

The Price of Decentralization in Block Building

Burak Öz*  

Flashbots, Berlin, Germany

Fei Wu*†  

King's College London, London, UK

Luis Correia  

Flashbots, London, UK

Sen Yang  

Yale University, New Haven, CT, USA

IC3, New York, NY, USA

Bruno Mazonra  

Flashbots, Barcelona, Spain

Stefanos Leonardos  

King's College London, London, UK

Abstract

Decentralized block building mechanisms replace the monopoly of a single proposer with multiple builders. However, their censorship-resistance and fair-access benefits depend not only on the number of builders, but also on where builders are geographically positioned to provide transaction coverage. We study this tension between builder location choice, user transaction coverage, and utility concentration by modeling decentralized block building as a stochastic coverage game. Builders choose regions, information sources emit transactions over a block construction round, and latency determines whether a transaction is received before the deadline.

We show that the game is an exact potential game and admits a pure Nash equilibrium. We prove an asymptotically tight factor-2 bound on the Price of Anarchy, which we interpret as the *price of decentralization* from uncoordinated builder placement. We also study builder utility concentration, showing that the lowest-utility builder earns at least half of the highest-utility builder's payoff, and the utility-share HHI is at most 12.5% above the egalitarian benchmark. We complement the theory with simulations under richer latency and source environments and show that slot times, builder participation, and reward-sharing rules are important protocol design choices that can shape the price of decentralization.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Computing methodologies → Modeling and simulation

Keywords and phrases decentralization, block building, potential game, price of anarchy

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.25

Related Version *Full Version*: <https://arxiv.org/abs/2606.01874> [37]

Supplementary Material *Software (Source Code)*: <https://github.com/boz1/db-sims> [1]
archived at `swb:1:dir:afa10f961c32d9645eb1b3fad7a89764a79699aa`

*These authors contributed equally to this work.

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



1 Introduction

Many blockchain protocols assign the block proposal right to a single leader in each slot. This temporary monopoly gives the proposer control over transaction inclusion and ordering, creating opportunities for Maximal Extractable Value (MEV) [8]. Moreover, proposers located in particular jurisdictions may face censorship pressures, such as those induced by OFAC sanctions in the U.S. [33], and geographic concentration can create potential liveness risks in the presence of regional outages. Recent multiple-proposer or decentralized block-building proposals seek to reduce this monopoly and improve censorship resistance by allowing multiple builders to construct candidate blocks or sub-blocks [16, 22].¹ If a user transaction can be included by multiple builders, who typically share the fees paid by the transaction, this increases the cost of censorship: excluding a transaction may require bribing multiple entities rather than a single proposer.

However, multiple builders do not automatically imply geographically broad transaction coverage. In permissionless blockchains, the consensus protocol typically does not mandate where builders should operate. Because latency affects whether transactions reach builders before the block construction deadline, self-interested builders may prefer regions that provide access to the same high-value transaction sources. Recent work shows that Ethereum’s block-building protocol design is not geographically neutral and can create co-locating incentives to minimize latency to payoff-critical parties [36]. If builders concentrate in a small number of regions, the benefits of having multiple builders can diminish, as users located farther away may face weaker inclusion guarantees due to latency effects. In auction-like settings, such latency differences can give colocated parties an informational advantage [24].

This raises a critical design question: if decentralized block building mechanisms rely on multiple self-interested builders who choose locations to maximize individual utility, will their location choices actually improve access for geographically distributed users and transaction sources? How inefficient and concentrated can the resulting placements be? These questions matter because the censorship-resistance and fair-access benefits of multi-builder designs depend not only on the number of builders, but also on whether their locations are distributed enough to provide timely coverage to users and transaction sources across regions over the existing Internet infrastructure with latency constraints.

In this paper, we model decentralized block building as a stochastic coverage game (cf. Section 2). Builders choose geographic regions, while information sources emit transactions over the course of a block-construction round. A builder’s region determines the probability that a transaction emitted by a given source reaches the builder before the block deadline. Since multiple builders may receive and include the same transaction, decentralized placement can create redundant coverage: several builders may cover the same high-value sources, while other sources remain weakly covered.

We study the welfare loss induced by this decentralized builder region game (cf. Section 3). Relative to a social planner who coordinates builder locations to maximize expected transaction coverage, selfish builders may over-concentrate in regions that provide access to the same valuable information sources. We measure this efficiency loss through the Price of Anarchy (PoA), which we interpret in this setting as the *price of decentralization*, and derive an asymptotically tight worst-case bound on it using the smoothness framework [31]. The key observation is that when a transaction is already covered by several builders, an additional

¹ In the paper, we use *builder* as the generic term for an entity that constructs and proposes a candidate (sub-)block.

builder receives only a reduced share of its value under the equal-split sharing rule, while not improving welfare. This allows us to relate unilateral deviations to aggregate welfare. We also prove the existence of a pure Nash equilibrium by constructing an exact potential function, and establish bounds on utility concentration among builders at equilibrium. In our model, welfare captures the total transaction value covered by at least one builder, while utility concentration captures how the economic benefits of block building are distributed.

The theoretical results characterize worst-case properties: whether stable placements exist, how inefficient they can be, and how concentrated builder utilities can become. To complement this analysis, we run simulations in richer latency and source environments. These simulations show when welfare losses arise in concrete instances, how they are driven by under-coverage of peripheral sources, and why geographic concentration and builder utility concentration need not coincide (cf. Sections 4 and 5). The results highlight how protocol designs can shape welfare in decentralized block building (cf. Section 6). We summarize our contributions as follows.

- We model decentralized block building as a latency-induced stochastic coverage game in which builders choose regions and transaction sources generate value over time. We analyze the equilibrium and efficiency properties of the game under the equal-split sharing rule. We show that the game admits a pure Nash equilibrium through an exact potential function and prove a tight factor-2 PoA bound.
- We study utility concentration among builders in equilibrium. We show that the lowest-utility builder earns at least half of the highest-utility builder's payoff, and that the resulting utility-share HHI is at most 12.5% above the egalitarian benchmark. Both bounds are tight.
- We use simulations [1] to study the game under richer latency and source environments. We find that welfare losses are largest in intermediate regimes where peripheral sources are reachable, but selfish incentives still favor high-value source regions. We also show that geographic concentration and utility concentration need not align: planner allocations may improve coverage by placing builders in lower-payoff peripheral regions, while equilibrium outcomes can be more geographically centralized but more utility-balanced.

1.1 Related work

Geographic (de)centralization. Prior measurement studies show that Ethereum nodes are geographically concentrated, with deployments repeatedly skewed toward a small number of countries, particularly in Europe and North America [6, 11, 19–21]. These studies establish geographic concentration as an empirical phenomenon, while recent work studies the incentives and mechanisms that shape geographic decentralization. Yang et al. [36] show that Ethereum's block building regimes create location-dependent validator rewards, allowing latency advantages and asymmetric access to information sources to reinforce concentration. Roeschlin et al. [30] study how to enforce geography-aware incentives when nodes may misreport their locations, using RTT-based localization to support truthful reporting. Our work is complementary to both lines. We study equilibrium outcomes in decentralized block building when builders strategically choose geographic regions, and analyze how latency-sensitive placement affects welfare, geographic concentration, and builder utility concentration.

Decentralized block building. Decentralized block building aims to reduce centralized control over transaction inclusion and ordering, especially under MEV-driven incentives, and to improve censorship resistance. Industry discussions have advocated distributed, parallel, and global block building architectures to counter latency moats and censorship risks [13, 23].

Recent work proposes Multiple Concurrent Proposers (MCP) designs to reduce the serial monopoly of a single proposer over inclusion and ordering [15, 16, 22]. Complementary work studies transaction assignment in such protocols, showing that stronger censorship resistance often requires more transaction duplication and hence lower throughput [10]. Our work complements these works through placement equilibria: we study how builders' strategic geographic choices affect welfare, geographic concentration, and builder utility concentration.

Coverage games and equilibrium inefficiency. Vetta introduced valid utility games, where agents' utilities are tied to a monotone submodular social objective and Nash equilibria admit constant welfare guarantees [32]. Gairing studied covering games, where players select subsets of weighted elements and covered value is allocated through utility-sharing rules [14]. Our equal-split rule corresponds to symmetric equal sharing, but coverage is stochastic and induced by heterogeneous source-to-builder reception probabilities.

