A Note on the Welfare Gap in Fair Ordering

Theo Diamandis Guillermo Angeris
tdiamand@mit.edu gangeris@baincapital.com

March 2023

Abstract

In this short note, we show a gap between the welfare of a traditionally 'fair' ordering, namely first-in-first-out (an ideal that a number of blockchain protocols strive to achieve), where the first transactions to arrive are the ones put into the block, and an 'optimal' inclusion that is, at least approximately, welfare-maximizing, such as choosing which transactions are included in a block via an auction. We show this gap is positive under a simple model with mild assumptions where we assume transactions are, roughly speaking, uniformly drawn from a reasonable distribution.

Introduction

When distorting a market, there is almost always a loss in efficiency. Typically, this loss manifests as a loss of social welfare: goods are not allocated to the individuals with the highest valuations (or utilities). In this short note, we examine ways in which first-in-first-out public blockchains might distort their block space markets and the resulting decrease in social welfare. This distortion is most commonly seen in the way in which blockchains allocate and charge for block space.

Public blockchains generally implement a *fee mechanism* to allocate finite computational resources across competing transactions. These transactions are broadcasted by users through the peer-to-peer network and collected in the mempool, which contains submitted transactions that have not yet been included on chain. A validator then selects which transactions from the mempool are included in the next block, possibly subject to an *inclusion rule*. (We note that this view is overly simplistic, but it suffices for the purposes of this note.)

Most protocols implement either a fixed fee per transaction, or a dynamic fee that fluctuates with the demand for block space (*e.g.*, the base fee on Ethereum). The majority of networks commonly used today implement an unconstrained inclusion rule: validators can include and order transactions in a block however they wish. However, many have proposed a first-in-first-out (FIFO) inclusion rule, where transactions are included in the order in which they arrive in the mempool, until the block is full. These methods (for example, those proposed in [KDK22, KDL$^+$21, ZSC$^+$20, CMSZ22, Kur20]) add extra rules to consensus to ensure, roughly speaking, that if a majority of validators receive transaction A before transaction B, then transaction A is guaranteed to be included in the block, if transaction B is included. (We assume this is achievable and ignore implementation difficulties of FIFO inclusion in this note.)

Externalities. The simplest case to consider is where the transaction fee is fixed and the transaction inclusion rule is (effectively) FIFO. The fixed (usually very low) transaction fee, coupled with FIFO inclusion, encourages users competing for specific transactions (*e.g.*, arbitrage trades or liquidations) to spam the network in hopes that their transaction is seen first and therefore included in the block. This spam may result in block space filled with reverting, 'garbage' transactions with low utility. A recent analysis of Solana transactions suggested that over 50% of transactions were failed arbitrage trades [Fou23]. Clearly, these users are not paying for the externality they are causing to the network, impacting other users of the network.

Ok, but let's play fair. A natural solution to spam is to charge a dynamic fee for block space; as the demand increases, the fee increases (see EIP-1559 [BCD$^+$19]). In this note we will see that, even in this case, FIFO inclusion still can dramatically reduce the social welfare generated by the network. As a simple example, liquidations (which are high-utility transactions) would not be prioritized, again encouraging searchers to spam the network to ensure their transaction is included rather than simply paying a higher fee. This type of challenge has been playing out in the Arbitrum ecosystem. Although there is a base fee, FIFO inclusion still encourages spamming the network for transaction prioritization [sno23]. This is the case we consider for the remainder of the note.

1 Building blocks

In this section, we briefly introduce the block building problem, which is a special case of the 0-1 knapsack problem [KPP04]. We then show a simple lower bound to the optimal value of this problem via a greedy heuristic, which can be computed efficiently. We use this lower bound in a later section to bound the gap between the 'FIFO' inclusion and the best possible transactions to include (to maximize social welfare, *i.e.*, net utility).

Transactions and resources. Following [DECA22], we consider a mempool with transactions $j = 1, \dots, n$. Each transaction j has utility $\tilde{q}_j \in \mathbf{R}_+$ and consumes some amount of a single resource, called *gas*, denoted by $a_j \in \mathbf{R}_+$. Gas has a per-unit cost $g \in \mathbf{R}_+$ in the same numéraire as the utility. We denote the gas limit per block by $b \in \mathbf{R}_+$.

