

Online Resource Sharing: Better Robust Guarantees via Randomized Strategies

David X. Lin, Daniel Hall, Giannis Fikioris, Siddhartha Banerjee, Éva Tardos

Cornell University

{dx12, deh275}@cornell.edu, gfikioris@cs.cornell.edu, {sbanerjee, eva.tardos}@cornell.edu

Abstract

We study the problem of fair online resource allocation via non-monetary mechanisms, where multiple agents repeatedly share a resource without monetary transfers. Previous work has shown that every agent can guarantee $1/2$ of their ideal utility (the highest achievable utility given their fair share of resources) robustly, i.e., under arbitrary behavior by the other agents. While this $1/2$ -robustness guarantee has now been established under very different mechanisms, including pseudo-markets and dynamic max-min allocation, improving on it has appeared difficult.

In this work, we obtain the first significant improvement on the robustness of online resource sharing. In more detail, we consider the widely-studied repeated first-price auction with artificial currencies. Our main contribution is to show that a simple randomized bidding strategy can guarantee each agent a $2 - \sqrt{2} \approx 0.59$ fraction of her ideal utility, irrespective of others' bids. Specifically, our strategy requires each agent with fair share α to use a uniformly distributed bid whenever her value is in the top α -quantile of her value distribution. Our work almost closes the gap to the known $1 - 1/e \approx 0.63$ hardness for robust resource sharing; we also show that any static (i.e., budget independent) bidding policy cannot guarantee more than a 0.6 -fraction of the ideal utility, showing our technique is almost tight.

1 Introduction

There are many settings where multiple self-interested agents share a resource controlled by a principal. Specifically, we consider the problem where the resource is repeatedly allocated to one of the agents over a long time horizon. For example, consider multiple scientific labs at a university sharing a computer cluster. Each lab wants to run experiments in every time slot, but only one lab can use the cluster at a time. The system administrator must decide who uses the cluster in each time slot. This has to be done in a “fair” way so that each lab is satisfied with the resulting allocation. In addition,

the allocation has to happen without requiring monetary payments from the labs.

The allocation model and mechanism that we study were first introduced by [Gorokh *et al.*, 2021] and subsequently also studied in [Fikioris *et al.*, 2023; Banerjee *et al.*, 2023]. Every agent is endowed with a *fair share* of the resource. Roughly speaking, an agent's fair share is their intrinsic right to enjoy a fraction of the resource. Letting different agents have different fair shares encapsulates many realistic real-world scenarios. For example, it may be natural for a larger research group within a university to get access to a computer cluster more often than a smaller group. As a benchmark, this line of work defines and uses *ideal utility*. Roughly speaking, an agent's ideal utility is the maximum per-round utility that she can obtain if she is restricted to obtaining only her fair share of the resource.

[Gorokh *et al.*, 2021] introduce a simple non-monetary mechanism in which each agent is endowed with an amount of artificial currency proportional to their fair share, and at each time, the item is allocated according to a first-price auction using the artificial currency. Assuming that each agent's value for the item is drawn independently from a fixed distribution each round, they show that in this mechanism, each agent can guarantee a $\frac{1}{2} - o(1)$ fraction of her ideal utility robustly, which means that each agent can make this guarantee regardless of the behavior of the other agents and their value distribution. [Fikioris *et al.*, 2023] study the same problem under a different mechanism and get the same result.

Our results We use the model and mechanism of [Gorokh *et al.*, 2021] and offer a much more detailed analysis of the achievable robust guarantees. Our main result is an improved bidding strategy that any agent can follow to guarantee a $2 - \sqrt{2} - o(1) \approx 0.59$ fraction of her ideal utility robustly (Theorem 1). We emphasize that since this result holds under arbitrary behavior by other agents, the same utility guarantee can be made under any equilibrium. This lower bound is close to the upper bound of $1 - \frac{1}{e} \approx 0.63$ that [Fikioris *et al.*, 2023] give. Our proposed strategy, *Randomized Robust Bidding*, is very simple. While budget suffices, an agent's bid is sampled from a certain uniform distribution if her value is in the top α -quantile of her value distribution, where α is her fair share.

This randomization in bidding is essential for improving

the $\frac{1}{2}$ ideal utility guarantee. For such a guarantee [Banerjee *et al.*, 2023; Gorokh *et al.*, 2021] suggest that an agent should bid a fixed constant whenever her value is sufficiently high. Under such simplistic bidding, the other agents know exactly how much they need to bid to beat agent i . It is not hard to show that no fixed bidding can robustly guarantee a better than $\frac{1}{2}$ of the ideal utility. The same holds for the $\frac{1}{2}$ guarantee of [Fikioris *et al.*, 2023]: their much simpler mechanism only allows requesting the resource or not, equating to a fixed bid for pseudo-markets.

Given the above observations about fixed bidding, a more involved strategy is essential for stronger results. While a uniform distribution is arguably the simplest continuous distribution, it enjoys the following property. It is the distribution that minimizes spending, subject to achieving the same fraction of ideal utility when the other agents are bidding b , for a range of values $b \geq 1$. Carefully optimizing this uniform distribution, we get the $(2 - \sqrt{2})$ guarantee of Theorem 1 when its support is $[0, 1 + \sqrt{2}]$.