Prior work studies the distributional properties of equilibria through notions such as utility uniformity [12] and fairness ratio for costs [3]. Our utility dispersion and HHI results complement this literature by studying reward concentration in a stochastic coverage game.

Beyond worst-case PoA guarantees, related work examines inefficiency in particular parameter regimes and among outcomes reached in practice or through specified response dynamics [4, 7, 26]. Our simulations take a similar instance-level perspective by evaluating equilibria selected by asynchronous better-response dynamics across sampled environments and initial conditions.

2 Model

2.1 Sources, regions, and probabilistic reception

We study a decentralized block-building network in which, in every round of a blockchain protocol, a new block containing transactions is appended. Each round has a duration $\Delta > 0$. Denote the set of builders by $\mathcal{B} := \{1, \dots, K\}$ and the set of geographical regions in the network by $\mathcal{R} := \{1, \dots, R\}$.

In each round, every builder constructs a sub-block from the transactions it receives before the round deadline. We assume that all received transactions are included and impose no block-capacity constraint. We discuss the implications of block-capacity constraints in Section 6.

Let $\mathcal{I}$ be a finite set of information sources. Each source $I \in \mathcal{I}$ is associated with a fixed region $r(I) \in \mathcal{R}$. For each source I , let $\mathcal{J}_I$ denote the random set of transactions emitted by source I during one round. Each transaction $j \in \mathcal{J}_I$ has an emission time $t_j \in [0, \Delta]$ and a positive value V_j .

We assume that source I has finite expected emitted value in a round. Specifically, there exist a scalar $\Lambda_I > 0$ and a density $\rho_I : [0, \Delta] \rightarrow \mathbb{R}_+$ with $\int_0^\Delta \rho_I(t) dt = 1$ such that, for every bounded function $h : [0, \Delta] \rightarrow \mathbb{R}$,

$$\mathbb{E} \left[\sum_{j \in \mathcal{J}_I} V_j h(t_j) \right] = \Lambda_I \int_0^\Delta h(t) \rho_I(t) dt.$$

Thus, Λ_I is the expected total value emitted by source I in one round, and ρ_I describes how this expected value is distributed over time within this round. This formulation lets us work directly with expected value flow rather than specifying the full arrival process of transactions and values.

For each ordered region pair $(i, k) \in \mathcal{R} \times \mathcal{R}$, let $D_{k \rightarrow i}$ denote the propagation latency from a source in region k to a builder located in region i , and let $F_{k \rightarrow i}$ be its cumulative distribution function. We allow the latency distribution to depend on the ordered region pair, so different source-builder pairs may exhibit different propagation characteristics. This captures persistent geographical differences in network connectivity while allowing for random latency fluctuations.

A strategy for builder b consists of selecting a region $s_b \in \mathcal{R}$, and a strategy profile is written as $\mathbf{s} := (s_1, \dots, s_K) \in \mathcal{R}^K$. Region selection is modeled as a one-shot simultaneous-move game: builders choose regions based on the known source environment, but not on the contemporaneous choices of the other builders.

Consider a transaction emitted at time $t \in [0, \Delta]$ by a source located in region k . If builder b is located in region $i = s_b$, then the transaction reaches builder b before the block deadline with probability

$$q_k(i, t) := \Pr(D_{k \rightarrow i} \leq \Delta - t) = F_{k \rightarrow i}(\Delta - t).$$

Thus, for a source I and a builder b , the corresponding reception probability is obtained by setting $k = r(I)$ and $i = s_b$, giving $q_{r(I)}(s_b, t)$. Since reception probabilities depend on builders only through their chosen regions, all builders located in the same region have identical reception probabilities.

Let $\mathcal{T}_b$ denote the random set of transactions received by builder b in a round, and write $\mathcal{T}_b^I := \mathcal{T}_b \cap \mathcal{I}_I$ for the transactions from source I received by b . Conditional on the realized transactions and their emission times, receptions are independent across builders. Thus, two builders located in the same region have the same reception probability for a given transaction, but need not have the same realized reception outcome.

2.2 Value sharing rule and welfare

After a round, the sub-blocks of all builders are aggregated into a final block. When multiple builders include the same transaction, its value is shared among them. For each transaction j , let

$$n_j(\mathbf{s}) := |\{b \in \mathcal{B} : j \in \mathcal{T}_b\}|$$

denote the number of builders that received and included transaction j under profile $\mathbf{s}$. When the profile is clear from context, we write n_j .

We assume an *equal-split sharing rule*: if $n_j \geq 1$, then each builder that included transaction j receives the fraction $1/n_j$ of its value. Builders that did not include j receive zero from it. Equivalently, the reward assigned to builder b from transaction j is

$$V_j \cdot \frac{\mathbf{1}\{j \in \mathcal{T}_b\}}{n_j},$$

with the convention that this term is zero when $n_j = 0$.

We use equal splitting as a simple symmetric and budget-balanced sharing rule: it depends only on the number of builders that include the transaction, treats all including builders equally, and distributes exactly V_j when the transaction is covered by at least one builder. Thus, the total reward distributed for transaction j is $V_j \mathbf{1}\{n_j \geq 1\}$. This rule is consistent with Shapley-style distribution designs for multi-proposer transaction fee mechanisms, where the reward from a transaction is shared among the proposers that include it [15]. We abstract from the bidding layer and treat V_j as the transaction payment to be shared; alternative reward-sharing rules are discussed in Section 6.

The realized reward of builder b under profile $\mathbf{s}$ is

$$\Pi_b(\mathbf{s}) = \sum_{I \in \mathcal{I}} \sum_{j \in \mathcal{J}_I} V_j \cdot \frac{\mathbf{1}\{j \in \mathcal{T}_b\}}{n_j},$$

again with the convention that the summand is zero when $n_j = 0$.

Since transaction generation, including values and emission times, and reception outcomes are random, $\Pi_b(\mathbf{s})$ is a random variable. The payoff-relevant quantity for builder b is its expected reward

$$u_b(\mathbf{s}) := \mathbb{E}[\Pi_b(\mathbf{s})],$$

where the expectation is taken over transaction generation and reception outcomes. Given the other builders' region choices, each builder chooses s_b to maximize its expected reward.

We define the *social welfare* of profile $\mathbf{s}$ as the expected total value of transactions received by at least one builder:

$$W(\mathbf{s}) := \mathbb{E} \left[\sum_{I \in \mathcal{I}} \sum_{j \in \mathcal{J}_I} V_j \mathbf{1}\{n_j \geq 1\} \right].$$

Thus, $W(\mathbf{s})$ measures the total user value realized by transactions that are successfully covered under profile $\mathbf{s}$. Following the standard welfare view in prior work [2], we treat transfers between users and builders as redistributing value rather than creating it. Therefore, in our stylized model, welfare equals the total value of included transactions.

The equal-split sharing rule implies a useful accounting identity: aggregate builder utility equals welfare. We state this formally in Lemma 5.

2.3 Equilibrium and social optimum

The preceding definitions induce a finite game in which builders are the players, regions are the actions, and expected rewards u_b are the payoffs. We refer to this game as the *builder region game*.

A pure strategy profile $\mathbf{s}$ is a Nash equilibrium of the builder region game if no builder can improve its expected reward through a unilateral deviation. That is, for every $b \in \mathcal{B}$ and every $s'_b \in \mathcal{R}$,

$$u_b(\mathbf{s}) \geq u_b(s'_b, \mathbf{s}_{-b}),$$

where $\mathbf{s}_{-b}$ denotes the strategies of all builders other than b .

A social planner chooses the entire profile $\mathbf{s}$ jointly to maximize welfare. We denote a welfare-maximizing profile by

$$\mathbf{s}^* \in \arg \max_{\mathbf{s} \in \mathcal{R}^K} W(\mathbf{s}),$$

and call $W(\mathbf{s}^*)$ the optimal welfare.

3 Equilibrium, Efficiency, and Utility Concentration

In this section, we analyze the equilibrium, efficiency, and utility-concentration properties of the builder region game. We first show that the equal-split sharing rule induces an exact potential function, which implies the existence of a pure Nash equilibrium. We then bound the welfare loss in equilibrium using the PoA. Finally, we study concentration in builder utilities through the Herfindahl–Hirschman Index (HHI). We provide the proofs in the appendix of the full version [37].