The block building problem. Using the above definitions, we then define the *block building problem* as

$$\begin{array}{ll} \text{maximize} & q^T x \\ \text{subject to} & a^T x \leq b \\ & x \in \{0,1\}^n, \end{array} \tag{1}$$

where $q = \tilde{q} - ga$ is the vector of net transaction utilities, after removing gas fees. Here, the variable $x \in \mathbf{R}^n$ indicates which transactions are included in the next block: $x_j = 1$ if transaction j is included in the block and $x_j = 0$ otherwise. In contrast to [DECA22], which considers arbitrarily complicated constraints, we only consider the constraint that the gas total of all included transactions must be lower than the upper limit. We assume that q is nonnegative (*i.e.*, users do not submit negative utility transactions) and, for the remainder of the note, denote the optimal value of problem (1) by $p^\star$.

Discussion. Problem (1) is an instance of the weighted 0-1 knapsack problem [KPP04]. Additionally, there are a number of important issues about how, exactly, to instantiate this problem and solve it in practice. The first is: how does one elicit the utilities q from users? This vector q can come from certain mechanisms (such as second-price auctions [Rou16]), though implementing these on chain remains an open research problem (see [CFK23, Fla22, CS21] and references therein). We do not deal with this question in this note and assume that there is some mechanism for receiving (or approximating) these utilities. There are other important issues to raise, such as the fact that this problem is NP-hard to solve [KPP04, App. A], but we note that this is a worst-case guarantee and that problems of this form are commonly solved to optimality in practice (using software such as Gurobi [Gur23]). However, we will use a relatively simple bound for the optimal value of this problem, which only depends on a heuristic solution that can always be efficiently computed. All results apply to this heuristic and therefore to any heuristics that do at least as well.

FIFO inclusion. In FIFO inclusion, on the other hand, transactions are included 'as they arrive', until the block reaches capacity. We assume that the transactions $j = 1, \dots, n$ are ordered by arrival time, from earliest to latest. (Whether FIFO inclusion is exactly possible in a decentralized setting has been a subject of great debate [KZGJ20, Theorem 1.2], but we treat is as an 'ideal' scenario that many protocols are trying to achieve.) In this scenario, the vector x must have the form $x = (\mathbf{1}_k, 0)$ for some k. We can define the utility from the FIFO transaction explicitly as the highest utility over over all vectors of this form:

$$p^{\text{FIFO}} = \max\{q^T(\mathbf{1}_k, 0) \mid a^T(\mathbf{1}_k, 0) \leq b,\ k = 0, 1, \dots, n\}.$$

A natural question is 'what is difference between the total utility of an optimal block packing (*i.e.*, a solution to problem (1)) and that of a FIFO packing?'. In other words, we want to find a bound on the gap

$$\Gamma = p^\star - p^{\text{FIFO}}.$$

Note, of course, that $\Gamma \geq 0$, since any FIFO solution is feasible for the original problem (1), though, as we will later show, the quantity Γ will be nonzero (and potentially large) 'in expectation' under certain weak assumptions about the distribution of the transactions.

1.1 A heuristic for block building

Problem (1) is a weighted knapsack problem, which has been well studied (see [KPP04] and references therein) and can be approximately solved by a number of simple heuristics with relatively tight guarantees. We present a very basic overview of one important heuristic and its corresponding proof of tightness here. We will then use this heuristic, and corresponding bound, to approximate $p^\star$ and show that Γ is large in a number of common scenarios.