A main component of our utility lower bound (and subsequent results) is Lemma 1, which greatly simplifies the problem of getting robust guarantees. Specifically, we show that for any $\beta \leq 1$, if an agent can guarantee a β fraction of her ideal utility when her value distribution is Bernoulli with mean equal to her fair share, then she can guarantee the same for any value distribution. In addition, our reduction works for any mechanism, showing that this Bernoulli value distribution is the worst-case for this problem.

We show that our robust utility guarantee is almost tight in the following sense. Suppose the agent does have the worst-case Bernoulli distribution. If the agent uses any strategy that involves bidding from the same fixed distribution every round in which she has value 1, she cannot robustly guarantee more than a 0.6 fraction of their ideal utility (Theorem 2). This result also showcases the strength of the uniform distribution we used before. If the probability that the agent bids more than b when she has value 1 is not at least $(1 - \beta)b$ at every point $b \geq 1$ (i.e., above a certain uniform distribution), then the other agents can bid b to make the agent get less than β fraction of her ideal utility. In other words, we prove that if the agent’s bidding is not aggressive enough, she cannot guarantee a β fraction of her ideal utility. On the other hand, by bidding too much, the agent will run out of budget. We combine these bounds to show that an agent cannot obtain a guarantee of more than a 0.6 fraction of their ideal utility regardless of the choice of bidding distribution.

Our final theoretical result is a bidding strategy the agents other than i can follow that ensures agent i receives at most a $(1 - \frac{1}{e})$ fraction of her ideal utility under the previous worst-case Bernoulli value distribution (Theorem 3). Unlike the previous result, this is under arbitrary strategies that agent i can follow, such as time-varying strategies. Due to [Fikioris *et al.*, 2023], it is already known that there is no mechanism can guarantee every agent a $(1 - \frac{1}{e})$ fraction of her ideal utility but this is the first explicit strategy that does this. We believe that the bidding distribution we provide for the other agents might be of interest to close the gap between the upper and lower bounds.

Finally, in Section 5, we provide empirical evidence of how our Randomized Robust Bidding strategy is superior to the fixed bidding of previous work. In particular, we show that when each of n agents with equal fair shares follows the Randomized Robust Bidding strategy, then every agent ends up with a $1 - (1 - 1/n)^n$ fraction of their ideal utility (note that $1 - (1 - 1/n)^n \rightarrow 1 - 1/e$ as $n \rightarrow \infty$). This fraction is the theoretical maximum *any allocation* can guarantee when there are n agents, even if it has knowledge of the agents’ realized values. This is much higher than the $1/2$ fraction of ideal utility that agents get when following the fixed bidding strategy of [Gorokh *et al.*, 2017; Banerjee *et al.*, 2023].

Related Work Our work follows in a long line of works that consider the problem of repeated allocation of resources without money. The interest in studying such mechanisms stems from its application in many real-world settings, and indeed, the theory has benefited from and influenced successful deployments for course allocation [Budish *et al.*, 2017], food banks [Walsh, 2014; Prendergast, 2022] and cloud computing [Dawson *et al.*, 2013; Vasudevan *et al.*, 2016].

The particular model we consider, with a single indivisible item per round, and agents with random valuations across rounds, was first considered in the work of [Guo and Conitzer, 2010]; however the core idea of ‘linking’ multiple allocations to incentivize truthful reporting without money goes back to the seminal work of [Jackson and Sonnenschein, 2007]. These mechanisms, and a long line of follow-up work [Cavallo, 2014; Gorokh *et al.*, 2017; Balseiro *et al.*, 2019; Blanchard and Jaillet, 2024], provide only Bayes-Nash equilibrium guarantees, and moreover, the mechanisms need to know the value distributions beforehand. None of these mechanisms, however, can provide any guarantee under non-equilibrium actions by other agents.

A more recent line of work, starting from [Gorokh *et al.*, 2021], considers the same setting, but focuses on *robust individual-level* guarantees: the aim now is to guarantee each agent some minimum utility irrespective of how other agents behave. [Gorokh *et al.*, 2021] propose the repeated first-price pseudo-market that we also use in our work, and show that every agent can guarantee a $1/2$ fraction of her ideal utility robustly. Since then, this $1/2$ -robustness guarantee has been re-obtained using very different mechanisms: [Banerjee *et al.*, 2023] give a simple argument to get this guarantee using a repeated first-price auction with a reserve (their main focus is to extend the robustness guarantees to reusable resources, i.e., resources that an agent might want for multiple consecutive rounds), and [Fikioris *et al.*, 2023] show how to get it using a very different non-market based mechanism called Dynamic Max-Min Fairness (while also showing how to get robust guarantees that are distribution-specific and extend to values that are correlated across time). Both these latter works also suggest that the $1/2$ -robustness guarantee is essentially tight under their respective approaches, which is far below the best upper-bound of $1 - 1/e$ we discuss above. The question of whether one can obtain stronger robustness guarantees for the basic single-item setting has, however, remained open until this work.

