or chaotic adjustment paths rather than convergence.

Our environment thus inherits two foundational IO insights. First, from contestability: when arbitrarily small entrants can profitably "hit and run," no static price equilibrium exists. Second, from dynamic price competition: when payoffs exhibit such discontinuities, the natural dynamic adjustment process generates *cycles*, not convergence. In our context, fee competition among EAs would therefore not settle at a decentralized steady state; instead, incumbents would repeatedly attempt to raise fees, entrants would repeatedly undercut to capture an $\bar{\varepsilon}$–foothold, and the market would cycle endlessly.

The mathematical parallel is direct. Fee competition among EAs is continuous; the ε–mass seeding trader plays the role of a small foothold segment of demand; and violation of the no–solo entry condition corresponds to the absence of fixed costs or minimum–scale constraints. Under these conditions, any positive fee margin invites profitable seeding by an entrant. As in the contestability and Edgeworth literatures summarized in Tirole (1988), the combination of continuous pricing, free small-scale entry, and payoff discontinuities induced by best-response behavior generically leads to nonexistence or dynamic cycling rather than a stable decentralized market structure.

When is Active Market-Sharing an Equilibrium?

Theorem 4 (Market-sharing SPNE with nonatomic traders: existence and characterization)**.** *Fix $m \geq 2$ and a nonempty set A_m of m EAs. There exists a subgame-perfect Nash equilibrium (SPNE) of the two-stage fee–routing game in which precisely the EAs in A_m are active and* every trader routes to all $i \in A_m$ *if and only if there exists a routing profile q with $q_i > 0$ for $i \in A_m$ and $q_j = 0$ for $j \notin A_m$ such that:*

(F1) ***Feasible pricing:*** *$p_i(q) \geq T_m$ for all $i \in A_m$ (so $r_i = 1 - T_m/p_i(q) \in [0,1)$ is well defined),*

(F2) ***No profitable*** *$(m{+}1)$**-st route:*** *$p_j(q) \leq \Delta\ell(m{+}1)$ for all $j \notin A_m$.*

When (F1)–(F2) *hold, the fee profile $r_i^* = 1 - T_m/p_i(q)$ for $i \in A_m$ and arbitrary $r_j \in [0,1)$ for $j \notin A_m$ supports a market-sharing SPNE in which each trader routes to all $i \in A_m$ and to none in I_m.*

Proof. (*If.*) Given (F1), set $r_i^* = 1 - T_m/p_i(q)$ for $i \in A_m$ so that $\Phi_i = T_m$; by Lemma 4, this is revenue-maximizing for active EAs given traders' continuation payoffs.

Consider now the Stage–2 routing subgame following $\mathbf{r}^*$. Given q and r^*, Lemma 4 implies each trader is (weakly) indifferent at every step $k \leq m$ and strictly against adding an $(m{+}1)$-st EA by (F2). Since traders are nonatomic, a unilateral deviation does not move q (thus not Φ), so routing to all $i \in A_m$ is a best reply. Hence, for the subgame following $\mathbf{r}^*$, q is a Nash equilibrium of the Stage–2 game, and (q, r^*) is obtained by backward induction. Therefore (q, r^*) constitutes a SPNE of the two-stage game with market-sharing.

(*Only if.*) In any such SPNE with market-sharing, Lemma 4 forces $\Phi_i = T_m$ on all active EAs, implying $p_i(q) \geq T_m$ (else $r_i \notin [0,1)$), giving (F1). If (F2) failed for some $j \notin A_m$, adding j as the $(m{+}1)$-st route would strictly increase a trader's payoff in the Stage–2 subgame, contradicting that the induced routing profile is a Nash equilibrium. Thus (F2) is necessary. □

Costless multihoming as a special case. If multihoming is costless, $\Delta\ell(m) \equiv 0$, then $T_m = \Phi_M$ for all m. The feasibility conditions reduce to $p_i(q) \geq \Phi_M$ for actives and $p_j(q) \leq 0$ for inactives (automatically true). Hence we obtain a large set of market-sharing equilibria:

Corollary 2 (Continuum of market-sharing SPNE when multihoming is free)**.** *If $\Delta\ell(m) \equiv 0$, then $T_m = \Phi_M$ for all m. For any nonempty active set A_m and any routing profile q with $p_i(q) \geq \Phi_M$ for all $i \in A_m$, there is a SPNE of the two-stage game supported by fees $r_i^* = 1 - \Phi_M/p_i(q)$ in which every trader routes to all $i \in A_m$. These equilibria are non-strict (traders are indifferent at the cutoff) and form a continuum (any split q across A_m works).*

When multihoming is literally costless ($\Delta\ell(m) = 0$), traders face no penalty for adding additional EAs. Their routing problem collapses to a simple rule: include every EA whose surplus share $\Phi_i \equiv p_i(1 - r_i)$ weakly exceeds the mempool outside option Φ_M. Because each active EA prices exactly to Φ_M, all active EAs provide exactly the same surplus, and traders are indifferent across how many of them to include.

This produces a *degenerate knife-edge* structure:

- Any split of flows q across the active set A_m works, as long as every active EA satisfies $p_i(q) \geq \Phi_M$.
- Any such q supports an SPNE.
- But the equilibrium is *not strict*: traders are exactly indifferent at $\Phi_i = \Phi_M$.

Hence, even arbitrarily small aggregate shocks to q (or to $p_i(q)$) can violate the condition $p_i \geq \Phi_M$ for some $i \in A_m$, causing that EA to be dropped and the system to "tip" toward a smaller active set. Equilibria exist, but none are trembling-hand perfect.

When $\Delta\ell(2) > 0$, Theorem 4 permits market-sharing SPNE provided (F1)–(F2) hold. However, they are also *non-strict* (indifference at $\Phi = T_m$) and not trembling-hand perfect: any arbitrarily small aggregate perturbation of q changes $p_i(q)$, violates $\Phi_{(2)} \geq T_2$, and triggers pruning to single-homing.

Comparative statics.

1. If $\Delta\ell(2) = 0$, multi-homing can be sustained in a knife-edge equilibrium where all traders send to all EAs with $S_i > 0$. Yet any infinitesimal shock in S_i breaks indifference and tips the system toward single-homing.
2. As $\Delta\ell(2)$ increases, the threshold for including a second EA rises, further strengthening the tipping effect toward monopoly POF.

3. Traders' option to use the public mempool (with surplus share $S_M = 1 - \tau_M$) bounds the monopolist's extraction rate but does not prevent concentration; the dominant EA needs only to satisfy $S_1 \geq S_M + \Delta\ell(2)$.

Total surplus decomposition. For a representative trade routed to m EAs in equilibrium,

$$\underbrace{v \sum_{i \in A_m} \Phi_i}_{\text{gross preserved value}} - \underbrace{v\,\ell(m)}_{\text{duplication loss}} = \underbrace{v\,m\,T_m}_{\text{trader-side value delivered}} - v\,\ell(m),$$

while the per-EA fee is $r_i^* = 1 - T_m/p_i(q_i)$, so expected EA extraction (per included EA) is $v\,[p_i - T_m]$, rising in p_i and in T_m. Thus, increases in marginal duplication costs $\Delta\ell(m)$ are *shared*: they directly reduce trader value via $\ell(m)$ and simultaneously raise r_i^* through T_m, transferring additional surplus to active EAs.

Implications. (i) Active EAs *price to the marginal inclusion constraint* T_m (duplication margin or mempool)—a multi-homing analogue of pricing to the participation constraint. (ii) For given inclusion probabilities, fees are *increasing* in the degree of multi-homing m whenever the duplication margin binds. (iii) Because p_i itself rises with the flow directed to i, the system exhibits a feedback loop: small increases in q_i allow both higher p_i and higher r_i^*, reinforcing the leader's advantage and undermining the stability of multiplicity unless duplication is literally costless.

Robustness to heterogeneous transaction sizes. A natural question is whether heterogeneity in traders' transaction values affects our concentration results. In our baseline environment, allowing traders to differ in value v_t does not alter any equilibrium characterization. Because both public and private routing payoffs scale linearly in v_t, the ranking of execution agents depends only on their surplus shares $p_i(1 - r_i)$, not on the trader's value. All traders therefore make identical routing decisions, and the effective block rewards R_i remain proportional across agents, leaving the comparative statics governing concentration unchanged.

Indeed, relaxing this linearity tends to *reinforce* our results. If execution agents can condition fees on v_t, or if duplication costs or extraction opportunities scale more than proportionally in v_t, higher-value traders are disproportionately attracted to the leading agent, amplifying reward dispersion and thereby strengthening the odds-magnification and sorting mechanisms that drive concentration. Heterogeneity in trade sizes is therefore not only innocuous but, under richer microfoundations, further exacerbates the centralizing forces we

identify.