3.1 Pure Nash equilibrium via an exact potential

We first show that the builder region game admits a pure Nash equilibrium by constructing an exact potential function. Under equal sharing, welfare itself is not a potential function: if a transaction is already covered by other builders, an additional builder can still obtain positive utility by also receiving it, while welfare does not increase. The key observation is instead that the marginal payoff from joining a transaction already received by m other builders is exactly $V_j/(m+1)$, which corresponds to the increment of a harmonic potential.

► **Proposition 1** (Exact potential under equal sharing). *Consider the builder region game under the equal-split sharing rule. For $m \geq 1$, let*

$$H_m := \sum_{\ell=1}^m \frac{1}{\ell}, \quad H_0 := 0$$

denote the harmonic numbers. For a realized transaction j , let

$$n_j(\mathbf{s}) := |\{b \in \mathcal{B} : j \in \mathcal{T}_b\}|$$

denote the number of builders that receive j under profile $\mathbf{s}$. Define

$$\Phi(\mathbf{s}) := \mathbb{E} \left[\sum_{I \in \mathcal{I}} \sum_{j \in \mathcal{J}_I} V_j H_{n_j(\mathbf{s})} \right].$$

Then Φ is an exact potential for the expected-utility game. That is, for every builder $b \in \mathcal{B}$, every profile $\mathbf{s}$, and every deviation $s'_b \in \mathcal{R}$,

$$\Phi(s'_b, \mathbf{s}_{-b}) - \Phi(\mathbf{s}) = u_b(s'_b, \mathbf{s}_{-b}) - u_b(\mathbf{s}).$$

Hence, under equal sharing, the builder region game is a finite exact potential game.

► **Corollary 2** (Pure equilibria and better-response convergence). *Under the equal-split sharing rule, the builder region game admits a pure Nash equilibrium. Moreover, every asynchronous sequence of strict better responses terminates in finite time at a pure Nash equilibrium.*

Proof. By Proposition 1, the game is a finite exact potential game. Hence, any maximizer of Φ over the finite pure profile space $\mathcal{R}^K$ is a pure Nash equilibrium [25].

Now consider an asynchronous sequence of strictly better responses. At each step, a single builder changes its region in a way that strictly increases its expected utility. By exact potentiality, this strictly increases the potential Φ . Since the pure profile space $\mathcal{R}^K$ is finite, such strict improvements cannot continue indefinitely. The process must therefore terminate at a profile with no profitable unilateral deviation, i.e., a pure Nash equilibrium. ◀

In the deterministic-reception special case, the game admits a direct congestion-game representation, where transaction sources are resources and a source of value Λ_I gives payoff Λ_I/m to each of its m covering builders, who dilute each others rewards. With probabilistic reception, expected payoff depends on the full vector of heterogeneous probabilities, so this compact source-level representation is less direct. Abstractly, every finite exact potential game can be mapped to a congestion game [25].

3.2 Price of Anarchy in decentralized block building

In the builder region game, builders may over-concentrate in regions that provide access to valuable information sources. This can lead multiple builders to include the same transactions, creating redundant coverage, while other sources remain weakly covered. The resulting mismatch between private incentives and aggregate coverage can reduce welfare. We use PoA to measure the welfare loss caused by selfish region selection relative to the welfare-maximizing profile chosen by a social planner.

Let $\mathcal{E}$ denote the set of pure Nash equilibria of the builder region game, and let $\mathbf{s}^*$ be a welfare-maximizing profile. We define PoA as

$$\text{PoA} := \frac{W(\mathbf{s}^*)}{\min_{\mathbf{s} \in \mathcal{E}} W(\mathbf{s})}.$$

The bound we derive follows the standard PoA argument for valid-utility games [32], a class that includes coverage and location-game variants, expressed through the smoothness framework [31]. The smoothness framework bounds inefficiency by comparing the total payoff players would obtain by unilaterally deviating from a current profile toward a benchmark profile with the welfare of the two profiles. Our contribution is to verify that the builder region game satisfies the required structure after taking expectations over transaction value flows and latency-induced probabilistic reception. Specifically, we show that welfare admits a source-level coverage representation, that equal-split sharing makes aggregate builder utility equal welfare, and that deviation payoffs can be lower-bounded by exclusive coverage. These properties allow the standard smoothness proof to carry over to our stochastic block-building setting.

A payoff-maximization game is (λ, μ) -smooth if, for every pair of profiles $\mathbf{s}$ and $\mathbf{s}^*$,

$$\sum_{b \in \mathcal{B}} u_b(s_b^*, \mathbf{s}_{-b}) \geq \lambda W(\mathbf{s}^*) - \mu W(\mathbf{s}).$$

If a game is (λ, μ) -smooth and aggregate utility equals welfare, then every pure Nash equilibrium has welfare at least a $\lambda/(1 + \mu)$ fraction of the optimal welfare. Equivalently,

$$\text{PoA} \leq \frac{1 + \mu}{\lambda}.$$

► **Definition 3 (Coverage).** *A transaction from source I emitted at time t is covered under profile $\mathbf{s}$ if at least one builder receives it before the deadline. The corresponding coverage probability is*

$$f_I(t; \mathbf{s}) := 1 - \prod_{b \in \mathcal{B}} (1 - q_{r(I)}(s_b, t)).$$

The expected coverage of source I under profile $\mathbf{s}$, averaged with respect to the source's emission-time profile ρ_I , is

$$\bar{f}_I(\mathbf{s}) := \int_0^\Delta \rho_I(t) f_I(t; \mathbf{s}) dt.$$

► **Lemma 4 (Coverage representation of welfare).** *For every profile $\mathbf{s}$, total expected welfare satisfies*

$$W(\mathbf{s}) = \sum_{I \in \mathcal{I}} \Lambda_I \bar{f}_I(\mathbf{s}).$$

The proof applies the expected value-flow assumption with $h(t) = f_I(t; \mathbf{s})$ and then sums over sources.

Interpretation. The lemma shows that welfare depends only on whether each transaction is covered by at least one builder. Additional builders receiving the same transaction do not create additional social value. Thus, welfare losses arise from inefficient region allocation: if too many builders cover the same high-value sources, other sources may remain weakly covered relative to the social optimum.

► **Lemma 5** (Aggregate utility identity). *Under the equal-split sharing rule, aggregate builder utility equals welfare. That is, for every profile $\mathbf{s}$,*

$$\sum_{b \in \mathcal{B}} u_b(\mathbf{s}) = W(\mathbf{s}).$$

This follows because a covered transaction distributes exactly its full value across the builders that include it, while an uncovered transaction distributes zero.

► **Lemma 6** (Smoothness inequality). *The builder region game is $(1, 1)$ -smooth. That is, for all profiles $\mathbf{s}$ and $\mathbf{s}^*$,*

$$\sum_{b \in \mathcal{B}} u_b(s_b^*, \mathbf{s}_{-b}) \geq W(\mathbf{s}^*) - W(\mathbf{s}).$$

The proof lower-bounds each builder's deviation payoff by the value it would cover exclusively under the deviation. Summing these exclusive-coverage terms across builders and applying a union-bound argument shows that they cover the welfare gap $W(\mathbf{s}^*) - W(\mathbf{s})$.

► **Theorem 7** (Price of Anarchy bound). *The PoA of the builder region game is at most 2. Equivalently, for every pure Nash equilibrium $\mathbf{s} \in \mathcal{E}$ and every welfare-maximizing profile $\mathbf{s}^*$,*

$$W(\mathbf{s}) \geq \frac{1}{2} W(\mathbf{s}^*).$$

► **Remark 8** (Broader equilibrium notions). The factor-2 guarantee follows from the standard smoothness argument for valid utility games [32]. The block-building-specific content is the verification that a latency-induced stochastic region-selection game preserves this structure in expectation. Although Theorem 7 is stated for pure Nash equilibria, the same smoothness argument implies the factor-2 welfare guarantee for mixed Nash, correlated, and coarse correlated equilibria, with welfare evaluated in expectation [31].

We next construct an example to show that the factor-2 bound is asymptotically tight. The worst case in this game arises when all builders crowd into a region with a single highly valuable source that is already fully covered. A unilateral move to an uncovered secondary source's region would improve social coverage but reduce builder's utility. In contrast, a social planner would keep only one builder at the dominant source and allocate the remaining builders to uncovered source regions.