Further afield, there are problems where agents’ values are known in advance. [Babaioff *et al.*, 2021; Babaioff *et al.*, 2022] study fair resource allocation when the agents’ values are adversarially picked have to be allocated simultaneously. Due to this information structure, their utility benchmarks are much weaker than ours. For example, for an agent who has positive value 1 for only T/n out of T items, they guarantee $\Theta(T/n^2)$ utility as an adversary can “block” the agent for only items of positive value. In contrast, for an agent with Bernoulli($1/n$) values, we guarantee $\approx 0.59T/n$ utility. [Avni *et al.*, 2018; Lazarus *et al.*, 1999] study ‘Poor-man games’, which are games involving bidding with artificial currencies. However, unlike our setting, players have full-information, making them fundamentally different, similar to [Babaioff *et al.*, 2021; Babaioff *et al.*, 2022].

2 Preliminaries

2.1 Model and Ideal Utility

We consider repeated online allocation of a single, indivisible resource via repeated first-price auctions using artificial currencies as introduced in [Gorokh *et al.*, 2021]. There are n agents, $1, 2, \dots, n$. At each time $t = 1, 2, \dots, T$, a principal decides which agent, if any, to receive the resource. Each agent i has a nonnegative value $V_i[t]$ for the item at time t and aims to maximize her total utility, which is the sum of the values of the items she got allocated. We assume the values $V_i[t]$ are drawn independently across both agents i and times t . Specifically, each agent has a time-invariant value distribution \mathcal{F}_i , and the values $V_i[t]$ are drawn from \mathcal{F}_i each round t . We make no assumptions about \mathcal{F}_i beyond non-negativity. The values $V_i[t]$ are private and are not known to the other agents or the principal.

Each agent i has some exogenously defined fair share α_i , where each $\alpha_i \geq 0$ and $\sum_{i=1}^n \alpha_i = 1$. An agent’s fair share measures the exogenously defined fraction of allocated items they should receive in a fair world. A fair principal should consider mechanisms that favor agents with higher fair shares in some way.

As in [Gorokh *et al.*, 2021; Banerjee *et al.*, 2023; Fikioris *et al.*, 2023], to evaluate an agent’s resulting utility, we use the benchmark of *ideal utility*. Intuitively, the ideal utility v_i^* of agent i is the maximum long-term time-average utility the agent can get if allocated an α_i fraction of the rounds. Formally, v_i^* is defined as the maximum expected utility they could achieve from a single round if they can only obtain the item with probability at most α_i :

$$v_i^* = \max_{\rho: [0, \infty) \rightarrow [0, 1]} \mathbb{E}_{V_i \sim \mathcal{F}_i} [V_i \rho(V_i)] \quad (1)$$

s.t. $\mathbb{E}_{V_i \sim \mathcal{F}_i} [\rho(V_i)] \leq \alpha_i$

where $\rho(V_i)$ denotes the probability of obtaining the item conditioned on the value V_i .

We will be interested in *robust* strategies used by the agents. These strategies approximate an agent’s ideal utility regardless of the other agents’ behavior, even if they behave adversarially. We give a formal definition below.

Definition 1. A policy used by an agent i is β_i -*robust* if when using the policy, regardless of the behavior of the other agents

Mechanism 1 Repeated first-price auction with artificial currency

Input: Number of rounds T , number of agents n , and fair shares $\alpha_1, \dots, \alpha_n$

- 1: Endow each agent i with $B_i[1] = \alpha_i T$ tokens of artificial currency.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Agents submit bids $b_i[t]$ where each $b_i[t] \leq B_i[t]$.
 - 4: Select the winner $i^* = \arg \max_i b_i[t]$ (ties broken arbitrarily).
 - 5: Set the payments as $P_i[t] = b_i[t] \mathbf{1}\{i = i^*\}$.
 - 6: Update budgets $B_i[t+1] = B_i[t] - P_i[t]$.
 - 7: Agents get utility $U_i[t] = V_i[t] \mathbf{1}\{i = i^*\}$.
 - 8: **end for**
-

$j \neq i$, the agent’s per-round expected utility is at least β_i fraction of her ideal utility, i.e.,

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

Robust Strategies, Equilibria and Price of Anarchy As also mentioned in previous work, we point out an additional benefit of Definition 1. If every agent has a β -robust policy, then under any equilibrium, every agent achieves a β fraction of her ideal utility. In addition, this implies Price of Anarchy guarantees when the agents’ fair shares are equal, $\alpha_i = 1/n$. In this case, the social welfare is upper-bounded by the sum of the agents’ ideal utilities, implying that, if every agent gets an β fraction of her ideal utility, the resulting social welfare is a β fraction of the optimal one. This means that the Price of Anarchy is at most $1/\beta$.

2.2 Pseudo-Market Mechanism

In this section, we introduce the mechanism we use to allocate the resource, a repeated first-price auction with artificial currency. We note again that this is the same mechanism as proposed in [Gorokh *et al.*, 2021].

At the beginning of time, each agent is endowed with a budget of $B_i[1] = \alpha_i T$ credits of artificial currency. At each time t , each agent submits a bid $b_i[t]$ no more than their current budget $B_i[t]$. The principal selects the agent i^* with the highest bid (ties broken arbitrarily) to allocate the item to. The winning agent i^* pays her bid in artificial currency, and no other agent pays. We denote the payment by agent i as $P_i[t] = b_i[t] \mathbf{1}\{i = i^*\}$. The budgets get updated as $B_i[t+1] = B_i[t] - P_i[t]$. The mechanism is summarized in Mechanism 1.

