

Dynamic Fair Division with Partial Information*

Gerdus Benadè

Boston University, Boston, MA benade@bu.edu

Daniel Halpern

Harvard University, Cambridge, MA, dhalpern@g.harvard.edu

Alexandros Psomas

Purdue University, West Lafayette, IN, apsomas@cs.purdue.edu

We consider the fundamental problem of fairly and efficiently allocating T indivisible items among n agents with additive preferences. Items become available over a sequence of rounds, and every item must be allocated immediately and irrevocably before the next one arrives. Previous work shows that when the agents' valuations for the items are drawn from known distributions, it is possible (under mild assumptions) to find allocations that are envy-free with high probability and Pareto efficient ex-post. However, this requires that agents accurately report their values to the algorithm, which rarely happens in practice.

We study a *partial-information* setting, where true item values are hidden from the algorithm and it is only possible to elicit ordinal information in the form of a ranking or pairwise comparison relative to prior items. When values are drawn from i.i.d. distributions, or correlated distributions consisting of a shared common value for each item with i.i.d. noise, we give an algorithm that is envy-free and $(1 - \epsilon)$ -welfare-maximizing with high probability. We provide similar guarantees (envy-freeness and a constant approximation to welfare with high probability) even with minimally expressive queries that ask for a comparison to a single previous item. For independent but non-identical agents, we obtain envy-freeness and a constant approximation to Pareto efficiency with high probability. Our results are asymptotically tight. A computational study on synthetic and real-world value distributions shows that envy-freeness and efficiency can be achieved on practical time-horizons.

Key words: Online allocation, fairness, online learning

1. Introduction

Motivated by operations in food rescue services, we consider the following fundamental fair division problem. A set of T indivisible items, arriving one at a time, must be allocated among a set of n agents with additive preferences. The value $v_{i,t}$ that agent i has for the item in round t is realized once the item arrives. Each item is allocated immediately and irrevocably upon arrival, and we ask that the overall allocation is both *fair* and *efficient*.

As fairness measure, we study *envy-freeness*, a prominent notion of fairness which requires that every agent prefers their allocation over the allocation of any other agent. Previous work shows that,

* A preliminary version of this paper appeared at NeurIPS'22

despite the uncertainty about future items, one can achieve simultaneous fairness and efficiency when agents’ values are stochastic. Specifically, when each $v_{i,t}$ is drawn i.i.d. from a distribution D , the simple algorithm that maximizes welfare — each item is allocated to the agent with the highest value — is envy-free with high probability and (obviously) ex-post Pareto efficient (Dickerson et al. 2014, Kurokawa et al. 2016). The same guarantee holds for independent and non-identical agents (where $v_{i,t}$ is drawn from an agent-specific distribution D_i) for the algorithm that maximizes weighted welfare (Bai and Gözl 2022). Even when agents’ valuations for an item are correlated (but items are independent), Pareto efficiency ex-post is compatible with strong fairness guarantees (Zeng and Psomas 2020).

Despite the computational simplicity of (most of) the aforementioned algorithms, an unappealing aspect, especially from a practical perspective, is the requirement that agents report an exact numerical value for each item. There are rare organizations that are able to elicit such fine-grained valuations, for example, Feeding America manages their allocations with a market-based mechanism in which recipients bid daily on available donations (Prendergast 2022). However, in the vast majority of cases, eliciting expressive additive valuations is impractical, e.g. due to low volume, the unpredictability of donations arriving, and agents’ cognitive limitations. For example, Shi et al. (2021) describes that at *412 Food Rescue*, based in Pittsburgh, PA, a dispatcher matches a donation to a recipient on an ad-hoc basis and gives them the opportunity to claim it; *MEANS database*, a nonprofit matching donors to food shelves and soup kitchens in 50 states, announces an available donation and assigns it to the first recipient who expresses interest (MEANS database 2023). In both cases eliciting values is limited to getting a binary signal of interest from a potential recipient, a far cry from knowing their exact value for the item.

In this paper, we study the power and limits of eliciting ordinal information in dynamic fair division. The value $v_{i,t}$ of agent i for item t is drawn from an *unknown* distribution upon arrival. Instead of this value, the algorithm is provided only partial ordinal information about the item, e.g., its rank relative to a subset of the past items allocated to this agent, or even just a pairwise comparison to a single previous item (a binary signal). Does this give up too much in an attempt to simplify elicitation? Or, can we learn the unknown distribution sufficiently accurately to simultaneously guarantee fairness and efficiency?

1.1. Our Contribution

We start by establishing a separation between the cardinal setting and our ordinal one. Pareto efficiency alone is trivial (allocate all goods to the same agent) and, in the cardinal setting, Pareto efficiency ex-post is compatible with envy-freeness with high probability as long as agents are independent. We prove in Theorem 1 that in our setting, even for the case of two i.i.d. agents

and any *known* distribution, envy-freeness with high probability is incompatible with even a very mild notion of exact Pareto efficiency, one-swap-Pareto efficiency, which requires that there is no beneficial one-to-one trade of items between agents but allows for improvements via many-to-many trades of items.

We proceed to give an essentially matching positive result. For any number of i.i.d. agents and an unknown value distribution D , there exists an algorithm (Algorithm 1) that is envy-free with high probability and guarantees a $(1 - \varepsilon)$ -approximation to the optimal utilitarian social welfare (the sum of utilities), for all $\varepsilon > 0$, with high probability (Theorem 2). When an item arrives, the algorithm learns for each agent i its relative rank compared to a subset of prior items allocated to agent i , but otherwise knows nothing about the underlying numerical valuation nor the value distribution. Our algorithm works in epochs. Each epoch has an exploration/sampling phase, where each agent gets a pre-determined number of items, followed by an exploitation/ranking phase, where each fresh item is given to an agent whose empirical quantile is largest. The goal is to make a sublinear number of errors compared to the “ideal” algorithm that allocates each item to the agent with the highest true quantile. The algorithm has to balance the need for sampling, which leads to more accurate empirical quantiles, against the inefficient allocations made while sampling. A significant technical barrier is that we cannot fix a target accuracy with which to approximate values because the underlying distribution is unknown. That is, for every fixed accuracy of the empirical quantiles, there exists a distribution for which this accuracy is not good enough for even a constant approximation to the optimal welfare. Instead, we need to make our epochs progressively longer, thereby guaranteeing progressively better bounds.

Given this strong positive result, we explore the limits of what we can achieve when further restricting the amount of information available. What if each agent can remember only a *single* item previously allocated to them, and the fresh item is compared to just this one item?¹ Surprisingly, the aforementioned positive result can almost be recovered even in this very restrictive setting. We prove that there exists an algorithm (Algorithm 2) that is envy-free with high probability and guarantees a $(1 - 1/e)^2 - \varepsilon$ approximation to the optimal welfare with high probability, for all $\varepsilon > 0$ (Theorem 4). Our algorithm again proceeds in epochs with gradually increasing exploration and exploitation phases; this time the goal is a sublinear number of differences compared to allocating each item to a uniformly random agent with quantile at least $1 - 1/n$, which we prove is envy-free and approximately efficient (Lemmas 1 and 2). When exploring, the algorithm puts a new item in memory, estimates its quantile, and rejects it if not sufficiently close to $1 - 1/n$. We need to sample enough to ensure high confidence in the estimated quantile, but also account for the

¹ So, the algorithm only learns if the new item is better or worse than the item in memory and may, at that time, choose to replace the item in memory.

additional sampling since an item’s quantile value might be far from $1 - 1/n$ to begin with; several technical details need to be accounted for. We give a near-matching lower bound: no algorithm with a memory of one item can achieve a 0.999-approximation to the social welfare with high probability; therefore a constant approximation like Algorithm 2 is all we can hope for.

Next, we relax the i.i.d. assumption and show that Algorithms 1 and 2 are still effective when agents are correlated or non-identical. First, we consider agents that agree on a noisy common valuation of each item, so $v_{it} = v_t + \epsilon_{it}$ for $v_t \sim D^{com}, v_{it} \sim D^{noise}$. Algorithm 1 (with slightly longer epochs) is enough to guarantee envy-freeness and a $1 - \epsilon$ approximation to welfare with high probability (Theorem 5). Second, when each agent i ’s values are drawn from an unknown distribution D_i , we show that it is impossible to get a $\frac{1+\sqrt{5}}{4} \approx .809$ approximation to Pareto efficiency with probability more than $2/3$, even for two agents and unbounded memory (Theorem 6). At the same time, Algorithms 1 and 2 are envy-free and $1/e$ approximately Pareto efficient with high probability! Note that, though both algorithms give the same formal guarantees for non-identical agents and Algorithm 2 elicits strictly less information, one might still prefer to use Algorithm 1 since it has significantly shorter exploration phases.

We conclude with a computational study on both synthetic value distributions and values learned from real data with a collaborative filtering algorithm. Though our theoretical results only guarantee that Algorithm 1 is attractive on an infinite time horizon, we find that on the vast majority of instances we evaluate, after 1000 items, the allocations are envy-free and provide more than 90% of the optimal welfare achievable with full information. Our results show that more correlated and skewed distributions are harder to learn. Finally, we consider variants of Algorithm 1 aimed at smoothing the relatively long periods of poor performance during sampling phases — these perform essentially as well as Algorithm 1, as long as the structure of resetting epochs is retained.

1.2. Related Work

Motivated by the reality that eliciting cardinal valuations is often impractical and prone to errors, a growing body of work in AI studies what can be achieved by algorithms that only elicit preferences of limited expressiveness. This idea originates from Procaccia and Rosenschein (2006), who defined the notion of distortion to measure the worst-case deterioration of an aggregate cardinal objective (e.g., utilitarian social welfare) due to only having access ordinal information. Recent works prove bounds on the distortion for many problems in social choice, including matching (Filos-Ratsikas et al. 2014), voting (Caragiannis et al. 2017, Goel et al. 2017, Mandal et al. 2020, Munagala and Wang 2019, Kempe 2020, Gkatzelis et al. 2020), and participatory budgeting (Benade et al. 2021); see (Anshelevich et al. 2021) for a recent survey. Beyond ordinal inputs, identical elicitation concerns inspired the study of abstractions, consisting of partial or coarsened information, for computing

market equilibria in Kroer et al. (2022). We are motivated by the same elicitation constraints but where distortion measures the *worst-case* loss over all instances, we assume values are stochastic, as a result, we can guarantee multiple attractive properties simultaneously with high probability.

Several papers study fair division in static settings under ordinal preferences, e.g., (Aziz et al. 2015, Bouveret et al. 2010, Baumeister et al. 2017, Nguyen et al. 2017), but often these models do not assume an underlying cardinal model and work directly on the ordinal preferences. Amanatidis et al. (2016) assume underlying cardinal information and, among other results, bound the approximation ratio of truthful mechanisms that elicit rankings. Closer to our work, Halpern and Shah (2021) study rules that have access to the ranking of the top- k items of each agent and bound the ratio of the social welfare of the allocation returned by a rule in the worst case. They also characterize the value of k needed to achieve prominent notions of fairness, namely envy-freeness up to one item (EF1) and approximate maximin share guarantee (MMS), and bound the loss in efficiency incurred due to fairness constraints in this setting.

Our work contributes to the growing literature in dynamic fair division (Kash et al. 2014, Aleksandrov et al. 2015, Friedman et al. 2015, 2017, Benade et al. 2018, He et al. 2019, Zeng and Psomas 2020, Gkatzelis et al. 2021, Barman et al. 2022, Gorokh et al. 2021, Vardi et al. 2022) and we note that the welfare-maximizing algorithms of Dickerson et al. (2014), Kurokawa et al. (2016), Bai and Gözl (2022) work in the dynamic setting, even though their settings are not explicitly dynamic. Bogomolnaia et al. (2022) study proportionality and envy-freeness and characterize undominated allocation rules for both goods and bads in a model which can be interpreted as online with potentially correlated stochastic valuations from unknown distributions, with additional access to the mean of each distribution. We make much stronger assumptions about valuations (i.e., they are either independent or correlated in specific way) but also have access to less information about the arriving item, Bogomolnaia et al. (2022) observe the vector of values in addition to the distribution means while we observe only ordinal information. Beyond stochastic valuations, (Benade et al. 2018) show that it is possible to achieve sublinear envy by randomly allocating every item when agents' valuations are adversarially generated (and this is optimal); however, sublinear envy is incompatible with non-trivial efficiency even in the cardinal setting (Zeng and Psomas 2020). To the best of our knowledge, we are the first to study imperfect expressivity in dynamic fair division.

Further afield, our paper is related to the vast literature on online learning (surveyed in Hoi et al. (2021)). In a classical setting, there are T days and on each day the algorithm follows the advice of one of n experts. The algorithm receives reward equal to value of the expert chosen of that day (in the full feedback variant), and objective is to minimize the difference in reward between the algorithm and the best expert in hindsight. In contrast, we allocate items to agents without knowing their values and minimize the difference in bundle values (envy). There are several

variants of online learning with partial information (or bandit algorithms) (see. e.g., [Cesa-Bianchi and Lugosi \(2006\)](#)) but we are not aware of technical connections. Our setting, where hidden values are drawn from unknown distributions, also reminds of prior-independent auctions ([Devanur et al. 2011](#)), where the task is to design mechanisms that perform well in the worst case even compared to the tailor-made mechanism which knows the distributions.

2. Preliminaries

A set of T indivisible items/goods, labeled by $\mathcal{G} = \{1, 2, \dots, T\}$, needs to be allocated to a set of n agents, labeled by $\mathcal{N} = \{1, \dots, n\}$. Agent $i \in \mathcal{N}$ assigns a value $v_{i,t} \in [0, 1]$ to item $t \in \mathcal{G}$. We assume agents have *additive* valuation functions, so $v_i(S) = \sum_{t \in S} v_{i,t}$ for $S \subseteq \mathcal{G}$. An allocation A is a partition of the items into bundles A_1, \dots, A_n , where A_i is the set of items assigned to agent $i \in \mathcal{N}$. Each allocation has an associated utility profile $v(A) = (v_1(A_1), \dots, v_n(A_n))$.

Items arrive online, one per round. The agents' valuations for the item in round t (the t -th item) are realized when the item arrives. Every item is allocated immediately and irrevocably before moving on to the next round. We write $\mathcal{G}^t = \{1, 2, \dots, t\}$ for the set of items that arrived in the first t rounds, and A_i^t for the allocation of agent i after the t -th item was allocated.

We consider three different models which specify how values are generated. In the **i.i.d. model**, the value of agent i for item t is independently drawn from an *unknown* distribution D with CDF F , i.e., $v_{i,t} \sim D$ for all $i \in \mathcal{N}$ and $t \in \mathcal{G}$. In the **correlated model**, the value of agent i for item t is $v_{i,t} = v_t^{com} + \varepsilon_{i,t}$, where $v_t^{com} \sim D^{com}$ is a common value drawn from an unknown value distribution with CDF F^{com} , and each agent draws independent noise $\varepsilon_{i,t}$ from an unknown noise distribution D^{noise} . For given item, agent values are now correlated, though they are still independent over time. In the **non-i.i.d. model**, the value of item t for agent i is independently drawn from an unknown, agent-dependent distribution D_i with CDF F_i , i.e., $v_{i,t} \sim D_i$ for all $i \in \mathcal{N}$ and $t \in \mathcal{G}$.

We write X_i for the random variable for i 's valuation, and $X_{i,t}$ for the random variable for i 's valuation for item t . It is often convenient to work directly with the quantile of an agent's value rather than the value itself; let $Q_i = F_i(X_i)$ and $Q_{i,t} = F_i(X_{i,t})$ respectively be the random variable denoting the quantile of agent i the associated item. Note that all Q_i and $Q_{i,t}$ are i.i.d. and follow a $\text{Unif}[0, 1]$ distribution. Unless explicitly stated otherwise, we assume all distributions are continuous (i.e., do not have point masses).

Ordinal Information. We assume the realizations $v_{i,t}$ are not available. Instead, our algorithms have access to *ordinal* information. Specifically, given current item t , the algorithm can access each agent's *ranking* for $S = \{t\} \cup M$, $M \subseteq \mathcal{G}^{t-1}$. The size of M , which we will informally refer to as the *memory size*, determines the complexity of eliciting information from each agent. In one extreme, agent i compares a new item t to a single item they had previously received, i.e., $M \subseteq A_i^{t-1}$, $|M| \leq 1$.

In the other extreme, t is compared to all previous items she received, so $M = A_i^{t-1}$. We write $\sigma_i(S)$ for the ranking of agent i for a subset S of the items, and $\sigma_i^{-1}(S, t)$ for the position of item $t \in S$ with respect to a subset S according to agent i . The highest value item is in position 1 and the lowest in position $|S|$. For example, if $S = \{1, 4\}$, $v_{i,1} = 1$ and $v_{i,4} = 0.1$, $\sigma_i(S) = (1 \succ 4)$, $\sigma_i^{-1}(S, 1) = 1$ and $\sigma_i^{-1}(S, 4) = 2$.