► **Proposition 9** (Tightness of the factor-2 bound). *For every $K \geq 2$ and every $\varepsilon \in (0, 1)$, there exists an instance of the builder region game with K builders, such that some pure Nash equilibrium $\mathbf{s}$ satisfies*

$$\frac{W(\mathbf{s}^*)}{W(\mathbf{s})} = 2 - \frac{1}{K} - \frac{K-1}{K} \varepsilon.$$

In particular, both the PoA and PoS can be arbitrarily close to 2.

► **Remark 10** (Price of Stability). The same construction also implies that the best and worst equilibria coincide, and the game admits a unique pure Nash equilibrium. Therefore, the factor-2 bound is asymptotically tight for the Price of Stability.

3.3 Builder utility dispersion and concentration

We next study how dispersed builder utilities can be at equilibrium. Our utility-dispersion measure coincides with the notion of *utility uniformity* studied in [12]. We then use this dispersion bound to quantify utility concentration among builders under decentralized block building with the equal-split sharing rule.

► **Lemma 11** (Equilibrium utility dispersion). *Under the equal-split sharing rule, let $\mathbf{s}$ be a pure Nash equilibrium. Define*

$$u_{\min}(\mathbf{s}) := \min_{b \in \mathcal{B}} u_b(\mathbf{s}), \quad u_{\max}(\mathbf{s}) := \max_{b \in \mathcal{B}} u_b(\mathbf{s}).$$

Then

$$u_{\min}(\mathbf{s}) \geq \frac{1}{2} u_{\max}(\mathbf{s}).$$

► **Proposition 12** (Tightness of the utility dispersion bound). *The factor $1/2$ in Lemma 11 is tight: there exists an instance and a pure Nash equilibrium $\mathbf{s}$ such that*

$$u_{\min}(\mathbf{s}) = \frac{1}{2} u_{\max}(\mathbf{s}).$$

► **Definition 13** (Utility shares and concentration). *For any profile $\mathbf{s}$ with $\sum_{b \in \mathcal{B}} u_b(\mathbf{s}) > 0$, define the utility share of builder b by*

$$x_b(\mathbf{s}) := \frac{u_b(\mathbf{s})}{\sum_{c \in \mathcal{B}} u_c(\mathbf{s})}.$$

The Herfindahl–Hirschman Index (HHI) of the utility profile is

$$\text{HHI}(\mathbf{s}) := \sum_{b \in \mathcal{B}} x_b(\mathbf{s})^2.$$

Its minimum possible value over K builders is attained at equal utilities:

$$\text{HHI}^{\text{egal}} = \frac{1}{K}.$$

► **Proposition 14** (Equilibrium bound on utility concentration). *Assume the equal-split sharing rule, and let $\mathbf{s}$ be a pure Nash equilibrium with $\sum_{b \in \mathcal{B}} u_b(\mathbf{s}) > 0$. Then*

$$\text{HHI}(\mathbf{s}) \leq \frac{9}{8K} = \frac{9}{8} \text{HHI}^{\text{egal}}.$$

In addition, the factor $9/8$ is tight as a universal bound.

Interpretation. This proposition shows that, under decentralized block building with the equal-split sharing rule, utility shares among builders cannot be much more concentrated than the egalitarian benchmark. At any pure Nash equilibrium, the utility-share HHI is at most 12.5% above the value attained when all builders earn equal utility. This follows from the equilibrium dispersion bound: the lowest-utility builder must earn at least half of the highest-utility builder’s payoff.

4 Simulation Setup

The theoretical results provide worst-case guarantees that hold uniformly across model instances. Simulations allow us to ask a complementary, instance-level question: how do welfare, transaction coverage, and concentration vary with concrete source-value, latency, and builder-participation configurations? To answer this question, we compare self-interested equilibrium outcomes with centralized planner benchmarks across a set of controlled simulation environments. For reproducibility, we open-source our simulator [1].

Regions and information sources. We instantiate the region set $\mathcal{R}$ using 24 regions from Google Cloud Platform (GCP), partitioned into two source-location pools (see the appendix of the full version [37]). The high-value pool contains regions near major financial or technology centers, such as North America, Western Europe, and East Asia². The peripheral pool contains regions that are geographically further from these centers. For each simulation instance, we sample the same number of source regions from the high-value and peripheral pools, without replacement, and place one information source in each selected region. Source locations remain fixed throughout a run. We denote the high-value and peripheral source clusters by $\mathcal{I}_H$ and $\mathcal{I}_P$, respectively. We normalize the total expected emitted value to $\sum_{I \in \mathcal{I}} \Lambda_I = 10$ in all reported experiments. Source-value asymmetry is parameterized by the high-to-low value share ratio γ , defined as the ratio between the total expected value emitted by $\mathcal{I}_H$ and $\mathcal{I}_P$. Thus, the two clusters receive total expected values $10 \cdot \frac{\gamma}{1+\gamma}$ and $10 \cdot \frac{1}{1+\gamma}$, respectively, divided evenly within each cluster.

Transaction and latency. To generate realized transactions, the simulator uses a Poisson arrival process with emission times sampled uniformly in $[0, \Delta]$. In the reported experiments, transaction values are deterministic and chosen so that each source has the desired expected emitted value Λ_I . As in previous literature [34, 36], the simulator also supports lognormal transaction values, but we set the value noise to zero to isolate latency and source-value effects. When sweeping Δ , we hold each Λ_I fixed, so slot-duration experiments isolate the deadline effect rather than changing total emitted value per slot.

We calibrate the propagation delays using an empirical GCP pairwise latency matrix. For a source region k and builder region i , latency $D_{k \rightarrow i}$ is modeled as lognormal with mean equal to the calibrated empirical latency and standard deviation 15% of the mean, with a small lower bound to avoid degeneracy. The corresponding reception probability is $q_k(i, t) = \Pr(D_{k \rightarrow i} \leq \Delta - t)$.

Equilibrium profiles. Builders choose regions from the same 24-region set and are initialized uniformly at random. We use asynchronous better-response dynamics (ABR), which is guaranteed to converge to a pure Nash equilibrium (PNE) under the exact-potential structure of the game (cf. Corollary 2). Builders update one at a time in round-robin order. When builder b is selected, all other builders' regions are fixed; candidate regions are considered in random order, and b moves to the first region that strictly improves its expected utility. If no improving region exists, it stays.

In the simulator, utilities are evaluated numerically, and convergence is declared when a full round-robin sweep completes without any builder changing regions. In all reported experiments, every simulation run converged to a profile with no utility-improving unilateral deviation within 250 builder updates. We denote the equilibrium profile by $\mathbf{s}^{\text{PNE}}$.

² This classification follows recent observations that Ethereum latency competition depends on relay geography and price-signal paths across North America, Europe, and East Asia [9].

To compute a builder's utility at a candidate region, we use the equal-split sharing rule directly. If builder b considers moving to region $i \in \mathcal{R}$, then for a transaction from source I emitted at time t , builder b receives the transaction with probability $q_{r(I)}(i, t)$. If $X_{I,t}^{-b}$ other builders also receive it, then builder b 's share is $\frac{1}{1+X_{I,t}^{-b}}$. Thus,

$$u_b(i, \mathbf{s}_{-b}) = \sum_{I \in \mathcal{I}} \Lambda_I \int_0^\Delta \rho_I(t) q_{r(I)}(i, t) \mathbb{E} \left[\frac{1}{1 + X_{I,t}^{-b}} \right] dt.$$

Here, $X_{I,t}^{-b}$ is the random number of builders other than b that receive a transaction from source I emitted at time t , under the profile $\mathbf{s}_{-b}$. Since receptions are independent across builders, $X_{I,t}^{-b}$ follows a Poisson-binomial distribution with parameters $\{q_{r(I)}(s_{b'}, t)\}_{b' \neq b}$. We compute the probability mass function by dynamic programming, evaluate $\mathbb{E}[1/(1 + X_{I,t}^{-b})]$, and approximate the time integral using a uniform grid of points. For each parameter setting, we sample five random source layouts. For each source layout, we simulate three independently sampled initial builder profiles, giving 15 runs per parameter setting.