Agents have no intrinsic value for the artificial currency and simply aim to maximize their total received value. Letting $W_i[t]$ be the indicator for whether agent i won the item at time t , we denote an agent’s utility gained at time t by $U_i[t]$, defined as $U_i[t] = V_i[t] W_i[t]$. Each agent seeks to maximize their total utility, $\sum_{t=1}^T U_i[t]$.

2.3 Formulation as a Two-Player Zero-Sum Game

When analyzing the robustness of a strategy for a particular agent i , we can think of the other $n - 1$ agents as one combined adversary. Specifically, we can think of a single player

with budget the sum of the other agents' budget and her bids are the maximum of their bids. Formally, at each time t , we let $B'[t] = \sum_{j \neq i} B_j[t]$ and $b'[t] = \max_{j \neq i} b_j[t]$. From the perspective of agent i , participating in the mechanism is the same as playing against a single adversary whose budget and bids at round t are $B'[t]$ and $b'[t]$. This gives a reduction from the problem of competing against multiple other players to the problem of only competing against a single adversary, i.e., the one adversary can always simulate the behavior of $n - 1$ other players. Since we are studying bounds on the achievable robustness of agent i 's strategies, we can think of this as a zero-sum game: agent i is trying to maximize her total utility $\sum_{t=1}^T U_i[t]$ and the other agents, thought of as a single combined adversary, is trying to minimize this.

This two-player zero-sum game is the perspective we will take in what follows. We will fix an agent i with fair share α and refer to her simply as "the agent" or "the player," and drop the i subscript from our notation. We will refer to the other players, playing as a single adversary, as just the "adversary" with a fair share $1 - \alpha$. In notation, we will use $'s$ to denote the adversary's quantities, e.g., $B'[t]$ is the adversary's budget, $P'[t]$ is the adversary's payment, etc.

We will note that while we have a two-person zero-sum game formulation, this game is extremely complicated. The strategy space for each player is extremely large, including all possible history-dependent bidding policies over time. Therefore, it is difficult to analyze the equilibrium behavior of this two-person zero-sum game.

3 Robust Strategy Lower Bound of $2 - \sqrt{2}$

In this section, we give an $(2 - \sqrt{2} - O(\sqrt{\log T/T}))$ -robust strategy for the agent. This means an agent can guarantee approximately a 0.59 fraction of their ideal utility under arbitrary behavior by the other agents. This is the first robust bound that breaks the $1/2$ barrier and shows any agent can guarantee more than half her ideal utility regardless of the behavior of the other agents.

Before presenting our lower bound, we simplify the problem by a reduction. Specifically, we reduce the problem of finding robust strategies for any value distribution \mathcal{F} to finding robust strategies for the Bernoulli(α) values. We prove that any β -robust strategy for the Bernoulli case can be converted to a β -robust strategy for arbitrary value distributions.

3.1 Reduction to Bernoulli(α) Value Distributions

In this section, we reduce the problem of arbitrary value distributions to Bernoulli value distributions. This reduction also implies that having a Bernoulli(α) value distribution is the worst case for the agent trying to achieve a high robustness factor. Our reduction is for an arbitrary mechanism, not just the pseudo-market we consider in the rest of the paper.

Lemma 1. *Fix an arbitrary mechanism and an agent with fair share α . Assume there is a β -robust policy $\hat{\pi}$ for that agent when she has a $\hat{\mathcal{F}} = \text{Bernoulli}(\alpha)$ value distribution. Then for any value distribution \mathcal{F} , we can construct a β -robust policy π for an agent with value distribution \mathcal{F} .*

The lemma's proof is based on the following observation. Using the definition of ideal utility, Eq. (1), we can map the values $V[t]$ of any distribution \mathcal{F} to Bernoulli values $\hat{V}[t]$ where $\hat{V}[t] = 1$ corresponds to the agent getting the resource when realizing her ideal utility. The β -robust policy $\hat{\pi}$ can guarantee the item a β fraction of the rounds when $\hat{V}[t] = 1$. In other words, $\hat{\pi}$ can guarantee a β fraction of the rounds that achieve v^*T utility in expectation.

Proof of Lemma 1. Let ρ^* be the maximizer of (1) with value distribution \mathcal{F} . We construct the policy π to "simulate" $\hat{\pi}$. Let π be the policy that, at each time t , defines $\hat{V}[t]$ to be 1 with probability $\rho^*(V[t])$ and 0 otherwise, and behaves the same as $\hat{\pi}$ in the mechanism when using the values $\hat{V}[t]$. It samples the probabilities independently so that $\hat{V}[t]$ is independent of $\hat{V}[t']$ conditioned on $V[t]$ and $V[t']$ for $t \neq t'$.

