Robust Resource Allocation via Competitive Subsidies

David X. Lin
Cornell University, Ithaca, NY, USA

Giannis Fikioris
Cornell University, Ithaca, NY, USA

Siddhartha Banerjee
Cornell University, Ithaca, NY, USA

Éva Tardos
Cornell University, Ithaca, NY, USA

Abstract

A canonical setting for non-monetary online resource allocation is one where agents compete over multiple rounds for a single item per round, with i.i.d. valuations and additive utilities across rounds. With n symmetric agents, a natural benchmark for each agent is the utility realized by her favorite $1/n$-fraction of rounds; a line of work has demonstrated one can *robustly* guarantee each agent a constant fraction of this ideal utility, irrespective of how other agents behave. In particular, several mechanisms have been shown to be 1/2-robust, and recent work established that repeated first-price auctions based on artificial credits have a robustness factor of 0.59, which cannot be improved beyond 0.6 using first-price and simple strategies. In contrast, even without strategic considerations, the best achievable factor is $1 - 1/e \approx 0.63$.

In this work, we break the 0.6 first-price barrier to get a new 0.625-robust mechanism, which almost closes the gap to the non-strategic robustness bound. Surprisingly, we do so via a simple auction, where in each round, bidders decide if they ask for the item, and we allocate uniformly at random among those who ask. The main new ingredient is the idea of *competitive subsidies*, wherein we charge the winning agent an amount in artificial credits that decreases when fewer agents are bidding (specifically, when k agents bid, then the winner pays proportional to $k/(k+1)$, varying the payment by a factor of 2 depending on the competition). Moreover, we show how it can be modified to get an equilibrium strategy with a slightly weaker robust guarantee of $5/(3e) \approx 0.61$ (and the optimal $1 - 1/e$ factor at equilibrium). Finally, we show that our mechanism gives the best possible bound under a wide class of auction-based mechanisms.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases Online Resource Allocation, Non-Monetary Mechanisms

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.96

Related Version *Full Version*: `https://arxiv.org/abs/2511.09934` [12]

Supplementary Material *Software (Source Code)*: `https://github.com/davidxlin/optimization-problem-for-robustness`, archived at `swh:1:dir:6152717a4eb16c8b4fddc4a049a1fbd060409386`

Funding *David X. Lin*: Supported by AFOSR grant FA9550-231-0068, NSF grant ECCS-1847393 and through a SPROUT Award from Cornell Engineering.
Giannis Fikioris: Supported in part by the Google PhD Fellowship, the Onassis Foundation – Scholarship ID: F ZS 068-1/2022-2023, and ONR MURI grant N000142412742.
Siddhartha Banerjee: Supported by AFOSR grant FA9550-231-0068, NSF grant ECCS-1847393 and through a SPROUT Award from Cornell Engineering.
Éva Tardos: Supported in part by AFOSR grant FA9550-23-1-0410, AFOSR grant FA9550-231-0068, ONR MURI grant N000142412742, and through a SPROUT Award from Cornell Engineering.

1 Introduction

Consider an indivisible resource shared between multiple selfish agents over time, e.g., a telescope shared by different research labs. Over multiple rounds, a principal needs to decide which agent gets allocated, aiming to allocate to each agent when their value is the highest. The mechanism used by the principal should be: (i) non-monetary, as the resource is being shared, and no one agent owns it, (ii) fair, in that each agent is individually satisfied with the utility they get, (iii) simple and understandable, so that participating in the mechanism is straightforward, (iv) robust, i.e., each individual has utility guarantees that hold independent of the behavior of others. Under these properties, any agent following a simple strategy is able to guarantee high utility regardless of how the other agents behave. Standard techniques in mechanism design, like the VCG mechanism, cannot be applied due to the lack of money. In addition, such techniques often ignore fairness in allocation and instead focus on efficiency, i.e., maximizing the total allocated value, which in this setting is ill-defined: the lack of money makes agents' values incomparable.

We study this problem via a canonical model, first introduced by [9]: n agents compete over T rounds for a single indivisible resource in each round. Agent i has (random) private values $V_i[t]$ for the resource in round t; these values are i.i.d. across rounds and independent across agents. The principal can award the item in each round to a single agent, or not award at all. Early work on this model looked at welfare approximations, with symmetric agents or known value distributions. More recent work on this setting has focused on the design of *share-based* mechanisms, where each agent i is endowed with a "fair share" α_i (with $\sum_i \alpha_i = 1$) indicating the nominal fraction of items they should be allocated. This was first suggested by [8], who showed that under a repeated first-price auction with artificial credits, each agent can *robustly* guarantee at least $1/2$ of her ideal utility (i.e., her maximum utility under her nominal share of the resource – see Definition 1), under arbitrary behavior by other agents. Recently, using the same mechanism, [13] improved this to $2 - \sqrt{2} \approx 0.59$ fraction of ideal utility, using a more complicated strategy where the agent has to use a randomized bid; they also show that for the first-price auction, no static bidding strategy can be better than 0.6 robust. In contrast, a natural upper bound on the robustness factor is $1 - 1/e$, which follows from ignoring strategic considerations – consider n agents with equal fair shares $1/n$ and Bernoulli($1/n$) values, wherein each agent should ideally get T/n rounds, but there are at most $(1 - (1 - 1/n)^n)T$ rounds in which any agent has non-0 utility. Thus, it appears fundamentally new ideas are required to make progress towards this bound.

1.1 Our Results

Our main result presents a new mechanism where any agent can robustly guarantee at least a $5/8 = 0.625$ fraction of her ideal utility, thus almost closing the gap to the upper bound of $1 - 1/e \approx 0.63$. In addition, our mechanism and strategies turn out to be simpler than the ones used in prior work on first-price mechanisms [8, 13], where each agent needs to bid in a carefully chosen (and potentially randomized) way, so their spending is not too low (to be aggressive enough) and not too high (to conserve budget). Instead, our proposed Competitive Subsidy Mechanism, abstracts this complexity away from the agent. As in earlier works, each agent is first endowed with a budget of artificial credits proportional to their fair share. However, in each round, instead of allowing arbitrary non-negative bids, we require each agent to either request the item or not, and allocate uniformly at random among requesting agents. Finally, the mechanism charges the winning agent an amount that is based on the competition in that round, subsidizing when fewer agents request. Specifically, if k agents request, the winner is charged proportionally to $\frac{k}{k+1}$. This is intuitive, as winning when more other agents request causes more externalities, and hence higher payments. Our 0.625 ideal-utility guarantee is now achieved when an agent with fair share α requests whenever her value is in the top α-quantile.

A natural question is whether the payment scheme $\frac{k}{k+1}$ is the optimal one. In Section 4, assuming static bidding policies, we show how to bound the robustness factor of any payment scheme (i.e., any function p_k charged to the winner when k agents request) via an optimization problem. Our numerical results indicate that we cannot improve over our 0.625 result, strongly suggesting that our simple payment rule is in fact optimal over any comparable payment function in such a mechanism.