Greedy heuristic. The simplest (and most common) heuristic is to relax the integrality constraint ($x \in \{0,1\}^n$) in problem (1) to an interval constraint, to get the linear programming problem:

$$\begin{array}{ll} \text{maximize} & q^T x \\ \text{subject to} & a^T x \leq b \\ & 0 \leq x \leq \mathbf{1}, \end{array} \tag{2}$$

with variable $x \in \mathbf{R}^n$ and the same problem data as (1). We write the optimal value of this relaxed problem as $r^\star$. Note that, since every x that is feasible for (1) is feasible for its relaxation (2), we have that

$$r^\star \geq p^\star \geq p^{\mathrm{FIFO}}.$$

Solution. Problem (2) is a linear programming problem that is easy to (computationally) solve in practice. In fact, it is possible to write a closed-form solution $x^\star$ to (2). To see this, start with the equivalent problem,

$$\begin{array}{ll} \text{maximize} & \bar{q}^T y \\ \text{subject to} & \mathbf{1}^T y \leq 1 \\ & 0 \leq y_i \leq b/a_i \quad i = 1, \dots, n. \end{array}$$

Here, the problem data are the *efficiencies* $\bar{q}_i = bq_i/a_i$ for each transaction $i = 1, \dots, n$, while the rest of the definitions are identical to the original relaxation (2). In this problem, we have simply done a variable substitution $y_i = a_i x_i / b$ in the relaxed problem (2). We can think of the $\bar{q}_i$ as a measure of the 'welfare-per-unit-resource' for transaction i.

A solution to this problem is then very simple to construct: sort the $\bar{q}_i$ in nondecreasing order (say, with indices $\tau_1, \dots, \tau_n$) then set $y^\star_{\tau_i} = b/a_{\tau_i}$, in order, for each i. If there is no index at which the bound is met, an optimal solution is to set $y^\star_{\tau_i} = b/a_{\tau_i}$ for $i = 1, \dots, n$. (It is *the* optimal solution if $\bar{q} > 0$.) If there is an index at which the bound is met or surpassed, say, index τ_k, then we choose $y^\star_{\tau_k}$ to be the largest possible value $\leq b/a_{\tau_k}$ such that the bound is met at equality. (There may be many solutions if there are many entries of $\bar{q}$ with value $\bar{q}_{\tau_k}$, in which case any entry suffices.) We can then recover a solution $x^\star$ for problem (2) from this optimal $y^\star$ by using the substitution provided above, $x_i^\star = by_i^\star/a_i$ for $i = 1, \dots, n$. Note that this solution, $x^\star$, is fractional (*i.e.*, has $0 < x_i^\star < 1$ for some i) in at most one entry.

Discussion. The solution above suggests a reasonable heuristic: since at most one entry of an optimal solution $x^\star$ to (2) (provided above) is fractional, we can simply round the (at most one) fractional entry down to zero, to receive, say, $x^0 \in \{0,1\}^n$. Note that x^0 is feasible for the original problem (1), and we denote its objective value as $p^0 = q^T x^0$. By definition, we have that $p^0 \le p^\star$, since x^0 is feasible. Surprisingly, we also have that $p^0 \ge Cp^\star$ for some constant $C \le 1$. In other words, this heuristic gives a feasible point with objective value p^0, which is actually 'close' to the true optimal value $p^\star$.

Bound. If the transactions all consume gas at most $a_i \le b/m$ for some integer $m > 1$, then we have that

$$\frac{m}{m-1} p^0 \ge p^\star.$$

To see this, note that the solution to the relaxation (2) described above satisfies $a^T x^\star = b$. By definition, there is at most one nonintegral entry, with gas $a_i \le b/m$ so the rounded solution must satisfy $a^T x^0 \ge (m-1)b/m$. Let i be the entry with largest q_i such that $x_i^0 = 0$ (*i.e.*, let i be the highest-utility transaction not included in the heuristic solution), then

$$\frac{q_i}{a_i} \le \frac{q^T x^0}{a^T x^0} = \frac{p^0}{a^T x^0} \le \frac{m}{(m-1)b} p^0.$$

The first inequality follows from the fact that, for positive t, u, v, w we have

$$\min\left\{\frac{t}{v}, \frac{u}{w}\right\} \le \frac{t+u}{v+w},$$

and q_i/a_i is no larger than the entries j with $x_j^0 = 1$, by construction of x^0. The equality follows by definition, and the last inequality comes from the previous discussion. Finally, using the fact that $a_i \le b/m$ again and the above, we have that

$$\frac{m}{b} q_i \le \frac{q_i}{a_i} \le \frac{m}{(m-1)b} p^0,$$

or, simplifying,

$$q_i \le \frac{p^0}{m-1}. \tag{3}$$

Since there is at most one transaction partially included by the relaxation, and this transaction has utility no larger than q_i (by definition of i), then

$$p^0 + q_i \ge r^\star \ge p^\star \ge p^0,$$

where, from before, $r^\star$ is the optimal value of the relaxation (2). Combining this with (3), we then have that

$$p^0 \le p^\star \le \frac{m}{m-1} p^0.$$

In other words, the heuristic solution, with optimal value p^0, is very close to the optimal value of exactly solving problem (1), which is NP-hard, whenever m is somewhat large. For more information on similar approximations, see [KPP04, §6].