By the feasibility and optimality of ρ^* in (1), the values $\hat{V}[t]$ are indeed i.i.d. Bernoulli(α). Recall that $W[t]$ denotes the indicator that the agent wins the item in round t . Then,

$$\sum_{t=1}^T \mathbb{E}[\hat{V}[t]W[t]] \geq \beta\alpha T \quad (2)$$

by the β -robustness of $\hat{\pi}$. When the agent uses the strategy π , everything in the game is independent of the actual values $V[t]$ conditioned on the Bernoulli values $\hat{V}[t]$. Therefore,

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[V[t]W[t]] &= \sum_{t=1}^T \mathbb{E}[\mathbb{E}[V[t]W[t] \mid \hat{V}[t]]] \\ &= \sum_{t=1}^T \mathbb{E}[\mathbb{E}[V[t] \mid \hat{V}[t]]\mathbb{E}[W[t] \mid \hat{V}[t]]] \\ &= \sum_{t=1}^T \mathbb{E}[\mathbb{E}[\mathbb{E}[V[t] \mid \hat{V}[t] = 1]\hat{V}[t]W[t] \mid \hat{V}[t]]] \\ &= \sum_{t=1}^T \frac{v^*}{\Pr(\hat{V}[t] = 1)} \mathbb{E}[\hat{V}[t]W[t]] \\ &= \frac{v^*}{\alpha} \sum_{t=1}^T \mathbb{E}[\hat{V}[t]W[t]] \\ &\geq \frac{v^*}{\alpha} \beta\alpha T = \beta v^* T. \end{aligned}$$

Above, the second line follows from the independence of $V[t]$ and $W[t]$ conditioned on $\hat{V}[t]$. The third line follows from the fact that $\mathbb{E}[V[t] \mid \hat{V}[t]] = \mathbb{E}[V[t] \mid \hat{V}[t] = 1]\hat{V}[t]$ since $\hat{V}[t]$ is Bernoulli. The fourth line uses $\mathbb{E}[V[t] \mid \hat{V}[t] = 1] = v^*/\Pr(\hat{V}[t] = 1)$ which follows from the definition of ideal utility. The fifth line uses the fact that $\Pr(\hat{V}[t] = 1) = \alpha$. The inequality comes from (2). \square

3.2 The Randomized Robust Bidding Strategy

Assuming a Bernoulli(α) value distribution, our proposed strategy for the agent is as follows. In rounds where the agent

has a value of 1, she bids using a sample from a specific uniform distribution (assuming she has enough budget).

We can see that using a proper distribution instead of a fixed bid (as done in previous work) is necessary to developing better than $1/2$ -robust strategies as follows. Assume the agent bids a fixed amount b every round she has value 1, while budget lasts. If $b \geq 2$, then with a budget of αT , she can only win at most $\alpha T/b$ times, yielding at most a $1/b \leq 1/2$ fraction of their ideal utility. If $b \leq 2$, then the adversary can bid $b + \epsilon$ to block the agent for approximately T/b rounds, leaving only $(1 - 1/b)T \leq T/2$ rounds for the agent.

Contrarily, against a continuous bid distribution, the adversary cannot predict the agent's exact bid. When both the agent and the adversary bid, the adversary has to either bid too high to ensure the agent does not get the item or bid lower but allow the agent to get the item with some probability.

We now describe our new Randomized Robust Bidding strategy for an arbitrary value distribution: the basic idea is that agent i bids *whenever* she has a high value (irrespective of everything else), but in a way such that her total bidding probability is exactly α_i . We formalize this below for the case where the value distribution of agent i has a unique α_i -quantile, as is the case for any absolutely continuous distribution; if the value distribution has atoms, then we can appropriately randomize at the cutoff to make sure the agent bids with probability exactly α_i .

Randomized Robust Bidding (RRB)

If $V[t]$ is in the top α -quantile of the value distribution, bid $b[t] \sim \text{Uniform}([0, \bar{b}])$, subject to the remaining budget.

More formally, the Randomized Robust Bidding proceeds as follows: Let τ be the last round in which the agent has at least \bar{b} budget left, $\tau = \max\{t \geq 1 : B[t] \geq \bar{b}\}$. Let ρ^* be the optimal solution that realized the ideal utility in (1). For each time t , if $t \leq \tau$, sample $\hat{V}[t] \sim \text{Bernoulli}(\rho^*(V[t]))$. If $\hat{V}[t] = 1$, bid $b[t] \sim \text{Uniform}([0, \bar{b}])$.

Theorem 1. *The RRB strategy with $\bar{b} = 1 + \sqrt{2}$ is β -robust for*

$$\beta = \left(2 - \sqrt{2} - O\left(\sqrt{\frac{\log T}{T}}\right) \right)$$

for any distribution \mathcal{F} that the agent has.

The uniform bidding distribution used in the RRB strategy stems from the following simplified analysis, which we also use for an upper bound in Section 4.1 (see Lemma 2). Consider the agent's bids are sampled from a distribution with CDF $F(\cdot)$. Also assume that her fair share α is small. Since the agent's behavior is the same round, consider the adversary also acting the same every round. Specifically, consider the adversary bidding b' , where $b' \geq 1$ (as otherwise she would have leftover budget). While the adversary's budget lasts, they pay b' each round in at least the $1 - \alpha$ fraction of the rounds that the agent bids 0. If α is small, the adversary's spending on the α fraction of the rounds that the agent has value is negligible. Hence, with a total budget of

$(1 - \alpha)T \approx T$, the adversary runs out of budget at time about T/b' . For the first T/b' rounds, the adversary still has budget remaining, and when the agent bids, she wins with probability $1 - F(b')$. For the remaining $T(1 - \frac{1}{b'})$ rounds, the adversary has no leftover budget, so the agent wins every time she bids. Assuming the agent does not run out of budget (which we ensure later), the resulting fraction of rounds she wins out of the ones she bids in is

$$\frac{1}{b'}(1 - F(b')) + \left(1 - \frac{1}{b'}\right) = 1 - \frac{F(b')}{b'}. \quad (3)$$

Under the above calculation, the agent wants to maximize the above quantity for any b' by setting $F(b') = \lambda b'$, i.e., using a uniform distribution, and then optimizing over λ . Specifically, λ is maximized while ensuring the assumption we made above: the agent does not run out of budget for any bid b' by the adversary. This results in $\lambda = \frac{1}{1 + \sqrt{2}}$, which is the uniform distribution we used above.