A notable benefit of the simplicity of our mechanism is that it is easy to modify to address metrics beyond robustness. To demonstrate this, we next consider a question raised in [14, 11] as to whether robust strategies can also be realized as equilibrium actions. Our proposed strategy of requesting when agent i's value is in her top α_i-quantile turns out not to be a best response if all other agents use this strategy (instead of acting adversarially, as we consider in Section 3). In fact, when all agents use this strategy, they do not utilize their entire budget. However, this is easy to fix by modifying our mechanism to have higher prices. Specifically, in Section 5, we show that when agents have equal fair shares, then increasing our payments by a factor of approximately $\frac{3}{8}e \approx 1.02$ makes the earlier strategy an equilibrium, albeit with a slightly worse robustness guarantee. In particular, every agent now enjoys a $(1 - 1/e) \approx 0.63$ fraction of her ideal utility at equilibrium and $\frac{5}{3e} \approx 0.61$ robustly. In Section 6, we modify this to realize an equilibrium strategy with the same robustness guarantee for arbitrary fair shares.

The innovations in our mechanism are twofold: First, our congestion-aware pricing is key to getting robustness better than 1/2. In fact, [11], whose pricing is invariant of the number of bidders, show that one cannot do better than 1/2 with static bidding under this rule. Secondly, simplifying the action space of the agents is key to our improved guarantees. Allowing agents to bid any real number in a first-price auction gives an adversary too much freedom and leads to a deterioration of robustness results, as shown by the robustness upper bound of 0.6 in [13]. Even for our equilibrium guarantees, ensuring that a strategy profile results in an equilibrium is very complicated when the action space is too large.

1.2 Related work

There is a long line of work on online resource allocation without money, building on the model of [9]. Earlier work focused on emulating outcomes of monetary mechanisms without using money [9, 5, 7]. These culminated in the black-box reduction of [7], which, building on the "linking decisions" idea of [10], showed how repeated all-pay auctions can emulate the equilibrium outcome of any monetary mechanism with vanishing efficiency loss. However, these approaches assume full knowledge of value distributions and provide no guarantees under off-equilibrium behavior.

Our work builds on a more recent line [8, 4, 6, 14, 13, 11], that considers the same model, but shifts the focus to distribution agnostic mechanisms, both achieving robust, individual-level guarantees, as well as analyzing equilibrium outcomes. [8] first showed that in a first-price auction with artificial currency, each agent has a 1/2-robust strategy when using an appropriate fixed bid. [4] extend their results to reusable resources (where an agent might require the resource for multiple consecutive rounds) using the same mechanism with a reserve price; they also show that in this setting, the 1/2 factor is tight. [6] use the Dynamic Max-min Fair (DMMF) mechanism (that allocates to the bidding agents with the least (normalized) number of wins) to get 1/2-robustness in the worst-case; they also get $(1 - o(1))$-robustness under assumptions on the agent's value distribution. [13] improve the robustness of the first-price mechanism to $2 - \sqrt{2} \approx 0.59$ by using randomization, where the agent's bid is uniformly distributed. In fact, our payment scheme and the uniform distribution used by [13] share a connection: if k bidders bid using a uniform $[0, 1]$ bid, then the expected payment of the winner is $\frac{k}{k+1}$. However, [13] allow arbitrary bidding by the other agents, giving an adversary too much freedom: they prove with this freedom of the adversary and static bidding cannot do better than a 0.6 fraction of ideal utility.

The issue with the above works is that the robust strategies considered do not result in an equilibrium if used by all players. In fact, [14] proves that even in very simple scenarios, the suggested strategy in the DMMF mechanism does not result in an equilibrium, in fact, DMMF does not have equilibria with static player strategies. [11] design a mechanism that remedies this issue. They limit each agent i with fair share α_i to at most $\alpha_i T$ requests and, using a complicated randomized allocation scheme, prove that requesting with probability α_i results in 1/2-robustness and a $\big(1-\prod_i(1-\alpha_i)\big)$-good equilibrium. In comparison to our work, their robustness factor $\lambda_{\texttt{ROB}} = 1/2$ is lower but their equilibrium factor $\lambda_{\texttt{NASH}} = 1 - \prod_i(1 - \alpha_i)$ can be greater for asymmetric fair shares $(\alpha_i)_i$. They prove that $\lambda_{\texttt{NASH}} = 1 - \prod_j(1 - \alpha_j)$ is optimal in that no mechanism, even if the principal can see all the values upfront, can guarantee each agent a greater fraction of their ideal utility with the worst-case value distributions. In contrast, our mechanism's $\lambda_{\texttt{NASH}} \geq 1 - 1/e$ is the best factor that does not depend on the fair shares α_i, but also is always worse than the factor of [11], since $\inf_{(\alpha_i)_i}(1 - \prod_i(1 - \alpha_i)) = 1 - 1/e$.

The ideal utility benchmark falls within a wider class of *share-based* approaches to fair allocation problems. A notable parallel notion is that of the AnyPrice Share (APS) [2, 3], defined as the value an agent with a given budget can guarantee under any choice of *normalized* item prices (i.e., where the sum of budgets equals the sum of prices). Constant-factor approximations of the APS have been characterized for a variety of full-information one-shot allocation problems; in contrast, we focus on repeated allocation with stochastic private valuations and strategic bidding. The worst-case nature of the APS also makes it much weaker than the ideal utility; for example, an agent with value 1 for a $1/n$-fraction of items can get a $\Theta(1/n)$-fraction of these in the worst case as all other agents may desire the same items; in contrast, under i.i.d. Bernoulli$(1/n)$ values, she can get ≈ 0.63 fraction of these items.

2 Preliminaries

We consider the following canonical setting for repeated online allocation, introduced by [9]. There are T rounds, and in each round t, a single indivisible item is available for allocation among n agents. At the start of round t, each agent i realizes a private value $V_i[t]$ for the item, drawn from a fixed distribution $\mathcal{F}_i$, independently across both agents and time. The value $V_i[t]$ becomes known to agent i at the start of round t, but is unknown to the principal and other agents. We assume the value distribution $\mathcal{F}_i$ is nonnegative and bounded by some constant not depending on T (which we take to be 1 without loss of generality).

At the end of each round t, following some mechanism, the principal chooses an agent to allocate the item to, or to not allocate. Let $W_i[t]$ be an indicator of whether agent i wins the round-t item or not, resulting in utility $U_i[t] = V_i[t]W_i[t]$. Each agent seeks to maximize their average per-round utility, $\frac{1}{T}\sum_{t=1}^{T} U_i[t]$.