Algorithms. An algorithm \mathcal{A} , in each step t , queries each agent for ordinal information with respect to some subset M and then makes a (possibly randomized) allocation decision based on this ordinal information and the history so far. An instance of our problem is parameterized by the number of agents n and the (unknown) value distributions D_1, \dots, D_n . Let $\mathcal{E}_P(t)$ be the event that some algorithm satisfies property P (e.g., envy-freeness or PO or ε -welfare) at time t . We are interested in the probability that an algorithm satisfies certain properties in the limit, as the number of rounds tends to infinity, where the randomness is over the random choices of the algorithm as well as the randomness in the valuations.

DEFINITION 1. An algorithm satisfies P with high probability if $\lim_{t \rightarrow \infty} \Pr[\mathcal{E}_P(t)] = 1$.

Note that this definition of high probability allows for dependency on n and the underlying distributions (i.e., they are treated as constants).

Efficiency notions. An allocation A Pareto dominates an allocation A' , denoted $A \succ A'$, when $v_i(A_i) \geq v_i(A'_i)$ for all $i \in \mathcal{N}$ and there exists $j \in \mathcal{N}$ with $v_j(A_j) > v_j(A'_j)$. An allocation A is *Pareto efficient* or *Pareto optimal* (PO) if there is no feasible (integral) allocation that Pareto dominates it. An allocation A' is in the (one) swap-neighborhood of A when it can be created from A with a single exchange of items between one pair of agents. Formally, there exist $i, j \in \mathcal{N}$ and items $z_j \in A_j$ and $z_i \in A_i$ so that $A'_i = (A_i \setminus \{z_i\}) \cup \{z_j\}$, $A'_j = (A_j \setminus \{z_j\}) \cup \{z_i\}$, and $A'_k = A_k$ for all other agents $k \neq i, j$. An allocation A is *one-swap Pareto optimal* (SPO) if it is undominated by any allocation in its swap-neighborhood. We use a notion of approximate efficiency defined by (Ruhe and Fruhwirth 1990) according to which an allocation A is α -Pareto efficient when $v(A)/\alpha$ is undominated.

The social welfare of an allocation A is $\text{sw}(A) = \sum_{i \in \mathcal{N}} v_i(A_i)$. Let allocation A^* denote a (social) welfare optimal allocation for which $\text{sw}(A^*) \geq \text{sw}(A)$ for all feasible allocations A . An allocation provides an α -approximation to welfare if $\text{sw}(A) \geq \alpha \cdot \text{sw}(A^*)$.

Fairness notions. We focus on a prominent notion of fairness called *envy-freeness*. An allocation $A^T = (A_1^T, \dots, A_n^T)$ of T items is *envy-free* (EF) when $v_i(A_i^T) \geq v_i(A_j^T)$ for all $i, j \in \mathcal{N}$, and *c-strongly-envy-free* (*c-strong-EF*) when $v_i(A_i) \geq v_i(V_j) + cT$.

3. Ideal Quantile-based Algorithms.

For our analysis, it will be useful to compare our algorithms with ideal algorithms that know *exact* quantile values for every item (and, in fact, several of our lower bounds apply to these stronger

algorithms too). Two ideal algorithms of interest are (1) quantile maximization, which allocates each item to the agent with the highest quantile value for it, and (2) “ q -threshold,” which allocates each item uniformly at random among agents whose quantile is at least q (and uniformly at random over all agents, if all quantile values are less than q).

In the i.i.d. model, quantile maximization is the same as value maximization, and thus envy-free with high probability and ex-post welfare optimal. The property we will use is c -strong envy-freeness, for some distribution-dependent constant c , which we state as Lemma 1. This was essentially proved by Dickerson et al. (2014); we provide an alternate proof that also works, essentially unchanged, for the $\frac{n-1}{n}$ -threshold algorithm; it can be found in Section EC.1.1.

LEMMA 1. [Essentially Dickerson et al. (2014).] *In the i.i.d. and non-i.i.d. models, the quantile maximization algorithm and the $\frac{n-1}{n}$ -threshold algorithm are c -strongly-envy-free, with high probability, where the constant $c = \min_{i \in \mathcal{N}} (\mathbb{E}[X_i \mid Q_i \geq 1/2] - \mathbb{E}[X_i]) / (4n)$.*

Note that c is strictly positive since our distributions are continuous. In the i.i.d. model, we show that the $\frac{n-1}{n}$ -threshold algorithm gives a $(1 - \frac{1}{e})^2 - \varepsilon$ approximation to welfare (Lemma 2) with high probability.

LEMMA 2. *In the i.i.d. model, the $\frac{n-1}{n}$ -threshold algorithm guarantees a $((1 - \frac{1}{e})^2 - \varepsilon)$ -approximation to welfare, with high probability, for all $\varepsilon > 0$.*

Proof. Let F be the CDF of an arbitrary continuous distribution. The expected contribution of an item to the welfare of the threshold algorithm is at least

$$\mathbb{E}_{Q \sim \text{Unif}[0,1]} \left[F^{-1}(Q) \mid Q \geq \frac{n-1}{n} \right] \cdot \Pr_{\vec{Q} \sim \text{Unif}[0,1]^n} \left[\max_{i \in \mathcal{N}} Q_i \geq \frac{n-1}{n} \right].$$

For the first term we have

$$\begin{aligned} \mathbb{E}_{Q \sim \text{Unif}[0,1]} \left[F^{-1}(Q) \mid Q \geq \frac{n-1}{n} \right] &= \mathbb{E}_{Q \sim \text{Unif}[\frac{n-1}{n}, 1]} [F^{-1}(Q)] \\ &= \left(\int_{\frac{n-1}{n}}^1 F^{-1}(q) \cdot f_{\text{Unif}[\frac{n-1}{n}, 1]}(q) dq \right) \\ &= \left(\int_{\frac{n-1}{n}}^1 F^{-1}(q) \cdot n dq \right) \\ &\geq (f_{\text{Beta}[n,1]}(x) = nx^{n-1}) \left(\int_{\frac{n-1}{n}}^1 F^{-1}(q) \cdot f_{\text{Beta}[n,1]}(q) dq \right) \\ &= \mathbb{E}_{Q \sim \text{Beta}[n,1]} \left[F^{-1}(Q) \mid Q \geq \frac{n-1}{n} \right] \cdot \Pr_{Q \sim \text{Beta}[n,1]} \left[Q \geq \frac{n-1}{n} \right] \\ &\geq \mathbb{E}_{Q \sim \text{Beta}[n,1]} [F^{-1}(Q)] \cdot \Pr_{Q \sim \text{Beta}[n,1]} \left[Q \geq \frac{n-1}{n} \right] \end{aligned}$$

$$= \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} \left[F^{-1}(\max_{i \in \mathcal{N}} Q_i) \right] \cdot \Pr_{\vec{Q} \sim \text{Unif}[0,1]^n} \left[\max_{i \in \mathcal{N}} Q_i \geq \frac{n-1}{n} \right],$$

where we used the fact that the maximum of n draws from $U[0, 1]$ follows a $\text{Beta}(n, 1)$. The expected contribution to the welfare is thus at least

$$\mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} \left[F^{-1}(\max_{i \in \mathcal{N}} Q_i) \right] \left(1 - \left(1 - \frac{1}{n} \right)^n \right)^2 \geq \left(1 - \frac{1}{e} \right)^2 \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} \left[F^{-1}(\max_{i \in \mathcal{N}} Q_i) \right]$$

Finally, for any fixed $\varepsilon > 0$, standard Chernoff bounds tell us that with high probability, the optimal welfare of T items is at most $T \cdot (1 + \varepsilon/2) \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} [F^{-1}(\max_{i \in \mathcal{N}} Q_i)]$ while the welfare of the threshold algorithm is at least $T \cdot (1 - \varepsilon/2) \left(1 - \frac{1}{e} \right)^2 \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} [F^{-1}(\max_{i \in \mathcal{N}} Q_i)]$. Indeed, the expected optimal welfare is equal to $T \cdot \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} [F^{-1}(\max_{i \in \mathcal{N}} Q_i)]$, the sum of T i.i.d. random variables with expectation $\mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} [F^{-1}(\max_{i \in \mathcal{N}} Q_i)]$. The standard multiplicative Chernoff bound says that the sum of i.i.d. variables exceeds $(1 + \varepsilon/2)$ times its expectation μ is at most $\exp(-\mu\varepsilon^2/12)$. Plugging in $\mu = T \cdot \mathbb{E}_{\vec{Q} \sim \text{Unif}[0,1]^n} [F^{-1}(\max_{i \in \mathcal{N}} Q_i)]$, we get the desired statement. The statement about the welfare of the threshold algorithm follows similarly. Thus, the algorithm is a

$$\left(1 - \frac{1}{e} \right)^2 \cdot (1 - \varepsilon/2)/(1 + \varepsilon/2) \geq \left(1 - \frac{1}{e} \right)^2 (1 - \varepsilon) \geq \left(1 - \frac{1}{e} \right)^2 - \varepsilon$$

approximation to welfare, with high probability. \square

Next, we prove that both ideal algorithms are approximately efficient. Let \mathcal{P}^* be the following property of an allocation: all items such that exactly one agent has quantile values at least $1 - 1/n$ are in the bundle of this agent. Both ideal algorithms (quantile maximization and $1 - 1/n$ -threshold) satisfy \mathcal{P}^* . We prove that, in the non-i.i.d. model, \mathcal{P}^* implies an almost $1/e$ approximation to efficiency. Our proof uses the fact that there is a (roughly) $1/e$ probability that exactly one agent has the high quantile, so the value of an agent's bundle in an algorithm that satisfies \mathcal{P}^* is, with high probability, a $1/e$ approximation to their value for their T/n most valuable items. Therefore, when considering an alternate allocation A' , the agent in A' that gets at most T/n items cannot be improved upon by more than a $1/e$ factor.

LEMMA 3. *In the non-i.i.d. model, every algorithm whose allocations satisfy \mathcal{P}^* is $(1/e - \varepsilon)$ -Pareto optimal, with high probability, for all $\varepsilon > 0$.*

Proof. Fix an $\varepsilon \in (0, 1)$, and choose ε' such that $\frac{1-\varepsilon'}{(1+\varepsilon')^2} \cdot \frac{1}{e} > \frac{1}{e} - \varepsilon$ (using $\varepsilon' = \varepsilon/3$ will do). Fix distributions with CDFs F_1, \dots, F_n for each agent $i \in \mathcal{N}$, and a time T . Suppressing the superscript, for ease of notation, let $A_i = A_i^T$ be the bundle allocated at time T to each agent i by an algorithm that satisfies \mathcal{P}^* . Let A_i^{top} be the set of the T/n most valuable items for each agent i . Let $A_i^{\text{high}} = \{t \in \mathcal{G}^T \mid F_i(v_{i,t}) \geq 1 - \frac{1+\varepsilon'}{n}\}$ be the set of items that agent i has “high” value for, in the sense that they come from the top $\frac{1+\varepsilon'}{n}$ portion of their distribution. We show the following $3n$ events, \mathcal{E}_{ij} for $i \in \mathcal{N}$ and $j \in \{1, 2, 3\}$, occur simultaneously with high probability (in T).

1. \mathcal{E}_{i1} : $v_i(A_i^{\text{top}}) \leq v_i(A_i^{\text{high}})$.
2. \mathcal{E}_{i2} : $v_i(A_i^{\text{high}}) \leq T \cdot \frac{(1+\varepsilon')^2}{n} \mathbb{E}_{Q \sim \text{Unif}[1-1/n, 1]}[F^{-1}(Q)]$.
3. \mathcal{E}_{i3} : $v_i(A_i) \geq T \cdot \frac{1-\varepsilon'}{en} \mathbb{E}_{Q \sim \text{Unif}[1-1/n, 1]}[F^{-1}(Q)]$.

Each of these individually will follow from a straightforward application of Hoeffding's inequality or Chernoff bounds, showing they each individually occur with probability exponentially close to 1 in T . This implies that they all occur simultaneously with high probability. Finally, we will show that conditioned on all $3n$ occurring, the allocation is $(1/e - \varepsilon)$ -PO.

Let us begin with \mathcal{E}_{i1} for each agent i . The event occurs when there are at least T/n items $t \in \mathcal{G}^T$ such that $F_i(v_{i,t}) \geq 1 - \frac{1+\varepsilon'}{n}$. Each item independently satisfies this property ($F_i(v_{i,t}) \geq 1 - \frac{1+\varepsilon'}{n}$) with probability $\frac{1+\varepsilon'}{n}$. Hence the probability this does not occur is at most $2 \exp(-2\varepsilon'^2 T)$.

Next, consider \mathcal{E}_{i2} for each agent i . The expected contribution of each item to $v_i(A_i^{\text{high}})$ is

$$\begin{aligned} \mathbb{E}_{Q \sim \text{Unif}[0, 1]} \left[F_i^{-1}(Q) \cdot \mathbb{I} \left[Q \geq 1 - \frac{1+\varepsilon'}{n} \right] \right] &= \frac{1+\varepsilon'}{n} \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1+\varepsilon'}{n}, 1]} [F_i^{-1}(Q)] \\ &\leq \frac{1+\varepsilon'}{n} \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]} [F_i^{-1}(Q)]. \end{aligned}$$

We now use the following multiplicative version of the Chernoff bound,

$$\Pr \left[\sum_i X_i \geq (1+\delta) \sum_i \mathbb{E}[X_i] \right] \leq \exp \left(-\frac{\delta^2}{3} \sum_i \mathbb{E}[X_i] \right),$$

to conclude that the probability that $v_i(A_i^{\text{high}})$ exceeds $T \cdot \frac{(1+\varepsilon')^2}{n} \mathbb{E}_{Q \sim \text{Unif}[1-1/n, 1]}[F^{-1}(Q)] \geq (1+\varepsilon') \cdot \mathbb{E}[v_i(A_i^{\text{high}})]$ is at most $\exp \left(-\frac{\varepsilon'^2 (1+\varepsilon') \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)]}{3n} \cdot T \right)$.

Finally, consider \mathcal{E}_{i3} for each agent i . We will show that the expected contribution of each item to $v_i(A_i)$ is at least $\frac{1}{en} \cdot \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)]$. Indeed, consider an item such that the quantile for agent i is $Q_i > 1 - 1/n$ while $Q_j < 1 - 1/n$ for all agents $j \neq i$. This occurs with probability $\frac{1}{n} \cdot (1 - \frac{1}{n})^{n-1} \geq \frac{1}{en}$, and when this occurs, since the algorithm satisfies \mathcal{P}^* , it must allocate the item to i . Further, when this does occur, the expected value of such an item is $\mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)]$, since it is independent of the other agent's values. Hence the expectation is at least $\frac{1}{en} \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)]$. Finally, we again use a multiplicative Chernoff bound to show that

$$\Pr \left[v_i(A_i) \leq (1-\varepsilon') \cdot \frac{T}{en} \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)] \right] \leq \exp \left(-\frac{\varepsilon'^2 \mathbb{E}_{Q \sim \text{Unif}[1-\frac{1}{n}, 1]}[F_i^{-1}(Q)]}{2en} \cdot T \right).$$

Now, suppose that \mathcal{E}_{ij} hold for all $i \in \mathcal{N}$ and $j \in \{1, 2, 3\}$. We show that this implies the allocation A_1, \dots, A_n is $(1/e - \varepsilon)$ -PO. Fix an arbitrary allocation A'_1, \dots, A'_n . We show there exists an agent $i \in \mathcal{N}$ such that $v_i(A'_i) < \frac{v_i(A_i)}{1/e - \varepsilon}$. First, there must be some agent i such that $|A'_i| \leq T/n$. Since A'_i can be at most as valuable as the most-valuable T/n items, we have

$$v_i(A'_i) \leq v_i(A_i^{\text{top}})$$

$$\begin{aligned}
&\leq^{(\mathcal{E}_{i1})} v_i(A_i^{\text{high}}) \\
&\leq^{(\mathcal{E}_{i2})} T \cdot \frac{(1 + \varepsilon')^2}{n} \mathbb{E}_{Q \sim \text{Unif}[1-1/n, 1]}[F^{-1}(Q)] \\
&\leq^{(\mathcal{E}_{i3})} \frac{(1 + \varepsilon')^2}{(1 - \varepsilon')(1/e)} v_i(A_i) \\
&< \frac{1}{1/e - \varepsilon} v_i(A_i),
\end{aligned}$$

as needed. \square

4. Unbounded memory in the i.i.d. model

We explore some fundamental limits of our setting. Efficiency by itself is easy: allocate all items to the same agent. However, in contrast to the cardinal setting, we find one-swap Pareto efficiency is incompatible with envy-freeness with high probability, even for two i.i.d. agents, and even when the underlying distribution is *known*.

THEOREM 1. *In the i.i.d. model, even for $n = 2$ agents, there does not exist an algorithm \mathcal{A} which is one-swap Pareto efficient and envy-free with high probability, even when values are sampled according to D , for any continuous, bounded and known value distribution D .*

Proof. Fix an arbitrary, continuous value distribution D and an algorithm \mathcal{A} .