9 Empirical Evidence

The mechanisms derived in Sections 5–8 correspond closely to empirical patterns observed in deployed blockchain systems. This section summarizes the evidence documenting heterogeneous rewards, sorting dynamics, contest-driven concentration, and the consolidation of private order flow, all of which align with the theoretical forces identified above.

A growing empirical literature documents the market patterns that motivate and support the mechanisms in our model. Across multiple independent datasets—including transaction-level records from Ethereum's execution layer, relay-provided MEV-Boost logs, and high-frequency data from centralized exchanges—a consistent picture emerges: transparency generates extractable value; this value is captured by a small number of specialized intermediaries; these intermediaries rely increasingly on private or exclusive routing arrangements rather than the public mempool; and block-building under proposer–builder separation (PBS) is already highly concentrated. Together, these findings provide the empirical foundation for our claim that transparent and decentralized trading are structurally incompatible in latency-bounded blockchain environments.

Early work on mempool dynamics by Daian et al. (2019) established the basic causal link between transparency and extractable value. Using detailed mempool traces, they showed that publicly visible transactions often reveal profitable opportunities that latency-advantaged actors can capture through frontrunning, copy-trading, and reordering. Their evidence demonstrated that transparency in transaction submission reliably leaks information about trader intentions, enabling predictable and contestable rents. This foundational mechanism—where transparency exposes profitable opportunities that speed-sensitive actors can systematically exploit—is the empirical starting point for the heterogeneity in effective rewards that appears in our model, even when base protocol rewards are symmetric.

Subsequent empirical studies show that the extraction of this value is highly concentrated. Heimbach, Pahari, and Schertenleib (2024), in the first comprehensive measurement of non-atomic (CEX–DEX) arbitrage, find that more than one-fourth of all volume on Ethereum's major DEXes can be attributed to this category of MEV. Strikingly, only eleven searchers execute more than eighty percent of non-atomic arbitrage volume, and the two largest entities alone account for nearly half. Their analysis shows that these opportunities are not dispersed among many traders, but rather are captured by a small and persistent set of specialized actors whose scale, execution infrastructure, and low-latency access place them beyond the competitive reach of ordinary traders. Wu et al. (2025) reach a similar conclusion in a dataset

covering nineteen months of CEX–DEX activity, documenting approximately $233 million in extracted value, three-quarters of which is captured by three searchers. They further show that searcher profitability depends heavily on relationships with specific builders and that vertically integrated searcher–builder pairs consistently outperform their non-integrated competitors.

A parallel empirical literature documents the growing role of private orderflow and exclusive routing. Thiery (2023), using a large sample of relay and mempool data, shows that between one-quarter and one-third of all transactions included by builders never appear in the public mempool. Although these exclusive transactions constitute a small fraction of total transaction count, they regularly account for the majority of value transferred to builders, often exceeding eighty percent of builder revenue in MEV-Boost auctions. Oz, Sui, Thiery, and Matthes (2024) corroborate this pattern by showing that the most successful builders rely disproportionately on exclusive transactions and coinbase transfers from searchers. In their data, builders' market shares are strongly correlated with access to exclusive orderflow, suggesting that block-building competition under PBS is largely a competition for private, not public, orderflow.

Exclusive routing is similarly pervasive on the searcher side. Wu et al. (2025) demonstrate that many of the largest CEX–DEX searchers send the majority of their trades—often between fifty and one hundred percent—to a single builder. These exclusive relationships are not incidental: the authors show that searcher volume and builder market share mutually reinforce each other over time. For example, during a four-month period in 2024, the daily share of CEX–DEX trades directed to Titan by the searcher Kayle is strongly correlated with Titan?s daily builder market share, both contemporaneously and with a one-day lag. This dynamic illustrates the feedback mechanism implied by our contest model: heterogeneous rewards create persistent advantages for particular execution agents, which in turn induce traders to route their transactions to those same agents, further amplifying concentration.

Recent empirical work by Pahari and Canidio (2025) provides some of the clearest evidence that high-value orderflow on Ethereum is overwhelmingly exclusive. Using a dataset of both winning and non-winning proposed blocks submitted to relays, they show that between 77% and 84% of all builder revenue comes from transactions that are exclusive to a single builder. Public-mempool transactions, while numerous, contribute only a small fraction of block value. Their analysis further demonstrates that exclusive transactions are the primary determinant of which builder wins a MEV-Boost auction, confirming that the decisive component of rewards is not protocol-level compensation but exclusive, privately routed orderflow. In related work, Canidio and Pahari (2025) use the contents of proposed blocks to document persistent builder-level differences in transaction inclusion and ordering. Their

results show that builders execute systematically different MEV strategies, include different sets of transactions even when facing the same public mempool, and extract different amounts of value. These builder ?fingerprints?? indicate that the heterogeneous effective rewards underlying our contest model are not only present but constitute the dominant economic force in contemporary block production. Together, these studies show that exclusive orderflow is both pervasive and predictive, and they provide direct empirical support for the mechanisms in our model: transparency induces heterogeneous rewards; these heterogeneities persist across blocks; and in a latency-bounded contest environment they naturally lead to increasing concentration.

Empirical evidence from block-building markets reinforces the conclusion that PBS has not produced a competitive or decentralized execution layer. Thiery (2023) and Wu et al. (2024) show that blocks built via MEV-Boost are heavily concentrated among a handful of builders such as beaverbuild, Titan, and rsyncbuilder. These builders not only construct the majority of high-value blocks but also capture a disproportionate share of all MEV-related transactions. The concentration intensifies during periods of high price volatility. In the Wu et al. (2025) dataset, blocks in the ninety-ninth percentile of intrablock volatility routinely derive more than eighty percent of their value from exclusive non-atomic arbitrage transactions. During such periods, integrated builders nearly always win the MEV-Boost auction, suggesting that volatility magnifies the advantage conferred by privileged access to private orderflow. Similar patterns appear in analyses of proposer–builder separation by Heimbach et al. (2023), who observe that builder concentration rises as the value of MEV spikes, and that exclusive orderflow and private bundles are now the dominant drivers of builder profitability.

Finally, although our theoretical focus is on block-production contests rather than market microstructure, related empirical work on automated market makers reinforces the informational logic embedded in our model. Milionis, Moallemi, Roughgarden, and Zhang (2024) quantify the losses incurred by liquidity providers due to stale quotes and show that arbitrageurs consistently capture predictable profits arising from delays in price adjustment. Their estimates of "loss-versus-rebalancing" match closely the observed arbitrage revenue in major AMM pools, further illustrating that predictable, transparency-driven value extraction is robust across DeFi trading venues. These findings are consistent with our mechanism in which transparency and latency jointly create systematic rents for informed or fast actors.

Taken together, the empirical evidence demonstrates that transparent trading environments generate substantial and predictable extractable value; that this value is captured almost entirely by a small set of highly specialized intermediaries; that private and exclusive routing arrangements now dominate the realization of high-value opportunities; and that

block-building under PBS is increasingly concentrated in the hands of builders who secure privileged access to this orderflow. The resulting equilibrium patterns—high concentration, vertically integrated searcher–builder pairs, and the dominance of private routing—parallel precisely the mechanisms identified by our theoretical analysis. In this sense, the empirical facts not only motivate the model but also illustrate the broader claim of this paper: under transparency and latency-bounded settlement, decentralized and competitive trading equilibria are economically unstable, and centralization emerges as the natural and robust outcome.

Having shown that the mechanisms predicted by the model are present and strong in practice, we now combine them into a formal impossibility result.

10 An Impossibility Result

The preceding sections establish two complementary forces. First, transparency and latency-bounded settlement generate heterogeneous rewards R_i through MEV. Even small heterogeneities trigger strict sorting (Theorems 1 and 2) and increasing concentration as contests become more decisive (Proposition 4) or rewards become more dispersed (Proposition 5). Second, once traders can route privately, their best responses feature strong coordination complementarities: larger private pools raise inclusion probabilities, reinforcing the leader's advantage (Lemma 3). These dynamics induce a tipping logic in which market-sharing allocations are either unstable (Theorem 3 and proposition 8) or remain only on a knife-edge relying on exact indifference conditions, which may be satisfied for all non-atomic traders or free multihoming (Theorem 4 and corollary 2). Any perturbation in fees or routing probabilities breaks the required equalities in surplus shares and tips best responses toward a single EA or the public mempool. Monopoly, however, is always stable.