With symmetric agents (i.e., with identical distributions $\mathcal{F}_i$ and equal importance), a natural aim for the principal is to allocate to the agent with the highest value. To extend this to heterogeneous agents in the absence of money, we use the benchmark of *ideal utility* introduced by [8]. Roughly speaking, the ideal utility is the highest expected per-round utility an agent could obtain if they are restricted to winning the item only for a pre-specified fraction of the rounds. Formally, we assume each agent has an exogenously given fair share $\alpha_i > 0$, where $\sum_i \alpha_i = 1$, and define their ideal utility as:

▶ **Definition 1** (Ideal Utility). *The ideal utility v_i^* of agent i is the value of the following maximization problem over measurable $\rho : [0,\infty) \to [0,1]$.*

$$\max_{\rho} \mathop{\mathbb{E}}_{V_i \sim \mathcal{F}_i}[V_i \rho(V_i)] \quad \textit{subject to} \quad \mathop{\mathbb{E}}_{V_i \sim \mathcal{F}_i}[\rho(V_i)] \leq \alpha_i \tag{1}$$

An agent's fair share measures an exogenously defined importance of this agent; symmetric fair shares $\alpha_i = 1/n$ mean each agent is equally important. If $\mathcal{F}_i$ has an absolutely continuous CDF F_i, then v_i^* is just the portion of the expectation of $V_i \sim \mathcal{F}_i$ that comes from the top α_i-quantile of $\mathcal{F}_i$, i.e., $v_i^* = \mathbb{E}_{V_i \sim \mathcal{F}_i}[V_i \mathbb{1}_{V_i \geq F_i^{-1}(1-\alpha_i)}]$. This is thus a natural benchmark for what an agent can hope to receive. Moreover, with identical $\mathcal{F}_i$ and equal shares, summing ideal utilities gives the so-called *ex-ante welfare* [1], which is an upper bound for overall welfare widely used in approximate mechanism design.

Fix a mechanism used to allocate the item. As in prior work, [8, 4, 6, 13, 11], we are interested in *robust* strategies, strategies that guarantee a certain fraction of the agent's ideal utility, regardless of the other agents' strategies (even if the other agents adversarially collude).

▶ **Definition 2** (λ-robust). *Fix a mechanism and an agent i. A strategy π_i used by agent i in the mechanism is λ-robust if regardless of the strategies of agents $j \neq i$,*

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{E}[U_i[t]] \geq \lambda v_i^*.$$

Robust strategies are nice in that they guarantee utility for the agent without assumptions the behavior of other agents. In addition, if all agents have a λ-robust strategy, then at any equilibrium, they must obtain at least a λ-fraction of their ideal utility; if not, they can deviate to the λ-robust strategy to gain more utility. In particular, for identical $\mathcal{F}_i$ and equal shares $\alpha_i = 1/n$, if each agent has a λ-robust strategy, then the ratio of the optimal social walfare and the achieved welfare (price of anarchy) is at most $1/\lambda$.

Our primary focus in this work is on getting mechanisms with good robustness bounds. A secondary goal is to realize such robustness bounds under equilibrium strategies. Unfortunately, robust strategies are not guaranteed to form an equilibrium, as it is possible that agents may want to deviate to an even higher payoff strategy. Indeed, for natural mechanisms, known robust strategies do not admit any equilibrium [14]. In such cases, a λ-robust strategy may not help in understanding a mechanism's performance under real agent behavior. To this end, [11] offers a mechanism where $\frac{1}{2}$-robust strategies do form an equilibrium, which motivates the following definition.

▶ **Definition 3** ($\lambda_{\texttt{ROB}}$-robust $\lambda_{\texttt{NASH}}$-good approximate-equilibrium). *A profile of strategies $(\pi_1, \ldots, \pi_n)$ is a $\lambda_{\texttt{ROB}}$-robust $\lambda_{\texttt{NASH}}$-good approximate-equilibrium if the following hold.*

1. *Each strategy π_i is $\lambda_{\texttt{ROB}}$-robust.*

2. *For some $\epsilon(T) = o(1)$, the profile of strategies $(\pi_1, \dots, \pi_n)$ is an $\epsilon(T)$-equilibrium: no agent can deviate from the strategy profile and gain more than an additive $\epsilon(T)$ in expected per-round utility.*[1]
3. *In the strategy profile $(\pi_1, \dots, \pi_n)$, each agent obtains at least a $\lambda_{\texttt{NASH}}$-fraction of their ideal utility in expectation.*

By definition, a $\lambda_{\texttt{ROB}}$-robust $\lambda_{\texttt{NASH}}$-good approximate-equilibrium has $\lambda_{\texttt{NASH}} \geq \lambda_{\texttt{ROB}} - o(1)$, but $\lambda_{\texttt{NASH}}$ could potentially be substantially higher. In Sections 5 and 6 we improve the result of [11] by offering a mechanism where 0.61-robust strategies form a $(1 - 1/e)$-good equilibrium.

3 0.625-robustness

In our mechanism, Competitive Subsidy Mechanism, each agent i is endowed with $\alpha_i T$ tokens of artificial credits. At each time t, each agent must either *bid* or not. As we mention in the introduction, having a binary action space limits the possible adversarial behavior of other agents, which is crucial for the improved guarantee over [13]. The item is allocated uniformly at random among bidding agents. The winning agent must pay $\bar{b}k/(1+k)$ artificial tokens where k is the number of bidding agents, and $\bar{b}$ is a scale factor, chosen later. As we show in Section 4, this payment scheme is optimal, in that we cannot get better robustness guarantees. When an agent's budget of artificial tokens becomes non-positive, she is barred from bidding in future rounds. We formally specify our mechanism in Algorithm 1.

Algorithm 1 Competitive Subsidy Mechanism.

Input: Fair shares $(\alpha_i)_{i\in[n]}$, number of rounds T, payment constant $\bar{b}$
Endow each agent with a budget $B_i[1] = \alpha_i T$ of bidding credits;
for $t = 1, 2, \dots, T$ **do**
  Agents either request to bid or not (let $r_i[t]$ be the indicator that agent i requests to bid);
  Enforce budgets: $r_i[t] \leftarrow 0$ for each i such that $B_i[t] \leq 0$;
  Define $S[t] = \{i : r_i[t] = 1\}$ to be the set of bidding agents;
  Select a winner uniformly at random from $S[t]$ (let $W_i[t]$ be the indicator agent i wins);
  Set payments $P_i[t] = \bar{b} \cdot \frac{|S[t]|}{1+|S[t]|} \cdot W_i[t]$ (note only the winner pays);
  Update budgets: $B_i[t+1] \leftarrow B_i[t] - P_i[t]$
end

Next, we describe our proposed robust strategy. We adopt the notion of an α-aggressive strategy from [6], whereby an agent bids only for the values that realize her ideal utility.