2 What's the gap?

Given the discussion above, it makes sense to consider a model where transactions are drawn from some distribution, which we will characterize in terms of the efficiencies. We will also assume that there are minimum and maximum gas limits for any transaction, denoted by B^- and B^+ respectively, so that $B^- \le a_i \le B^+$ for $i = 1, \dots, n$. Clearly, the minimum and maximum transaction sizes imply a maximum and minimum number of transactions we can include in each block, given by b/B^- and b/B^+, respectively. From this, we can construct a lower bound of the optimal block utility and an upper bound of the FIFO-inclusion block utility. In this section, we will construct these bounds and discuss when there is a strictly positive gap between the utilities.

Lower bound. First, we will lower bound the utility of a block that was packed using the greedy heuristic (which we know is close to optimal). Denote the number of transactions included in the block using the greedy heuristic by $\bar{k}$. We define the average utility of these transactions by

$$q^+ = \frac{1}{\bar{k}} \sum_{i=1}^{\bar{k}} q_{\tau_i}.$$

Since $\bar{k} \ge b/B^+$, a lower bound for the utility of the greedily packed block is

$$L = \bar{k} q^+ \ge \frac{b}{B^+} q^+.$$

Upper bound. Now we will upper bound the expected utility of a FIFO-inclusion block, assuming the arrival time of the transactions is random; *i.e.*, we get a uniformly random permutation as the order of the arrival of the transactions. We know that at most b/B^- transactions can be included in a block and that the n transactions are uniformly randomly permuted, which means that the expected utility of FIFO, which we will call U, is no larger than

$$U \le \frac{b}{B^-} \frac{1}{n} \sum_{i=1}^{n} q_i.$$

If we define q^- as the average utility for the transactions not included by the greedy heuristic,

$$q^- = \frac{1}{n - \bar{k}} \sum_{i=\bar{k}+1}^{n} q_{\tau_i},$$

then we can write

$$\frac{1}{n}\sum_{i=1}^{n} q_i = \frac{\bar{k}}{n}q^+ + \left(1 - \frac{\bar{k}}{n}\right) q^-.$$

(We can view q^- as, roughly speaking, the average utility of the 'tail' of transactions, as the efficiencies get small.) This means the average FIFO utility is bounded from above by

$$U \le \frac{b}{B^-}\left(\frac{\bar{k}}{n}q^+ + \left(1 - \frac{\bar{k}}{n}\right) q^-\right).$$

What's the gap? We now seek to characterize the gap between these two bounds, which gives us a lower bound on the gap between FIFO inclusion and the optimal, since

$$\Gamma \ge L - U \ge \frac{b}{B^+}q^+ - \frac{b}{B^-}\left((q^+ - q^-)\frac{\bar{k}}{n} + q^-\right).$$

We now look for a basic condition for when the right-hand-side of this inequality is positive, which would imply that the gap $\Gamma > 0$. Rearranging, it is easy to see that the right hand side is positive whenever

$$q^+\left(1 - \frac{\bar{k}\eta}{n}\right) > \eta q^-\left(1 - \frac{\bar{k}}{n}\right), \tag{4}$$

where $\eta = B^+/B^-$ is the ratio between the largest and smallest possible transaction. If

$$q^+ > \eta q^-, \tag{5}$$

then, as the number of outstanding transactions n becomes large relative to the number of greedily-chosen transactions, $\bar{k}$, while the average utilities stay roughly constant, we get that $\Gamma > 0$. We may also wish to consider the ratio of optimal to FIFO block utility. From the preceding discussion, we have that

$$\frac{p^\star}{p^{\text{FIFO}}} \ge \frac{L}{U} \ge \frac{q^+/\eta}{(q^+ - q^-)\frac{\bar{k}}{n} + q^-},$$

which is strictly greater than one under the positive gap condition (4).