Taken together, these mechanisms imply a formal impossibility: under transparency, any equilibrium that is stable with respect to small perturbations must be centralized. Attempts to sustain decentralized private routing either collapse to monopoly or eliminate equilibrium existence altogether. We now formally define what we mean by decentralized trading and state the impossibility theorem.

Definition 1 (Transparent and decentralized trading equilibrium)**.** *An equilibrium of the fee–routing environment is:*

- transparent *if transactions submitted to the public mempool are observable before settlement, generating strictly positive MEV (Section 3);*

- decentralized *if at least two execution agents (EAs) receive strictly positive expected private order flow and strictly positive expected payoffs, and if the allocation is bounded away from monopoly: there exists* $\varepsilon > 0$ *such that the Herfindahl index* $H \leq 1 - \varepsilon$ *for either block-production shares or POF shares.*

Theorem 5 (Impossibility of transparent and decentralized DeFi trading)**.** *Consider the fee–routing environment with transparent public mempool, latency-bounded settlement, and contest environment* Γ *satisfying Assumptions 1–4. Suppose traders may route privately (single- or multi-homing) with duplication costs* $\Delta\ell(m) \geq 0$*, and may always fall back to the public mempool.*

Then transparency generates heterogeneous effective rewards R_i *via MEV extraction (Section 7). In any equilibrium:*

(a) ***The contest layer becomes concentrated.*** *Heterogeneous rewards imply strict sorting (Theorems 1 and 2) and increasing concentration as decisiveness or reward dispersion rise (Propositions 4 and 5).*

(b) ***Single-homing POF admits only monopoly.*** *Positive feedback in surplus shares makes all market-sharing allocations unstable, and the only stable POF allocation is monopoly (Proposition 6).*

(c) ***Costly multi-homing with some atomic traders yields monopoly or nonexistence.*** *If duplication costs satisfy* $\Delta\ell(2) > 0$*, no market-sharing SPNE exists (Proposition 7). Under the no–solo–entry condition, the unique SPNE is monopoly POF (Theorem 3); when some traders are large enough that this condition fails, no SPNE with positive POF exists (Proposition 8).*

(d) ***Multi-homing SPNE exists with non-atomic traders or costless multi-homing, but is knife-edge and unstable.*** *Market-sharing SPNE exist only when routing profiles satisfy the knife-edge conditions (F1)–(F2) in Theorem 4 or multi-homing is free (Corollary 2). These equilibria are non-strict and not trembling-hand perfect: any small aggregate shock collapses them to monopoly.*

Therefore, for an open set of primitives — including any environment with strictly positive MEV and strictly positive duplication costs — there exists no *equilibrium that is simultaneously transparent and decentralized. All structurally stable equilibria are centralized (monopoly POF and concentrated block production), while attempts to sustain decentralization either collapse to monopoly or eliminate equilibrium existence. Transparent and latency-bounded DeFi trading cannot remain decentralized in equilibrium.*

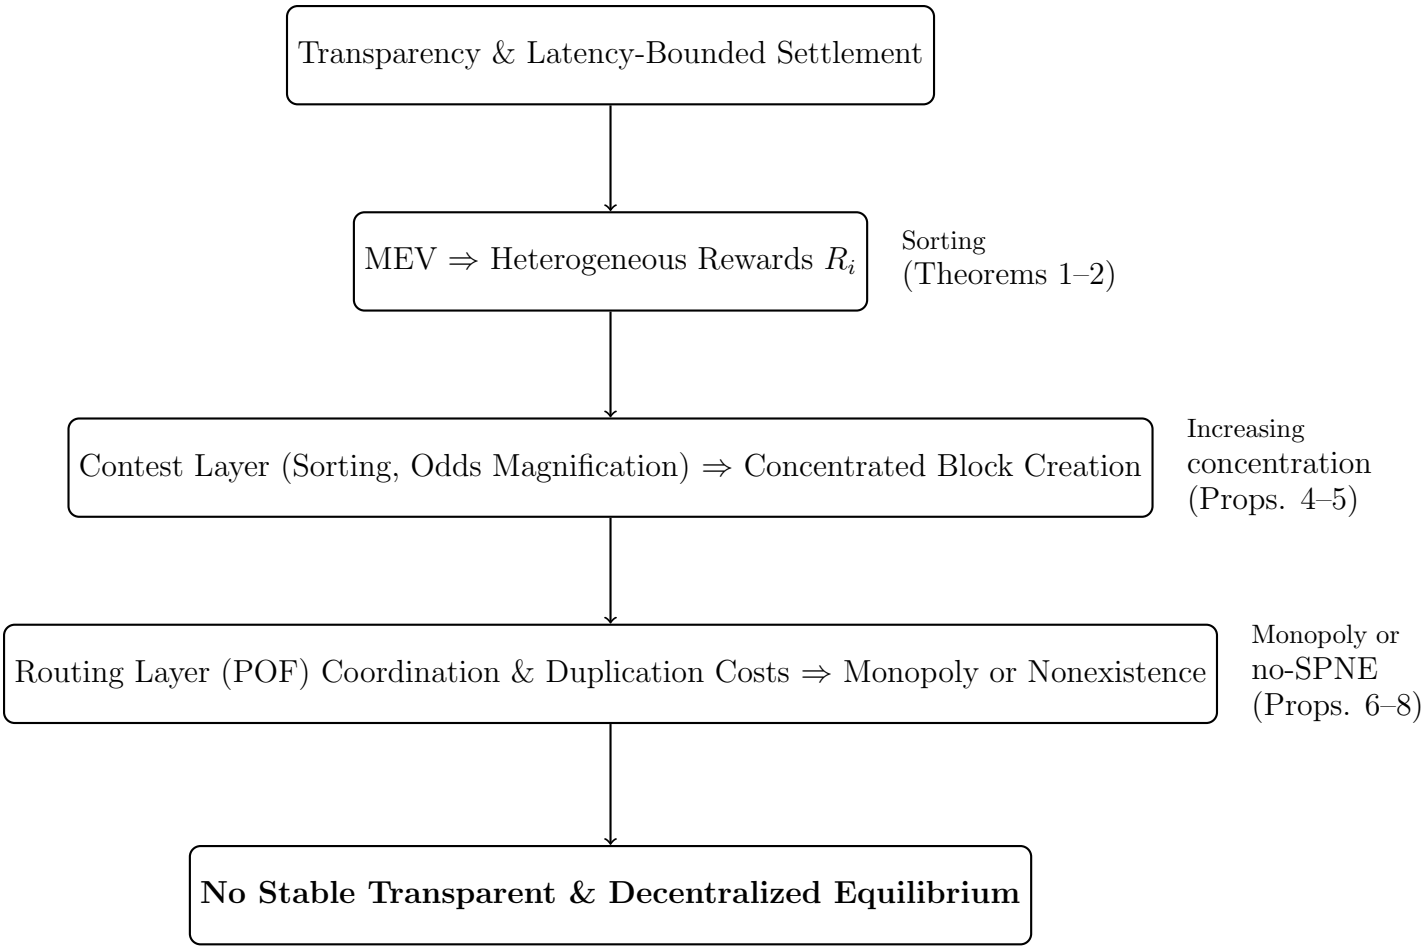


Figure 1: Logical structure behind the impossibility result. Transparency induces MEV-driven heterogeneity, which is amplified by the contest layer and by network complementarities in private routing. All structurally stable equilibria are centralized; market-sharing requires knife-edge conditions or fails to exist.

The theorem highlights a structural tension between the architectural ideals of blockchain platforms and the equilibrium consequences of transparency. Transparency ensures openness and verifiability but also exposes traders' intentions and MEV opportunities. This induces heterogeneous rewards that are amplified by the contest structure of block production. Traders' attempts to escape extraction through POF only reinforce these asymmetries through network complementarities in routing. The forces of sorting, odds magnification, and coordination all point in the same direction: transparent markets generate the very concentrated intermediaries they were designed to eliminate.

11 Conclusion

Blockchain-based finance rests on a foundational idea: that transparency, open participation, and permissionless entry can jointly sustain a decentralized market structure. Yet despite an architecture that decentralizes access, verification, and participation, equilibrium outcomes have centralized.

This paper explains why transparency and decentralization are structurally in tension. Our central insight is that transparency—widely viewed as a fairness guarantee—creates predictable, actionable, and extractable opportunities in the settlement layer. In any latency-bounded blockchain, transactions are visible before settlement, enabling strategic reordering and Maximal Extractable Value (MEV). This induces heterogeneity in the effective rewards R_i of the block creators. Even small differences in MEV-extraction skill become persistent asymmetries in expected rewards.