As the agents are a priori identical, we can assume without loss of generality that \mathcal{A} gives the first item to agent 1. We will show that, with a positive probability, this decision becomes an irrevocable “mistake,” in the sense that agent 2 really liked the item and agent 1 did not. This mistake will make envy-freeness and one-swap PO incompatible.

First, we find values to make this mistake sufficiently bad. Let $g : [0, 1] \rightarrow [0, 1]$ be the function $g(q) = \mathbb{E}[X \mid X \leq F^{-1}(q)] / \mathbb{E}[X]$, which maps a quantile q to the ratio of the expected value of an item below quantile q to the expected value of an arbitrary item. g is a continuous increasing function with $g(1) = 1$, so there is some quantile $\hat{q} < 1$ such that $g(\hat{q}) \geq 0.9$. Let $q_2^* = \max(\hat{q}, 0.9)$. Since g is increasing, $g(q_2^*) \geq g(\hat{q}) \geq 0.9$. Let $q_1^* = 0.1$, $v_1^* = F^{-1}(q_1^*)$ and $v_2^* = F^{-1}(q_2^*)$. Let $\mathcal{E}^{\text{mistake}}$ be the event that $X_{1,1} < v_1^*$ and $X_{2,1} > v_2^*$. Define $c := \Pr[\mathcal{E}^{\text{mistake}}] = (1 - q_2^*) \cdot q_1^*$ to be the probability that $\mathcal{E}^{\text{mistake}}$ occurs. D is continuous, so $c > 0$. Our lower bound on the probability that the allocation at step t violates either envy-freeness or one-swap PO will only depend on c .

Let \mathcal{E}_j be the event that for item j we have that both $X_{1,j} \geq v_1^*$ and $X_{2,j} \leq v_2^*$. If $\mathcal{E}^{\text{mistake}}$ occurs, the only way to maintain one-swap Pareto efficiency is to allocate item j to agent 1 every time \mathcal{E}_j occurs; otherwise, swapping items 1 and j between the two agents yields a Pareto improvement. This constraint will make envy-freeness unlikely.

Let $\mathcal{E}^{\text{manyhigh}}(t)$ be the event $\sum_{j=2}^t X_{2,j} \cdot \mathbb{I}[\mathcal{E}_j] \geq (t - 1) \cdot 0.7 \cdot \mathbb{E}[X]$. In other words, $\mathcal{E}^{\text{manyhigh}}(t)$ occurs when agent 2 has a high value for items j , $2 \leq j \leq t$, for which \mathcal{E}_j occurs (i.e., the items

that must be given to agent 1 in order to satisfy one-swap PO). Let $\mathcal{E}^{\text{normalval}}(t)$ denote the event that $\sum_{j=2}^t X_{2,j} \leq (t-1) \cdot 1.1 \cdot \mathbb{E}[X]$. We first show that for sufficiently large t , the probability that both $\mathcal{E}^{\text{manyhigh}}(t)$ and $\mathcal{E}^{\text{normalval}}(t)$ occur is at least $1/2$. To do so, we prove each event occurs with probability at least $3/4$, and then apply a union bound.

First, since each $X_{1,j}$ and $X_{2,j}$ are independent, $\Pr[\mathcal{E}_j] \geq 0.9 \cdot 0.9 = 0.81$, and $\mathbb{E}[X_{2,j} | \mathcal{E}_j] = \mathbb{E}[X_{2,j} | X_{2,j} \leq v_2^*]$. Also, from the definition of $g(\hat{q})$ and the choice of q_2^* , $\mathbb{E}[X_{2,j} | X_{2,j} \leq v_2^*] \geq 0.9 \cdot \mathbb{E}[X]$. It follows that $\mathbb{E}[X_{2,j} \cdot \mathbb{I}[\mathcal{E}_j]] = \mathbb{E}[X_{2,j} | \mathcal{E}_j] \cdot \Pr[\mathcal{E}_j] \geq 0.729 \cdot \mathbb{E}[X]$. A straightforward Chernoff bound establishes that $\Pr[\mathcal{E}^{\text{manyhigh}}(t)] \geq 3/4$ for t at least $\frac{6}{\mathbb{E}[X]}$.

Let $Y_j = X_{2,j} \cdot \mathbb{I}[\mathcal{E}_j]$ for all j . Then, $\mathbb{E}[Y_j] \geq 0.729 \cdot \mathbb{E}[X]$, and $\mathbb{E}[\sum_{j=2}^t Y_j] \geq (t-1) \cdot 0.729 \cdot \mathbb{E}[X]$. We are interested in the probability that $\sum_{j=2}^t Y_j$ is at least $(t-1) \cdot 0.7 \cdot \mathbb{E}[X]$, i.e., the probability that $\sum_{j=2}^t Y_j$ is at least $\frac{0.7}{0.729}$ its expectation.

We use the following Chernoff bound: Let Y_1, \dots, Y_n be independent random variables that take values in $[0, 1]$, and let Y be their sum. Then, for all $\delta \in [0, 1]$, $\Pr[Y \leq (1 - \delta) \mathbb{E}[Y]] \leq e^{-\frac{\mathbb{E}[Y]\delta^2}{2}}$.

Continuing our derivation:

$$\begin{aligned} \Pr\left[\sum_{j=2}^t Y_j \geq (t-1) \cdot 0.7 \cdot \mathbb{E}[X]\right] &= \Pr\left[\sum_{j=2}^t Y_j \geq \frac{0.7}{0.79} \mathbb{E}\left[\sum_{j=2}^t Y_j\right]\right] \\ &= 1 - \Pr\left[\sum_{j=2}^t Y_j < \frac{0.7}{0.79} \mathbb{E}\left[\sum_{j=2}^t Y_j\right]\right] \\ &\geq 1 - \Pr\left[\sum_{j=2}^t Y_j \leq 0.89 \mathbb{E}\left[\sum_{j=2}^t Y_j\right]\right] \\ &\geq 1 - \exp\left(-\frac{\mathbb{E}[\sum_{j=2}^t Y_j](0.89)^2}{2}\right), \end{aligned}$$

which is at least $3/4$ when $\frac{\mathbb{E}[\sum_{j=2}^t Y_j](0.89)^2}{2}$ is at least $\ln(4)$, or, equivalently, if $t \geq 1 + \frac{2\ln(4)}{0.7 \cdot (0.89)^2 \cdot \mathbb{E}[X]}$. Since $\frac{2\ln(4)}{0.7 \cdot (0.89)^2} < 5$ and $\mathbb{E}[X] < 1$, so $t \geq \frac{6}{\mathbb{E}[X]}$ suffices. $\Pr[\mathcal{E}^{\text{normalval}}(t)] \geq 3/4$ follows similarly.

Next, observe that $\mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)$ is independent of $\mathcal{E}^{\text{mistake}}$, since the two events depend on disjoint sets of independent random variables. Therefore, $\Pr[\mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] = \Pr[\mathcal{E}^{\text{mistake}}] \cdot \Pr[\mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] \geq c \cdot 1/2$ for $t \geq 6/\mathbb{E}[X]$.

Let $\mathcal{E}_{\text{SPO}}(t)$ and $\mathcal{E}_{\text{EF}}(t)$ be the events that the allocation at step t is one-swap PO, and envy-free, respectively. When $\mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)$ occur, the allocation cannot be both one-swap PO and envy-free, i.e. $\Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cap \mathcal{E}_{\text{EF}}(t)} | \mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] = 1$. To see this, notice that first, due to $\mathcal{E}^{\text{mistake}}$, the only way to remain one-swap PO is to give each item j to agent 1 every time \mathcal{E}_j occurs. Second, $\mathcal{E}^{\text{manyhigh}}(t)$ ensures that agent 2's value for these items, and hence agent 2's value for agent 1's bundle, is at least $0.7 \cdot (t-1) \cdot \mathbb{E}[X] + v_{2,1}$. Third, $\mathcal{E}^{\text{normalval}}(t)$ ensures that agent 2's value for all items is at most $1.1 \cdot (t-1) \cdot \mathbb{E}[X] + v_{2,1}$, which is strictly less

than twice her value for agent 1's bundle. We conclude that the allocation at step t cannot be proportional, and is hence not envy-free. Overall, we have that

$$\begin{aligned}
\Pr[\overline{\mathcal{E}_{\text{SPO}}(t)}] + \Pr[\overline{\mathcal{E}_{\text{EF}}(t)}] &\geq \Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cup \mathcal{E}_{\text{EF}}(t)}] \\
&= \Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cap \mathcal{E}_{\text{EF}}(t)}] \\
&\geq \Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cap \mathcal{E}_{\text{EF}}(t)} \cap \mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] \\
&= \Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cap \mathcal{E}_{\text{EF}}(t)} \mid \mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] \\
&\quad \cdot \Pr[\mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] \\
&\geq c/2.
\end{aligned}$$

Therefore, for $t \geq 6/\mathbb{E}[X]$, at least one of $\Pr[\overline{\mathcal{E}_{\text{SPO}}(t)}]$ and $\Pr[\overline{\mathcal{E}_{\text{EF}}(t)}]$ is at least $c/4$. We conclude that no algorithm can be both envy-free and one-swap PO with high probability. \square

Theorem 1 implies that when we have access to only ordinal information, we need to settle for *some* approximation to envy-freeness and efficiency. Our main positive result for this section is an algorithm that essentially matches the aforementioned lower bound.

THEOREM 2. *In the i.i.d. model, Algorithm 1 achieves envy-freeness and a $(1 - \varepsilon)$ approximation to welfare, with high probability, for all $\varepsilon > 0$.*

Algorithm 1 works in epochs: each epoch k has an exploration/sampling phase, where each agent i receives a pre-determined set of items, denoted G_i^k , irrespective of their valuation. This is followed by an exploitation/ranking phase, where each item is given to the agent with the highest empirical quantile (with respect to items received in the preceding exploration phase, i.e. G_i^k).

Algorithm 1: EF + $(1 - \varepsilon)$ -Welfare

```

for epoch  $k = 1 \dots$  do
    Sampling Phase: ( $n \cdot k^4$  items)
        Give the  $j$ -th item in this phase to agent  $j(\bmod n)$ .
    Ranking Phase: ( $k^8$  items)
        for each item  $g$  in this phase do
            Elicit  $\sigma_i^{-1}(G_i^k \cup \{g\}, g)$  for all  $i \in \mathcal{N}$ .
            Allocate  $g$  to an agent  $j \in \arg \min_{i \in \mathcal{N}} \sigma_i^{-1}(G_i^k \cup \{g\}, g)$ .

```

We start with a technical lemma, which gives us a bound on the length of the exploration period we need in each epoch. The following definition will be useful.

DEFINITION 2. A sample of $n \cdot m$ items (where each agent is allocated exactly m items) is ε -accurate if, with probability at least $1 - \varepsilon$, the relative rank of a fresh item (with respect to the sample) is highest for the agent with highest quantile value.

LEMMA 4. If $\varepsilon, \delta \in (0, 1)$, and $m \in \mathbb{Z}^+$ are such that $\varepsilon > 2n\sqrt{\frac{\ln(2n/\delta)}{2m}}$, then giving m samples to each agent is ε -accurate with probability at least $1 - \delta$.

Proof. We will use the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality (Dvoretzky et al. 1956, Massart 1990) to show the empirical CDF of sampled quantiles is reasonably close to a uniform distribution with probability $1 - \delta$. We then show this is sufficient to guarantee ε -accuracy for the chosen ε . Let \hat{F}_i be the empirical CDF of the sampled *quantiles* for agent i , i.e., $\hat{F}_i(q)$ for $q \in [0, 1]$ is a random variable that describes the proportion of sampled items with quantile at most q . Note that \hat{F}_i exactly captures agent i 's ranking for a new item: if a fresh item has quantile q_i for agent i and q_j for agent j , then i ranks it higher than j exactly when $\hat{F}_i(q_i) > \hat{F}_j(q_j)$.

Noting that the CDF for the actual quantile distribution (i.e., the uniform distribution) is the identity on $[0, 1]$, the DKW inequality states that for all $\gamma > 0$, $\Pr\left[\sup_{q \in [0, 1]} |\hat{F}_i(q) - q| > \gamma\right] \leq 2e^{-2m\gamma^2}$. We want this condition to hold for all n agents, simultaneously, with probability at least $1 - \delta$, so we pick γ such that $2e^{-2m\gamma^2} \leq \delta/n$ and apply a union bound; it suffices to choose $\gamma = \sqrt{\frac{\ln(2n/\delta)}{2m}}$.

We now show that the DKW condition ($\sup_{q \in [0, 1]} |\hat{F}_i(q) - q| \leq \gamma$) being satisfied for all agents i is sufficient to guarantee ε -accuracy. Consider sampling quantiles Q_1, \dots, Q_n for a fresh item. Let $i^{\max} \in \arg \max_{i \in \mathcal{N}} Q_i$ be a quantile-maximizing agent (technically a random variable). Our goal is to show that with probability at least $1 - \varepsilon$ (with respect to the samples of Q_1, \dots, Q_n) $\hat{F}_{i^{\max}}(Q_{i^{\max}}) > \hat{F}_j(Q_j)$ for all $j \neq i^{\max}$. This ensures that i^{\max} has the highest empirical rank, and hence receives the item. Let $Q_{(1)}, \dots, Q_{(n)}$ be the respective order statistics. A key observation is that $Q_{(n)} - Q_{(n-1)} \sim \text{Beta}[1, n]$ (Gentle 2019). The PDF of a $\text{Beta}[1, n]$ distribution is $f(x) = nx^{n-1}$ for $x \in [0, 1]$. Since $f(x) \leq n$, $\Pr[Q_{(n)} - Q_{(n-1)} < \rho] < n\rho$ for all $\rho > 0$. Plugging in $\rho = 2\gamma$, we have $\Pr[Q_{(n)} - Q_{(n-1)} \leq 2\gamma] < 2n\gamma$. We will show that as long as $\varepsilon > 2n\gamma$, ε -accuracy holds. First, we have $\Pr[Q_{(n)} - Q_{(n-1)} > 2\gamma] > 1 - \varepsilon$. Conditioned on $Q_{(n)} - Q_{(n-1)} > 2\gamma$, the item is given to i^{\max} . To see why, observe $Q_{i^{\max}} = Q_{(n)}$ and $Q_j \leq Q_{(n-1)}$ for all $j \neq i^{\max}$, by definition. Using the DKW inequality condition, it follows that $\hat{F}_{i^{\max}}(Q_{i^{\max}}) \geq Q_{i^{\max}} - \gamma > Q_j + \gamma \geq \hat{F}_j(Q_j)$. We conclude that for $\varepsilon > 2n\sqrt{\frac{\ln(2n/\delta)}{2m}}$, ε -accuracy is satisfied with probability at least $1 - \delta$. \square

Using Lemma 4, we can get, for each epoch, a bound on the number of decisions where Algorithm 1 differs from the quantile maximization algorithm.

LEMMA 5. The allocation of Algorithm 1 differs from that of the quantile maximization algorithm after T steps by at most $f(T)$ items with high probability, where $f(T) \in O(\text{poly}(n) \cdot T^{15/16})$.

Proof. We start by bounding the accuracy of Algorithm 1 in each epoch k . In epoch k , each agent receives k^4 items during the sampling phase. We claim that the sample in epoch k for $k \geq 3n$ is ε_k -accurate for $\varepsilon_k := 3n/k^{3/2}$ with probability at least $1 - \delta_k$, for $\delta_k := 2n/e^{2k}$. Indeed, first note that by the choice of k , we have that $\varepsilon_k, \delta_k \in (0, 1)$. Hence, we just need to show that these values satisfy the inequality of Lemma 4. We have that

$$\varepsilon_k = \frac{3n}{k^{3/2}} > \frac{2n}{k^{3/2}} = 2n\sqrt{\frac{1}{k^3}} = 2n\sqrt{\frac{\ln(e^{2k})}{2k^4}} = 2n\sqrt{\frac{\ln(2n/\delta_k)}{2k^4}}.$$

Next, fix a time T . Slightly abusing notation, let $k(t) = \min\{K \in \mathbb{N} \mid \sum_{k=1}^K nk^4 + k^8 \geq t\}$ be the function that given an item t returns the epoch item t is in. Notice that $T \geq \sum_{k=1}^{k(T)-1} nk^4 + k^8 \geq (k(T) - 1)^8$, and therefore $k(T) \leq 2T^{1/8}$. In any run of the algorithm, we can classify every item $t \leq T$ into one of four categories.

1. Item t was allocated in one of the first $3n - 1$ epochs, that is, $k(t) < 3n$.
2. Item t was allocated in the sampling phase of epoch $k(t) \geq 3n$.
3. Item t was allocated in the ranking phase of epoch $k(t) \geq 3n$; the epoch was $\varepsilon_{k(t)}$ -accurate.
4. Item t was allocated in the ranking phase of epoch $k(t) \geq 3n$; the epoch was not $\varepsilon_{k(t)}$ -accurate.