Discussion. We note that the bounds derived in (4) and (5) are actually fairly loose in practice (as we will see next). The reason such bounds are loose is primarily due to the worst-case assumptions made in the lower bound, where we only use the fact that $\bar{k} \ge b/B^+$ and the upper bound, where we use the fact that b/B^- bounds the maximum possible number of transactions that can be included by FIFO. Both of these cases can be quite loose if there are only a few transactions close to the bounds, relative to the rest of the distribution. As this is just a note, we only seek 'qualitatively reasonable' behavior from the bounds in order to gain intuition, but not tight constants. An interesting avenue for future research would be to either tighten the bounds given here or generalize them to the multidimensional setting considered in [DECA22].

2.1 Special cases and experiments

From inequalities (4) and (5), we can deduce some conditions under which we are guaranteed to have a large gap between the utilities of the optimal and FIFO-inclusion blocks. We also show some basic numerical experiments which suggest that, in practice, the gap may be far larger than the one suggested by our bounds.

Sharp distribution. First, if the distribution over the efficiencies is sufficiently sharp (*e.g.*, if are a small number of transactions with very large utility and equal or lower gas relative to the others) then $q^+ \gg q^-$ for reasonable block sizes. This situation is common in practice when many similar transactions are submitted but only one can be executed profitably, as is the case in many MEV opportunities, including liquidations and DEX arbitrage. (As discussed in the introduction, this distribution is not rare, even in practice.) Since only one of these transactions can be executed profitably, the others will revert and, therefore, have non-positive utility. Note that this situation corresponds to these efficiencies being drawn from a heavy-tailed probability distribution.

Equal size transactions. Another important case is if the transactions have roughly the same size, *i.e.*, $\eta \approx 1$. From bound (4), we see that any distribution of utilities that is not flat will create a gap between the optimal and FIFO-inclusion blocks. Intuitively, this is straightforward to see: there is always a chance that high-utility transactions will not be included in the block, and, since all transactions consume roughly the same gas, this is a strict loss in total utility.

Simple experiments. We plot our bound on the utility gap and the realized gap for several transaction utility distributions, shown in figure 1 (we provide definitions in appendix A). All transaction sizes are drawn uniformly at random from the interval $[1, 3]$. We generate 1000 transactions in the mempool and vary the block size from 20 to 2000 gas (for an average of 10 to all 1000 transactions per block). We run 100 trials for each block size. All experiments are done in the Julia programming language [BEKS17] and we use the `Distributions.jl` package [BPA+21]. Code is available at

`https://github.com/bcc-research/fifo-note`

The empirical gaps for flat distributions (*i.e.*, those where we do not expect transactions with very large utility, relative to the other transactions) are shown in figure 2. We see that when the block size is small relative to the number of transactions, there is a significant gap between the utilities of the optimal and FIFO-inclusion blocks. Furthermore, this gap is quite a bit worse than what is predicted by our bound (4). (This difference is not surprising, given the discussion above.) For sharp distributions, the gap, shown in figure 3, is quite a bit larger. These distributions more closely resemble competitive MEV opportunities.