▶ **Definition 4** (α-aggressive strategy). *Agent i follows a α_i-aggressive strategy if she bids whenever her budget is positive and her value $V_i[t]$ is in the top α_i-quantile of her value distribution.*[2]

[1] We allow this additive $o(1)$ to account for stochastic deviation.
[2] If the CDF $\mathcal{F}_i$ has jumps and the top α_i-quantile is not uniquely-defined, then the agent randomizes appropriately at the cutoff to bid with probability exactly α_i. Formally, the agent bids at time t with probability $\rho(V_i[t])$ where ρ solves Eqn. (1).

Our main result of this section is as follows, which says that an α_i-aggressive strategy is $(0.625 - o(1))$-robust, i.e., each agent i guarantees that fraction of ideal utility regardless of other agents' behavior.

▶ **Theorem 5.** *When running Competitive Subsidy Mechanism with $\bar{b} = 8/3$, an α_i-aggressive strategy is λ_i-robust for some $\lambda_i \geq \frac{5}{8} - O\left(\sqrt{\frac{\log T}{T}}\right)$.*

We prove the theorem formally in the full version of the paper and give a proof sketch here.

Proof Sketch. When using an α_i-aggressive strategy, the expected value of $V_i[t]$ when conditioned on requesting is the expected value of $V_i[t]$ conditioned on $V_i[t]$ being in the top α_i-quantile. This is exactly v_i^*/α_i by the definition of the ideal utility v_i^*. Therefore, to show λ_i-robustness, we must show that agent i wins at least an $\approx \lambda_i\alpha_i$ fraction of the rounds.

For this sketch, we will only work with expectations, the full proof, that uses standard probability concentration bounds to convert these to high probability statements, is given in the full version of the paper. At a high level, we need to consider two cases: either agent i does not use up all her budget (i.e., $\sum_{t=1}^{T} P_i[t] < \alpha_i T$), or she runs out of budget (i.e., $\sum_{t=1}^{T} P_i[t] \geq \alpha_i T$). In the first case, we argue that she gets utility $\frac{1}{T}\sum_{t=1}^{T} W_i[t] \geq \alpha_i(1-1/\bar{b})$; on the other hand, if she does use up all of her budget, we argue that she gets utility $\frac{1}{T}\sum_{t=1}^{T} W_i[t] \geq 5\alpha_i/(3\bar{b})$. By setting $\bar{b} = 8/3$, we get that the fraction of rounds won is at least $5\alpha_i/8$ in either case.

First, let us consider the case that agent i does not run out of budget. At each time t, some number k of agents $j \neq i$ bid. Let x_k be the fraction of times t that there are k agents $j \neq i$ bidding,

$$x_k = \frac{1}{T} \cdot \#\{t : \#\{j \neq i : r_j[t] = 1\} = k\}.$$

By the independence of agent i's bid $r_i[t]$ from the others' bids $r_j[t]$ for $j \neq i$, if we restrict to only times t in which agent i bids, x_k is also the fraction of these times t that both agent i bids and there are k agents $j \neq i$ bidding at time t. Then focusing on expectations[3] we have,

$$\frac{1}{T}\sum_{t=1}^{T} W_i[t] \approx \alpha_i \sum_{k=0}^{n-1} \frac{x_k}{1+k}, \tag{2}$$

since at each time t, agent i bids with probability α_i, and if k others bid, there will be $1+k$ total bidders, and agent i will win with probability $1/(1+k)$ since the mechanism uniformly at random allocates among bidding agents. Agents' $j \neq i$ budget constraint imply (again focusing on expectations and ignoring $o(1)$ terms)

$$\sum_{k=1}^{n-1} \left((1-\alpha_i) \cdot \frac{\bar{b}k}{1+k} + \alpha_i \cdot \frac{\bar{b}k}{2+k} \right) x_k \lessapprox 1 - \alpha_i. \tag{3}$$

This is because at any time, with probability $1-\alpha_i$, if k other agents $j \neq i$ bid, they pay in total $\frac{\bar{b}k}{1+k}$ by the allocation rule, and with probability α_i, there are $1+k$ total bidding agents, so some agent $j \neq i$ wins with probability $\frac{k}{1+k}$ in which case they pay $\bar{b} \cdot \frac{1+k}{2+k}$. Agents' $j \neq i$

[3] Formally, we write $f(T) \approx g(T)$ if $|f(T) - g(T)| \leq o(1)$ with probability at least $1 - o(1)$. Similarly, $f(T) \lessapprox g(T)$ if $f(T) \leq g(T) + o(1)$ with probability at least $1 - o(1)$.

total budget is $\sum_{j \neq i} \alpha_j T = (1 - \alpha_i)T$, so their total expected per-round payment has to be at most $1 - \alpha_i$. Using (2) and (3), it suffices to lower bound the value of the following linear program.

$$\begin{aligned} \min_{(x_k)_{k=0}^{n-1}} \quad & \alpha_i \sum_{k=0}^{n-1} \frac{x_k}{1+k} \\ \text{s.t.} \quad & \bar{b} \sum_{k=1}^{n-1} \left((1-\alpha_i) \cdot \frac{k}{1+k} + \alpha_i \cdot \frac{k}{2+k} \right) x_k \leq 1 - \alpha_i \\ & \sum_{k=0}^{n-1} x_k = 1 \\ & x_k \geq 0 \qquad \forall k \end{aligned} \tag{4}$$

Since this is a linear program with two constraints, its minimum is achieved at some (x_k^*) that only has two nonzero coordinates. It is not hard to show that of these nonzero coordinates must be x_0^* if $\bar{b} \geq 2$ for (4) to be satisfied. Letting $x_{k^*}^*$ be the other nonzero coordinate, one can work out that the objective value is at least

$$\alpha_i \sum_{k=0}^{n-1} \frac{x_k^*}{1+k} \geq \alpha_i \left(1 - \frac{(2+k^*)(1-\alpha_i)}{\bar{b}(2+k^*-\alpha_i)} \right) \geq \alpha_i \left(1 - \frac{3(1-\alpha_i)}{3\bar{b} - \alpha_i \bar{b}} \right) \geq \alpha_i \left(1 - \frac{1}{\bar{b}} \right).$$

Now let us handle the case where agent i does run out of budget. The argument is similar. Let τ be the time at which the agent uses up all of her budget. Let x_k be the fraction of times $t \leq \tau$ that there are k agents $j \neq i$ bidding. Agent i's number of wins is as before, but we multiply by τ/T to account for the fact that agent i runs out of budget early:

$$\frac{1}{T} \sum_{t=1}^{T} W_i[t] \approx \frac{\tau}{T} \cdot \alpha_i \sum_{k=0}^{n-1} \frac{x_k}{1+k}. \tag{5}$$

The budget constraint on agents $j \neq i$ works similarly:

$$\frac{\tau}{T} \sum_{k=1}^{n-1} \left((1-\alpha_i) \cdot \frac{\bar{b}k}{1+k} + \alpha_i \cdot \frac{\bar{b}k}{2+k} \right) x_k \lessapprox 1 - \alpha_i. \tag{6}$$