We develop a unified and protocol-agnostic model capturing competition among payoff-maximizing execution agents (miners, validators, builders). The model nests Proof-of-Work, Proof-of-Stake, and Proposer-Builder Separation (PBS) within a single generalized contest environment Γ. We show that transparent block production systematically amplifies reward heterogeneity, pushing the system toward concentration even when agents are ex ante identical and entry is fully free. Three mechanisms underpin this result.

1. Transparency $\rightarrow$ MEV $\rightarrow$ Heterogeneous Rewards R_i from Block Creation. Transparency and latency-bounded settlement convert the mempool into a predictable source of extractable value. Any agent that wins the block can reorder transactions to capture MEV, producing heterogeneity in effective rewards R_i based on differences in technical skill or infrastructure. As Section 3 shows, this heterogeneity is endogenous to transparency itself.

2. The Contest Layer Strictly Sorts and Amplifies Small Advantages. In environment Γ, contest success is determined by a generalized Tullock CSF with decisiveness parameter a. We show:

Strict Sorting: any heterogeneity in R_i leads to strictly ordered equilibrium efforts and block shares (Theorem 1 and Theorem 2).

Odds Magnification: more decisive contests (higher a, as under PBS) disproportionately increase the dominance of high-R_i agents (Proposition 4).

Reward Dispersion $\rightarrow$ Concentration: widening the leading agent's advantage strictly increases concentration (Proposition 5).

Thus the contest layer creates a structural tendency toward concentration, independent of protocol-level intentions.

3. Trader Response → POF → Strong Network Effects → Monopoly. Aware of extraction in the public mempool, traders try to protect themselves by routing privately to a block creator. This generates:

network complementarities: more traders joining an EA's private pool raise its inclusion probability and thus increase the value of routing there (Lemma 3);

vertical differentiation: EAs differ in reliability p_i, so low-fee entrants cannot discipline the dominant agent;

tipping: only monopoly POF is stable (Proposition 6).

Under costly multihoming—or even the slightest friction—POF equilibria collapse to monopoly (Theorem 3). Under conditions permitting "profitable seeding," no equilibrium with positive POF exists at all (Proposition 8). When multihoming is costless or traders are nonatomic, market-sharing equilibria exist but lie on a knife-edge (Theorem 4 and Corollary 2): any perturbation collapses them back to monopoly.

Altogether, the model shows that transparent and decentralized trading are jointly untenable. Transparency produces MEV, which induces reward heterogeneity. Contest dynamics amplify that heterogeneity into concentrated block production. Trader behavior, intended to escape MEV, produces network complementarities that tip private routing into monopoly. Any equilibrium robust to small perturbations must therefore be centralized.

The main result (Theorem 5) is an impossibility theorem:

> In any latency-bounded and transparent blockchain environment with MEV and free entry, no equilibrium that is stable in the sense of tremble-free best responses can remain decentralized. Stable equilibria are inevitably centralized, while decentralized ones are knife-edge or fail to exist.

This paper unifies empirical facts about MEV, POF consolidation, and builder concentration into a single theoretical logic. It shows that DeFi's centralization is not the outcome of mismanagement or imperfect designs—it is a structural consequence of transparency in a latency-bounded environment. Radical transparency, intended to eliminate intermediation, re-creates it.

Appendix

A Limit of the Tullock CSF as $a \to \infty$

Lemma 5 (Limit of the Tullock CSF as $a \to \infty$)**.** *For any effort profile* $\mathbf{e} = (e_1, \ldots, e_N)$ *with* $\sum_i e_i > 0$*, define the generalized Tullock contest success function*

$$p_i^{(a)}(\mathbf{e}) \;=\; \frac{e_i^a}{\sum_{k=1}^N e_k^a}.$$

Then

$$\lim_{a\to\infty} p_i^{(a)}(\mathbf{e}) \;=\; \begin{cases} \dfrac{1}{|\arg\max_k e_k|}, & \text{if } i \in \arg\max_k e_k \text{ and } \max_k e_k > 0, \\ 0, & \text{otherwise.} \end{cases}$$

Proof. Let $e_{\max} = \max_k e_k$ and $M = \{\, k : e_k = e_{\max} \,\}$ denote the set of maximizers, with cardinality $|M| = m \geq 1$.

For each i, define the normalized effort $x_i = e_i/e_{\max} \in [0, 1]$. Then

$$p_i^{(a)}(\mathbf{e}) = \frac{x_i^a}{\sum_{k=1}^N x_k^a} = \frac{x_i^a}{m + \sum_{k\notin I} x_k^a}.$$

If $i \notin M$, then $x_i < 1$ and hence $x_i^a \to 0$ as $a \to \infty$, so $p_i^{(a)} \to 0$.

If $i \in M$, then $x_i = 1$, and for all $k \notin M$, $x_k^a \to 0$. Therefore

$$\lim_{a\to\infty} p_i^{(a)}(\mathbf{e}) = \frac{1}{m + \sum_{k\notin I} 0} = \frac{1}{|M|}.$$

When all efforts are zero ($e_{\max} = 0$), we define $p_i^{(\infty)}(\mathbf{e}) = 0$ by convention. □

B First-order structure and interiority

Write, statewise, $A_{-i} := \sum_{k\neq i} e_k^a$, so $A = e^a + A_{-i}$ and

$$p_i((e, e_{-i})) = \frac{e^a}{e^a + A_{-i}}, \qquad \frac{\partial p_i}{\partial e} = \frac{a e^{a-1} A_{-i}}{(e^a + A_{-i})^2}.$$

Lemma 6 (Statewise derivative and common operator)**.** *Under Assumptions 1–3, for $e > 0$,*

$$\frac{\partial}{\partial e}\Phi_i(e; R_i \,|\, \sigma_{-i}) = \mathbb{E}\left[\frac{a e^{a-1} A_{-i}}{(e^a + A_{-i})^2}\big(R_i - \beta(e) - \tilde{\gamma}(S)\big) - \frac{e^a}{e^a + A_{-i}}\beta'(e) - \frac{e^a}{e^a + A_{-i}}\tilde{\gamma}'(S)\,\frac{\partial S}{\partial e_i} - \alpha'(e)\right]$$
$$=: \ \mathcal{H}(e; R_i \,|\, \sigma_{-i}).$$

Proof. Differentiate inside the expectation by Assumption 3. Product rule yields the three bracketed terms from $p_i(\cdot)\big(R_i - \beta(e) - \tilde{\gamma}(S)\big)$ plus $-\alpha'(e)$. □

Lemma 7 (Strictly positive profit ⇒ interior and degenerate)**.** *If $\mathbb{E}[u_i(\sigma)] > 0$ in equilibrium, then $e_i > 0$ almost surely and (by Assumption 4) player i plays a unique interior effort $e_i^* > 0$ (i.e., no mixing).*

Proof. At $e = 0$, $p_i((0, e_{-i})) = 0$ so $\Phi_i(0; R_i \,|\, \sigma_{-i}) = -\alpha(0) \le 0$. Strict profitability implies the support cannot put mass on 0. Uniqueness then rules out mixing. □

C Proof of Proposition 3

Proof. Let $\mathcal{H}$ be as in Lemma 6. By Assumption 4, for each R the interior best reply is the unique $e > 0$ solving $\mathcal{H}(e; R \,|\, \sigma_{-i}) = 0$, and $e \mapsto \mathcal{H}(e; R \,|\, \sigma_{-i})$ is strictly decreasing on $(0, \infty)$. For $R' > R$ and $e^\star := e_i^*(R' \,|\, \sigma_{-i})$,

$$\mathcal{H}(e^\star; R \,|\, \sigma_{-i}) = \mathcal{H}(e^\star; R' \,|\, \sigma_{-i}) - \int_R^{R'} \frac{\partial}{\partial t}\mathcal{H}(e^\star; t \,|\, \sigma_{-i})\, dt = -(R' - R)\cdot\mathbb{E}\left[\frac{a\, e^{\star a-1} A_{-i}}{((e^\star)^a + A_{-i})^2}\right] \le 0,$$

with strict inequality if $\Pr[A_{-i} > 0] > 0$. Since $e \mapsto \mathcal{H}(e; R \,|\, \sigma_{-i})$ is strictly decreasing, its unique zero lies (strictly) to the left of $e^\star$. □

D Proof of Theorem 1

Proof. By Lemma 7, strictly profitable players choose unique interior (nonrandom) efforts $e_i^*, e_j^* > 0$. Because $\mathcal{H}$ has the same functional form for all players under Assumption 2, their FOCs satisfy

$$\mathcal{H}(e_i^*; R_i \,|\, \sigma) = 0 = \mathcal{H}(e_j^*; R_j \,|\, \sigma).$$