4.1 Benchmarks and metrics

Planner benchmark. For each simulation instance, we compare the equilibrium profile with a centralized social-planner benchmark. When exact optimization is tractable, the planner benchmark is $\mathbf{s}^* \in \arg \max_{\mathbf{s} \in \mathcal{R}^K} W(\mathbf{s})$, and $W^* := W(\mathbf{s}^*)$. We assume builders are symmetric, so exact search can be performed over occupancy vectors rather than over all labeled profiles in $\mathcal{R}^K$. The number of such occupancy vectors is $\binom{|\mathcal{R}|+K-1}{K}$. We use exact exhaustive search in the experiments with fixed number of builders $K = 5$. For experiments that sweep over larger builder counts, exact search becomes computationally expensive, so we use a greedy welfare-maximization routine as the planner benchmark. Across all tested instances, the worst observed greedy-to-exact welfare ratio is 0.999129 (see the appendix of the full version [37]). In experiments using the greedy benchmark, welfare ratios should be interpreted relative to the greedy planner benchmark rather than the exact optimum.

Metrics. The primary efficiency metric is *welfare ratio*, calculated as $WR := W_{\text{PNE}}/\widehat{W}^*$, where $\widehat{W}^* = W^*$ for exact planner search and equals the greedy planner welfare when the greedy benchmark is used. When $\widehat{W}^* = W^*$, the theoretical PoA ≤ 2 bound corresponds to $WR \geq 1/2$.

We report utility concentration with *utility HHI* following Definition 13. This metric measures how concentrated expected builder utilities are at the selected PNE.

Geographical concentration of builder locations is measured with *geographic HHI*. Let $g_i(\mathbf{s}) := |\{b \in \mathcal{B} : s_b = i\}|/K$ be the fraction of builders located in region $i \in \mathcal{R}$. We define

$$\text{HHI}_{\text{geo}}(\mathbf{s}) = \sum_{i \in \mathcal{R}} g_i(\mathbf{s})^2.$$

When $K \leq |\mathcal{R}|$, its minimum value is $1/K$, attained when all builders occupy distinct regions. Its maximum value is 1, attained when all builders co-locate in a single region.

Finally, for a source cluster $G \subseteq \mathcal{I}$, we define *value-weighted cluster coverage* by

$$\text{Coverage}(G; \mathbf{s}) = \frac{\sum_{I \in G} \Lambda_I \bar{f}_I(\mathbf{s})}{\sum_{I \in G} \Lambda_I}.$$

We report this metric for $G = \mathcal{I}_H$ and $G = \mathcal{I}_P$. Higher coverage means that transactions from that cluster are more likely to reach at least one builder before the deadline.

5 Simulation Results

We now study how the equilibria behave across concrete source-value and latency environments. The goal is to understand how self-interested builder location choice affects welfare, transaction coverage, geographic centralization, and builder utility concentration in equilibrium, and how these outcomes compare with those in the welfare-optimal planner benchmark. The experiments vary the two main forces in the model. Source-value asymmetry changes builders' private incentives by making some sources more attractive than others. The slot duration controls how strongly latency affects whether transactions arrive before the deadline.

Experiment 1: source-value asymmetry. We isolate the effect of source-value asymmetry by sweeping the high-to-low source value ratio at a fixed slot duration and builder count.

Experiment 2: latency deadline effect. We isolate the latency deadline effect by sweeping the slot duration while holding value asymmetry fixed.

Experiment 3: interaction effect. We study the interaction between source-value asymmetry and latency sensitivity by jointly varying the value ratio and the slot duration.

Experiment 4: builder participation. We vary the number of builders to test whether additional builders expand coverage or mainly duplicate coverage around attractive sources.

The main figures in these experiments report aggregate metrics across randomized source layouts and initializations. These summaries are useful for comparing welfare, coverage, and concentration, but they do not show the underlying builder placements in individual simulation runs. To make the placement mechanism more concrete, we visualize a representative builder placement in equilibrium and the planner benchmark for each experiment in the appendix of the full version [37].

5.1 Experiment 1: source-value asymmetry effect


This experiment isolates the effect of source-value asymmetry. We test whether increasing the value gap between high-value and peripheral sources causes builders to concentrate near high-value sources, creating redundant coverage and reducing peripheral-source coverage.

Setup. We fix $K = 5$ builders and $\Delta = 50\text{ms}$ slot duration.³ Each simulation instance contains five high-value sources and five peripheral sources, sampled from their respective region pools. We vary the high-to-low source value ratio γ from 1 to 20, while keeping the total expected emitted value fixed. Thus, increasing the ratio reallocates value toward the high-value sources without changing the total value generated.

Results. In the top-left panel of Figure 1, the welfare in equilibrium approaches that in the planner benchmark as the high-to-low value ratio increases. The welfare ratio rises toward one, indicating that builder placement in equilibrium recovers most of the optimum welfare achieved by the planner when high-value sources dominate the total value.

This improvement is not driven by region-level co-location. The bottom-left panel shows that geographic HHI remains essentially at the uniform baseline $1/K$ throughout the value-ratio sweep. In this experiment, there are exact $K = 5$ builders and five high-value sources, so builders can cover the high-value cluster in equilibrium by spreading across those

³ We use $K = 5$ as a baseline, motivated by the fact that the current Ethereum builder market is concentrated among a few builders [28, 29, 35]. 50ms slot duration is motivated by Solana Constellation's configuration, in which proposers submit transactions and the leader assembles a batch in each cycle [22].



■ **Figure 1** *Effect of source-value asymmetry.* Lines show medians across runs; shaded bands show inter-quartile ranges. As the high-to-low source value ratio increases, equilibrium welfare approaches the optimal welfare while geographic HHI remains exactly $1/K$. The main effect is coverage reallocation: at intermediate value ratios, coverage already concentrates on high-value sources in equilibrium, whereas the planner maintains peripheral source coverage and induces higher utility concentration.


regions rather than co-locating. The main effect of source-value asymmetry is therefore the reallocation of coverage. In the top-right panel, as the high-to-low value ratio increases, builders become increasingly incentivized to cover high-value sources in equilibrium: high-value cluster coverage quickly approaches one, while peripheral cluster coverage drops sharply. In the planner benchmark, the planner keeps peripheral coverage higher at intermediate value ratios, but this comes with a higher builder utility HHI (bottom-right panel), indicating that welfare-improving coverage can require allocating builders to lower-payoff regions.

Overall, in this balanced-capacity setting, source-value asymmetry reallocates coverage toward high-value sources without inducing region-level co-location: with as many high-value source regions as builders, builders can spread across distinct high-value regions. This should not be interpreted as saying that value asymmetry cannot create geographic concentration. If high-value sources are themselves geographically clustered, the same incentive to cover high-value sources may instead lead builders to concentrate in a small number of regions.

5.2 Experiment 2: slot-duration effect

Next, we fix the source-value distribution and vary the slot duration, isolating the other main force in the model: deadline-induced latency sensitivity. This experiment tests how the block-construction round duration changes latency-sensitive builder placement. When the duration is short, propagation latency has a stronger effect on whether transactions reach builders before the deadline. When the duration is long, more source-builder pairs can receive transactions in time, and geographic placement should matter less for coverage.

Setup. We fix $K = 5$ builders and high-to-low value ratio $\gamma = 10$. Each instance contains five high-value sources and five peripheral sources. We vary slot durations from 10ms to 12s. The lower end is tighter than the proposed 50ms Solana Constellation cycle [22], while the upper end matches the 12s slot duration used by Ethereum Proof-of-Stake [5]. When



■ **Figure 2** *Effect of slot duration.* Larger slot durations relax the latency deadline and increase source-cluster coverage. The welfare gap is largest at intermediate durations, where peripheral coverage is feasible in the planner benchmark but not reliably supplied in equilibrium. At long durations, coverage is nearly saturated, so geographic concentration has limited welfare cost even when the builder placement in equilibrium remains more concentrated than in the planner benchmark.

sweeping slot durations Δ , we hold each source’s expected emitted value Λ_I fixed, so the experiment isolates the effect of the deadline on transaction reception rather than changing the amount of value generated per round.

Results. Figure 2 shows a non-monotone effect of slot duration. In the top-left panel, when slots are very short, the welfare ratio is close to one, but this does not indicate broad coverage. Rather, the deadline is so tight that builders can effectively capture value only from nearby sources. Since high-value sources dominate the value distribution in this experiment, both the planner and the equilibrium incentive place builders near high-value source regions, leaving peripheral sources mostly uncovered.