To calculate τ, we calculate that

$$\frac{1}{T} \sum_{t=1}^{\tau} P_i[t] \approx \alpha_i \bar{b} \sum_{k=0}^{n-1} \frac{x_k}{2+k},$$

using the payment rule, since at each time t, agent i bids with probability α_i, and if k other agents $j \neq i$ bid, there are $1+k$ bidding agents total, so agent i wins with probability $\frac{1}{1+k}$ in which case they pay $\bar{b} \cdot \frac{1+k}{2+k}$. Since agent i has $\alpha_i T$ tokens total,

$$\frac{\tau}{T} \approx \frac{\alpha_i T}{\sum_{t=1}^{\tau} P_i[t]} \approx \frac{1}{\bar{b} \sum_{k=0}^{n-1} \frac{x_k}{2+k}}. \tag{7}$$

Substitute (7) into (5) and (6) to see that we now must lower bound the value of the following minimization problem.

$$\begin{aligned}\min_{(x_k)_{k=0}^n} \quad & \alpha_i \cdot \frac{\sum_{k=0}^{n-1} \frac{x_k}{1+k}}{\bar{b}\sum_{k=0}^{n-1} \frac{x_k}{2+k}}. \\ \text{s.t.} \quad & (1-\alpha_i)\sum_{k=0}^{n-1} \frac{x_k}{2+k} \geq \sum_{k=1}^{n-1} \left((1-\alpha_i)\cdot \frac{k}{1+k} + \alpha_i \cdot \frac{k}{2+k}\right) x_k \\ & \sum_{k=0}^{n-1} x_k = 1 \\ & x_k \geq 0 \qquad \forall k\end{aligned}$$

The rest of the argument is similar to the previous case. We argue that this is the minimization of a linear-fractional function with positive denominator, and thus quasi-concave function, over a polytope, so there is an optimal solution x_k^* with only two nonzero coordinates. We then argue that one of these nonzero coordinates must be x_0^*, and letting $x_{k^*}^*$ be the other nonzero coordinates, we show the value of this minimization problem is lower bounded by

$$\alpha_i \cdot \frac{3+2k^* - \alpha_i}{\bar{b}(2+k^*-\alpha_i)} \geq \alpha_i \cdot \frac{5-\alpha_i}{\bar{b}(3-\alpha_i)} \geq \frac{5\alpha_i}{3\bar{b}},$$

as desired. ◀

4 Upper Bound on the Robustness under an Arbitrary Payment Rule

In this section, we consider payment schemes different than the ones in Section 3, and show that the one used in our Competitive Subsidy Mechanism achieves optimal robustness under static strategies. We come up with an optimization problem that bounds this robustness factor and then solve it numerically to get the desired upper bound.

Consider the generalization of the Competitive Subsidy Mechanism where agents pay some $p_k \geq 0$ when there are k bidding agents (with our mechanism being a special case where $p_k = \bar{b} \cdot k/(1+k)$). We call this generalization General Cost Mechanism and bound its performance under any payment scheme.

We consider an agent i following the α_i-aggressive strategy and all other agents (with a combined budget $(1-\alpha_i)T$) coordinating so that at each time, exactly k agents other than i bid. Under such simple strategies, we can explicitly calculate how many rounds agent i wins in expectation, which is equal to the fraction of ideal utility she receives by the definitions of ideal utility and α_i-aggressive strategy. Given the commitment to these strategies, agents $j \neq i$ are picking k to minimize i's utility, while the mechanism picks the payment scheme to maximize it. The harder part is limiting the search space of this optimization problem. We consider agents $j \neq i$ only bidding 1 or 2 at a time, making only the payments p_1, p_2, p_3 relevant for the mechanism designer. While in principle considering more agents bidding at the same time could harm agent i more, we show below that even with this strategy by the other players, we can argue that the 0.625 bound is best possible. Next, we prove upper bounds on these three payments, given that having them too high can make agents run out of budget even without adversarial competition.

This process is shown in the next theorem, where $\mu(p_1, p_2, p_3, k)$ is the number of rounds agent i wins, which is minimized over k and maximized over p_1, p_2, p_3. In the computation of $\mu(\cdot)$, γ is the fraction of rounds where agents $j \neq i$ still have budget remaining out of all the rounds that agent i still has budget remaining. In the full version of the paper, we work out and prove the above computation in detail, which considers agents with equal fair shares $1/n$ and $n \to \infty$.

▶ **Theorem 6.** *Suppose λ is greater than the value of the following optimization problem.*

$$\max_{\substack{0<p_1\le 2e\\ 0\le p_2\le 4e\\ 0\le p_3\le 12e}} \min_{k\in\{1,2\}} \mu(p_1,p_2,p_3,k)$$

where

$$\begin{aligned}\mu(p_1,p_2,p_3,k) := \max_{\gamma} & \left(\left(\frac{\gamma}{1+k}+(1-\gamma)\right)\min\left\{1, \frac{1}{\frac{p_{k+1}}{k+1}\cdot\gamma+p_1(1-\gamma)}\right\}\right)\\ \text{s.t.} \quad & 1\ge\gamma\ge\min\left\{1,\max\left\{\frac{1}{p_k}, \frac{p_1}{p_k-\frac{p_{k+1}}{k+1}+p_1}\right\}\right\}\\ & \gamma\in\left\{\min\left\{1,\max\left\{\frac{1}{p_k}, \frac{p_1}{p_k-\frac{p_{k+1}}{k+1}+p_1}\right\}\right\}, \frac{p_1}{\frac{p_{k+1}}{k+1}-p_1}\right\}.\end{aligned}$$

Then, there exists a number of players n such that with equal fair shares $\alpha_i = 1/n$, an $1/n$-aggressive strategy has a robustness factor at most $\lambda + O\left(\sqrt{\log T/T}\right)$ in General Cost Mechanism no matter the choice of the payment scheme $(p_k)_k$.

We numerically brute-force optimized the optimization problem in Theorem 6 by discretizing the space of (p_1,p_2,p_3) and evaluating the objective function at each (p_1,p_2,p_3) in the discretized space. In our code, we found no (p_1,p_2,p_3) such that the objective is more than 0.625, giving numerical evidence that the value of the optimization problem is 0.625.[4]

▶ **Numerical Result.** *By numeric calculations, the expression of Theorem 6 is at most* 0.625*, which indicates that there is no payment rule $(p_k)_k$ that makes the α_i-aggressive strategy better than $\left(0.625+\omega\left(\sqrt{\log T/T}\right)\right)$-robust for all n.*

5 Making the Robust Strategies Form an Equilibrium

While robust strategies like in the previous section are nice, they do not always form an equilibrium. Previous work [8, 6, 4, 13] proposed robust strategies in different mechanisms for this setting that do not form an equilibrium. On the other hand, [14] designed strategies that form an equilibrium but offer no robustness. [11] achieve both types of results by designing a mechanism that has a 1/2-robust $\left(1-\prod_{j\in[n]}(1-\alpha_j)\right)$-good equilibrium (see Definition 3). They also note that the $1-\prod_j(1-\alpha_j)$ factor is optimal in that even if the principal is able to see all the values, they cannot guarantee a higher fraction of ideal utility than $1-\prod_j(1-\alpha_j)$ to every agent in expectation.