By Proposition 3, the unique interior best reply is (strictly) increasing in R; hence $R_i > R_j$ implies $e_i^* > e_j^*$. Since $p_i(\omega) - p_j(\omega) = \frac{e_i^{*\,a} - e_j^{*\,a}}{A(\omega)}$ for each state with $A(\omega) > 0$, taking expectations gives $\mathbb{E}[p_i] > \mathbb{E}[p_j]$. □

E Proof of Theorem 2

Proof. **(Monotone best reply).** Fix σ_{-i} and write agent i's expected objective as

$$\Phi(e;R) \;=\; \mathbb{E}_{\sigma_{-i}}\big[\,p_i^{\infty}(e,e_{-i})\,\big]\cdot R \;-\; \mathbb{E}_{\sigma_{-i}}\big[\,p_i^{\infty}(e,e_{-i})\big(\beta(e)+\gamma(e,e_{-i})\big)\big] \;-\; \alpha(e).$$

For $R' > R$, their difference is

$$\Phi(e;R') - \Phi(e;R) \;=\; (R'-R)\,\mathbb{E}_{\sigma_{-i}}\big[p_i^{\infty}(e,e_{-i})\big].$$

The map $e \mapsto \mathbb{E}_{\sigma_{-i}}[p_i^{\infty}(e,e_{-i})]$ is nondecreasing (raising e weakly raises the event that i is among the maximizers), hence the difference $\Phi(\cdot;R') - \Phi(\cdot;R)$ is nondecreasing in e. This is exactly the *single–crossing* (increasing differences) property in (e,R).

By the Milgrom–Shannon monotone selection theorem, any selection from the argmax correspondence of a single–crossing objective is nondecreasing in the parameter. Since the best reply is unique on the profitable region by assumption, it is a (single–valued) nondecreasing function: $R \uparrow \Rightarrow e_i^*(R) \uparrow$. This proves the Monotone Best Reply property.

(Strictness). Suppose $\Pr[\max_{k\neq i} e_k > 0] > 0$. Then there exists e such that $\mathbb{E}[p_i^{\infty}(e,e_{-i})]$ is *strictly* increasing at some points: intuitively, near the current optimizer $e_i^*(R)$, a small upward move in e strictly raises the chance of being the unique maximizer on the events where some rival exerts positive effort. Formally, there is $\bar{e}$ and $\delta > 0$ with

$$\mathbb{E}[p_i^{\infty}(\bar{e}+\delta,e_{-i})] \;>\; \mathbb{E}[p_i^{\infty}(\bar{e},e_{-i})].$$

Hence $\Phi(\bar{e}+\delta;R') - \Phi(\bar{e}+\delta;R) \;>\; \Phi(\bar{e};R') - \Phi(\bar{e};R)$ for all $R' > R$. This is the *strict* single–crossing property on a neighborhood of the optimizer. By the strict version of the monotone selection argument (or by contradiction using the uniqueness of the maximizer and upper hemicontinuity of the argmax), the unique optimizer must move strictly upward when R increases. This proves the Strictness property.

(Strict sorting in equilibrium). Let σ be a Nash equilibrium and take i,j with $\mathbb{E}[u_i] > 0$ and $\mathbb{E}[u_j] > 0$. By (i), best replies are weakly increasing in R, so $R_i > R_j$ implies $e_i^* \geq e_j^*$. If $e_i^* = e_j^*$, Lemma 2 rules out a profitable tie with positive probability; but in the hard–max environment equal efforts among two profitable players would induce a positive-probability tie at the top on the set where both are maximal, yielding a profitable ε-deviation. Hence $e_i^* > e_j^*$.

Finally, since $p_i^{\infty}(\mathbf{e})$ is (weakly) increasing in e_i and strictly increasing at the top, $e_i^* > e_j^*$ implies $\mathbb{E}_{\sigma}[p_i^{\infty}] > \mathbb{E}_{\sigma}[p_j^{\infty}]$. This proves Strict Sorting in Equilibrium. □

F Proof of Proposition 4

Let's start by formally restating Proposition 4.

Proposition. Consider the environment Γ described above. For $a \geq 1$, let $\sigma^*(a)$ be a Nash equilibrium (possibly mixed), and let

$$p_i^*(a) \;:=\; \mathbb{E}_{\sigma^*(a)}[\, p_i(\mathbf{e})\,], \qquad \Pi_i^*(a) \;:=\; \mathbb{E}_{\sigma^*(a)}[\, u_i(\mathbf{e}; R_i)\,].$$

Assume there exists a nondegenerate interval $A \subseteq [1, \infty)$ such that:

(B1) (*Locally fixed profitable set*) The set of profitable players

$$P(a) \;:=\; \{\, i : \; \Pi_i^*(a) > 0 \,\}$$

is constant on A; denote it by P. Moreover, each $i \in P$ is interior in equilibrium: $\Pr_{\sigma^*(a)}(e_i > 0) = 1$ for all $a \in A$.

(B2) (*Single-valued best responses for profitable players*) For every $i \in P$, the equilibrium induces a unique action almost surely: $\Pr_{\sigma^*(a)}(e_i = e_i^*(a)) = 1$ with a single effort level $e_i^*(a) > 0$.

(B3) (*Label-neutral externality*) The map $\gamma(\cdot)$ is label-neutral and independent of e_i.

Then:

(i) (*Odds magnification*) For any $i, j \in P$,

$$\frac{p_i^*(a)}{p_j^*(a)} \;=\; \Big(\frac{e_i^*(a)}{e_j^*(a)}\Big)^a.$$

Moreover, if $e_i^*(a) > e_j^*(a)$ at some $a \in A$, then $\frac{p_i^*(a)}{p_j^*(a)}$ is strictly increasing in a on A.

(ii) (*Concentration increases in a*) Let $p^*(a)$ denote the vector $(p_1^*(a), \ldots, p_N^*(a))$ sorted in nonincreasing order. Then $p^*(a)$ becomes more concentrated as a increases on A:

- the top share $p_1^*(a)$ is (weakly) increasing in a;
- every Schur-convex concentration index (e.g., the Herfindahl index $H(a) = \sum_i (p_i^*(a))^2$) is (weakly) increasing in a.

Furthermore, this increase is strict for any strictly Schur-convex index (like $H(a)$) if and only if the equilibrium efforts of the profitable players are not all identical (i.e., if $e_i^*(a) \neq e_j^*(a)$ for at least one pair $i, j \in P$).

Proof. Let $\mathbf{e}^*(a) = (e_i^*(a))_{i=1}^N$ be the vector of equilibrium efforts. By (B1) and (B2), players $i \in P$ choose deterministic efforts $e_i^*(a) > 0$. Players $i \notin P$ must have $e_i^*(a) = 0$, otherwise $\Pi_i^*(a) \leq 0$ would be violated by positive costs.

The Tullock Contest Success Function (CSF) is $p_i(\mathbf{e}) = e_i^a / \sum_{k=1}^N e_k^a$. For any two profitable players $i, j \in P$, their probability ratio is:

$$\frac{p_i^*(a)}{p_j^*(a)} = \frac{p_i(\mathbf{e}^*(a))}{p_j(\mathbf{e}^*(a))} = \frac{e_i^*(a)^a / \sum_{k\in P} e_k^*(a)^a}{e_j^*(a)^a / \sum_{k\in P} e_k^*(a)^a} = \frac{e_i^*(a)^a}{e_j^*(a)^a} = \left(\frac{e_i^*(a)}{e_j^*(a)}\right)^a.$$

This proves the first part of (i).

To prove the second part of (i), we analyze how this ratio changes with a. Let $r_{ij}(a) := e_i^*(a)/e_j^*(a)$ be the effort ratio. We compute the logarithmic derivative of the odds ratio with respect to a:

$$\begin{aligned}\frac{d}{da}\log\left(\frac{p_i^*(a)}{p_j^*(a)}\right) &= \frac{d}{da}\log\left([r_{ij}(a)]^a\right) \\ &= \frac{d}{da}\left[a \cdot \log r_{ij}(a)\right] \\ &= \log r_{ij}(a) + a \cdot \frac{r'_{ij}(a)}{r_{ij}(a)}.\end{aligned}$$

Assume $e_i^*(a) > e_j^*(a)$. This implies $r_{ij}(a) > 1$, so the first term $\log r_{ij}(a)$ is **strictly positive**.

Now, we analyze the second term. We appeal to Lemma 8, which establishes (via First-Order Conditions) that if player i has an advantage over player j, their relative effort $r_{ij}(a)$ is weakly increasing in a. That is, $\mathbf{r'_{ij}(a) \geq 0}$.