We say an item t was a mistake if it was given to an agent with a non-maximum quantile for it. We show that the number of mistakes in each category are bounded by $3^{10}n^9, 2nT^{5/8}, 9nT^{15/16}$, and $158n \ln(T)$ respectively, with high probability. This implies that the total number of mistakes is at most the sum of these quantities, which is $O(\text{poly}(n) \cdot T^{15/16})$, with high probability, via a union bound.

The number of items in the first category is at most

$$\sum_{k=1}^{3n-1} k^4 n + k^8 \leq \sum_{k=1}^{3n} (3n)^4 n + (3n)^8 \leq (3n)^5 n + (3n)^9 \leq 3^{10} n^9.$$

Hence, the number of mistakes in the first category is also at most $3^{10}n^9$.

For the second category, since $k(T) \leq 2T^{1/8}$, we have that the total number of items in the sampling phase is (with probability 1) upper bounded by

$$\sum_{k=1}^{k(T)} nk^4 \leq nk(T)^5 \leq 2nT^{5/8}.$$

Each item t in the third category has probability $\varepsilon_{k(t)}$ of being a mistake. The expected number of mistakes is therefore at most $\sum_{k=3n}^{k(T)} \varepsilon_k k^8 = \sum_{k=3n}^{k(T)} 3nk^{13/2} \leq 3nk(T)^{15/2} \leq 8nT^{15/16}$. Using Hoeffding's inequality we get that with high probability the number of mistakes is at most $9nT^{15/16}$, since a deviation of $nT^{15/16}$ occurs with probability at most $\exp(-2n^2T^{15/8}/T) = \exp(-2n^2T^{7/8})$.

For the fourth category, the expected number of items is at most $\sum_{k=3n}^{k(T)} \delta_k k^8 = 2n \sum_{k=3n}^{k(T)} \frac{k^8}{e^{2k}} \leq 2n \sum_{k=1}^{\infty} \frac{k^8}{e^{2k}} \leq 158n$. Using Markov's inequality we have that the number of mistakes is at most $158n \ln(T)$ with probability at least $1 - 1/\ln(T)$, i.e., with high probability. \square

Finally, we can prove Theorem 2 as a relatively straightforward consequence of Lemma 5, since the ideal quantile maximization algorithm satisfies nice properties (e.g., Lemma 1).

Proof of Theorem 2 Fix a distribution D with CDF F and let X be a random variable with distribution D . Fix some ε to be $(1 - \varepsilon)$ -welfare-maximizing. Let \mathcal{E}_1^T be the event that the maximum social welfare at time T is at least $1/2 \cdot \mathbb{E}[X] \cdot T$, let \mathcal{E}_2^T be the event that quantile maximization is c -strongly-EF for $c = \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])}{4n}$, and let \mathcal{E}_3^T be the event that Algorithm 1 differs from quantile maximization on at most $f(T)$ items from Lemma 5. We first claim that $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T$ occurs with high probability in T . Note that Lemmas 1 and 5 tell us \mathcal{E}_2^T and \mathcal{E}_3^T each occur with high probability, respectively. For \mathcal{E}_1^T , the maximum value for each item is in expectation at least the expected value for a single agent $\mathbb{E}[X]$. Hence, a Chernoff bound tells us \mathcal{E}_1^T occurs with probability at least $1 - \exp\left(\frac{-\mathbb{E}[X]T}{8}\right)$, i.e., with high probability. The claim holds because the intersection of a finite number of high probability events occurs with high probability.

Next, note that for sufficiently large T , since $f(T) \in o(T)$, $f(T) \leq \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])}{8n} \cdot T$ and $f(T) \leq \varepsilon/2 \cdot \mathbb{E}[X] \cdot T$ (for any fixed ε that does not depend on T). Fix such a sufficiently large T . We show that, conditioned on $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T$, both EF and $(1 - \varepsilon)$ -welfare hold. Let $A^{QM} = (A_1^{QM}, \dots, A_n^{QM})$ be the allocation of quantile maximization and $A = (A_1, \dots, A_n)$ be the allocation of Algorithm 1. Beginning with envy-freeness, we have that for all pairs of agents i and j ,

$$\begin{aligned} v_i(A_i) &\stackrel{(\mathcal{E}_3^T)}{\geq} v_i(A_i^{QM}) - f(T) \\ &\stackrel{(\mathcal{E}_2^T)}{\geq} v_i(A_j^{QM}) - f(T) + \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])T}{4n} \\ &\stackrel{(\mathcal{E}_3^T)}{\geq} v_i(A_j) - 2f(T) + \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])T}{4n} \\ &\geq v_i(A_j), \end{aligned}$$

so the allocation is envy-free. Further, noting that $\text{sw}(A^{QM})$ is the maximum social welfare, we have the welfare approximation is at least

$$\begin{aligned} \frac{\text{sw}(A)}{\text{sw}(A^{QM})} &= \frac{\text{sw}(A^{QM}) - (\text{sw}(A^{QM}) - \text{sw}(A))}{\text{sw}(A^{QM})} \\ &\stackrel{(\mathcal{E}_3^T)}{\geq} \frac{\text{sw}(A^{QM}) - f(T)}{\text{sw}(A^{QM})} \\ &= 1 - \frac{f(T)}{\text{sw}(A^{QM})} \\ &\stackrel{(\mathcal{E}_1^T)}{\geq} 1 - \frac{f(T)}{1/2 \cdot \mathbb{E}[X] \cdot T} \\ &\stackrel{(\mathcal{E}_3^T)}{\geq} 1 - \frac{\varepsilon/2 \cdot \mathbb{E}[X] \cdot T}{1/2 \cdot \mathbb{E}[X] \cdot T} \\ &= 1 - \varepsilon, \end{aligned}$$

as needed. \square

5. Bounded memory in the i.i.d. model

In this section, we are interested in the more ambitious problem of designing dynamic algorithms with even more limited partial information: each agent is allowed to “remember” only a single item. We first show that, in this case, we need to settle for constant approximations of welfare.

THEOREM 3. *In the i.i.d. model, given a memory of one item per agent, there is no algorithm \mathcal{A} that is .999-welfare maximizing with high probability for all continuous and bounded value distributions.*

Proof We prove that this negative result holds even for an even stronger class of algorithms in which, at each step t , the algorithm *selects* quantile thresholds $q_1^t, \dots, q_n^t \in [0, 1]$ for each agent, and once an item arrives the algorithm observes, for each agent, whether the quantile of their sampled value $Q_{i,t}$ is above or below the threshold q_i^t . Note that this provides at least as much information about the fresh item as comparing it to any single prior item, since there is some uncertainty about the values and quantiles of all prior items.

We first focus on the algorithm for a single time-step and show there is a distribution of values such that, regardless of the quantile thresholds selected and allocations made, it cannot do well.

Fix a number of agents n and assume $n \geq 3$. We handle the special case of $n = 2$ at the end of this proof, as it requires a different distribution. For simplicity we consider a distribution that takes values larger than 1; re-scaling (specifically, dividing all values by $2 + \varepsilon$) gives a distribution upped bounded by 1 and does not affect any of our arguments. Consider the value distribution X , with

$$X \sim \begin{cases} \text{Unif}[0, \varepsilon] & \text{with probability } 1 - \frac{1}{n}, \\ \text{Unif}[1, 1 + \varepsilon] & \text{with probability } \frac{2}{3n}, \text{ and} \\ \text{Unif}[2, 2 + \varepsilon] & \text{with probability } \frac{1}{3n} \end{cases}$$

for some small $\varepsilon > 0$ to be fixed later. Intuitively, X is a continuous version of a discrete distribution which takes low value (near 0) with probability $1 - \frac{1}{n}$, medium value (near 1) with probability $\frac{2}{3n}$, and high value (near 2) with probability $\frac{1}{3n}$. Let F_X be its CDF. Trivially, the maximum social welfare of T items when all agents have this value distribution is at most $T \cdot (2 + \varepsilon)$.

We show that regardless of what quantile thresholds the algorithm chooses at step t and which decision it makes given the resulting signals, the expected value of the agent receiving item t is at least $(1 - \varepsilon) \cdot \frac{1}{144e}$ away from optimal. To that end, fix arbitrary thresholds q_1, \dots, q_n . First, we partition the agents depending on whether their quantile q_i is above or below $1 - \frac{2n}{3}$. We let $N^{\text{below}} = \{i \in [n] \mid q_i < 1 - \frac{2n}{3}\}$ and $N^{\text{above}} = \{i \in [n] \mid q_i \geq 1 - \frac{2n}{3}\}$. Either $|N^{\text{below}}| \geq \lceil n/2 \rceil$ or $|N^{\text{above}}| \geq \lceil n/2 \rceil$; we analyse each case separately. Since $n \geq 3$, we have $\lceil n/2 \rceil \geq 2$.

Case I: $|N^{\text{below}}| \geq \lceil n/2 \rceil$. In this case, it will be difficult for the algorithm to distinguish between agents in N^{below} with medium value and those with high value. Consider the event \mathcal{E} that one

agent $i^{\max} \in N^{\text{below}}$ has quantile $Q_{i^{\max}} > 1 - \frac{1}{3n}$, one agent $i^{\text{smax}} \in N^{\text{below}}$ has quantile $Q_{i^{\text{smax}}} \in (1 - \frac{2}{3n}, 1 - \frac{1}{3n})$, and all other agents $i \in \mathcal{N} \setminus \{i^{\max}, i^{\text{smax}}\}$ have quantile $Q_i < 1 - \frac{1}{n}$. First, we show that $\Pr[\mathcal{E}] \geq \frac{1}{72e}$, a constant. To compute this probability, note that there are at least $\lceil n/2 \rceil \cdot (\lceil n/2 \rceil - 1)$ choices of i^{\max} and i^{smax} . Once these have been selected, the probability of \mathcal{E} occurring for this pair of agents is

$$\frac{1}{3n} \cdot \frac{1}{3n} \cdot \left(1 - \frac{1}{n}\right)^{n-2} \geq_{(n \geq 3)} \frac{1}{9n^2} \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{9en^2}.$$

Since $\lceil n/2 \rceil \cdot (\lceil n/2 \rceil - 1) \geq n^2/8$, we can then conclude $\Pr[\mathcal{E}] \geq \frac{1}{72e}$. Conditioned on \mathcal{E} occurring, i^{\max} has high value, i^{smax} has medium value, and all other agents have low value. However, from the perspective of the algorithm, two agents (i^{\max} and i^{smax}) give a high signal, and it's equally likely that each of them is the agent with the high value (note that we condition on \mathcal{E}). The algorithm must therefore allocate the item to an agent with at most medium value (upper bounded by $1 + \varepsilon$) with probability at least $1/2$, even though an agent with value at least 2 exists. Hence, in this timestep, the algorithm has an additive error (compared to the optimum welfare) of at least $(1 - \varepsilon)$ with probability at least $\frac{1}{144e}$.

Case II: $|N^{\text{above}}| \geq \lceil n/2 \rceil$. In this case, it will be difficult for the algorithm to distinguish between agents in N^{above} that have medium value and those with low value. Consider the event \mathcal{E} that one agent $i^{\max} \in N^{\text{above}}$ has quantile $Q_{i^{\max}} \in (1 - \frac{1}{n}, 1 - \frac{2}{3n})$ and all other agents $i \in \mathcal{N} \setminus \{i^{\max}\}$ have quantile $Q_i < 1 - \frac{1}{n}$. First, we show that $\Pr[\mathcal{E}] \geq \frac{1}{6e}$. Indeed, there are at least $n/2$ choices for i^{\max} . For a fixed choice of i^{\max} , the probability of \mathcal{E} occurring is $\frac{1}{3n} \cdot \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{3en}$, and there are at least $n/2$ choices for i^{\max} , so $\Pr[\mathcal{E}] \geq \frac{1}{6e}$. Agent i^{\max} and the other members of N^{above} (there is at least one more) are indistinguishable to the algorithm as they all have a low signal, so the algorithm must give it to an agent with value at most ε with probability at least $1/2$ even though an agent with value at least 1 exists. Hence, in this timestep, the algorithm has an additive error (compared to the optimum welfare) of at least $(1 - \varepsilon)$ with probability at least $\frac{1}{12e}$.

In either case, for every time step, the algorithm has an additive error of at least $(1 - \varepsilon)$ with probability at least $\frac{1}{144e}$, irrespective of the past allocations. As time steps are independent, standard tail bounds give that, for sufficiently small $\varepsilon > 0$, the error is at least $\frac{1-\varepsilon}{1000}T$ with high probability. The optimal social welfare is at most $(2 + \varepsilon) \cdot T$; we conclude the algorithm can be no more than an 0.999-approximation to welfare.

The case of two agents. Finally, we handle the case of two agents. Assume values are drawn from a $\text{Unif}[0, 1]$ distribution. Let q_1, q_2 be the quantile thresholds selected by the algorithm and, without loss of generality, suppose that $0 \leq q_1 \leq q_2 \leq 1$. At least one of the differences $q_1 - 0, q_2 - q_1, 1 - q_2$ must be at least $1/3$. Suppose $q_2 - q_1 \geq 1/3$ (the other cases are symmetric). We investigate the event that both agents have $Q_i \in [q_1, q_2]$, so that agent 1 signals high and agent 2 signals low,

Algorithm 2: Bounded Memory

```

for Epoch  $k = 1 \dots$  do
  Sampling Phase: ( $k^9$  items)
  NOTWITHINERROR  $\leftarrow \mathcal{N}$ 
  for  $trial = 1, \dots, k^3$  do
    for  $i \in \text{NOTWITHINERROR}$  do
      | Allocate the next item to agent  $i$ , and update her memory
    Test  $k^6 - |\text{NOTWITHINERROR}|$  number of items (for each agent)
    for  $i \in \text{NOTWITHINERROR}$  do
      | if Proportion of test items for agent  $i$  is within  $\pm 1/k^2$  of  $(n-1)/n$  then
        | NOTWITHINERROR  $\leftarrow \text{NOTWITHINERROR} \setminus \{i\}$ 
  Ranking Phase: ( $k^{18}$  items)
  for each item  $g$  in this phase do
    if Some agent  $i$  has high signal then
      | Give  $g$  to a (uniformly) random such agent
    else
      | Give  $g$  to an agent uniformly at random

```

which occurs with probability at least $1/9$. Conditioned on this event, the signals do not provide any additional information, so the algorithm chooses the agent with smaller value at least half of the time. In this case, the expected difference between the larger and smaller values is $1/9$. Hence, the expected difference of the value from the algorithm versus the maximum social welfare is at least $\frac{1}{9} \cdot \frac{1}{2} \cdot \frac{1}{9} = 1/162$ on each item. The maximum social welfare is at most T , and we expect the difference to be at least $T/1000$ due to concentration, so the algorithm cannot guarantee more than a .999 approximation, as needed. \square

Our positive result matches this lower bound up to a constant.

THEOREM 4. *In the i.i.d. model, given a memory of one item per agent, Algorithm 2 achieves envy-freeness and a $(1 - 1/e)^2 - \varepsilon$ approximation to welfare, with high probability, for all $\varepsilon > 0$.*

Algorithm 2 works in epochs, similar to Algorithm 1. In each epoch's exploration/sampling phase, it tries to find an item whose quantile is close to the $\frac{n-1}{n}$ -threshold algorithm. Epoch k makes k^3 such attempts, and each candidate item is tested against k^6 fresh items to get an estimated quantile. If everything is within the error we can tolerate, the algorithm remembers this item for this epoch; otherwise, the agent has an arbitrary item in memory during this epoch. During the exploitation/ranking phase, Algorithm 2 tries to mimic the $\frac{n-1}{n}$ -threshold algorithm (instead of the quantile maximization algorithm as Algorithm 1 did), and, in fact, inherits its approximation factor (Lemma 2) exactly.

Our first technical lemma, Lemma 6, gives necessary bounds on the various variables of Algorithm 2 for a sample to be ε -accurate with respect to the ideal threshold algorithm; see Definition 3. Its proof can be found in Section EC.1.2.

DEFINITION 3. A set of n items in memory, one for each agent, is ε -accurate with respect to q^* if with probability at least $1 - \varepsilon$, when a fresh item is sampled, the agents with true quantile above q^* are exactly those that value the fresh item more than their item in memory.

LEMMA 6. For all $\varepsilon, \delta \in (0, 1)$, if (1) at least τ trials are done with $\tau \geq \frac{\ln(2n/\delta)}{\varepsilon/(3n)}$, and (2) at least ℓ test items are used per trial for $\ell \geq \frac{18n^2}{\varepsilon^2} \ln\left(\frac{4\tau n}{\delta}\right)$, and (3) the tolerance for accepting an item is $\varepsilon/(3n)$, then the items in memory are ε -accurate (for all agents, simultaneously) with respect to $q^* = \frac{n-1}{n}$, with probability at least $1 - \delta$.

Though Lemmas 4 and 6 resemble each other (and are used in analogous ways), the proofs require different techniques, as the sampling processes are very different. Next, we prove an analogue to Lemma 5: the number of disagreements between Algorithm 2 and the ideal threshold algorithm is sublinear.² The proofs of Lemmas 5 and 7 are similar, precisely because Lemma 4 matches Lemma 6. Theorem 4 follows from Lemma 7 as in the i.i.d. case. The proofs of Lemma 7 and Theorem 4 can be found in Sections EC.1.3 and EC.1.4 respectively.