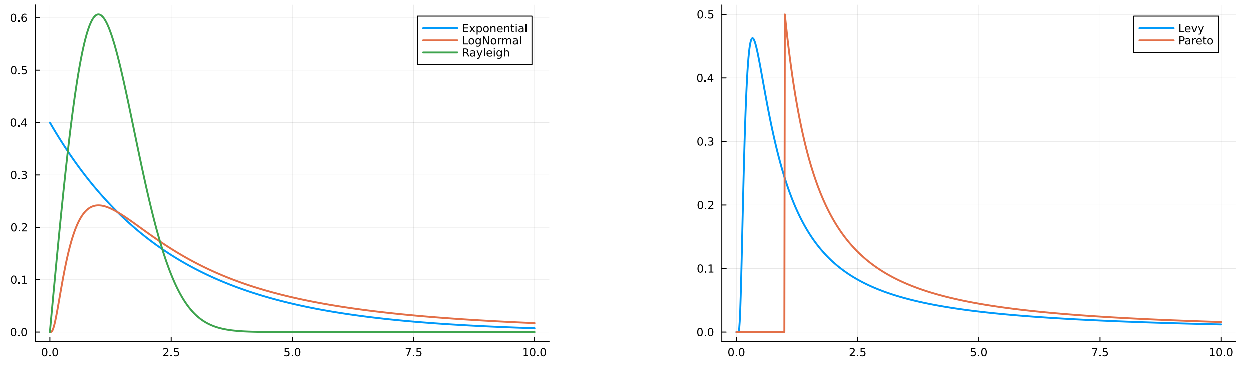

(a) Light-tailed utility distributions.

(b) Heavy-tailed utility distributions.

Figure 1: Distributions used for experiments. Light-tailed utility distributions lead to ‘flat’ efficiency distributions, while heavy-tailed ones lead to ‘sharp’ efficiency distributions.

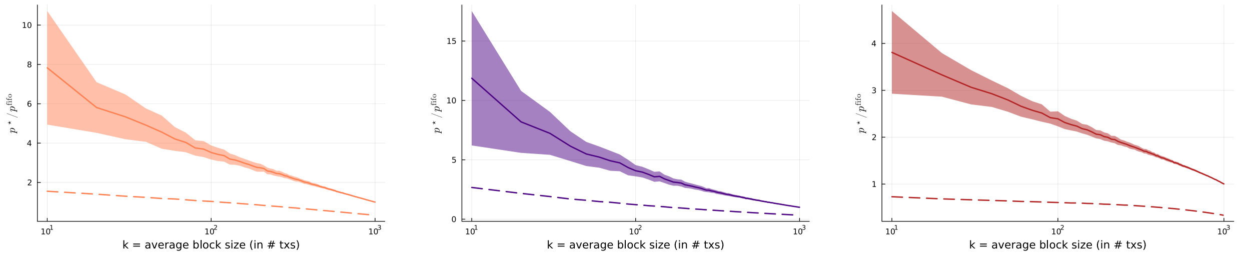

(a) $q_i \sim \text{Exponential}(2.5)$

(b) $q_i \sim \text{LogNormal}(1, 1)$

(c) $q_i \sim \text{Rayleigh}(1)$

Figure 2: Gap between FIFO and optimal block packing, as a function of block size, for flat (light-tailed) utility distributions. The solid line indicates the mean over 100 trials, the ribbon indicates the ±σ range, and the dashed line indicates the bound (4).

References

[BCD+19] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. Eip-1559: Fee market change for eth 1.0 chain, 2019.

[BEKS17] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

[BPA+21] Mathieu Besançon, Theodore Papamarkou, David Anthoff, Alex Arslan, Simon Byrne, Dahua Lin, and John Pearson. Distributions.jl: Definition and modeling of probability distributions in the juliastats ecosystem. *Journal of Statistical Software*, 98(16):1–30, 2021.

[CFK23] Tarun Chitra, Matheus VX Ferreira, and Kshitij Kulkarni. Credible, optimal auctions via blockchains. *arXiv preprint arXiv:2301.12532*, 2023.