Since $r_{ij}(a) > 0$, $a \geq 1$, and $r'_{ij}(a) \geq 0$, the second term is non-negative:

$$a \cdot \frac{r'_{ij}(a)}{r_{ij}(a)} \geq 0.$$

Combining these results, the logarithmic derivative is strictly positive:

$$\frac{d}{da}\log\left(\frac{p_i^*(a)}{p_j^*(a)}\right) = \underbrace{\log r_{ij}(a)}_{>0} + \underbrace{a \cdot \frac{r'_{ij}(a)}{r_{ij}(a)}}_{\geq 0} > 0.$$

Since the logarithmic derivative is positive, the function $\frac{p_i^*(a)}{p_j^*(a)}$ is strictly increasing in a. This completes the proof of (i).

Proof of (ii): Part (i) shows that for any pair $i, j \in P$, if $p_i^*(a) > p_j^*(a)$ (which is equivalent to $e_i^*(a) > e_j^*(a)$), then their probability ratio p_i^*/p_j^* is **strictly increasing** with a. This means the distribution of probabilities $p(a) = (p_1^*(a), \ldots, p_N^*(a))$ becomes more "spread out" or "unequal" as a increases.

Formally, this implies that if $a_2 > a_1$ (for $a_1, a_2 \in A$), the vector $p^*(a_2)$ (the sorted vector of probabilities) **majorizes** the vector $p^*(a_1)$.

A concentration index $C(\cdot)$ is **Schur-convex** if $C(\mathbf{x}) \geq C(\mathbf{y})$ whenever $\mathbf{x}$ majorizes $\mathbf{y}$. Since $p^*(a)$ increases in the majorization order as a increases, it follows directly that $p_1^*(a)$, $H(a)$, and any other Schur-convex index are (weakly) increasing in a.

For the strict case: If the equilibrium is not perfectly symmetric, there exists at least one pair $i, j \in P$ with $e_i^*(a) \neq e_j^*(a)$. By part (i), their probability ratio p_i^*/p_j^* is strictly increasing. This is sufficient to imply that $p^*(a_2)$ **strictly majorizes** $p^*(a_1)$. An index $C(\cdot)$ is **strictly Schur-convex** (which the Herfindahl index $H(a) = \sum_i (p_i^*(a))^2$ is) if $C(\mathbf{x}) > C(\mathbf{y})$ whenever $\mathbf{x}$ strictly majorizes $\mathbf{y}$ (and $\mathbf{x}$ is not just a permutation of $\mathbf{y}$). Therefore, if the equilibrium efforts are not all identical, any strictly Schur-convex index $H(a)$ is **strictly increasing** in a. This proves (ii). □

Lemma 8 (FOC and Non-decreasing Effort Ratios)**.** *Under the assumptions of the proposition, the First-Order Condition (FOC) for any $i \in P$ at an interior solution $e_i^* > 0$ is:*

$$\partial_{e_i} u_i(\mathbf{e}^*) = \frac{a}{e_i^*} p_i(1-p_i)\left(R_i - \beta(e_i^*) - \gamma(e_{-i}^*)\right) \; - \; p_i\,\beta'(e_i^*) \; - \; \alpha'(e_i^*) \; = \; 0.$$

Furthermore, under standard regularity conditions (e.g., $R_i > R_j$ with identical cost functions α, β), a comparative statics analysis of the system of FOCs shows that the effort ratio $r_{ij}(a) = e_i^*(a)/e_j^*(a)$ is weakly increasing in a (i.e., $\mathbf{r'_{ij}}(\mathbf{a}) \geq \mathbf{0}$).

This result is a standard finding in the contest theory literature.

Proof sketch of Lemma 8. The utility is $u_i = p_i(\mathbf{e}) \cdot \left(R_i - \beta(e_i) - \gamma(e_{-i})\right) - \alpha(e_i)$. The partial derivative of the Tullock CSF $p_i(\mathbf{e}) = e_i^a / \sum_k e_k^a$ is:

$$\begin{aligned}\frac{\partial p_i}{\partial e_i} &= \frac{(a e_i^{a-1})(\sum_k e_k^a) - (e_i^a)(a e_i^{a-1})}{(\sum_k e_k^a)^2} \\ &= \frac{a e_i^{a-1}(\sum_{k \neq i} e_k^a)}{(\sum_k e_k^a)^2}\end{aligned}$$

We factor this into the $p_i(1-p_i)$ form:

$$\frac{\partial p_i}{\partial e_i} = \frac{a}{e_i} \cdot \frac{e_i^a}{\sum_k e_k^a} \cdot \frac{\sum_{k \neq i} e_k^a}{\sum_k e_k^a} = \frac{a}{e_i} p_i(1-p_i).$$

Applying the product rule to $\partial_{e_i} u_i = (\partial_{e_i} p_i) \cdot (\text{Net Prize}) + p_i \cdot (\partial_{e_i} \text{Net Prize}) - \alpha'(e_i) = 0$ yields the FOC stated in the lemma. □

G Proof of Proposition 5

Let's start by formally restating Proposition 5.

Proposition. Fix $a > 0$. For a valuation vector $\mathbf{R}$, let $\sigma(\mathbf{R})$ be a Nash equilibrium (pure or mixed) of game Γ. Define the set of *profitable players*

$$P(\mathbf{R}) \;:=\; \{\, i : \; \mathbb{E}_{\sigma(\mathbf{R})}[u_i] > 0 \,\}.$$

Assume that in a neighborhood of $\mathbf{R}$:

(B1) (*Locally fixed profitable set*) $P(\cdot)$ is constant and, for every $i \in P(\mathbf{R})$, $p_i(\sigma(\mathbf{R})) > 0$.

(B2) (*Single-valued efforts for profitable players*) For each $i \in P(\mathbf{R})$, the equilibrium effort $e_i^*(\cdot)$ induced by $\sigma(\cdot)$ is uniquely determined and differentiable on that neighborhood.

(B3) (*Positive own effect within P*) $\frac{\partial e_i^*}{\partial R_i} > 0$ for all $i \in P(\mathbf{R})$.

(B4) (*Strategic substitutes within P*) $\frac{\partial e_j^*}{\partial R_i} \leq 0$ for all distinct $i, j \in P(\mathbf{R})$.

Let player $1 \in P(\mathbf{R})$ satisfy $R_1 > R_j$ for all $j \in P(\mathbf{R}) \setminus \{1\}$. Then, for every $j \in P(\mathbf{R}) \setminus \{1\}$,

$$\frac{d}{dR_1}\left(\frac{p_1}{p_j}\right) \;>\; 0.$$

Consequently, the *within-P* Herfindahl index

$$\widetilde{H}_P(\mathbf{p}) \;:=\; \sum_{k \in P(\mathbf{R})} \left(\frac{p_k}{\sum_{\ell \in P(\mathbf{R})} p_\ell}\right)^2$$

strictly increases with R_1 locally.

Moreover, if in addition

(B5) $S_P(\mathbf{R}) := \sum_{k \in P(\mathbf{R})} p_k(\sigma(\mathbf{R}))$ is (weakly) nondecreasing in R_1,

then the (unnormalized) index $H_P(\mathbf{p}) := \sum_{k\in P(\mathbf{R})} p_k^2$ strictly increases with R_1 locally.

Proof. For $i \in P$, the contest success function gives $p_i = e_i^a / \sum_m e_m^a$ with $e_i > 0$. Hence for $j \in P \setminus \{1\}$,

$$\frac{p_1}{p_j} = \Big(\frac{e_1}{e_j}\Big)^a, \quad \frac{d}{dR_1}\log\Big(\frac{p_1}{p_j}\Big) = a\Big(\frac{1}{e_1}\frac{de_1^*}{dR_1} - \frac{1}{e_j}\frac{de_j^*}{dR_1}\Big).$$

By (B3) and (B4), the right-hand side is strictly positive, establishing $\frac{d}{dR_1}(p_1/p_j) > 0$.

Write $S_P := \sum_{k\in P} p_k$ and the normalized within-P shares $\tilde{p}_k := p_k/S_P$. The strict increase of all pairwise ratios p_1/p_j for $j \neq 1$ implies a mean-preserving reallocation within P toward player 1, i.e., the vector $(\tilde{p}_k)_{k\in P}$ becomes strictly more concentrated in the sense of the Lorenz order. Therefore

$$\frac{d}{dR_1}\widetilde{H}_P \;=\; \frac{d}{dR_1}\sum_{k\in P}\tilde{p}_k^2 \;>\; 0.$$

Finally, $H_P = S_P^2\,\widetilde{H}_P$. If (B5) holds, then $dH_P/dR_1 = 2S_P\widetilde{H}_P\,\frac{dS_P}{dR_1} + S_P^2\,\frac{d\widetilde{H}_P}{dR_1} > 0$. □

H Proof of Lemma 3

Proof. Consider a marginal reallocation in which one trader moves from some $j \neq i$ to i, keeping the total informed mass α fixed. Write $T = (T_1, \ldots, T_N)$ for routing masses and let the effective reward map satisfy $R_i(T) = R + \rho_i(T_i)$ with $\rho_i'(\cdot) > 0$.[4] Let $p^*(T)$ denote the inclusion probabilities induced in Γ at the corresponding equilibrium.