In this section, we explore if our mechanism achieves similar results. First, we shall argue that in our mechanism, if we choose $\bar{b}$ as suggested by the previous section, each player using an α_i-aggressive strategy does not form an equilibrium. However, if we choose $\bar{b}$ slightly differently and make a slight modification of the mechanism by forcing players to bid at least an α_i fraction of the time, we can make the α_i-aggressive strategies form an equilibrium at the expense of sacrificing the robustness factor a little bit. Here, we demonstrate how to make this modification to $\bar{b}$ in the symmetric case where $\alpha_i = 1/n$. In Section 6, we give a reduction from the general α_i case to the symmetric case.

[4] Our code can be found at `https://github.com/davidxlin/optimization-problem-for-robustness`.

First, let us argue why when choosing the payment constant $\bar{b} = 8/3$ as in the previous section, everyone playing the robust strategy is not necessarily an equilibrium. Assume every agent is playing the robust strategy, so they each bid with probability $1/n$ independently across agents and time, so long as they have budget. Then, at any given time that everyone has budget remaining, for any agent i, the number of other agents $j \neq i$ that bid is distributed as $X \sim \text{Binomial}(n-1, 1/n)$. Therefore, conditioned on agent i bidding, agent i's expected payment is $\mathbb{E}\left[\frac{\bar{b}}{2+X}\right]$ since conditioned on the number of other agents bidding X, agent i wins with probability $\frac{1}{1+X}$, in which case they pay $\bar{b} \cdot \frac{1+X}{2+X}$. By substituting in $\bar{b} = 8/3$ and computing $\mathbb{E}[1/(2+X)]$, which we show is equal to $(n+1)/(1+n(1-1/n)^{n+1}$ in the full version of the paper, we can show that this expected payment is less than 1 for $n \geq 3$. Since this is less than 1, agents have T/n budget, and are bidding with probability $1/n$, this means that agents will have $\Omega(T)$ budget remaining at the end of the mechanism with high probability. Clearly, for value distributions that are nonzero with probability greater than $1/n$, the agents are not best responding: they should bid more to use more of their budget.

However, the above calculation suggests the following idea. We should set $\bar{b} = 1/\mathbb{E}[1/(2+X)]$, and then agents will be spending their budget exactly in expectation. We still obtain high robustness beating all previous work with this choice of $\bar{b} = 1/\mathbb{E}[1/(2+X)] = (n+1)/(1+n(1-1/n)^{n+1}) \approx e$: the robustness factor (using the calculations in the proof of Theorem 5) becomes

$$\min\left\{1 - \frac{3(1-1/n)}{3\bar{b} - \bar{b}/n}, \frac{5-1/n}{\bar{b}(3-1/n)}\right\} \geq \frac{5}{3e} \approx 0.61.$$

Notice that the mechanism allocates the item so long as there is at least one bidder. Each bidder playing a $1/n$-aggressive strategy bids with probability $1/n$ as long as they have budget remaining. If everyone plays a $1/n$-aggressive strategy, by the fact that agents spend their budget exactly in expectation and by concentration inequalities, no agent runs out of budget too early with high probability. Then, there are $1-(1-1/n)^n$ rounds in which the item is allocated in expectation. By the strategies and symmetry, this will be the fraction of ideal utility each agent achieves, matching the optimal as shown in [11].

What is left to do is to argue that choosing $\bar{b}$ in this way is indeed an equilibrium. Intuitively, for a fixed agent i, they have two potential deviations: they could bid more often, or they could bid less often. By bidding more often, agent i spends more artificial currency. We can also see that agent i bidding more often decreases the spending of other agents as follows. In a given round, conditioned on there being k bidders, the expected payment of a bidder is $\bar{b}/(k+1)$: each bidder wins with probability $1/k$ in which case they pay $\bar{b}k/(k+1)$. Hence, fixing the strategies of agents $j \neq i$, agent i bidding more often increases the number of bidders k at each timestep, lowering the other agents' payments. Therefore, agent i will run out of money quicker with the same probability of winning each round (conditioned on bidding), so they will obtain more lower-valued rounds at the beginning instead of spacing their wins.

Whether or not agent i should bid less is slightly more nuanced. In fact, if we retain the same mechanism as in Competitive Subsidy Mechanism, agent i may want to bid less. For example, if agent i's value distribution were $\mathcal{F}_i = \text{Bernoulli}(1/(2n))$, then a $1/n$-aggressive strategy would imply that agent i is bidding sometimes when they have 0 value. Instead, they should not bid when they have 0 value, which would cause others' payments to go up (using the previous reasoning about the other agents' spending), and then when other agents run out of budget, agent i will have a higher probability of winning on rounds that they actually have value 1.

To accommodate this, we modify the mechanism to enforce that at each time t, each agent must have bid at least $t/n - o(T)$ times; otherwise, the principal will force them to bid. Then, attempting to underbid is not a helpful strategy. We describe the mechanism formally in Algorithm 2, where the only difference from Algorithm 1 is the enforcement of the minimum bidding requirement.

Algorithm 2 Competitive Subsidy Mechanism with Bidding Minimum.

Input: Number of rounds T, payment constant $\bar{b}$, Underbidding allowance $\epsilon = o(T)$
Endow each agent with a budget $B_i[1] = T/n$ of artificial credits;
for $t = 1, 2, \ldots, T$ **do**
  Endow each agent with a budget $B_i[1] = \alpha_i T$ of bidding credits;
  for $t = 1, 2, \ldots, T$ **do**
    Agents either request to bid or not (let $r_i[t]$ be the indicator that agent i requests to bid);
    Enforce bidding minimums: $\boldsymbol{r_i[t] \leftarrow 1}$ **for each** $\boldsymbol{i}$ **such that** $\boldsymbol{\sum_{s=1}^{t} r_i[s] \leq \frac{t}{n} - \epsilon}$;
    Enforce budgets: $r_i[t] \leftarrow 0$ for each i such that $B_i[t] \leq 0$;
    Define $S[t] = \{i : r_i[t] = 1\}$ to be the set of bidding agents;
    Select a winner uniformly at random from $S[t]$ (let $W_i[t]$ be the indicator agent i wins);
    Set payments $P_i[t] = \bar{b} \cdot \frac{|S[t]|}{1+|S[t]|} \cdot W_i[t]$ (note only the winner pays);
    Update budgets: $B_i[t+1] \leftarrow B_i[t] - P_i[t]$
  end
end

Our main result of this section is as follows, proved formally in the full version of the paper.