LEMMA 7. The allocation of Algorithm 2 differs from that of the $\frac{n-1}{n}$ -threshold algorithm after T steps by at most $f(T)$ items with high probability, where $f(T) \in O(\text{poly}(n) \cdot T^{17/18})$.

6. Agents with correlated values

Recall that $v_{i,t} = v_t^{\text{com}} + \varepsilon_{i,t}$, with common value v_t^{com} drawn from a common distribution D^{com} and agent specific noise $\varepsilon_{i,t}$ drawn from noise distribution D^{noise} . This class of valuations was captured in a more general class considered by Dickerson et al. (2014), who show that welfare maximization is still EF with high probability (and, by definition, a 1-approximation to welfare). However, it is unclear whether these results carry over when only given partial information since the correlation can make it harder to “learn” agents’ relative values during sampling.

In this section we show that, at least under mild restrictions on D^{com} and D^{noise} , we can still devise algorithms that are able to well-approximate the ideal welfare maximizing algorithm. The restrictions are as follows: (i) *Interval support*: the support of each of these distributions are some intervals $[a^{\text{com}}, b^{\text{com}}]$ and $[a^{\text{noise}}, b^{\text{noise}}]$. (ii) *PDF-boundedness*: there are constants $0 < p \leq q$ such that the probability density functions of D^{com} and D^{noise} are bounded between p and q on their

² Note these are randomized algorithms, so by “differ on a item” here we mean that the distributions over agents receiving the item differ.

support. These assumptions are required only in this section and are quite common in the distributional fair division literature — they are the exact assumptions of [Bai and Gözl \(2022\)](#) and weaker than the those of [Manurangsi and Suksompong \(2021\)](#), who also require the support to be $[0, 1]$. We call this method of generating values the *common-noise model* and establish the following.

THEOREM 5. *In the common-noise model, running Algorithm 1 with sampling phases per agent of length k^6 and exploiting phases of length k^{12} achieves envy-freeness and a $(1 - \varepsilon)$ approximation to welfare, with high probability, for all $\varepsilon > 0$.*

The proof of this theorem is similar to that of Theorem 2. The main difference is that Lemma 4 is not valid for correlated values. Nonetheless, using new techniques, we show the following analog.

LEMMA 8. *For all (p, q) -bounded common and noise distributions D^{com} and D^{noise} supported on $[a^{com}, b^{com}]$ and $[a^{noise}, b^{noise}]$, if $\varepsilon, \delta \in (0, 1)$ and $m \in \mathbb{Z}^+$ are such that*

$$\varepsilon > \frac{2nq}{\min\left(\frac{b^{com}-a^{com}}{b^{noise}-a^{noise}}, 1\right) \cdot p} \cdot \left(\frac{\ln(2n/\delta)}{2m}\right)^{1/4},$$

then giving m samples to each agent is ε -accurate with probability at least $1 - \delta$.

Proof. It is without loss of generality to assume that the supports of D^{com} and D^{noise} are translated to start at 0, i.e., are of the form $[0, b^{com} - a^{com}]$ and $[0, b^{noise} - a^{noise}]$. Indeed, translating the values does not change whether or not an item goes to the correct agent. For convenience we assume the supports are $[0, b^{com}]$ and $[0, b^{noise}]$ throughout the proof, then translate the distributions back for the final bound by replacing b^{com} by $b^{com} - a^{com}$ and b^{noise} by $b^{noise} - a^{noise}$.

Let D^{sum} be the distribution obtained by adding independent samples from D^{com} and D^{noise} . Note that D^{sum} is the marginal distribution of agent values. Let F^{com} , F^{noise} , and F^{sum} and f^{com} , f^{noise} , and f^{sum} be the CDFs and PDFs of the distributions D^{com} , D^{noise} , and D^{sum} , respectively. Additionally, since D^{sum} is the sum of independent samples of D^{com} and D^{noise} , it is well known that $f^{sum}(x) = \int_{-\infty}^{\infty} f^{com}(t)f^{noise}(x-t)dt$, the convolution of the summand densities. Let \hat{F}_i be the empirical CDF of agent i 's values after m samples. (Note that unlike in Lemma 4, we are working with values instead of quantiles.) Using the DKW inequality, it is still the case that for all $\gamma > 0$, $\Pr\left[\sup_v |\hat{F}_i(v) - F^{sum}(v)| > \gamma\right] \leq 2e^{-2m\gamma^2}$, and specifically for $\gamma = \sqrt{\frac{\ln(2n/\delta)}{2m}}$, this holds for all n agents simultaneously with probability $1 - \delta$.

We again condition on the DKW event, that $\sup_v |\hat{F}_i(v) - F^{sum}(v)| > \gamma$ for all agents i . Let V^{com} and $V_1^{noise}, \dots, V_n^{noise}$ be fresh samples of common and noise values. Let $V_i = V^{com} + V_i^{noise}$ be the total value of each agent i . We would like to show that if i^{max} is the agent with the highest value, then $\hat{F}_{i^{max}}(V_{i^{max}}) > \hat{F}_j(V_j)$ for all $j \neq i^{max}$. This ensures that i^{max} receives the item. A sufficient

condition for this to occur is that $F^{sum}(V_{i^{max}}) - F^{sum}(V_j) > 2\gamma$ for all $j \neq i^{max}$, because of the DKW condition.

To get a handle on conditions to ensure this difference in quantiles is sufficiently large, we will begin by proving the following inequality.

LEMMA 9. *For all $v_2 \geq v_1 \in [0, b^{com} + b^{noise}]$ such that $v_2 - v_1 \leq b^{noise}$,*

$$F^{sum}(v_2) - F^{sum}(v_1) \geq \left(\min \left(\frac{b^{com}}{b^{noise}}, 1 \right) p(v_2 - v_1) \right)^2 / 2. \quad (1)$$

Proof of Lemma 9. Fix such a v_1 and v_2 . Writing this out more explicitly, we have that

$$F^{sum}(v_2) - F^{sum}(v_1) = \int_{v_1}^{v_2} f^{sum}(x) dx.$$

To lower bound this integral, we will first lower bound $f^{sum}(x)$. Fix an x in the support of D^{com} , so $x \in [0, b^{com} + b^{noise}]$. We have that

$$\begin{aligned} f^{sum}(x) &= \int_{-\infty}^{\infty} f^{com}(t) f^{noise}(x-t) dt \\ &\geq \int_{-\infty}^{\infty} (p \cdot \mathbb{I}[t \in [0, b^{com}]] \cdot (p \cdot \mathbb{I}[x-t \in [0, b^{noise}]]) dt \\ &= \int_{-\infty}^{\infty} (p \cdot \mathbb{I}[t \in [0, b^{com}]] \cdot (p \cdot \mathbb{I}[t \in [x-b^{noise}, x]]) dt \\ &= p^2 \int_{-\infty}^{\infty} \mathbb{I}[t \in [0, b^{com}] \wedge t \in [x-b^{noise}, x]] dt. \end{aligned}$$

Since $0 \leq x \leq b^{com} + b^{noise}$, rearranging shows that $x - b^{noise} \leq b^{com}$, and (trivially), $x \geq 0$. Therefore, $t \in [0, b^{com}] \wedge t \in [x - b^{noise}, x]$ reduces to $t \in [\max(0, x - b^{noise}), \min(b^{com}, x)]$, and hence

$$\begin{aligned} \int_{-\infty}^{\infty} \mathbb{I}[t \in [0, b^{com}] \wedge t \in [x - b^{noise}, x]] dt &= \min(b^{com}, x) - \max(0, x - b^{noise}) \\ &= \min(b^{com}, x) + \min(0, b^{noise} - x) \\ &= \min(b^{com}, b^{noise}, x, (b^{noise} + b^{com}) - x). \end{aligned}$$

Putting this together, we have that for $x \in \text{supp}(D^{sum})$,

$$f^{sum}(x) \geq p^2 \min(b^{com}, b^{noise}, x, (b^{noise} + b^{com}) - x).$$

Let $g(x) = p^2 \min(b^{com}, b^{noise}, x, (b^{noise} + b^{com}) - x)$ and let us now consider the shape of $g(x)$. A plot of $g(x)$ can be found in Figure 1. It increases linearly (with a slope of p^2) from 0 until $\min(b^{com}, b^{noise})$, stays constant until $\max(b^{com}, b^{noise})$, and then decreases (with a slope of $-p^2$) until $b^{com} + b^{noise}$. Note that both the nonconstant intervals are of length $\min(b^{com}, b^{noise})$. For our purposes, the shape of $g(x)$ means that integrating over any interval $[c, d] \subseteq [0, b^{com} + b^{noise}]$




Figure 1 Plot of lower bound $g(x)$.

of length $\ell := d - c$ is at least as large as integrating $[0, \ell]$, i.e., $\int_c^d g(x) dx \geq \int_0^\ell g(x) dx$. Further, as long as $\ell \leq \min(b^{com}, b^{noise})$, the area under the curve of $g(x)$ on this interval is simply a triangle, and it has area $p^2 \ell^2 / 2$. These facts together imply that as long as $v_2 - v_1 \leq \min(b^{com}, b^{noise})$, then

$$F^{sum}(v_2) - F^{sum}(v_1) \geq (p(v_2 - v_1))^2 / 2. \quad (2)$$

To extend this to the case needed for Inequality (1) with the only constraint being $v_2 - v_1 \leq b^{noise}$, let $v'_2 = v_1 + (v_2 - v_1) \cdot \min(\frac{b^{com}}{b^{noise}}, 1)$. Note that $v'_2 \leq v_2$, and, in addition,

$$v'_2 - v_1 \leq \min\left(\frac{b^{com}}{b^{noise}}, 1\right)(v_2 - v_1) \leq \min(b^{com}, b^{noise}).$$

It follows from Ineq. (2), $v'_2 \leq v_2$ and the definition of v'_2 that, as required,

$$F^{sum}(v_2) - F^{sum}(v_1) \geq F^{sum}(v'_2) - F^{sum}(v_1) \geq p^2 (v'_2 - v_1)^2 / 2 = \left(\min\left(\frac{b^{com}}{b^{noise}}, 1\right) p(v_2 - v_1)\right)^2 / 2. \quad \square$$

Having established Lemma 9, we continue with the proof of Lemma 8 and now consider what constraints on $V_{i^{max}}$ and V_j ensure that $F^{sum}(V_{i^{max}}) - F^{sum}(V_j) \geq 2\gamma$ for all $j \neq i^{max}$. Since $V_{i^{max}}$ and V_j can differ by at most b^{noise} , we immediately get that $F^{sum}(V_{i^{max}}) - F^{sum}(V_j) \geq (\min(\frac{b^{com}}{b^{noise}}, 1) p(V_{i^{max}} - V_j))^2 / 2$. A sufficient condition for this to be at least 2γ is that

$$(V_{i^{max}} - V_j)^2 \geq \frac{4\gamma}{(\min(\frac{b^{com}}{b^{noise}}, 1) p)^2},$$

and, equivalently,

$$V_{i^{max}} - V_j \geq \frac{2\sqrt{\gamma}}{\min(\frac{b^{com}}{b^{noise}}, 1) p}.$$

Recall that $V_{i^{max}} = V^{com} + V_{i^{max}}^{noise}$ and $V_j = V^{com} + V_j^{noise}$, so $V_{i^{max}} - V_j = V_{i^{max}}^{noise} - V_j^{noise}$. Additionally, since f^{noise} is upperbounded by q ,

$$F^{noise}(V_{i^{max}}^{noise}) - F^{noise}(V_j^{noise}) = \int_{V_j^{noise}}^{V_{i^{max}}^{noise}} f^{noise}(x) dx \leq q(V_{i^{max}}^{noise} - V_j^{noise}).$$

Hence, as long as $F^{noise}(V_{i^{max}}^{noise}) - F^{noise}(V_j^{noise}) \geq \frac{2q\sqrt{\gamma}}{\min(\frac{b^{com}}{b^{noise}}, 1)p}$, then $V_{i^{max}}^{noise} - V_j^{noise} \geq \frac{2\sqrt{\gamma}}{\min(\frac{b^{com}}{b^{noise}}, 1)p}$. $F^{noise}(V_j^{noise})$ is distributed uniformly on $[0, 1]$ for each i , and these random variables are independent across agents. $F^{noise}(V_{i^{max}}^{noise})$ will be the largest of n such draws. Hence, we have reduced this to the case handled by Lemma 4 — the Beta $[1, n]$ analysis from that lemma shows that $\Pr[\forall j \neq i^{max}, F^{noise}(V_{i^{max}}^{noise}) - F^{noise}(V_j^{noise}) > \rho] \geq 1 - n\rho$. Plugging in $\rho = \frac{2q\sqrt{\gamma}}{\min(\frac{b^{com}}{b^{noise}}, 1)p}$, we get that this holds with probability at least $1 - \frac{2nq\sqrt{\gamma}}{\min(\frac{b^{com}}{b^{noise}}, 1)p}$. Therefore, as long as $\varepsilon > \frac{2nq}{\min(\frac{b^{com}}{b^{noise}}, 1)p} \cdot \left(\frac{\ln(2n/\delta)}{2m}\right)^{1/4}$, ε -accuracy holds. \square

We are now in position to prove Theorem 5. The proof structure is nearly identical to Theorem 2, we give an overall sketch and describe the differences.

Proof sketch of Theorem 5. Fix (p, q) -bounded distributions D^{com} and D^{noise} supported on $[a^{com}, b^{com}]$ and $[a^{noise}, b^{noise}]$. The proof is nearly identical to Theorem 2, we primarily describe the differences here, while giving an overall sketch.

An analog of Lemma 1 continues to hold, with slightly different constants that depend on the distribution. Hence, all we need to show is that running Algorithm 1 with a sampling length of k^6 per agent and an exploiting length of k^{12} differs from value maximization by at most some sublinear number of items with high probability. We will show there are $O(T^{47/48})$ errors; note the asymptotic notation is hiding constants that depend on n and the common and noise distributions.

Let $C = \frac{3nq}{\min(\frac{b^{com}-a^{com}}{b^{noise}-a^{noise}}, 1)p}$. For each epoch k , choose $\delta_k = 2n/e^{2k}$ and set $\varepsilon_k = C \cdot k^{-5/4}$. Now

$$\begin{aligned} \varepsilon_k &= C \cdot k^{-5/4} = C \cdot \left(\frac{2k}{2k^6}\right)^{1/4} = C \cdot \left(\frac{\ln(e^{2k})}{2k^6}\right)^{1/4} \\ &> \frac{2nq}{\min(\frac{b^{com}}{b^{noise}}, 1)p} \cdot \left(\frac{\ln(2n/\delta_k)}{2k^6}\right)^{1/4}. \end{aligned}$$

The constant C is unimportant, as it does not depend on T or k it will essentially disappear as the timestep grows large. The primary difference to Theorem 2 is that, due to the longer phase lengths (k^6 and k^{12} instead of k^4 and k^8), we need a dependence of $k^{-5/4}$ for ε_k (rather than $k^{-3/2}$).

We count the number of items on which Algorithm 1 differs from the welfare-maximizing algorithm, as in Lemma 5, and recall that $k(t)$ is the epoch in which item t arrives. Fix a timestep T . Due to the different sampling lengths, it now holds that $k(T) \leq 2T^{1/12}$.

The number of items allocated in sampling phases is now

$$\sum_{k=1}^{k(T)} nk^6 \leq nk(T)^7 \leq 2nT^{7/12},$$

which is still sublinear (with probability one).

The expected number mistakes in exploit phases of $\varepsilon_{k(t)}$ -accurate epochs is at most

$$\sum_{k=1}^{k(T)} \varepsilon_k k^{12} = \sum_{k=1}^{k(T)} C k^{43/4} = C k(T)^{47/4} = O(T^{47/48}).$$

Due to Hoeffding's inequality a deviation of more than $T^{47/48}$ does not occur with high probability, so the number of mistakes here is $O(T^{47/48})$ with high probability.

The expected number of items in non- ε_k -accurate exploit phases is again at most

$$\sum_{k=1}^{\infty} \delta_k k^{12} = 2n \sum_{k=1}^{\infty} \frac{k^{12}}{e^2 k} = O(1),$$

because the infinite series converges. As a result, the probability it is more than $\log(T)$ times this value is at most $O(1/\log(T))$.

With these modifications, the rest of Theorem 2 goes through essentially unchanged. \square

7. The non-i.i.d. model

In this section, we study the non-i.i.d. model. We first establish a strong lower bound for the non-i.i.d. model. The following negative result holds even for algorithms that know the associated quantile for every fresh item.