In the full proof in the full version of this paper, we have to account for the adversary using more complicated strategies, e.g., bidding as a function of time, remaining budget, etc. We show that there is little benefit to such strategies, which is where the $\sqrt{\log T/T}$ term comes from. This is done by a careful martingale analysis of three key quantities: the agent's total spending, the agent's total utility, and the adversary's spending. We show that with high probability, these three increase with the same rate under the adversary's optimal strategy. This essentially makes the simplified analysis given above correct.

4 Robust Strategy Upper Bounds

In this section, we provide upper bounds for the pseudo-market mechanism. First, we show that our RRB strategy for the $2 - \sqrt{2}$ lower bound in Section 3 is almost tight: for every fixed bidding distribution the agent might use, the adversary can bid to prevent the agent from achieving more than a $\frac{3}{5}$ fraction of her ideal utility. Second, we show that the adversary can bid so that the agent cannot guarantee more than a $1 - \frac{1}{e}$ fraction of her ideal utility under any strategy.

4.1 Fixed Bidding Distribution $\frac{3}{5}$ Upper Bound

In this section, we upper bound the robustness factor of an agent with $\text{Bernoulli}(\alpha)$ value distribution under the assumption that the agent must pick a fixed distribution to bid each round. Specifically, the agent chooses a distribution \mathcal{D} and bids $b[t] = r[t]V[t]$ at each time t where the $r[t]$ are i.i.d. with distribution \mathcal{D} . We allow any distribution \mathcal{D} , even ones that depend on T ; however, we assume the support of \mathcal{D} is in some range $[0, \bar{b}]$ not depending on T . Finally, as is standard with these upper bounds, we assume that α is small. The final result is the following, showing that an agent cannot hope for better than $\frac{3}{5}$ -robustness under such strategies.

Theorem 2. *Assume that an agent with fair share α and value distribution $\text{Bernoulli}(\alpha)$ bids according to distribution \mathcal{D} whenever her value is 1. Then this strategy cannot be β -robust for $\beta > \frac{3}{5}$ as $\alpha \rightarrow 0$.*

The theorem’s main lemma is the formal statement of our argument in Section 3.2. Specifically, we prove that if the agent’s bidding strategy is β -robust, then its CDF must be upper-bounded by the CDF of a certain uniform distribution. This implies that the agent must make high bids with decent probability. Next, we will leverage this to show that if the agent bids too high, she will run out of budget.

Lemma 2. *Let $F(\cdot)$ be the CDF of a bidding distribution that is β -robust, as described in Theorem 2. Then for $x \geq 1 - \alpha$,*

$$F(x-) \leq \frac{1 - \beta}{1 - \alpha} x$$

where $F(x-) = \lim_{y \rightarrow x-} F(y)$.

We prove the lemma by considering the adversary bidding x every round, similar to our argument in (3). Specifically, if $x \geq 1$ (consider the case when $\alpha \rightarrow 0$), we know that the adversary will run out of budget approximately on round $\frac{T}{x}$. While the adversary is bidding x , the agent gets a $1 - F(x-)$ fraction of the rounds she requests. This means that, even if she always wins the item after the adversary runs out of budget, she wins at most

$$\frac{T}{x} (1 - F(x-))\alpha + \left(T - \frac{T}{x}\right)\alpha$$

rounds where she has a value of 1. By robustness, this has to be at least $\beta T \alpha$, yielding the lemma. The full proof, in which we consider the interactions of the agent and the adversary more carefully, is deferred to the full version of this paper.

The final step for proving Theorem 2 is to show that if $\beta > \frac{3}{5}$, then the agent’s bidding is too high. Specifically, if the adversary uses a constant bid b' every round then the agent has to pay in expectation

$$\mathbb{E}_{b \sim \mathcal{D}} [b \mid b \geq b']$$

every time she wins the resource. Using Lemma 2, we can show that if $\beta > \frac{3}{5}$, then under the right b' , the above payment is too high, and the agent will run out of money before accumulating a β fraction of her ideal utility.

We defer the complete proof of Theorem 2 to the full version of this paper.

4.2 Constructive Proof for an Upper Bound of

$$1 - \frac{1}{e} + \frac{\alpha}{e} + o(1)$$

[Fikioris *et al.*, 2023] prove that no mechanism can guarantee every agent a $1 - \frac{1}{e}$ fraction of her ideal utility when agents have Bernoulli valuations. While this bound remains the best-known upper bound for robust guarantees, its proof is existential. In particular, this result does not provide insight into how the adversary can prevent an agent from obtaining more than such a fraction of her ideal utility. In this section, we provide a simple bidding strategy that the adversary can follow to ensure that the agent does not get more than a $1 - \frac{1}{e}$ fraction of her ideal utility. This result is stronger than the one in Theorem 2 in two ways. First, the adversary’s strategy is independent of the agent’s strategy. Second, the upper bound holds for any strategy the agent might follow, even time-varying ones. Specifically, we prove the following theorem.