At intermediate slot durations, peripheral sources become reachable for some regions, but coverage still depends strongly on builder placement. This is where the welfare ratio falls: the planner spreads builders across regions to improve coverage of different source clusters, while in equilibrium, builders are concentrated around regions with better access to high-value sources. The top-right cluster-coverage panel shows the resulting mechanism: high-value coverage is already near saturation, whereas peripheral coverage in equilibrium lags behind that in the planner benchmark. Thus, the welfare gap comes primarily from under-coverage of peripheral sources in equilibrium, induced by redundant coverage around already attractive high-value sources. The utility HHI panel indicates that this welfare gap is not primarily an effect of payoff inequality: builder utilities remain relatively balanced even when peripheral coverage is under-provided.

This transition is best interpreted relative to the latency scale of the underlying region graph (see the appendix of the full version [37]). The largest welfare losses arise in the intermediate regime, where Δ is comparable to cross-region latencies: peripheral coverage is feasible, but only for appropriately placed builders. When Δ is much smaller than typical cross-region propagation delays, builders can capture value only from nearby sources, so

builders have limited ability to cover peripheral sources in both the equilibrium and the planner benchmark. When Δ is much larger than these latencies, most source-builder pairs can communicate before the deadline, so transaction coverage becomes less sensitive to builder geography. In this long-slot regime, geographic centralization may persist, but it no longer produces a large welfare loss in this model because coverage is close to saturated.

We repeat this experiment with a lower value ratio of $\gamma = 2$, where peripheral sources contribute to a larger share of the total value. The qualitative patterns are similar, but the welfare gap is larger at intermediate slot durations, where peripheral sources are valuable enough to matter and reachable enough for the planner to cover. We provide the figure and discussion in the appendix of the full version [37].

5.3 Experiment 3: value-asymmetry–slot-duration interaction effect

Having isolated source-value asymmetry and slot duration separately, we now study their interaction. The goal is to identify which combinations produce inefficient welfare outcomes, weak transaction coverage, or concentrated builder placement in equilibrium.

Setup. Each simulation instance contains five builders, five high-value sources, and five peripheral sources, sampled from their respective region pools. We sweep the high-to-low source value ratio γ from 1 to 20 and the slot duration from 10ms to 12s.

Results. Figure 3 shows that source-value asymmetry and slot duration jointly determine how close the equilibrium outcome is to the welfare-optimal outcome, and reinforces our prior results. The welfare ratio is primarily non-monotone in the slot duration. At both extremes, welfare in equilibrium is close to that in the planner benchmark. The largest welfare loss occurs at intermediate slot durations, where peripheral sources are reachable from certain regions but coverage remains sensitive to builder placement. The value ratio affects how costly this peripheral under-coverage is. When the high-to-low value ratio is small, and the peripheral sources meaningfully contribute to the total emitted value, the welfare loss becomes larger in equilibrium. As the value ratio increases, high-value sources dominate the total emitted value, so the welfare penalty from leaving peripheral sources under-covered becomes smaller.

The coverage-gap surfaces identify the mechanism. A positive gap means the corresponding source cluster is better covered in the planner benchmark than in equilibrium. When the value ratio $\gamma = 1$, the two source clusters have equal expected value, so the labels “high-value” and “peripheral” distinguish geography rather than value. In this case, planner advantage in one cluster and equilibrium advantage in the other need not create a large welfare gap.

As the value ratio increases, the high-value cluster receives a larger share of total value, and builders in equilibrium have stronger incentives to prioritize regions with access to high-value sources. At intermediate slot durations, peripheral sources are still reachable from some regions, so the planner can recover additional welfare by allocating builders to improve peripheral coverage. By contrast, peripheral coverage is under-provided in equilibrium, and builder placement remains tilted toward high-value sources, contributing to a negative coverage gap in high-value source coverage, a positive gap in peripheral source coverage, and a corresponding welfare penalty.

The HHI surfaces are consistent with this mechanism. Utility concentration becomes higher in the planner benchmark because welfare-improving allocations yield lower individual payoffs for builders placed away from high-value sources. Conversely, the utility HHI in




Figure 3 *Joint effect of source-value asymmetry and slot duration.* Surfaces show medians across source layouts and initializations for fixed $K = 5$. Darker colors indicate higher absolute values. Coverage gaps are defined as planner coverage minus equilibrium coverage, so positive gaps indicate that the planner covers more of the corresponding source cluster. The HHI panels compare geographic and utility concentration in equilibrium and in the planner benchmark. Welfare losses are largest at intermediate slot durations and lower-to-moderate value ratios.


equilibrium remains balanced and close to the egalitarian baseline in much of the grid, but it comes with a more geographically centralized builder placement around the high-value sources.

5.4 Experiment 4: builder participation effect

Previous experiments isolate the effects of source-value asymmetry and slot duration, as well as their interaction. We now vary the number of builders to test whether larger builder sets expand source coverage or mainly duplicate coverage around already attractive sources. Adding builders increases the system’s potential coverage capacity, but when source values are asymmetric, additional builders may still prefer regions with strong access to high-value sources.

Setup. We vary the number of builders from $K = 3$ to $K = 12$. We fix the high-to-low value ratio at $\gamma = 10$ and the slot duration at $\Delta = 50\text{ms}$. For this builder-count sweep, we use the greedy welfare-maximization routine as the planner benchmark, since the cost of exact search grows rapidly with the builder count K .

Results. Figure 4 shows that increasing the number of builders does not automatically yield broader source coverage in equilibrium. In the top-left panel, the welfare ratio is non-monotone in K . When the number of builders is small, additional builders mainly improve high-value source coverage, so the welfare in equilibrium quickly moves closer to the optimal welfare. Once high-value coverage is nearly saturated, further welfare gains require better peripheral-source coverage.



■ **Figure 4** *Effect of builder participation.* As the number of builders K increases, high-value source coverage is improved in equilibrium, while peripheral coverage remains low, contributing to the welfare ratio first increasing then declining. Geographic HHI shows that the builder locations remain more spatially concentrated in equilibrium than the planner benchmark for larger K , while utility HHI remains comparatively close to the egalitarian baseline.

The top-right cluster-coverage panel shows that the planner uses additional builders to expand peripheral-source coverage. In equilibrium, however, they remain attracted to regions with strong access to high-value sources, even after those sources are nearly saturated. As a result, larger builder sets can generate redundant high-value coverage rather than broad source coverage, explaining why the welfare ratio first improves and then declines.

The HHI panels reveal a tradeoff between geographic decentralization and utility equality. As the number of builders increases, geographic HHI moves closer to the minimum $1/K$ in the planner benchmark, because additional builders expand coverage toward peripheral sources. Builder placement in equilibrium remains more geographically centralized: additional builders continue to locate in regions with stronger access to the high-value cluster, even after high-value coverage is nearly saturated.

The pattern is reversed for utility concentration. Utility HHI is closer to the egalitarian baseline in equilibrium, while it is higher in the planner benchmark. Improved peripheral source coverage in the planner benchmark achieves higher total welfare, but yields lower individual utility for those builders located in the peripheral regions. By contrast, in equilibrium, builders are incentivized to locate in high-value regions, where equal sharing can make payoffs more balanced. Thus, geographic decentralization and builder-utility equality need not coincide: improving source coverage and total welfare may require distributing builders into lower-payoff regions, while selfish behavior can achieve a more utility-balanced but geographically more centralized outcome in equilibrium.

We repeat this experiment with a lower value ratio of $\gamma = 2$, where peripheral sources contribute to a larger share of the total value. The qualitative patterns are similar, but the welfare gap is more pronounced. We provide the figure and discussion in the appendix of the full version [37].

6 Discussion

We now discuss the implications of our results for decentralized block-building design. The central message is that decentralization is not only a question of how many builders participate, but also of where self-interested builders choose to locate. In our model, the *price of decentralization* measures the welfare cost of decentralized, self-interested location decisions relative to a centrally coordinated placement that maximizes transaction coverage. Decentralized decisions may themselves produce a geographically concentrated outcome. We then discuss directions for future work and limitations.

6.1 Implications for decentralized block-building design

Our theory and simulations point to a structural feature of decentralized block-building or MCP designs. Multiple builders reduce the monopoly power of a single proposer, but if their location choices are uncoordinated, equilibrium placement can still duplicate coverage around privately attractive high-value sources, weakening the system’s censorship-resistance and fair-access guarantees.