THEOREM 6. *Even for 2 non-identical agents, there is no algorithm that is EF and c -PO with probability p , for $c > \frac{1+\sqrt{5}}{4} \approx .809$ and $p > 2/3$, for all continuous and bounded value distributions.*

Proof. Suppose for contradiction that there is an algorithm \mathcal{A} so that for all bounded continuous distributions (X_1, X_2) there exists a $T^* = T^*(X_1, X_2)$ where for all $t \geq T^*$, \mathcal{A} is envy-free and c -PO with probability p with $p > 2/3$ for some constant $c > \frac{1+\sqrt{5}}{4}$. Hence, there is some ε such that $p > 2/3 + \varepsilon$ and $1/c < \frac{4}{1+\sqrt{5}} - \varepsilon = \sqrt{5} - 1 - \varepsilon$.

Consider two distributions D_F and D_S ; we describe these later in the proof. Consider the three instances $I_0 = (D_F, D_F)$, $I_1 = (D_S, D_F)$ and $I_2 = (D_F, D_S)$.

Let $\mathcal{E}_j^{A,t}$ be the event that \mathcal{A} is envy-free and c -PO on instance I_j at time t for $j \in \{0, 1, 2\}$. By construction, $\Pr[\mathcal{E}_j^{A,t}] \geq 2/3 + \varepsilon$ for all $j \in \{0, 1, 2\}$ and $t \geq T^*$.

Let z be a parameter we will fix later in the proof, and let $Z_i^t = \mathbb{I}\{Q_{i,t} \geq 1 - z\}$ for $i = \{1, 2\}$. Observe that $Z_1^t \cdot Z_2^t$ is 1 with probability z^2 and 0 otherwise. The following events characterize a specific notion of a “nice” sample, in which the number of items with high quantiles for both agents is near its expectation: $\mathcal{E}_1^T = \mathbb{I}\{|\frac{1}{T} \sum_{t=1}^T Z_1^t \cdot Z_2^t - z^2| < \delta\}$, $\mathcal{E}_2^T = \mathbb{I}\{|\frac{1}{T} \sum_{t=1}^T Z_1^t - z| < \delta\}$, and $\mathcal{E}_3^T = \mathbb{I}\{|\frac{1}{T} \sum_{t=1}^T Z_2^t - z| < \delta\}$ for some $\delta > 0$. By Hoeffding's inequality, $\Pr[\bar{\mathcal{E}}_1^T] = \Pr\left[|\frac{1}{T} \sum_{t=1}^T Z_1^t \cdot Z_2^t - z^2| \geq \delta\right] \leq 2 \exp(-2T\delta^2)$. It follows that for $T \geq \log(2/\varepsilon)/(2\delta^2)$, $\Pr[\bar{\mathcal{E}}_1^T] \leq \varepsilon$.

Similarly, for $T \geq \log(2/\varepsilon)/(2\delta^2)$, it holds that $\Pr[\bar{\mathcal{E}}_2^T] \leq \varepsilon$, and $\Pr[\bar{\mathcal{E}}_3^T] \leq \varepsilon$. Consider an arbitrary $T > T_{\max} = \max\{T_0, T_1, T_2, \log(2/\varepsilon)/(2\delta^2)\}$. Applying a union bound,

$$\Pr[\bar{\mathcal{E}}_0^{A,T} \cup \bar{\mathcal{E}}_1^{A,T} \cup \bar{\mathcal{E}}_2^{A,T} \cup \bar{\mathcal{E}}_1^T \cup \bar{\mathcal{E}}_2^T \cup \bar{\mathcal{E}}_3^T] \leq \sum_{i=0}^2 \Pr[\bar{\mathcal{E}}_i^{A,T}] + \sum_{i=1}^3 \Pr[\bar{\mathcal{E}}_i^T] < 3 \cdot \left(\frac{1}{3} - \varepsilon\right) + 3\varepsilon = 1.$$

It follows that $\Pr[\mathcal{E}_0^{A,T} \cap \mathcal{E}_1^{A,T} \cap \mathcal{E}_2^{A,T} \cap \mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T] > 0$. Therefore, there must exist a sequence of T items whose quantiles satisfy all of \mathcal{E}_1^T , \mathcal{E}_2^T , and \mathcal{E}_3^T , and, since \mathcal{A} does not have access to the items' values, there must exist an allocation A^T for these T items (in the support of \mathcal{A}) that is EF and c -PO, no matter which of I_0 , I_1 or I_2 the values were taken from. Let $q^T = \{(q_1(t), q_2(t))\}_{t=1}^T$ be these items' quantiles. Let $H_B = \{t \in [T] : q_1(t) \geq 1 - z \text{ and } q_2(t) \geq 1 - z\}$ be the items for which $Z_1^t \cdot Z_2^t = 1$, and $H_1 = \{t \in [T] : q_1(t) \geq 1 - z\}$ the items for which $Z_1^t = 1$.

Set distributions $D_F = \text{Unif}[1 - w, 1]$ and D_S , under which each item is $\text{Unif}[0, w]$ with probability z and at $\text{Unif}[1 - w, 1]$ with probability $1 - z$, for some small positive w that we fix later in the proof.

We have that some agent receives at most half the items in H_B ; without loss of generality this is agent 2, i.e., $|A_2^T \cap H_B| \geq |H_B|/2$. We show that there exists a feasible more than $1/c$ Pareto-improvement under the values in I_1 . To that end, we compare A^T to the allocation \hat{A} where $\hat{A}_1 = H_1$ and $\hat{A}_2 = \bar{H}_1$.

We next bound the utilities of each agent under A^T and \hat{A} . Beginning with agent 1, we have

$$\begin{aligned} u_1(\hat{A}_1) &= u_1(H_1) \\ &\geq |H_1| \cdot (1 - w) \\ &\geq^{(\mathcal{E}_1^T)} T \cdot (z - \delta)(1 - w) \\ &= T(z - \delta - zw + \delta w) \\ &\geq T(z - \delta - w) \end{aligned}$$

and

$$\begin{aligned} u_1(A_1) &\leq w \cdot |A_1 \cap \bar{H}_1| + 1 \cdot |A_1 \cap H_1| \\ &\leq T \cdot w + |H_1| - |A_2 \cap H_1| \\ &\leq T \cdot w + |H_1| - |A_2 \cap H_B| \\ &\leq^{(\mathcal{E}_2^T)} T \cdot w + T(z + \delta) - |A_2 \cap H_B| \\ &\leq T \cdot w + T(z + \delta) - |H_B|/2 \\ &\leq^{(\mathcal{E}_1^T)} T \cdot w + T(z + \delta) - T(z^2 - \delta)/2 \\ &= T(z - z^2/2 + w + 3\delta/2). \end{aligned}$$

Together, these imply

$$\frac{u_1(\hat{A}_1)}{u_1(A_1^T)} \geq \frac{z - \delta - w}{z - z^2/2 + w + 3\delta/2} = \frac{2z - 2\delta - 2w}{2z - z^2 + 2w + 3\delta}.$$

Next, we consider agent 2. We have

$$\begin{aligned} u_2(\hat{A}_2) &= u_2(\bar{H}_1) \\ &\geq (1-w)|\bar{H}_1| \\ &= (1-w)(T - |H_1|) \\ &\geq^{(\varepsilon_2^T)} (1-w)T \cdot (1 - (z + \delta)) \\ &= T(1 - z - \delta - w + wz + w\delta) \\ &\geq T(1 - z - \delta - w). \end{aligned}$$

By \mathcal{E}_0^A , A^T is envy-free on I_0 . It follows that $|A_1^T| \geq (1-w)|A_2^T|$. Since $|A_1^T| + |A_2^T| = T$, we have that $|A_2^T| \leq \frac{1}{2-w}T$. Hence, $u_2(A_2^T) \leq |A_2^T| \leq \frac{1}{2-w}T$. Combining these, we have

$$\frac{u_2(\hat{A}_2)}{u_2(A_2^T)} = \frac{1 - z - \delta - w}{\frac{1}{2-w}} = 2 - 2z - 2\delta - 2w - w + wz + w\delta + w^2 \geq 2 - 2z - 2\delta - 3w.$$

Choose $z = \frac{3-\sqrt{5}}{2}$. Note that $z^2 = \frac{7-3\sqrt{5}}{2}$. Choose $\delta, w < \varepsilon/25$. We then have,

$$\begin{aligned} \frac{u_1(\hat{A}_1)}{u_1(A_1^T)} &> \frac{3 - \sqrt{5} - \varepsilon/5}{(\sqrt{5} - 1)/2 + \varepsilon/5} \\ &= \frac{3 - \sqrt{5}}{(\sqrt{5} - 1)/2 + \varepsilon/5} - \frac{\varepsilon/5}{(\sqrt{5} - 1)/2 + \varepsilon/5} \\ &> \frac{3 - \sqrt{5}}{(\sqrt{5} - 1)/2 + \varepsilon/5} - \frac{2\varepsilon}{5} && (\frac{\sqrt{5}-1}{2} + \frac{\varepsilon}{5} > 1/2) \\ &> \frac{3 - \sqrt{5}}{(\sqrt{5} - 1)/2 \cdot (1 + 2\varepsilon/5)} - \frac{2\varepsilon}{5} && (\sqrt{5} - 1 > 1) \\ &= (\sqrt{5} - 1) \cdot \frac{1}{1 + 2\varepsilon/5} - \frac{2\varepsilon}{5} \\ &> (\sqrt{5} - 1) \cdot (1 - 2\varepsilon/5) - \frac{2\varepsilon}{5} \\ &> (\sqrt{5} - 1) - \varepsilon/2 - \frac{2\varepsilon}{5} && ((\sqrt{5} - 1) \cdot 2/5 < 1/2) \\ &> \sqrt{5} - 1 - \varepsilon \\ &> 1/c \end{aligned}$$

and

$$\frac{u_2(\hat{A}_2)}{u_2(A_2^T)} > 2 - (3 - \sqrt{5}) - \varepsilon/5 > \sqrt{5} - 1 - \varepsilon > 1/c,$$

so this is more than a $1/c$ Pareto improvement. \square

Algorithms 1 and 2 are envy-free with high probability, even in the non-i.i.d. model, since envy-freeness is not an “inter-agent” property. Our last result shows that they also give a constant approximation to Pareto efficiency, by combining Lemma 3 with Lemmas 5 and 7. Its proof can be found in Section EC.1.5.

THEOREM 7. *In the non-i.i.d. model, both Algorithm 1 (unbounded memory) and Algorithm 2 (one-item memory) are EF and $(1/e - \varepsilon)$ -PO, with high probability, for all $\varepsilon > 0$.*

Interestingly, the guarantees for Algorithm 2 in the non-i.i.d. model are only marginally worse compared to the i.i.d. model; the approximation ratio decreases from $(1 - 1/e)^2 \approx 0.4$ to $1/e \approx 0.37$. Finally, we note that, even though the formal guarantees in Theorem 7 are the same for the two algorithms, and even though Algorithm 2 uses memory size of one, Algorithm 1 has the benefit of much shorter epoch lengths (in addition to better guarantees under the i.i.d. model).

8. Computational study

The purpose of this section is twofold. First, while our theoretical results ensure that algorithms such as Algorithm 1 satisfy desirable properties, these guarantees are in the only limit, so a priori it is possible that it may take an extremely long time for them to kick in. With this in mind, we verify that these properties are satisfied on a variety of generated values, as well as compare how qualitative shifts in the value distributions affect these convergence rate. Second, although Algorithm 1 was designed to be amenable for theoretical analysis, we compare it to variations that may perform better or be preferable for practical reasons.

8.1. Setup

The experiments are conducted as follows. For each setting, we sample 100,000 item values for five agents. We run our algorithms on this either two random agents or all five. We repeat this 100 times for each setting so as to get reasonable statistics about the performance.

We generate agent values both from distributions and from real-world data. From distributions, we first consider several instances of the beta distribution. The first set are of the form $\beta(1, x)$ and $\beta(x, 1)$ for different values of x . Recall that $\beta(1, 1)$ is the uniform distribution over $[0, 1]$ and, as x grows larger, the distributions skew left or right. This allows us to understand the effect of skew (are there a few items that are extremely valuable, or are most valuable except for a few duds?) on performance. The next set were of the form $\beta(x, x)$ for increasing values of x . As x increases, the distribution becomes more peaked while remaining centered around $1/2$. The density functions of these distributions are visualized in Figure 2.

Next, we investigate the effect of correlation on performance. We generate uniform common values $v_t \sim U(0, 1)$ and an agent-specific values $\varepsilon_{it} \sim U(0, 1)$ for each agent i and item t . We then set

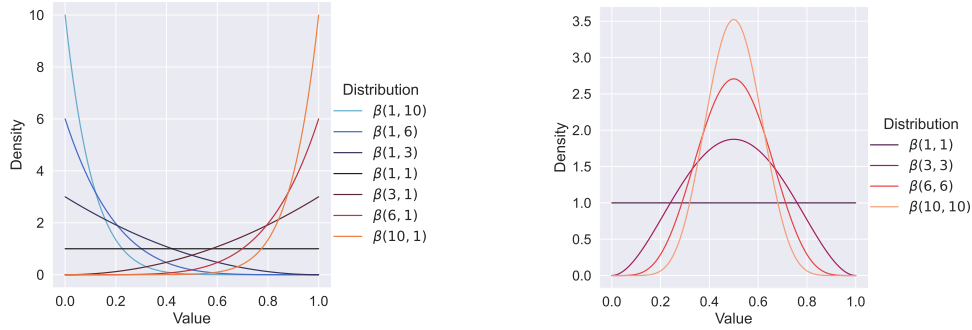


Figure 2 Density functions of the beta distributions we test on.

the agent value to $v_{it} = \alpha \cdot v_t + (1 - \alpha)\varepsilon_{it}$. Note that $\alpha = 0$ corresponds to independent $U(0, 1)$ values, $\alpha = 1$ corresponds to fully correlated identical values, and increasing α increases the correlation.

Finally, we test on real-world data obtained by training a matrix factorization based collaborative filtering algorithm on publicly available ratings data from Amazon,³ sample agents and use the predicted ratings as proxies for value.⁴

8.2. Results

All plots that appear in the main body are for experiments run with two agents. Additional plots with five agents can be found in Appendix EC.2 and are qualitatively similar.

We begin with an analysis on how quickly Algorithm 1 becomes envy-free. The corresponding plots for Algorithm 1, alongside the ideal welfare maximizing algorithm with full information, can be found in Figure 3. The main takeaways we find are that the more left skewed distributions tended to have faster convergence than more right skewed ones, less peaked distributions tended to converge faster than more peaked ones, and increasing the correlation, of course, made the problem more difficult. In general, the algorithm kept pace reasonably well alongside its “ideal” counter-part, and did not converge much slower. Instances where Algorithm 1 performed worse exactly corresponded to those where welfare/quantile maximization also struggled.

Next, we visualize the approximation to welfare in Figure 4. Here we find a complete reversal. Left skewed distributions had *worse* approximations than right skewed ones, less peaked distribution had worse approximations, and increasing correlation lead to better approximations. One possible explanation is that these trends exactly correspond to the ratio between the expected value of these distributions and the expected maximum of several draws. When this ratio is large (the expectation is quite close to the expected maximum) it means that giving an item to the “wrong” agent does not have too big of an effect on the welfare approximation. The one exception to this

³ All data are available at <https://cseweb.ucsd.edu/~jmcauley/datasets.html>.

⁴ Ratings measure customer satisfaction and is commonly used to represent utilities in the recommender systems literature. See, for example, Adomavicius and Tuzhilin (2005).

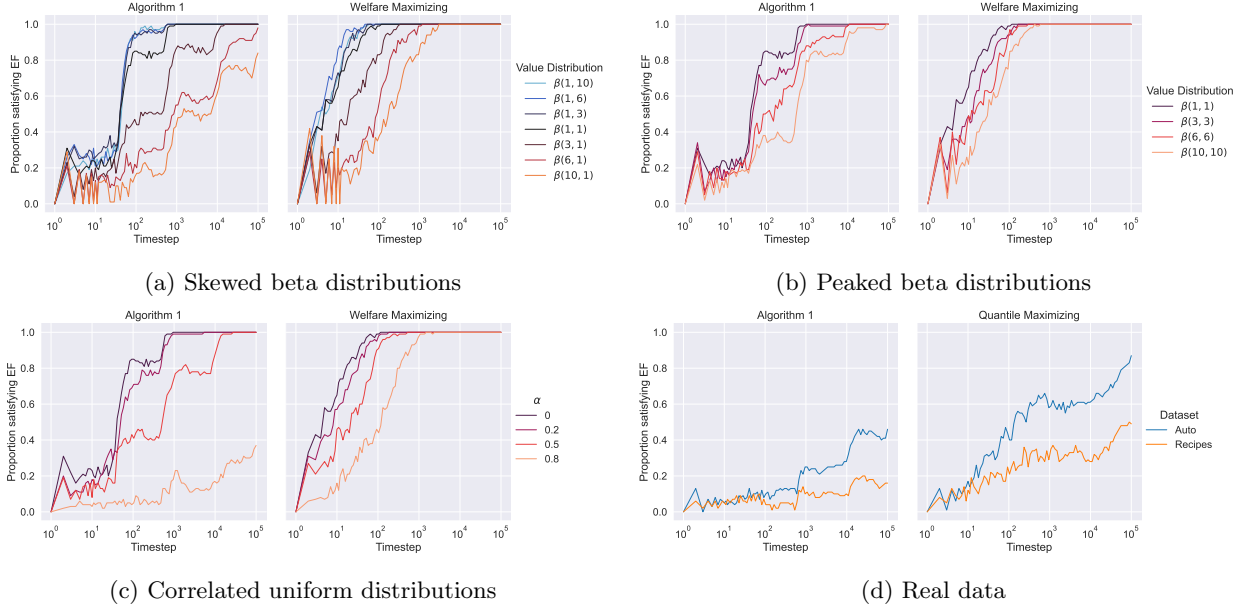


Figure 3 Proportion of envy-free runs on all datasets. For each dataset, the left graphs show the results for Algorithm 1, while the right shows the performance of the “ideal” welfare maximizing algorithm (or quantile maximizing in the case of real data where the underlying value distributions are heterogeneous).