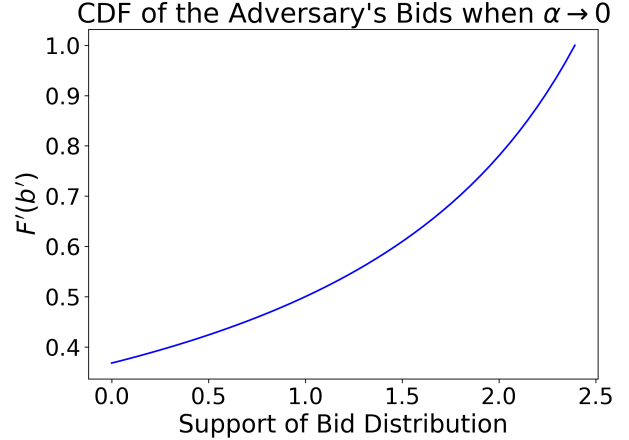


Figure 1: CDF of the adversary’s bid distribution used in Theorem 3 when $\alpha \rightarrow 0$. The expected bid under this distribution is 1, so the adversary will not run out of budget. The CDF is carefully chosen such that the agent cannot win more than $1 - 1/e$ fraction of the rounds regardless of strategy.

Theorem 3. *There exists a stationary bidding policy by the adversary such that an agent with fair share α and value distribution Bernoulli(α) cannot get expected utility more than a $1 - \frac{1}{e} + \frac{\alpha}{e} + O(\sqrt{\log T/T})$ fraction of her ideal utility, i.e.,*

$$\sum_{t=1}^T \mathbb{E}[U[t]] \leq \alpha T \left(1 - \frac{1}{e} + \frac{\alpha}{e}\right) + O(\sqrt{T \log T}).$$

The bidding policy we use for the theorem’s proof is roughly (for simplicity we assume the adversary’s budget constraint holds only in expectation) the following CDF (depicted also in Fig. 1)

$$F'(b') = \begin{cases} 0 & \text{if } b' < 0 \\ \frac{1-\alpha}{e-(e-2)b'} + \alpha & \text{if } 0 \leq b' \leq \frac{e-1}{e-2} \\ 1 & \text{if } b' > \frac{e-1}{e-2} \end{cases}.$$

A first observation from the above distribution is that the adversary’s expected bid is $1 - \alpha$. This implies that the adversary will not run out of budget in expectation, even if the agent never wins any rounds. The next key observation is that the highest bid the agent can use to not run out of budget in expectation is $b \stackrel{\alpha \rightarrow 0}{=} \frac{e}{e-1}$. In particular, this bid wins with probability $F'(b) = \frac{e-1+\alpha}{e}$, which is the claimed bound.

For the full proof that can be found in the full version of this paper, we present a more careful analysis of what happens when the agent follows *any* bidding strategy. In particular, we show that with high probability, the adversary will not run out of budget. Again with high probability, we show that any sequence of bids by the agent, even if they change according to the past, cannot guarantee more than the claimed fraction of ideal utility.

5 Experimental Evaluation of RRB

In our theoretical results, we investigated the worst-case utility guarantees that we could obtain under arbitrary behavior




Figure 2: Fraction of ideal utility that an agent obtains under differing strategy profiles. We compare the agents’ utility when they all use the previously best-known robust strategy from [Gorokh *et al.*, 2021], labeled Deterministic Robust Bidding, with the agents’ utility when they all use Randomized Robust Bidding. We also plot an agent’s utility when they use Randomized Robust Bidding but the other agents adversarially always bid 1 regardless of their values, labeled Randomized Robust Bidding against Adversary. When all agents use Randomized Robust Bidding, they achieve $\approx 1 - (1 - 1/n)^n$ fraction of their ideal utility, the theoretical maximum for any allocation procedure. When one agent uses Randomized Robust Bidding but the other agents behave adversarially, the agent using Randomized Robust Bidding achieves at least a $2 - \sqrt{2}$ fraction, the guarantee of Theorem 1.

by the other agents, which includes adversarial (and collusive) behavior that may not be realistic. In this section, we experimentally investigate the fraction of ideal utility an agent gets when all agents use robust strategies and show that our proposed strategy performs very well. Specifically, we compare the agents’ utilities under the following two strategies. First, all agents use the deterministic $(1/2 - o(1))$ -robust strategy given by [Gorokh *et al.*, 2021], where each agent bids 2 each time their value is in the top α_i -quantile of their value distribution. Second, all agents use our Randomized Robust Bidding strategy, where each agent bids according to a uniform distribution instead. To also illustrate our theoretical results where the other agents are behaving adversarially, we also run an experiment where one agent is using Randomized Robust Bidding but the other agents adversarially always bid 1 regardless of their values.

We consider the symmetric agent case, where each agent has fair share $\alpha_i = 1/n$. We consider each agent’s value distribution to be Bernoulli($1/n$). For each strategy, we compare the agents’ resulting utility for each number of players

$n \in \{2, 3, \dots, 30\}$. We ran the mechanism for $T = 100000$ time periods 10 times and recorded the average fraction of ideal utility that a particular agent obtained¹. We plot our results in Fig. 2.