The simulations identify where this price of decentralization is most visible. Welfare losses are largest in intermediate regimes where peripheral sources are reachable and still contribute meaningful value, but decentralized equilibrium incentives still favor regions with strong access to high-value sources. Thus, inefficiency is not driven simply by source-value asymmetry or by latency sensitivity alone, but by their interaction: peripheral coverage must be feasible enough to matter, while still requiring builders to locate away from the most privately attractive regions.

From a protocol-design perspective, the relevant timing parameter is not slot duration in isolation, but slot duration relative to transaction propagation latencies. Longer slot durations or more relaxed (sub-)block construction deadlines can make transaction coverage less geography-dependent in our model, but geographic centralization may persist when its welfare penalty becomes small. Conversely, shorter deadlines make source-builder proximity more important, while transaction coverage is restricted to be local, limiting coverage of distant sources. The fragile regime lies between these extremes, where cross-region coverage is possible but still placement-sensitive.

Increasing builder participation expands potential coverage capacity. However, in the finite regimes our simulations examine (up to 12 builders), additional builders may not cover peripheral sources immediately. When individual incentives favor high-value source access, extra builders can duplicate coverage around already attractive regions. The results also show that geographic and utility concentration capture different objectives. Planner allocations may be geographically more dispersed but utility-wise more unequal, because the planner places builders in lower-payoff peripheral regions to improve total coverage. Equilibrium profiles may instead be more geographically concentrated while keeping builder utilities relatively balanced under equal sharing.

Since source-value asymmetry is largely an environmental primitive, while timing, builder participation, and reward-sharing rules are at least partly protocol design choices, these parameters are natural levers for shaping transaction coverage incentives. Overall, protocol designers should distinguish between maximizing welfare and coverage, improving geographic decentralization, and equalizing builder rewards, since these objectives need not move together.

6.2 Future work

Location-market modeling. Our model treats decentralized block building as a stochastic coverage game. Transaction sources are exogenous: they emit transactions, builders choose regions, and latency determines whether each transaction reaches at least one builder before the deadline. This formulation is useful for studying source-side coverage and the welfare loss from redundant coverage of the same sources.

A richer model would endogenize transaction routing. In such a model, users or applications choose which builders to send transactions to, and the builder location affects which transactions a builder can serve with sufficiently low latency. Builder rewards would then depend not only on which sources are reachable, but also on how users route order flow and how builders compete for that flow. This would make the model closer to a location-market game than a pure coverage game.

Such a formulation may be important for protocols and applications in which latency-sensitive users actively choose among builders, such as the Solana ecosystem with Constellation [22]. High-value order-flow sources, such as exchanges or latency-sensitive application servers, may emerge as strategic hubs: co-locating builders may receive higher-quality order flow or offer stronger latency guarantees. We leave the analysis of endogenous transaction routing and builder competition for future work.

Alternative reward-sharing rules. Our analysis focuses on equal splitting among builders who receive and include a transaction. Future work could compare this rule with other reward-sharing mechanisms. Winner-takes-all or unique-inclusion rules may reduce duplicate inclusion, but can strengthen incentives for exclusivity, private routing, and reward concentration. Broader committee-level sharing [15] may encourage builders to maximize aggregate source coverage, but can create free-riding incentives if rewards are paid to builders that do not contribute to coverage.

Another natural extension is contribution-weighted or Tullock-style sharing. The equal-split rule can be viewed as the uniform-score special case of proportional sharing. More general rules could assign contribution scores, for example, based on the source-builder propagation latency, and allocate value proportionally to the score. Such rules interpolate between egalitarian sharing and winner-takes-all allocation, and relate to work on the blockspace distribution and contest-style mechanisms [17, 18, 27]. Studying their effects on welfare, geographic centralization, and reward equality is an important direction for future work.

6.3 Scope and limitations

We close by clarifying the scope of the model and simulations. The goal of the paper is not to provide a full protocol-level simulator, but to isolate how latency-sensitive transaction value shapes the region choices of self-interested builders in decentralized block-building.

First, our simulations use asynchronous better-response dynamics to select a pure Nash equilibrium. We use this dynamic because it is theoretically aligned with the exact-potential structure of the game under equal sharing and gives a stable equilibrium benchmark. Simultaneous updates or learning dynamics may cycle or induce non-stationary empirical distributions; while studying such out-of-equilibrium behavior is an interesting direction, it is outside the scope of this paper. We use ABR as an equilibrium-selection procedure, not as a literal description of how all builders relocate in practice. The underlying incentive tension remains the same: builders prefer regions with stronger access to certain high-value sources, while the system may benefit from broader source coverage.

Our exact-potential argument guarantees finite convergence of strict asynchronous better-response dynamics, but it does not provide a polynomial bound on the length of improvement paths. Thus, the complexity of computing a pure Nash equilibrium in the builder region game is a natural direction for future work. For finite, explicitly represented instances with polynomial-time evaluable utilities, pure Nash equilibrium computation can be viewed as a local-search problem over the potential function, suggesting that a formal PLS analysis may be appropriate. In our simulations, however, ABR served only as an equilibrium-selection procedure, and all reported runs converged to a no-improving-deviation profile within 250 builder updates. Moreover, the smoothness guarantee in Theorem 7 extends beyond pure Nash equilibria to weaker equilibrium notions, so the factor-2 efficiency bound is not tied only to the ABR selection procedure.

Second, we abstract from many communication details. We model source-to-builder latency directly and use a common deadline for transaction reception. In practice, protocols may involve peer-to-peer propagation, intermediaries (e.g., relays), committee attestation, and leader aggregation. These features would require a richer model of consensus-layer and network communication. However, they do not remove the central mechanism studied here: location affects timely access to payoff-relevant information. Additional latency-sensitive targets may, in fact, create more location incentives [36], reinforcing the same kind of concentration pressure.

Third, our simulations use stylized source locations and value asymmetries rather than direct empirical estimates of order-flow geography. Such calibration is difficult because high-value transaction sources and their value distribution are not directly observable. We therefore use empirical observations on geographic latency competition in Ethereum block building [9] to instantiate the region graph and sweep broad ranges of value asymmetry, slot duration, and builder participation. The simulation results should be read as comparative statics over plausible environments, rather than as a point estimate for a particular network.

Finally, we omit several strategic features that are important for full protocol design. We impose no block-capacity constraint, fix the builder set, and rule out Sybil behavior and collusion. Capacity constraints would couple transaction selection decisions and may strengthen the welfare-loss channel we identify, because redundant inclusion of already-covered high-value transactions could crowd out otherwise uncovered transactions. Entry and rotation would affect which builders are active in a given round and could change equilibrium selection. Sybil behavior and collusion would require a richer model of coalitional incentives, where coordinated builders may improve welfare by avoiding redundant coverage, but they could also undermine the censorship-resistance benefits of decentralized block building by behaving as a single economic entity. These extensions are therefore important for full protocol design, but complementary to the location-choice mechanism studied here, which focuses on the inefficiency of uncoordinated self-interested placement.

7 Conclusion

We studied location choice in decentralized block building by modeling it as a stochastic coverage game. Our results show that replacing a single proposer with multiple builders is not sufficient, by itself, for broad transaction coverage or welfare-efficient placement. Under equal sharing, the builder region game admits a pure Nash equilibrium, has an asymptotically tight factor-2 PoA bound, and satisfies tight bounds on builder utility concentration. Our simulations complement these worst-case guarantees by showing that welfare losses are most pronounced in intermediate regimes where peripheral sources are reachable and valuable, but

self-interested builders still prefer regions with strong access to high-value sources. Overall, the paper highlights that the benefits of decentralized block building depend not only on the number of builders, but also on where builders locate and how protocol rules shape their incentives. Timing parameters, participation rules, and reward-sharing mechanisms can therefore affect the price of decentralization, especially in networks with asymmetric source values and latency-sensitive transaction propagation.