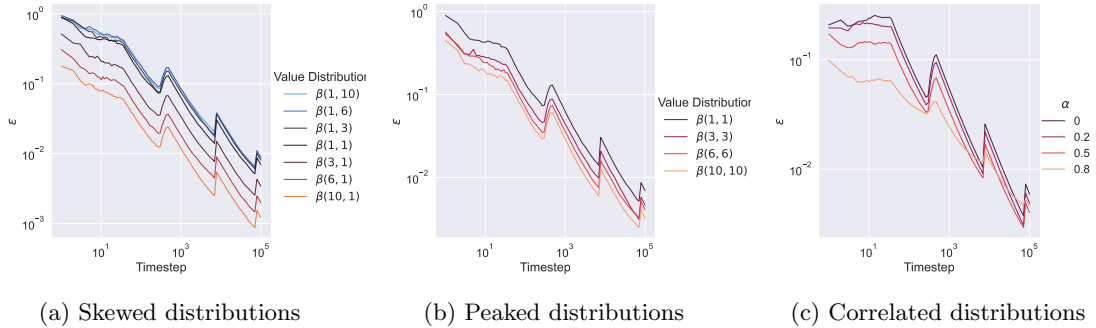


Figure 4 95th percentile welfare approximations. In words, if at timestep 1000 the algorithm has $\varepsilon = .02$, then on 95% of runs, at the thousandth time step, the algorithm achieved at least 98% of the optimal welfare.

is for extremely correlated distributions ($\alpha = 0.8$) where this improvement is counteracted by the fact that additional correlation makes it more difficult to learn agents’ values.

Finally, we compare Algorithm 1 to three variants in Figure 5. We first try an algorithm with shorter epochs: Rather than sampling phases of length k^4 and exploiting phases of length k^8 , we try lengths of k^2 and k^4 , respectively. The result is similar performance overall with additional smoothing: the alternation between degraded performance (during sampling) and good performance (during exploiting) is now more frequent. Next, we try not ending the sampling phase after just k^4 steps. Now items given during the exploit phase are still added to the sample, which should give the algorithm additional information for the rest of the epoch. This lead to a marginal boost

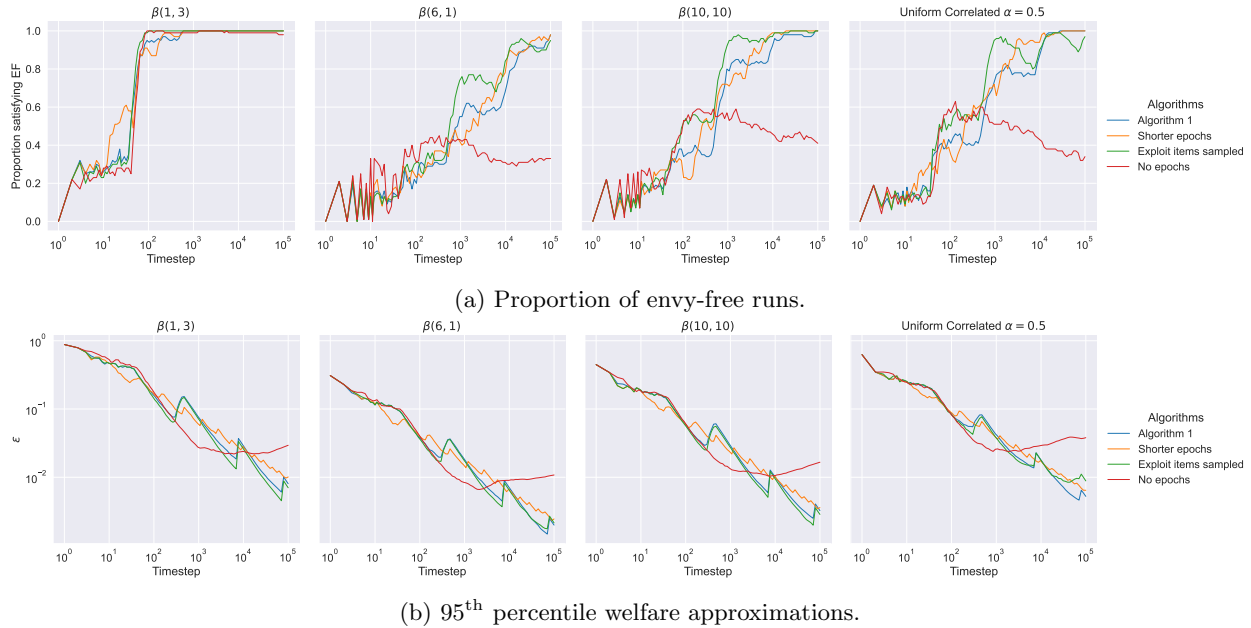


Figure 5 Comparison of alternative algorithms in various settings.

in performance. Finally, we try running the algorithm *without* resetting the sampled items at the end of every epoch. Namely, first agents were given 20 items each, and then all future items are simply added to this sample for comparison. This algorithm performed significantly worse and unfortunately seems to have an asymptote. When the initial 20-item sample is “good,” then the overall run may perform reasonable well, but if we are unlucky and the initial sample is not great, then there is no chance for a later reset. Hence, in many settings, we see that only about half of the runs would lead to EF allocations and welfare appears to cap out well below 99% of optimal without hope further improvements from additional items.

9. Conclusion

To conclude, we have analyzed the online fair division problem when agents only reveal partial information. In multiple settings, we show that ordinal information is enough to obtain strong fairness and efficiency guarantees, even when given as little as binary signals about agent preferences. For food rescue services who are already constrained to eliciting binary preferences this is good news, though we see that the asymptotically optimal algorithms require repeated sampling phases during which items are (purposefully) allocated suboptimally.

Building on this work, there are many other forms of partial information that may be practical to elicit in specific contexts and which may enable different guarantees. For example, if agents can compare small subsets of items, rather than single items, it may be possible to achieve stronger results such as arbitrarily good approximations to PO even in the non-i.i.d. setting. Another interesting direction is to ask what guarantees are possible given a limited time horizon or sample

budget. Now convergence rate matters and, for example, there is a reason to prefer Algorithm 1 over Algorithm 2 in the non-i.i.d. setting. We assume throughout agents are truthful, we leave the study of strategic agents to future work.

References

- Adomavicius G, Tuzhilin A (2005) Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering* 17(6):734–749.
- Aleksandrov MD, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2540–2546.
- Amanatidis G, Birmpas G, Markakis E (2016) On truthful mechanisms for maximin share allocations. *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 31–37.
- Anshelevich E, Filos-Ratsikas A, Shah N, Voudouris AA (2021) Distortion in social choice problems: the first 15 years and beyond. *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 4294–4301.
- Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
- Bai Y, Gözl P (2022) Envy-free and pareto-optimal allocations for agents with asymmetric random valuations. *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, forthcoming.
- Barman S, Khan A, Maiti A (2022) Universal and tight online algorithms for generalized-mean welfare. *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, forthcoming.
- Baumeister D, Bouveret S, Lang J, Nguyen NT, Nguyen TT, Rothe J, Saffidine A (2017) Positional scoring-based allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems* 31(3):628–655.
- Benade G, Kazachkov AM, Procaccia AD, Psomas CA (2018) How to make envy vanish over time. *Proceedings of the 10th ACM Conference on Economics and Computation (EC)*, 593–610.
- Benade G, Nath S, Procaccia AD, Shah N (2021) Preference elicitation for participatory budgeting. *Management Science* 67(5):2813–2827.
- Bogomolnaia A, Moulin H, Sandomirskiy F (2022) On the fair division of a random object. *Management Science* 68(2):1174–1194.
- Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, 387–392.
- Caragiannis I, Nath S, Procaccia AD, Shah N (2017) Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58:123–152.
- Cesa-Bianchi N, Lugosi G (2006) *Prediction, learning, and games* (Cambridge University Press).

- Devanur N, Hartline J, Karlin A, Nguyen T (2011) Prior-independent multi-parameter mechanism design 122–133.
- Dickerson JP, Goldman J, Karp J, Procaccia AD, Sandholm T (2014) The computational rise and fall of fairness. *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 1405–1411.
- Dvoretzky A, Kiefer J, Wolfowitz J (1956) Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics* 642–669.
- Filos-Ratsikas A, Frederiksen SKS, Zhang J (2014) Social welfare in one-sided matchings: Random priority and beyond. *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, 1–12.
- Friedman E, Psomas CA, Vardi S (2015) Dynamic fair division with minimal disruptions. *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, 697–713.
- Friedman E, Psomas CA, Vardi S (2017) Controlled dynamic fair division. *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 461–478.
- Gentle JE (2019) *Computational statistics* (Springer).
- Gkatzelis V, Halpern D, Shah N (2020) Resolving the optimal metric distortion conjecture. *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, 1427–1438.
- Gkatzelis V, Psomas A, Tan X (2021) Fair and efficient online allocations with normalized valuations. *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 5440–5447.
- Goel A, Krishnaswamy AK, Munagala K (2017) Metric distortion of social choice rules: Lower bounds and fairness properties. *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 287–304.
- Gorokh A, Banerjee S, Iyer K (2021) From monetary to nonmonetary mechanism design via artificial currencies. *Mathematics of Operations Research* 46(3):835–855.
- Halpern D, Shah N (2021) Fair and efficient resource allocation with partial information. *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 5456–5463.
- He J, Procaccia AD, Psomas A, Zeng D (2019) Achieving a fairer future by changing the past. *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)* 343–349.
- Hoi SC, Sahoo D, Lu J, Zhao P (2021) Online learning: A comprehensive survey. *Neurocomputing* 459:249–289.
- Kash I, Procaccia AD, Shah N (2014) No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research* 51:579–603.
- Kempe D (2020) Communication, distortion, and randomness in metric voting. *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2087–2094.

- Kroer C, Peysakhovich A, Sodomka E, Stier-Moses NE (2022) Computing large market equilibria using abstractions. *Operations research* (1):329–351.
- Kurokawa D, Procaccia AD, Wang J (2016) When can the maximin share guarantee be guaranteed? *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 523–529.
- Mandal D, Shah N, Woodruff DP (2020) Optimal communication-distortion tradeoff in voting. *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 795–813.
- Manurangsi P, Suksompong W (2021) Closing gaps in asymptotic fair division. *SIAM Journal on Discrete Mathematics* 35(2):668–706.
- Massart P (1990) The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The annals of Probability* 1269–1283.
- MEANS database (2023) Organization website and faqs. URL <https://meansdatabase.org/>.
- Munagala K, Wang K (2019) Improved metric distortion for deterministic social choice rules. *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, 245–262.
- Nguyen NT, Nguyen TT, Rothe J (2017) Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 262–271.
- Prendergast C (2022) The allocation of food to food banks. *Journal of Political Economy* 130(8):1993–2017.
- Procaccia AD, Rosenschein JS (2006) The distortion of cardinal preferences in voting. *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, 317–331.
- Ruhe G, Fruhwirth B (1990) ε -optimality for bicriteria programs and its application to minimum cost flows. *Computing* 44(1):21–34.
- Shi ZR, Lizarondo L, Fang F (2021) A recommender system for crowdsourcing food rescue platforms. *Proceedings of the Web Conference 2021*, 857–865.
- Vardi S, Psomas A, Friedman E (2022) Dynamic fair resource division. *Mathematics of Operations Research* 47(2):945–968.
- Zeng D, Psomas A (2020) Fairness-efficiency tradeoffs in dynamic fair division. *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 911–912.

Details omitted from the main body

EC.1. Proofs

EC.1.1. Proof of Lemma 1

We focus on quantile maximization. The same proof goes through essentially unchanged for the threshold algorithm; we explain the differences whenever it's appropriate.

Fix distributions D_1, \dots, D_n with CDFs F_1, \dots, F_n . Fix two agents i and j . We will show with high probability, i does not envy j (in a strong sense). Union bounding over all $\binom{n}{2}$ pairs yields the lemma statement.

As in [Dickerson et al. \(2014\)](#), we compare the expected contribution of an item t to i 's bundle and its expected contribution to j 's bundle. Let A be the random variable denoting the agent that received the item. We want to consider the difference $\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i]] - \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = j]]$. Let H_i be the event that $F_i(v_{i,t}) \geq \frac{n-1}{n}$, and L_i be the compliment. We split each of the two terms into conditional expectations depending on the signal, beginning with the first one:

$$\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i]] = \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i] \mid H_i] \cdot \Pr[H_i] + \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i] \mid L_i] \cdot \Pr[L_i].$$

Note that, under quantile maximization, $v_{i,t}$ is positively correlated with $\mathbb{I}[A = i]$: for any fixed value $v_{i,t}$, $A = i$ with probability $F(v_{i,t})^{n-1}$, which is increasing in $v_{i,t}$. Therefore, the expectation of the product is greater than or equal to the product of the expectations: $\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i] \mid H_i] \geq \mathbb{E}[v_{i,t} \mid H_i] \cdot \Pr[A = i \mid H_i]$ and $\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i] \mid L_i] \geq \mathbb{E}[v_{i,t} \mid L_i] \cdot \Pr[A = i \mid L_i]$. Therefore

$$\begin{aligned} \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i]] &\geq \mathbb{E}[v_{i,t} \mid H_i] \cdot \Pr[A = i \mid H_i] \cdot \Pr[H_i] + \mathbb{E}[v_{i,t} \mid L_i] \cdot \Pr[A = i \mid L_i] \cdot \Pr[L_i] \\ &= \mathbb{E}[v_{i,t} \mid H_i] \cdot \Pr[A = i \text{ and } H_i] + \mathbb{E}[v_{i,t} \mid L_i] \cdot \Pr[A = i \text{ and } L_i]. \end{aligned}$$

For the threshold algorithm, we have equality above, since conditioned on either H_i or L_i , $v_{i,t}$ is independent of $\mathbb{I}[A = i]$, as the allocation depends only on the high vs low signal.

On the other hand, $v_{i,t}$ is negatively correlated with $\mathbb{I}[A = j]$. Therefore

$$\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = j]] \leq \mathbb{E}[v_{i,t} \mid H_i] \cdot \Pr[A = j \text{ and } H_i] + \mathbb{E}[v_{i,t} \mid L_i] \cdot \Pr[A = j \text{ and } L_i].$$

Again, for the threshold algorithm, we have equality.

Combined, we have

$$\mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i]] - \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = j]] \geq \mathbb{E}[v_{i,t} \mid H_i] \cdot (\Pr[A = i \text{ and } H_i] - \Pr[A = j \text{ and } H_i]) \quad (\text{EC.1})$$

$$- \mathbb{E}[v_{i,t} \mid L_i] \cdot (\Pr[A = j \text{ and } L_i] - \Pr[A = i \text{ and } L_i]). \quad (\text{EC.2})$$

We analyze (EC.1), $\Pr[A = i \text{ and } H_i] - \Pr[A = j \text{ and } H_i]$. Let H_j be the event that $F_j(v_{j,t}) \geq \frac{n-1}{n}$. Let L_j be its complement. We have:

$$\begin{aligned} & (\Pr[A = i \text{ and } H_i \text{ and } L_j] + \Pr[A = i \text{ and } H_i \text{ and } H_j]) \\ & - (\Pr[A = j \text{ and } H_i \text{ and } L_j] + \Pr[A = j \text{ and } H_i \text{ and } H_j]). \end{aligned}$$

Notice that $\Pr[A = j \text{ and } H_i \text{ and } L_j] = 0$ because if agent i has a high quantile and j has a low quantile, j cannot receive the item (in either algorithm). Additionally, by symmetry, $\Pr[A = i \text{ and } H_i \text{ and } H_j] = \Pr[A = j \text{ and } H_i \text{ and } H_j]$. Therefore, (EC.1) simplifies to $\Pr[A = i \text{ and } H_i \text{ and } L_j]$. Finally, we note that $\Pr[A = i \text{ and } H_i \text{ and } L_j] \geq \frac{1}{n-1}$, again, for both algorithms.