We can see that when every agent uses the deterministic strategy, each agent gets $1/2$ of their ideal utility, similar to the theoretical guarantee. When each player plays our randomized strategy, they enjoy a higher utility. In particular, they achieve close to a $1 - (1 - 1/n)^n$ fraction of their ideal utility. This is the best we can hope for: no allocation procedure can guarantee each agent a greater fraction of their ideal utility when they have the aforementioned Bernoulli values. This is superior to our theoretical lower bound of $2 - \sqrt{2} \approx 0.59$ in Theorem 1 and shows the empirical performance of RRB can be even greater than the one under worst-case competition.

Acknowledgements

David X. Lin was supported by NSF grant ECCS-1847393. Giannis Fikioris was 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 was supported in part by AFOSR grant FA9550-23-1-0068, ARO MURI grant W911NF-19-1-0217, and NSF grants ECCS-1847393 and CNS-195599. Éva Tardos was supported in part by AFOSR grant FA9550-23-1-0410, AFOSR grant FA9550-231-0068, and ONR MURI grant N000142412742.

Contribution Statement

David X. Lin and Daniel Hall contributed equally to this work as co-first authors, and their order was determined randomly.

References

- [Avni *et al.*, 2018] Guy Avni, Thomas A Henzinger, and Rasmus Ibsen-Jensen. Infinite-duration poorman-bidding games. In *Web and Internet Economics: 14th International Conference, WINE 2018, Oxford, UK, December 15–17, 2018, Proceedings 14*, pages 21–36. Springer, 2018.
- [Babaioff *et al.*, 2021] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair-share allocations for agents with arbitrary entitlements. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 127–127, 2021.
- [Babaioff *et al.*, 2022] Moshe Babaioff, Tomer Ezra, and Uriel Feige. On best-of-both-worlds fair-share allocations. In *International Conference on Web and Internet Economics*, pages 237–255. Springer, 2022.
- [Balseiro *et al.*, 2019] Santiago R Balseiro, Huseyin Gurkan, and Peng Sun. Multiagent mechanism design without money. *Operations Research*, 67(5):1417–1436, 2019.
- [Banerjee *et al.*, 2023] Siddhartha Banerjee, Giannis Fikioris, and Eva Tardos. Robust pseudo-markets for reusable public resources. In *Proceedings of the 24th*

¹Our code can be found at <https://github.com/davidxlin/repeated-fisher-market-experiments>

- ACM Conference on Economics and Computation*, pages 241–241, 2023.
- [Blanchard and Jaillet, 2024] Moise Blanchard and Patrick Jaillet. Near-optimal mechanisms for resource allocation without monetary transfers. *arXiv preprint arXiv:2408.10066*, 2024.
- [Budish *et al.*, 2017] Eric Budish, Gérard P Cachon, Judd B Kessler, and Abraham Othman. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2):314–336, 2017.
- [Cavallo, 2014] 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 '14, Paris, France, May 5-9, 2014*, pages 1313–1320, Paris, France, 2014. IFAAMAS/ACM.
- [Dawson *et al.*, 2013] Christopher J Dawson, Vincenzo V DiLuoffo, Michael D Kendzierski, and James W Seaman. Reserving services within a cloud computing environment, 2013. US Patent 8,615,584.
- [Fikioris *et al.*, 2023] Giannis Fikioris, Siddhartha Banerjee, and Éva Tardos. Online resource sharing via dynamic max-min fairness: efficiency, robustness and non-stationarity. *arXiv preprint arXiv:2310.08881*, 2023.
- [Gorokh *et al.*, 2017] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. From monetary to non-monetary mechanism design via artificial currencies. In Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin, editors, *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 563–564, Cambridge, MA, USA, 2017. ACM.
- [Gorokh *et al.*, 2021] Artur Gorokh, Siddhartha Banerjee, and Krishnamurthy Iyer. The remarkable robustness of the repeated fisher market. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 562–562, 2021.
- [Guo and Conitzer, 2010] Mingyu Guo and Vincent Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 10-14, 2010, Volume 1-3*, pages 881–888, Toronto, Canada, 2010. IFAAMAS.
- [Jackson and Sonnenschein, 2007] Matthew O Jackson and Hugo F Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007.
- [Lazarus *et al.*, 1999] Andrew J Lazarus, Daniel E Loeb, James G Propp, Walter R Stromquist, and Daniel H Ullman. Combinatorial games under auction play. *Games and Economic Behavior*, 27(2):229–264, 1999.
- [Prendergast, 2022] Canice Prendergast. The allocation of food to food banks. *Journal of Political Economy*, 130(8):1993–2017, 2022.
- [Vasudevan *et al.*, 2016] Ramesh Vasudevan, Anjani Kalyan Prathipati, Pradeep Seetharam, and Gopalan Arun. Customizable model for throttling and prioritizing orders in a cloud environment, 2016. US Patent 9,253,113.
- [Walsh, 2014] Toby Walsh. Allocation in practice. In Carsten Lutz and Michael Thielscher, editors, *KI 2014: Advances in Artificial Intelligence - 37th Annual German Conference on AI, Stuttgart, Germany, September 22-26, 2014. Proceedings*, volume 8736 of *Lecture Notes in Computer Science*, pages 13–24, Stuttgart, Germany, 2014. Springer.