References

- 1 Authors. Decentralized Block Building Simulator, 2026. URL: <https://github.com/boz1/db-sims>.
- 2 Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. Transaction Fee Mechanism Design in a Post-MEV World. In Rainer Böhme and Lucianna Kiffer, editors, *6th Conference on Advances in Financial Technologies (AFT 2024)*, volume 316 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:24, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2024.29.
- 3 Davide Bilò, Tobias Friedrich, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. Fair tree connection games with topology-dependent edge cost. In Nitin Saxena and Sunil Simon, editors, *40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020)*, volume 182 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.FSTTCS.2020.15.
- 4 Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 138–144. AAAI Press, 2013. doi:10.1609/aaai.v27i1.8667.
- 5 Vitalik Buterin, Diego Hernandez, Thor Kamphofner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining ghost and casper. *arXiv preprint arXiv:2003.03052*, 2020. arXiv:2003.03052.
- 6 Cambridge Centre for Alternative Finance (CCAF). Ethereum network analytics. https://ccaf.io/cbnsi/ethereum/network_analytics, 2026. Accessed: 2026-05-09.
- 7 Riccardo Colini-Baldeschi, Roberto Cominetti, Panayotis Mertikopoulos, and Marco Scarsini. The asymptotic behavior of the price of anarchy. In *Web and Internet Economics: 13th International Conference, WINE 2017, Bangalore, India, December 17–20, 2017, Proceedings*, pages 133–145, Berlin, Heidelberg, 2017. Springer-Verlag. doi:10.1007/978-3-319-71924-5_10.
- 8 Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability. In *2020 IEEE symposium on security and privacy (SP)*, pages 910–927, Los Alamitos, CA, USA, 2020. IEEE Computer Society. doi:10.1109/SP40000.2020.00040.
- 9 Data Always. The Geography of Block Building. The Flashbots Collective, November 2025. Flashbots. URL: <https://collective.flashbots.net/t/the-geography-of-block-building/5367>.
- 10 Fatima Elsheimy, Ioannis Kaklamanis, Sarisht Wadhwa, Charalampos Papamanthou, and Fan Zhang. Censorship Resistance vs Throughput in Multi-Proposer BFT Protocols. *Cryptology ePrint Archive*, 2026.
- 11 ethernodes.org. Ethernodes – countries. <https://www.ethernodes.org/countries>, 2026. Accessed: 2026-05-09.
- 12 Michal Feldman, Kevin Lai, and Li Zhang. The proportional-share allocation market for computational resources. *IEEE Transactions on Parallel and Distributed Systems*, 20(8):1075–1088, 2009. doi:10.1109/TPDS.2008.168.

- 13 Flashbots. decentralized building: wat do? <https://writings.flashbots.net/decentralized-building-wat-do>, 2026. Accessed: 2026-05-09.
- 14 Martin Gairing. Covering games: Approximation through non-cooperation. In Stefano Leonardi, editor, *Internet and Network Economics*, volume 5929, pages 184–195. Springer Berlin Heidelberg, 2009. Series Title: Lecture Notes in Computer Science. doi:10.1007/978-3-642-10841-9_18.
- 15 Pranav Garimidi, Lioba Heimbach, and Tim Roughgarden. Transaction fee mechanism design for leaderless blockchain protocols. In *Financial Cryptography and Data Security: 29th International Conference, FC 2025, Miyakojima, Japan, April 14–18, 2025, Revised Selected Papers, Part II*, pages 20–35, Berlin, Heidelberg, 2025. Springer-Verlag. doi:10.1007/978-3-032-07035-7_2.
- 16 Pranav Garimidi, Joachim Neu, and Max Resnick. Multiple concurrent proposers: Why and how. Cryptology ePrint Archive, Paper 2025/1772, 2025. URL: <https://eprint.iacr.org/2025/1772>.
- 17 Pranav Garimidi, Michael Neuder, and Tim Roughgarden. Tullock contests in the wild: Applications in blockchains. *ACM SIGecom Exchanges*, 23(1):24–34, July 2025. doi:10.1145/3821553.3821558.
- 18 Pranav Garimidi, Michael Neuder, and Tim Roughgarden. Beyond winner-take-all procurement auctions. In *Financial Cryptography and Data Security: 30th International Conference, FC 2026*, 2026. URL: <https://ifca.ai/fc26/preproceedings/231.pdf>.
- 19 Lioba Heimbach, Yann Vonlanthen, Juan Villacis, Lucianna Kiffer, and Roger Wattenhofer. Deanonymizing Ethereum Validators: The P2P Network Has a Privacy Issue. In *34th USENIX Security Symposium*, Seattle, WA, USA, 2025. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity25/presentation/heimbach>.
- 20 Lucianna Kiffer, Lioba Heimbach, Dennis Trautwein, Yann Vonlanthen, and Oliver Gasser. Multiple Sides of 36 Coins: Measuring Peer-to-Peer Infrastructure Across Cryptocurrencies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 9(3):1–32, 2025. doi:10.1145/3771572.
- 21 Seoung Kyun Kim, Zane Ma, Siddharth Murali, Joshua Mason, Andrew Miller, and Michael Bailey. Measuring Ethereum Network Peers. In *Proceedings of the Internet Measurement Conference 2018*, pages 91–104, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3278532.3278542.
- 22 Quentin Knip, Max Resnick, Jakub Sliwinski, and Roger Wattenhofer. Solana constellation: Internet capital markets. White paper, Anza, 2026.
- 23 Robert Miller. Parallel block building. <https://writings.flashbots.net/parallel-builder>, 2024. Accessed: 2026-05-09.
- 24 Ciamac C. Moallemi, Mallesh M. Pai, and Dan Robinson. Latency advantages in common-value auctions, 2026. arXiv:2504.02077.
- 25 Dov Monderer and Lloyd S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996. doi:10.1006/game.1996.0044.
- 26 Barnabé Monnot, Francisco Benita, and Georgios Piliouras. How bad is selfish routing in practice? arXiv preprint arXiv:1703.01599, 2017. doi:10.48550/arXiv.1703.01599.
- 27 Michael Neuder, Pranav Garimidi, and Tim Roughgarden. On block-space distribution mechanisms, 2024. Accessed: 2026-05-25. URL: <https://ethresear.ch/t/on-block-space-distribution-mechanisms/19764>.
- 28 Burak Öz, Danning Sui, Thomas Thiery, and Florian Matthes. Who Wins Ethereum Block Building Auctions and Why? In Rainer Böhme and Lucianna Kiffer, editors, *6th Conference on Advances in Financial Technologies (AFT 2024)*, volume 316 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:25, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.AFT.2024.22.
- 29 relayscan.io. MEV-Boost Relay & Builder Stats. <https://www.relayscan.io/>, 2026. Accessed: 2026-05-18.

- 30 Marc Roeschlin, Evangelos Markakis, Raghav Bhaskar, and Aggelos Kiayias. Incentivizing Geographic Diversity for Decentralized Systems. *Cryptology ePrint Archive*, 2026.
- 31 Tim Roughgarden. Intrinsic Robustness of the Price of Anarchy. *J. ACM*, 62(5), November 2015. doi:10.1145/2806883.
- 32 Adrian Vetta. Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions . In *The 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 416–425, Los Alamitos, CA, USA, 2002. IEEE, IEEE Computer Society. doi:10.1109/SFCS.2002.1181966.
- 33 Anton Wahrstätter, Jens Ernstberger, Aviv Yaish, Liyi Zhou, Kaihua Qin, Taro Tsuchiya, Sebastian Steinhorst, Davor Svetinovic, Nicolas Christin, and Mikolaj Barczentewicz. Blockchain censorship. *arXiv preprint arXiv:2305.18545*, 2023. doi:10.48550/arXiv.2305.18545.
- 34 Fei Wu, Thomas Thiery, Stefanos Leonardos, and Carmine Ventre. Strategic bidding wars in on-chain auctions. In *2024 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 503–511, Dublin, Ireland, 2024. IEEE. doi:10.1109/ICBC59979.2024.10634354.
- 35 Sen Yang, Kartik Nayak, and Fan Zhang. Decentralization of Ethereum’s Builder Market. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 1456–1474, Los Alamitos, CA, USA, 2025. IEEE Computer Society. doi:10.1109/SP61157.2025.00157.
- 36 Sen Yang, Burak Öz, Fei Wu, and Fan Zhang. Geographical Centralization Resilience in Ethereum’s Block-Building Paradigms. *Proc. ACM Meas. Anal. Comput. Syst.*, 10(2), June 2026. doi:10.1145/3805637.
- 37 Burak Öz, Fei Wu, Luis Correia, Sen Yang, Bruno Mazorra, and Stefanos Leonardos. The price of decentralization in block building, 2026. doi:10.48550/arXiv.2606.01874.