▶ **Theorem 7.** *Consider Competitive Subsidy Mechanism with Bidding Minimum with payment constant* $\bar{b} = (n+1)/(1+n(1-1/n)^{n+1})$ *and underbidding allowance* $\epsilon = \sqrt{T \log T}$. *Then, when players have equal shares, every agent playing a* $1/n$*-aggressive strategy is a* $\lambda_{\texttt{ROB}}$*-robust* $\lambda_{\texttt{NASH}}$*-good approximate-equilibrium for some* $\lambda_{\texttt{ROB}}$ *and* $\lambda_{\texttt{NASH}}$ *satisfying*

$$\lambda_{\texttt{ROB}} \geq \frac{5}{3e} - O\left(\sqrt{\frac{\log T}{T}}\right), \quad \lambda_{\texttt{NASH}} \geq 1 - \left(1 - \frac{1}{n}\right)^n - O\left(\sqrt{\frac{\log T}{T}}\right).$$

6 Robust Equilibrium with Asymmetric Fair Shares

In this section, we generalize Section 5 to asymmetric fair shares. The issue with Competitive Subsidy Mechanism is that it allocates the item uniformly at random among the bidding agents. While this is not an issue for robustness claims (all such results of previous sections hold for arbitrary fair shares), it is an issue for the equilibrium claim. In fact, to get better than 1/2 ideal utility guarantees in equilibrium, we have to allocate the item asymmetrically. As before, to get equilibrium guarantees, we want a payment scheme such that agents use their budget exactly in expectation. Consider such a scheme and the following simplified setting: two agents with fair shares α and $1-\alpha$ that request the item with probabilities α and $1-\alpha$, respectively. Allocating uniformly at random makes the agent with fair share α win $\alpha(\alpha + (1-\alpha)/2) = \alpha(1+\alpha)/2$ fraction of the rounds. This results in 1/2 fraction of her ideal utility when α is small, which is smaller than the robustness guarantee when not setting $\bar{b}$ to obtain an equilibrium.

To remedy this issue and extend our mechanism, we simulate agents with different fair shares with multiple "small" agents. The basic idea is that if the fair shares α_i are all rational with common denominator m[5], we run Competitive Subsidy Mechanism with Bidding Minimum with m agents where agent i gets to control $k_i := \alpha_i m$ simulated agents in Competitive Subsidy Mechanism with Bidding Minimum.

Implementing this idea naively does not quite work to obtain the same equilibrium behavior, where each simulated agent is bidding independently across rounds with probability $1/m$. There is no particular reason why an agent i controlling multiple simulated agents should have the simulated agents bid independently of each other. Instead, we shall have the principal force some level of independence by only allowing an agent to request whether they want at least one simulated agent bidding or not. If agent i requests to bid at time t, then the principal shall sample requests to bid $(\widehat{r}_{(i,1)}[t], \dots, \widehat{r}_{(i,k_i)}[t])$ distributed as i.i.d. Bernoulli$(1/m)$ conditioned on at least one of them being nonzero to use as the requests in the simulated agents that i controls.

If agent i uses a $\left(1 - \left(1 - \frac{1}{m}\right)^{k_i}\right)$-aggressive strategy, where a β-aggressive strategy is as before, requesting whenever the value is in the top β-quantile of the value distribution, then the $\widehat{r}_{(i,i')}[t]$ will be i.i.d. Bernoulli$(1/m)$. This emulates each simulated agent playing a $1/m$-aggressive strategy. If each simulated agent wins at least λ_i/m rounds, then agent i will win $\lambda_i k_i/m$ rounds, which then implies robustness and equilibrium utility lower bounds. In particular, if we use $\bar{b} = (m+1)/(1+m(1-1/m)^{n+1})$ as in Section 5, a $\left(1 - \left(1 - \frac{1}{m}\right)^{k_i}\right)$-aggressive strategy is $(5/(3e) - o(1))$-robust, and if each agent i plays a $\left(1 - \left(1 - \frac{1}{m}\right)^{k_i}\right)$-aggressive strategy, each agent obtains a $1-\left(1-\frac{1}{m}\right)^m - o(1) \geq 1 - 1/e - o(1)$ fraction of their ideal utility.

Using the same arguments as in Section 5, if we put a minimum bidding constraint, each agent playing this $\left(1 - \left(1 - \frac{1}{m}\right)^{k_i}\right)$-aggressive strategy is an equilibrium. The way we have set the mechanism up, agents can only control the frequency at which they bid. They can't underbid by enforcement, and overbidding does not help because this only decreases the payments of others while winning worse rounds for the overbidder.

Our full mechanism for asymmetric fair shares is as follows. Given the fair shares $\alpha_i = k_i/m$, create a set of simulated agents $\widehat{N} = \{(i,i') : i \in [n], i' \in [k_i]\}$. Initialize an instance $\widehat{\mathcal{M}}$ of Competitive Subsidy Mechanism with the agent set $\widehat{N}$ and equal fair shares: $\widehat{\alpha}_{(i,i')} = 1/m$ for each $(i,i') \in \widehat{N}$. At each time t, each real agent i can request to bid or not. Let $r_i[t]$ denote the indicator that agent i requests to bid. We enforce a bidding minimum that $\sum_{s=1}^{t} r_i[s] \geq \left(1 - \left(1 - \frac{1}{m}\right)^{k_i}\right) t - o(T)$. Let $S[t] = \{i : r_i[t] = 1\}$ be the set of bidding agents. For $X_1, \dots, X_k$ i.i.d. Bernoulli$(1/m)$ random variables, let $\mathcal{D}_{k,m}$ be the distribution of $(X_1, \dots, X_k)$ conditioned on the event that at least one of the $X_i = 1$. For each bidding agent $i \in S[t]$, we sample $(\widehat{V}_{(i,1)}[t], \dots, \widehat{V}_{(i,k_i)}[t]) \sim \mathcal{D}_{k_i,m}$. Set $\widehat{S}[t] = \{(i,i') : i \in S[t], \widehat{V}_{(i,i')} = 1\}$, which we set as the set of requesting simulated agents at time t, and we simulate $\widehat{\mathcal{M}}$ at time with bidding agents $\widehat{S}[t]$. From $\widehat{\mathcal{M}}$, there is simulated agent (i,i') who won the item, and we give the item to the real agent i. (We fully simulate $\widehat{\mathcal{M}}$, so the budgets of the simulated agents also get updated within $\widehat{\mathcal{M}}$, and they are enforced in $\widehat{\mathcal{M}}$ when the simulated agents request to bid.)

[5] If the fair shares α_i are irrational, or if the common denominator is large, we can approximate the fair shares with rational fair shares with small denominators and obtain approximate guarantees. Specifically, we show in the full version of the paper that it suffices to approximate the fair shares with rational numbers with denominators at least $1/(2 \min_i \alpha_i \epsilon)$ to obtain a $(1-\epsilon)$-fraction of our guarantees on the achieved fraction of ideal utility.