Case A (no block creator drop-outs; locally fixed profitable set). Assume Assumptions 1–4 and, locally, that the profitable set is fixed and profitable block creators choose unique interior efforts. When one trader moves from j to i, we have $R_i \uparrow$ and $R_j \downarrow$. By Proposition 3 (monotone best reply in R), e_i^* strictly increases in R_i. Since the Tullock CSF $p_i(e) = e_i^a / \sum_k e_k^a$ is strictly increasing in e_i and weakly decreasing in rivals' efforts,

$$p_i^*(T \text{ with the extra trader at } i) \;>\; p_i^*(T).$$

Therefore $\Pi_{POF}(i) = p_i^*\, v(1 - r_i)$ strictly increases.

Case B (a block creator exits). Suppose the outflow causes some block crator ℓ to drop from the profitable set (its equilibrium effort becomes 0). Then, relative to the pre-move profile, (i) R_i increases, (ii) a competitor disappears from the denominator $\sum_k e_k^a$, and (iii) e_i^* increases by the same arguments as above. Thus p_i^* increases strictly (the denominator

[4] E.g., $R_i(T) = R + r_i\big((T_i + T_m)v\big)$.

shrinks while the numerator rises). Thus $\Pi_{POF}(i)$ is increasing in T_i when an exit occurs.

Combining the two cases yields the claim. □

I Proof of Proposition 6

Proof. (a) **Instability of Sharing:**
Consider a symmetric state where a set of $k \geq 2$ block creators offer the same maximal surplus share, S_{max}. The total mass of informed traders, α, is distributed equally among them. Each j of these block creators serves $n = \alpha/k$ traders.

The payoff for a trader at any of these block creators is $\Pi_j = p(n)v(1 - r_j)$, where $p(n)$ is the inclusion probability resulting from having n traders. Since all k block creators are symmetric, they offer the same r_i and have the same $p(n)$.

Consider a single trader who unilaterally deviates from a block creator i to another top-tier block crator, j. The number of traders at these block creators changes to $N_i' = n - 1$ and $N_j' = n + 1$.

The deviating trader's new payoff is $\Pi_{dev}' = p(n + 1)v(1 - r_j)$.

For the symmetric equilibrium to be stable, no trader must have an incentive to deviate. This requires $\Pi_i \geq \Pi_{dev}'$.

$$p(N)v(1 - r_i) \geq p(N + 1)v(1 - r_j)$$

However, since $r_i = r_j$ and the block creator's inclusion probability p is an increasing function of the number of traders N_i, $p(n + 1) > p(n)$, the stability condition is violated. The deviation is strictly profitable for the trader.

Because a single trader has an incentive to deviate, the symmetric sharing state is not a stable Nash Equilibrium.

(b) **Stability of Monopoly:**
Consider a monopolistic state where all α traders have coordinated on a single block creator, j^*, which offers the best surplus share S_{j^*}.

The payoff for every trader is $\Pi_{j^*} = p(\alpha)v(1 - r_{j^*})$. Any other competing block creator i has zero traders, $N_i = 0$, and thus a lower inclusion probability, $p(0) < p(\alpha)$.

Consider a single trader who unilaterally deviates from the monopolist j^* to an empty competitor, i. The number of traders becomes $N_{j^*}' = \alpha - 1$ and $N_i' = 1$.

The deviating trader's new payoff is $\Pi_{dev}' = p(1)v(1 - r_i)$.

The trader will only deviate if $\Pi'_{dev} > \Pi_{j^*}$, which requires:

$$p(1)v(1-r_i) > p(\alpha)v(1-r_{j^*}) \implies p(1)(1-r_i) > p(\alpha)(1-r_{j^*})$$

However, j^* by definition offers the best (or equal best) surplus share, so $p_{j^*}(1-r_{j^*}) \geq p_i(1-r_i)$ for any competing offer. Furthermore, for any significant number of traders α, the effect of network size on inclusion probability will be substantial, meaning $p(\alpha) \gg p(1)$. A single trader abandoning the dominant pool for a solo pool would face a drastically lower chance of inclusion. The deviation is not profitable.

Since no single trader has an incentive to deviate from the monopolistic state, it is a stable Nash Equilibrium.

(c) **No other Nash Equilibria:**
No other type of equilibrium is possible. Any block creator who is not offering the maximal S_j will be strictly less attractive to all traders and will therefore receive no order flow.

□

Nonexistence of SPNE with positive POF and Proof of Proposition 8

Recall the timing of the game Γ:

1. **Stage 1 (fees).** EAs simultaneously choose fees $r_i \in [0, 1)$.
2. **Stage 2 (routing).** Given a realized fee vector $\mathbf{r}$, traders simultaneously choose routing sets and associated flows, and we restrict attention to Nash equilibria of this routing subgame.

A SPNE of Γ is a profile of (possibly mixed) fee strategies for EAs and routing strategies for traders such that, for every realized $\mathbf{r}$, the induced routing profile is a Nash equilibrium of the Stage–2 subgame, and each EA's fee strategy is a best reply given the induced continuation payoffs.

We now formalize what happens when the no–solo–entry condition fails under strictly positive duplication cost $\Delta\ell(2) > 0$, and contest primitives stated earlier.

Let $\tilde{G}_t(\varepsilon, \mathbf{r}, r_k)$ denote the incremental expected net payoff for a trader t of mass ε from *adding* an inactive EA k as a second private route, given that the incumbent EA $i^\star$ currently receives all POF under fees $\mathbf{r}$ and the entrant charges a fee r_k. Violations of no–solo–entry mean that for some trader t and some EA k we can profitably seed the entrant at arbitrarily small scale.

Assumption 6 (Violation of no–solo–entry)**.** *There exist an EA k, a trader t, and $\bar{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}]$ there exists a fee $\hat{r}_k(\varepsilon) \in [0, 1)$ and an incumbent fee region $\mathcal{R}_{i^\star}(\varepsilon) \subset [0, 1)$ with the following properties:*

1. *For any incumbent fee $r_{i^\star} \in \mathcal{R}_{i^\star}(\varepsilon)$, when the routing subgame is played with the fee vector $(r_{i^\star}, \hat{r}_k(\varepsilon), r_{-i^\star,k})$, trader t strictly prefers to add k as a second private route:*

$$\tilde{G}_t(\varepsilon, (r_{i^\star}, r_{-i^\star}), \hat{r}_k(\varepsilon)) \;>\; \Delta\ell(2).$$

2. *EA k obtains strictly positive expected fee revenue from the seeded ε–mass flow at $(r_{i^\star}, \hat{r}_k(\varepsilon))$.*

Assumption 6 is a formalization of "profitable seeding": whenever the incumbent's fee lies in a nondegenerate region $\mathcal{R}_{i^\star}(\varepsilon)$, an inactive EA k can choose a seeding fee $\hat{r}_k(\varepsilon)$ that makes some ε–mass trader strictly better off by multihoming and generates strictly positive revenue for k.[5]

We can now restate Proposition 8.

Proposition. (No SPNE with positive POF under profitable seeding) Suppose $\Delta\ell(2) > 0$ and Assumption 6 holds. Then the two–stage fee–routing game Γ admits no subgame–perfect Nash equilibrium—pure or mixed—in which private orderflow is positive in some subgame. In any SPNE of Γ (if one exists), POF is never used in equilibrium.

Proof. Suppose, towards a contradiction, that there exists a SPNE of Γ in which POF is positive with positive probability. Let σ denote the profile of (possibly mixed) fee strategies for the EAs, and let $\rho(\cdot \mid \mathbf{r})$ denote the equilibrium routing profile in the Stage–2 subgame following fees $\mathbf{r}$.

By assumption, there exists a measurable set of fee realizations $\mathcal{S} \subset [0, 1)^N$ with $\Pr_\sigma(\mathcal{S}) > 0$ such that, for every $\mathbf{r} \in \mathcal{S}$, the routing equilibrium $\rho(\cdot \mid \mathbf{r})$ assigns strictly positive POF to some EA $i^\star$. Because duplication is costly ($\Delta\ell(2) > 0$) and the contest primitives stated earlier, any routing equilibrium for a fixed $\mathbf{r}$ has at most one EA with positive POF and all POF traders single–home on that EA (following directly from Lemma 3). Thus, for each $\mathbf{r} \in \mathcal{S}$, there is a *unique* incumbent EA $i^\star$ that receives all private orderflow in $\rho(\cdot \mid \mathbf{r})$.