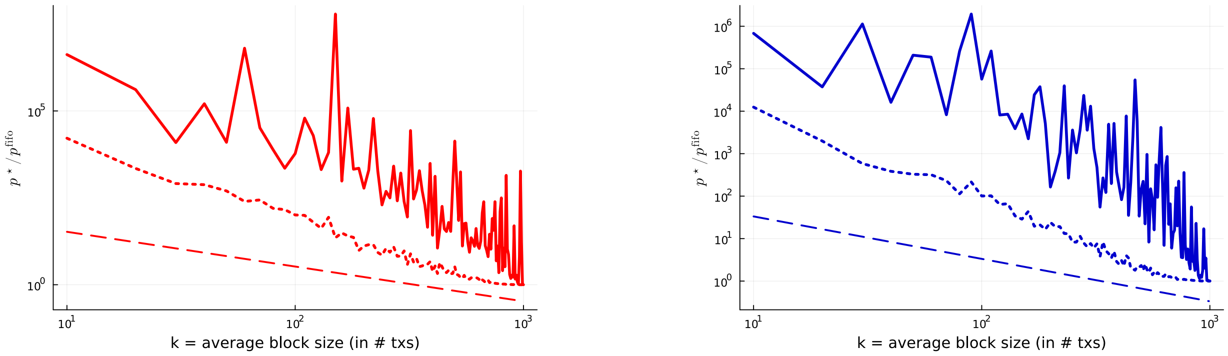

(a) $q_i \sim \text{Levy}(0, 1)$

(b) $q_i \sim \text{Pareto}(0.5)$

Figure 3: Gap between FIFO and optimal block packing, as a function of block size, for sharp (heavy-tailed) distributions. The solid line indicates the mean over 100 trials, the dotted line indicates the median, and the dashed line indicates the bound (4). Since these distributions have infinite variance (and mean), $\pm\sigma$ ranges are not displayed.

[CMSZ22] Christian Cachin, Jovana Mićić, Nathalie Steinhauer, and Luca Zanolini. Quick order fairness. In *International Conference on Financial Cryptography and Data Security*, pages 316–333. Springer, 2022.

[CS21] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. *arXiv preprint arXiv:2111.03151*, 2021.

[DECA22] Theo Diamandis, Alex Evans, Tarun Chitra, and Guillermo Angeris. Dynamic pricing for non-fungible resources. *arXiv preprint arXiv:2208.07919*, 2022.

[Fla22] Flashbots. The future of mev is suave: Flashbots, Nov 2022.

[Fou23] Jito Foundation. Solving the mev problem on solana: A guide for stakers, Feb 2023.

[Gur23] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.

[KDK22] Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. Order-fair consensus in the permissionless setting. In *Proceedings of the 9th ACM on ASIA Public-Key Cryptography Workshop*, pages 3–14, 2022.

[KDL+21] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *Cryptology ePrint Archive*, 2021.

[KPP04] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack problems*. Springer, 2004.

[Kur20] Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 25–36, 2020.

[KZGJ20] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In *Advances in Cryptology–CRYPTO 2020: 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17–21, 2020, Proceedings, Part III 40*, pages 451–480. Springer, 2020.

[Rou16] Tim Roughgarden. *Twenty lectures on algorithmic game theory.* Cambridge University Press, 2016.

[sno23] snoopy_mev. `https://twitter.com/snoopy_mev/status/1629283898453811200`, Feb 2023.

[ZSC+20] Yunhao Zhang, Srinath Setty, Qi Chen, Lidong Zhou, and Lorenzo Alvisi. Byzantine ordered consensus without byzantine oligarchy. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 633–649, 2020.

A Distribution definitions

Flat distributions. We consider three distributions which yield relatively flat efficiencies as they have a light (*i.e.*, sub-exponential or sub-Gaussian) tail. These distributions have a finite mean and variance, so we expect samples to cluster tightly together. The exponential distribution with parameter θ has probability density function

$$f(x;\theta) = \frac{1}{\theta}e^{-\frac{x}{\theta}}, \quad x > 0.$$

The log normal distribution with parameters μ and σ has probability density function

$$f(x;\mu,\sigma) = \frac{1}{x\sqrt{2\pi\sigma^2}}\exp\left(-\frac{(\log(x)-\mu)^2}{2\sigma^2}\right), \quad x > 0.$$

The Rayleigh distribution with parameter σ has probability density function

$$f(x;\sigma) = \frac{x}{\sigma^2}e^{-\frac{x^2}{2\sigma^2}}, \quad x > 0.$$

Sharp distributions. We consider two utility distributions which have heavy tails. These distributions have an infinite mean and variance, and we expect the presence of large outliers when sampling, yielding a relatively sharp distribution of the efficiencies. The Levy distribution with parameters μ and σ has probability density function

$$f(x;\mu,\sigma) = \sqrt{\frac{\sigma}{2\pi(x-\mu)^3}}\exp\left(-\frac{\sigma}{2(x-\mu)}\right), \quad x > \mu.$$

The Pareto distribution with parameter α has probability density function

$$f(x;\alpha) = \frac{\alpha}{x^{\alpha+1}}, \quad x \geq 1.$$