We formally describe this mechanism in Algorithm 3 and prove the following theorem in the full version of the paper.

Algorithm 3 Asymmetric Fair Share Mechanism.

Input: Fair shares $(\alpha_i)_{i\in[n]}$ where $\alpha_i = k_i/m$, number of rounds T, payment constant $\bar{b}$, underbidding allowance $\epsilon = o(T)$
Let $\widehat{N} = \{(i, i') : i \in [n], i' \in [k_i]\}$;
Initialize an instance $\widehat{\mathcal{M}}$ of Competitive Subsidy Mechanism with the agent set $\widehat{N}$, equal fair shares $\widehat{\alpha}_{(i,i')} = 1/m$, and payment constant $\bar{b}$;
for $t = 1, 2, \ldots, T$ **do**
  Agents either request to bid or not (let $r_i[t]$ be the indicator that agent i requests to bid);
  Enforce bidding minimums: $r_i[t] \leftarrow 1$ for each i such that $\sum_{s=1}^{t} r_i[s] \leq \left(1 - (1 - \frac{1}{m})^{k_i}\right) t - \epsilon$;
  Let $S[t] = \{i : r_i[t] = 1\}$ be the set of bidding agents;
  For each $i \in S[t]$, sample $(\widehat{V}_{(i,1)}[t], \ldots, \widehat{V}_{(i,k_i)}[t]) \sim \mathcal{D}_{k_i,m}$ (where $\mathcal{D}_{k,m}$ is the distribution of $(X_1, \ldots, X_k)$ conditioned on at least one $X_i = 1$ where the X_i are i.i.d. Bernoulli$(1/m)$) ;
  Let $\widehat{S}[t] = \left\{(i, i') : i \in S[t], \widehat{V}_{(i,i')} = 1\right\}$;
  Simulate $\widehat{\mathcal{M}}$ at time t with requesting agents $\widehat{S}[t]$ to determine a simulated winner in $\widehat{N}$;
  For the simulated agent (i, i') that won the item in $\widehat{\mathcal{M}}$, give the item to agent i.
end

▶ **Theorem 8.** *Consider Asymmetric Fair Share Mechanism with $\bar{b} = (m+1)/(1 + m(1 - 1/m)^{m+1})$ and $\epsilon = \sqrt{T \ln T}$. Then, each agent i playing a $\left(1 - (1 - \frac{1}{m})^{k_i}\right)$-aggressive strategy is a $\lambda_{\texttt{ROB}}$-robust $\lambda_{\texttt{NASH}}$-good approximate-equilibrium for some $\lambda_{\texttt{ROB}}$ and $\lambda_{\texttt{NASH}}$ satisfying*

$$\lambda_{\texttt{ROB}} \geq \frac{5}{3e} - O\left(\sqrt{\frac{\log T}{T}}\right), \lambda_{\texttt{NASH}} \geq 1 - \left(1 - \frac{1}{m}\right)^m - O\left(\sqrt{\frac{\log T}{T}}\right).$$

▶ Remark 9. At a high level, our techniques are similar to that of [11]: in the equal fair share case, their mechanism is just Competitive Subsidy Mechanism, but instead of agents paying an amount dependent on the number of bidders, each bidder always pays 1 each round they request. Our works also differ in how we handle the asymmetric fair share case. Our approach in reducing asymmetric fair shares to symmetric fair shares does not seem to be able to get the optimal $\lambda_{\texttt{NASH}} = 1 - \prod_j (1 - \alpha_j)$, since if there are m agents with equal fair shares, the best possible is $\lambda_{\texttt{NASH}} = 1 - (1 - 1/m)^m$, and if $1/m \leq \alpha_j$ for each j with the inequality strict for at least one j, then $1 - (1 - 1/m)^m < 1 - \prod_j (1 - \alpha_j)$. In contrast, [11] are able to obtain $\lambda_{\texttt{NASH}} = 1 - \prod_j (1 - \alpha_j)$. They do this by changing the uniformly random allocation rule. Specifically, if at some time a set of agents S bid, they allocate the item to an $i \in S$ according a probability distribution $(p_i^S)_{i\in S}$ that depends on S. These probability distributions $(p_i^S)_{i\in S}$ are extremely complicated. Much of their work is dedicated to only proving the existence of $(p_i^S)_{i\in S}$ that guarantee their robustness and equilibrium guarantees, and they do not have any formula for the p_i^S. In contrast, our mechanism, Asymmetric Fair Share Mechanism, is much simpler.

References

1 Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014. doi:10.1137/120878422.

2 Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair-share allocations for agents with arbitrary entitlements. *Mathematics of Operations Research*, 49(4):2180–2211, 2024. doi:10.1287/MOOR.2021.0199.

3 Moshe Babaioff and Uriel Feige. Share-based fairness for arbitrary entitlements. In *ACM Symposium on Theory of Computing (STOC)*, 2025.

4 Siddhartha Banerjee, Giannis Fikioris, and Eva Tardos. Robust pseudo-markets for reusable public resources. In *ACM Conference on Economics and Computation (EC)*, 2023.

5 Ruggiero Cavallo. Incentive compatible two-tiered resource allocation without money. In Ana L. C. Bazzan, Michael N. Huhns, Alessio Lomuscio, and Paul Scerri, editors, *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS*, Paris, France, 2014.

6 Giannis Fikioris, Siddhartha Banerjee, and Eva Tardos. Beyond worst-case online allocation via dynamic max-min fairness. In *ACM Conference on Economics and Computation (EC)*, 2025.

7 Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. From monetary to nonmonetary mechanism design via artificial currencies. *Mathematics of Operations Research*, 46(3):835–855, 2021. doi:10.1287/MOOR.2020.1098.

8 Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. The remarkable robustness of the repeated fisher market. In *ACM Conference on Economics and Computation (EC)*, pages 562–562, 2021.

9 Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Toronto, Canada, 2010.

10 Matthew O Jackson and Hugo F Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007.

11 David X Lin, Siddhartha Banerjee, Giannis Fikioris, and Éva Tardos. Robust equilibria in shared resource allocation via strengthening border's theorem. *arXiv preprint arXiv:2505.11431*, 2025. doi:10.48550/arXiv.2505.11431.

12 David X Lin, Giannis Fikioris, Siddhartha Banerjee, and Éva Tardos. Robust resource allocation via competitive subsidies. *arXiv preprint arXiv:2511.09934*, 2025.

13 David X Lin, Daniel Hall, Giannis Fikioris, Siddhartha Banerjee, and Éva Tardos. Online resource sharing: Better robust guarantees via randomized strategies. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2025.

14 Chido Onyeze, Siddhartha Banerjee, Giannis Fikioris, and Éva Tardos. Allocating public goods via dynamic max-min fairness: Long-run behavior and competitive equilibria. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS)*, 9(1):1–45, 2025. doi:10.1145/3711695.