Fix such a fee profile $\mathbf{r} \in \mathcal{S}$ and its incumbent $i^\star$. Consider a trader t and an entrant EA $k \neq i^\star$ as in Assumption 6, and choose $\varepsilon > 0$ sufficiently small and $r_{i^\star} \in \mathcal{R}_{i^\star}(\varepsilon)$. Under the

[5] Competition from the public mempool implies that the set of fee pairs $(r_{i^\star}, r_k)$ for which seeding is profitable is open in $[0, 1)^2$. Assumption 6 simply selects an open strip of incumbent fees and a best–response seeding fee within that region.

equilibrium routing $\rho(\cdot \mid \mathbf{r})$, trader t currently single–homes on $i^\star$ and pays the duplication cost only if they multihome. Assumption 6 then guarantees the existence of a seeding fee $\hat{r}_k(\varepsilon)$ such that, when the fee vector changes from $\mathbf{r}$ to $\mathbf{r}' := (\hat{r}_k(\varepsilon), r_{-k})$, trader t strictly prefers to add k as a second private route in any Nash equilibrium of the Stage–2 subgame following $\mathbf{r}'$, and EA k earns strictly positive fee revenue from the seeded ε–mass flow.

Because we are in a SPNE, EAs correctly anticipate that traders best–respond to any deviation in fees; therefore, deviating from $\mathbf{r}$ to $\mathbf{r}'$ is a profitable unilateral deviation for EA k in the Stage–1 fee game. This contradicts the assumption that σ is a Nash equilibrium in Stage 1. Hence there can be no SPNE in which there exists a fee realization $\mathbf{r}$ with positive probability under σ for which some EA receives positive POF in the continuation game.

It remains to argue that randomization over fees cannot restore equilibrium with positive POF. Suppose, again for contradiction, that there exists a SPNE σ with positive POF and define $\mathcal{S}$ as above. For each incumbent EA $i^\star$ that receives positive POF with positive probability, the support of its equilibrium fee distribution under σ must be such that:

(i) there is zero probability mass on all regions $\mathcal{R}_{i^\star}(\varepsilon)$ for which Assumption 6 yields a profitable seeding deviation for some entrant k, otherwise k would have a profitable deviation; and

(ii) given (i), $i^\star$ has no profitable upward deviation in fees.

However, because the set of fee pairs $(r_{i^\star}, r_k)$ that make seeding profitable for some ε is open by continuity of primitives, condition (i) requires that the support of $i^\star$'s fee distribution be contained in a (possibly lower–dimensional) set that avoids an open subset of $[0, 1)$. For generic primitives, any such restriction forces the incumbent to choose fees strictly below the limit price that would equate traders' net benefits to the duplication cost, thus leaving a positive profit margin on the inframarginal traders it serves. Then the incumbent has a profitable upward deviation in Stage 1, violating (ii). Therefore, no choice of mixed strategies over fees can simultaneously satisfy both (i) and (ii) while sustaining positive POF in any subgame.

We conclude that there is no SPNE—pure or mixed—in which private orderflow is positive in any subgame. Any SPNE of Γ (if it exists) must feature zero POF at every realized fee vector, i.e., all traders always use the public mempool and private routing is never active. □

Proposition 8 is not a knife–edge result. Profitable seeding is defined by an open region of fee pairs; mixed strategies cannot "average out" the deviation because, with positive probability, an incumbent fee will fall into a region where an entrant can profitably seed and, conversely, excluding all such regions forces the incumbent to leave rents on the table.

The failure of equilibrium is thus fundamental rather than an artefact of a particular tie–breaking rule or of restricting attention to pure strategies.

References

Adams, Hayden, Noah Zinsmeister, and Dan Robinson. 2020. *Uniswap V2.* https://uniswap.org/whitepape

Adams, Hayden, et al. 2021. *Uniswap V3: Concentrated Liquidity.* https://uniswap.org/whitepaper-v3.pdf.

Arnosti, Nick, and S. Matthew Weinberg. 2022. Bitcoin: A Natural Oligopoly. *Management Science* 68(1): 152–171.

Baumol, William, John C. Panzar, and Robert D. Willig. 1982. *Contestable Markets and the Theory of Industry Structure.* Harcourt.

Canidio, Andrea, and Vabuk Pahari. 2025. Becoming Immutable: How Ethereum Is Made. Working paper.

Capponi, Agostino, Ruizhe Jia, and Ye Wang. 2023. Blockchain Private Pools and Price Discovery. AEA Papers and Proceedings 2023, 113: 253–256

Capponi, Agostino, Jiaoyang Huang, and Sveinn Olafsson. 2024. Proposer-Builder Separation, Exclusive Order Flow, and Centralization in Blockchain. Working paper.

Catalini, Christian, and Joshua S. Gans. 2016. Some Simple Economics of the Blockchain. NBER Working Paper No. 22952.

Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2019. Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability. arXiv:1904.05234.

Dixit, Avinash. 1978. A Model of Duopoly Suggesting a Theory of Entry Barriers and Economies of Scale. *Bell Journal of Economics* 10(1): 20–32.

Farrell, Joseph, and Garth Saloner. 1986. Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation. *American Economic Review* 76(5): 940–955.

Gosselin, Stephane, and Ankit Chiplunkar. 2023. The Order Flow Auction (OFA) Design Space. Flashbots Research Report.

Gupta, Tivas, Mallesh M. Pai, and Max Resnick. 2023. The Centralizing Effects of Private Order Flow on Proposer?Builder Separation. In *Proceedings of the 5th Conference on Advances in Financial Technologies (AFT).*

Heimbach, Lioba, Vabuk Pahari, and Eric Schertenleib. 2024. Non-Atomic Arbitrage in Decentralized Finance. In *IEEE Symposium on Security and Privacy.*

Heimbach, Lioba, Lucianna Kiffer, Christof Ferreira Torres, and Roger Wattenhofer. 2023. Ethereum?s Proposer-Builder Separation: Promises and Realities. In *ACM Internet*

Measurement Conference (IMC).

Katz, Michael L., and Carl Shapiro. 1985. Network Externalities, Competition, and Compatibility. *American Economic Review* 75(3): 424–440.

Kilbourn, Philip. 2022. Order Flow, Auctions, and Centralisation — A Warning. Flashbots Writings (Research Blog).

Konrad, Kai. 2009. *Strategy and Dynamics in Contests*. Oxford University Press.

Maskin, Eric, and Jean Tirole. 1988. A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs. *Econometrica* 56(3): 549?569.

Milionis, Jason, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2024. Automated Market Making and Loss-Versus-Rebalancing. arXiv:2208.06046.

Nti, Kofi O. 1998. Effort and Performance in Group Contests. *European Journal of Political Economy* 14(4): 769–781.

Oz, Burak, Danning Sui, Thomas Thiery, and Florian Matthes. 2024. Who Wins Ethereum Block Building Auctions and Why? In *Proceedings of the 6th Conference on Advances in Financial Technologies (AFT)*.

Pahari, Vabuk, and Andrea Canidio. 2025. How Exclusive Are Ethereum Transactions? Working paper.

Rochet, Jean-Charles, and Jean Tirole. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1(4): 990–1029.

Spence, A. Michael. 1977. Entry, Capacity, Investment and Oligopolistic Pricing. *Bell Journal of Economics* 8(2): 534?544.

Thiery, Thomas. 2023. Empirical Analysis of Builders? Behavioral Profiles (BBPs). Ethereum Research Forum (August 10). https://ethresear.ch/t/empirical-analysis-of-builders-behavioral-profiles-bbps/16327.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. MIT Press.

Tullock, Gordon. 1980. Efficient Rent Seeking. In *Toward a Theory of the Rent-Seeking Society*, edited by James M. Buchanan, Robert D. Tollison, and Gordon Tullock, Texas A&M University Press.

Wahrstätter, Anton, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to Bribe: Measuring Block Construction Market. arXiv:2305.16468.

Wu, Fei, Danning Sui, Thomas Thiery, and Mallesh Pai. 2025. Measuring CEX–DEX Extracted Value and Searcher Profitability: The Darkest of the MEV Dark Forest. arXiv:2507.13023.

Yang, Sen, Kartik Nayak, and Fan Zhang. 2025. Decentralization of Ethereum?s Builder Market. In *IEEE Symposium on Security and Privacy*.