We analyze (EC.2), $\Pr[A = j \text{ and } L_i] - \Pr[A = i \text{ and } L_i]$. Let \mathcal{E}^{low} be the event that all agents other than i have quantile lower than $\frac{n-1}{n}$. Let $\overline{\mathcal{E}^{\text{low}}}$ be its complement, the probability that at least one agent other than i has a high quantile. We have:

$$\begin{aligned} & \left(\Pr[A = j \text{ and } L_i \text{ and } \mathcal{E}^{\text{low}}] + \Pr[A = j \text{ and } L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}] \right) \\ & - \left(\Pr[A = i \text{ and } L_i \text{ and } \mathcal{E}^{\text{low}}] + \Pr[A = i \text{ and } L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}] \right). \end{aligned}$$

Notice that $\Pr[A = i \text{ and } L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}] = 0$ because if agent i has a low quantile and at least one other agent has a high quantile, i cannot receive the item. Additionally, by symmetry,

$$\Pr[A = j \text{ and } L_i \text{ and } \mathcal{E}^{\text{low}}] = \Pr[A = i \text{ and } L_i \text{ and } \mathcal{E}^{\text{low}}],$$

since when all agents have low quantiles, i and j are equally likely to receive the item. Hence, the probability in (EC.2) simplifies to $\Pr[A = j \text{ and } L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}] = \Pr[A = j | L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}] \cdot \Pr[L_i \text{ and } \overline{\mathcal{E}^{\text{low}}}]$. The first term is equal to $1/(n-1)$, since j is equally likely to receive the item compared to any agent. The second term is $\frac{n-1}{n} \cdot \left(1 - \left(\frac{n-1}{n}\right)^{n-1}\right)$. Observing that $\left(\frac{n-1}{n}\right)^{n-1} = \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{e}$, we have that the probability in (EC.2) is at most $\frac{1}{n}(1 - \frac{1}{e})$.

Overall, we have shown that

$$\begin{aligned} \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = i]] - \mathbb{E}[v_{i,t} \cdot \mathbb{I}[A = j]] & \geq \mathbb{E}[v_{i,t} \mid H_i] \frac{1}{n-1} - \mathbb{E}[v_{i,t} \mid L_i] \frac{e-1}{en} \\ & \geq \frac{e-1}{en} (\mathbb{E}[v_{i,t} \mid H_i] - \mathbb{E}[v_{i,t} \mid L_i]) \\ & \geq \frac{1}{2n} (\mathbb{E}[v_{i,t} \mid H_i] - \mathbb{E}[v_{i,t} \mid L_i]) \\ & = \frac{1}{2n} (\mathbb{E}[X_i \mid H_i] - \mathbb{E}[X_i \mid L_i]) \end{aligned}$$

It remains to show that the value of i for A_i^T is at least her value for A_j^T plus $\frac{1}{4n} (\mathbb{E}[X_i \mid H_i] - \mathbb{E}[X_i \mid L_i])$ with high probability. Towards this, notice that the value of i for A_i^T

minus her value for A_j^T is the sum of T i.i.d. random variables, supported in $[-1, 1]$, whose expectation is at least $\frac{1}{2n} (\mathbb{E}[X_i | H_i] - \mathbb{E}[X_i | L_i])$, as we've established so far. Hoeffding's inequality then implies that the probability that this difference is less than $b = \frac{1}{4n} (\mathbb{E}[X_i | H_i] - \mathbb{E}[X_i | L_i])$ is at most $2 \exp\left(-\frac{b^2 T}{2}\right)$, i.e., exponentially small, since b is a constant. Observing that $\mathbb{E}[X_i | H_i] \geq \mathbb{E}[X_i | Q_i \geq 1/2]$ and $\mathbb{E}[X_i | L_i] \leq \mathbb{E}[X_i]$ concludes the proof. \square

EC.1.2. Proof of Lemma 6

Fix such an ε , δ , τ , and ℓ . We claim that a sufficient condition for ε -accuracy is that all agents accept an item with quantile within $q^* \pm \varepsilon/(2n)$. Indeed, note that any sampled quantile outside this range will be classified (as high vs low) correctly. With such an error tolerance, the probability a specific agent's quantile (for a fresh item) falls within this range is at most ε/n . Via a union bound over all n agents, the probability that *no* agent has a quantile (for a fresh item) within this range is at least $1 - \varepsilon$. Hence, all that needs to be shown is that with probability $1 - \delta$, all agents accept an item and the accepted item has quantile within the allowed range.

Since there are τ trials, there are at most $n\tau$ items tested across all agents. We show that ℓ is large enough such that with probability $1 - \delta/2$, all these tests are within $\pm\varepsilon/(6n)$ of the true value. Using Hoeffding's inequality, the probability any specific test fails is at most

$$2 \exp\left(-2 \left(\frac{\varepsilon}{6n}\right)^2 \cdot \ell\right) = 2 \exp\left(-\frac{\varepsilon^2}{18n^2} \cdot \ell\right) \leq 2 \exp\left(-\ln\left(\frac{4\tau n}{\delta}\right)\right) = \frac{\delta}{2\tau n},$$

a union bound over all $n\tau$ tests yields the required probability.

Note that under the condition that all the tests are this accurate, since the threshold for acceptance is $\pm\varepsilon/(3n)$, any accepted item will be within $\pm\varepsilon/(2n)$ of q^* , as needed. What remains to be shown is that each agent will, with reasonable probability, accept an item. To that end, we need to show that with probability $1 - \delta/2$, all agents will test an item that is within $\pm\varepsilon/(6n)$ of q^* . If such an item is tested and the test is accurate, the empirical estimate of its quantile is within $\pm\varepsilon/(3n)$, and the item would hence be accepted. A union bound will then tell us that both of these events would occur with probability $1 - \delta$.

Towards proving that each agent will test an item within $\pm\varepsilon/(6n)$ of q^* with probability $1 - \delta/2$, we use a union bound, showing that each agent individually will *not* sample such an item with probability at most $\delta/(2n)$. In each of τ trials, the probability such an item is sampled is $\varepsilon/(3n)$. Hence, the probability no such item is sampled is $\left(1 - \frac{\varepsilon}{3n}\right)^\tau$. We then have that

$$\begin{aligned} \left(1 - \frac{\varepsilon}{3n}\right)^\tau &= \left(1 - \frac{1}{\frac{3n}{\varepsilon}}\right)^\tau \\ &= \left(\left(1 - \frac{1}{\frac{3n}{\varepsilon}}\right)^{\frac{3n}{\varepsilon}}\right)^{\tau \cdot \frac{\varepsilon}{3n}} \end{aligned}$$

$$\begin{aligned}
&\leq (e^{-1})^{\tau \cdot \frac{\varepsilon}{3n}} && (1 - 1/x)^x \leq e^{-1} \text{ for all } x \geq 1 \\
&= e^{-\tau \cdot \frac{\varepsilon}{3n}} \\
&\leq e^{-\ln(2n/\delta)} \\
&= \frac{\delta}{2n},
\end{aligned}$$

as needed. \square

EC.1.3. Proof of Lemma 7

First, we prove that for all $k \geq 10n$, epoch k is ε_k -accurate with probability δ_k for $\varepsilon_k = 3n/k^2$ and $\delta_k = 2ne^{-k}$. Since $k > 3n$ these are valid values between 0 and 1. Hence, we simply need to check that the τ and ℓ inequalities hold for the number of trials and number of test items specified in Algorithm 2. For arbitrary epoch k ,

$$\frac{\ln(2n/\delta_k)}{\varepsilon_k/(3n)} = \ln(e^k)k^2 = k^3,$$

so the number of trials is sufficiently large. Further,

$$\begin{aligned}
\frac{18n^2}{\varepsilon_k^2} \ln\left(\frac{4k^3n}{\delta_k}\right) &= 2k^4 \cdot \ln(2k^3e^k) \\
&\leq 2k^4 \cdot \ln(k^4e^k) && (k \geq 2) \\
&= 2k^4 \cdot (k + 4 \ln k) \\
&\leq 2k^4 \cdot (k + 4k) && (\ln k < k) \\
&\leq 10k^5 \\
&\leq k^6. && (k \geq 10)
\end{aligned}$$

Recall that $k(t)$ is defined as the epoch of item t . As in the proof of Lemma 5, we characterize deviations from the ideal algorithm in four ways.

1. Item t was allocated in one of the first $10n - 1$ epochs; that is, $k(t) < 10n$.
2. Item t was allocated during the sampling phase of epoch $k(t) \geq 10n$.
3. Item t was allocated during the ranking phase of epoch $k(t) \geq 10n$, which was $\varepsilon_{k(t)}$ -accurate.
4. Item t was allocated during the ranking phase of epoch $k(t) \geq 10n$, which was not $\varepsilon_{k(t)}$ -accurate.

We say an item t is incorrect (incorrectly allocated) when it is given to an agent with non-maximum quantile for it. We show that the number of mistakes in each category are bounded by $10^{20}n^{19}$, $2T^{5/9}$, $7nT^{17/18}$ and $1.3 \cdot 10^{16}n$ respectively, with high probability. This implies, via a union bound, that the total number of mistakes is at most the sum of these quantities, or $O(\text{poly}(n) \cdot T^{17/18})$, with high probability.

The number of items in category 1, is at most

$$\sum_{k=1}^{10n} k^9 + k^{18} \leq (10n)^{10} + (10n)^{19} \leq 10^{20} n^{19}$$

Notice that $T \geq \sum_{k=1}^{k(T)-1} k^9 + k^{18} \geq (k(T) - 1)^{18}$, and therefore $k(T) \leq 2T^{1/18}$.

For the second category, since $k(T) \leq 2T^{1/18}$, the total number of items in the sampling phase is (with probability 1) upper bounded by

$$\sum_{k=1}^{k(T)} k^9 \leq k(T)^{10} \leq 2T^{5/9}.$$

For the third category, note that each item t in this category has probability $\varepsilon_{k(t)}$ of being incorrect. The expected number of mistakes is at most

$$\sum_{k=10n}^{k(T)} \varepsilon_{k(t)} k^{18} = \sum_{k=10n}^{k(T)} 3nk^{16} \leq 3nk(T)^{17} \leq 6nT^{17/18}.$$

Using Hoeffding's inequality we get that with high probability the number of mistakes is at most $7nT^{17/18}$, since a deviation of $nT^{17/18}$ occurs with probability at most $\exp(-2n^2T^{17/9}/T) = \exp(-2n^2T^{8/9})$.

For the fourth category, the expected number of items in this category is at most

$$\sum_{k=10n}^{k(T)} \delta_k k^{18} = 2n \sum_{k=10n}^{k(T)} \frac{k^{18}}{e^k} \leq 2n \sum_{k=1}^{\infty} \frac{k^{18}}{e^k} \leq 1.3 \cdot 10^{16} n.$$

Using Markov's inequality we have that the number of mistakes is at most $1.3 \cdot 10^{16} n \ln(T)$ with probability at least $1 - \ln(T)$, i.e., with high probability. \square

EC.1.4. Proof of Theorem 4

The proof is nearly identical to the proof of Theorem 2. Fix a distribution D with CDF F and let X be a random variable with distribution D . Fix some ε to be $(1 - 1/e) - \varepsilon$ welfare maximizing. Let \mathcal{E}_1^T be the event that the maximum social welfare at time T is at least $1/2 \cdot \mathbb{E}[X] \cdot T$, let \mathcal{E}_2^T be the event the ideal threshold algorithm is c -strongly-EF for $c = \frac{(\mathbb{E}[X] \mid F(X) \geq 1/2] - \mathbb{E}[X])}{4n}$, let \mathcal{E}_3^T be the event that the ideal threshold algorithm is a $(1 - 1/e)^2 - \varepsilon/2$ approximation to welfare, and let \mathcal{E}_4^T be the event that Algorithm 2 differs from the ideal threshold algorithm on at most $f(T)$ items from Lemma 7. We first claim that $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T \cap \mathcal{E}_4^T$ occurs with high probability in T . Note that Lemmas 2, 1, and 7 tell us each of \mathcal{E}_2^T , \mathcal{E}_3^T , and \mathcal{E}_4^T occur with high probability. For \mathcal{E}_1^T , the maximum value for each item is in expectation at least the expected value for a single agent $\mathbb{E}[X]$. Hence, a Chernoff bound tells us \mathcal{E}_1^T occurs with probability at least $1 - \exp\left(\frac{-\mathbb{E}[X]T}{8}\right)$, i.e., with high probability. The claim holds because the intersection of a finite number of high probability events occurs with high probability.

Next, note that for sufficiently large T , since $f(T) \in o(T)$, $f(T) \leq \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])}{8n} \cdot T$ and $f(T) \leq \varepsilon/4 \cdot \mathbb{E}[X] \cdot T$. Fix such a sufficiently large T . We show that conditioned on $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T \cap \mathcal{E}_4^T$, both EF and $((1 - 1/e)^2 - \varepsilon)$ -welfare hold. Recall that a “difference” between Algorithm 2 and the ideal threshold algorithm refers to different distributions over the agents that get some item (i.e., a different randomized allocation). In order to make statements about envy-freeness and efficiency we need a way to argue about the differences between the algorithms ex-post. However, notice that without loss of generality we can couple the decision made by the two algorithms when randomized allocation is the same; that is, when Algorithm 2 does not differ from the ideal threshold algorithm we can assume without loss of generality that the agent who gets the item is the same. Let $A^{IT} = (A_1^{IT}, \dots, A_n^{IT})$ be the allocation of the ideal threshold algorithm and $A = (A_1, \dots, A_n)$ be the allocation of Algorithm 2. Beginning with envy-freeness, we have that for all pairs of agents i and j ,

$$\begin{aligned} v_i(A_i) &\stackrel{(\mathcal{E}_4^T)}{\geq} v_i(A_i^{IT}) - f(T) \\ &\stackrel{(\mathcal{E}_2^T)}{\geq} v_i(A_j^{IT}) - f(T) + \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])T}{4n} \\ &\stackrel{(\mathcal{E}_4^T)}{\geq} v_i(A_j) - 2f(T) + \frac{(\mathbb{E}[X \mid F(X) \geq 1/2] - \mathbb{E}[X])T}{4n} \\ &\geq v_i(A_j), \end{aligned}$$

so the allocation is envy-free. Let A^* be a welfare-maximizing algorithm. For the welfare approximation, we then have

$$\begin{aligned} \frac{\text{sw}(A)}{\text{sw}(A^*)} &= \frac{\text{sw}(A^{IT}) - (\text{sw}(A^{IT}) - \text{sw}(A))}{\text{sw}(A^*)} \\ &\stackrel{(\mathcal{E}_4^T)}{\geq} \frac{\text{sw}(A^{IT}) - f(T)}{\text{sw}(A^*)} \\ &= \frac{\text{sw}(A^{IT})}{\text{sw}(A^*)} - \frac{f(T)}{\text{sw}(A^*)} \\ &\stackrel{(\mathcal{E}_3^T)}{\geq} (1 - 1/e)^2 - \varepsilon/2 - \frac{f(T)}{\text{sw}(A^*)} \\ &\stackrel{(\mathcal{E}_1^T)}{\geq} (1 - 1/e)^2 - \varepsilon/2 - \frac{f(T)}{1/2 \cdot \mathbb{E}[X] \cdot T} \\ &\geq (1 - 1/e)^2 - \varepsilon/2 - \frac{\varepsilon/4 \cdot \mathbb{E}[X] \cdot T}{1/2 \cdot \mathbb{E}[X] \cdot T} \\ &= (1 - 1/e)^2 - \varepsilon, \end{aligned}$$

as needed. \square

EC.1.5. Proof of Theorem 7

The proof of envy-freeness for each algorithm is nearly identical to Theorems 2 and 4 respectively; we show it here for completeness. We focus on Algorithm 1. The proof for Algorithm 2 goes through

identically with all occurrences of quantile maximization replaced with the ideal threshold algorithm and all occurrences of Lemma 5 replaced with Lemma 7.

Fix a distributions D_1, \dots, D_n with CDFs F_1, \dots, F_n and let X_i be a random variable with distribution D_i . Fix some ε to be $(1/e - \varepsilon)$ -PO. Let $E = \min_{i \in \mathcal{N}} \mathbb{E}[X_i]$ be the minimum expected value for all agents. Let \mathcal{E}_1^T be the event that each agent i 's value for their bundle at time T is at least $1/(2n) \cdot E \cdot T$, let \mathcal{E}_2^T be the event that quantile maximization is c -strongly-EF for $c = \min_{i \in \mathcal{N}} \frac{(\mathbb{E}[X_i \mid F_i(X_i) \geq 1/2] - \mathbb{E}[X_i])}{4n}$, let \mathcal{E}_3^T be the event that quantile maximization is a $(1/e - \varepsilon/2)$ -PO, and let \mathcal{E}_4^T be the event that Algorithm 1 differs from quantile maximization on at most $f(T)$ items from Lemma 5. We first claim that $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T \cap \mathcal{E}_4^T$ occurs with high probability in T . Note that Lemmas 1, 3, and 5 tell us $\mathcal{E}_2^T, \mathcal{E}_3^T$, and \mathcal{E}_4^T each occur with high probability, respectively. For \mathcal{E}_1^T , note that under quantile maximization, the probability each agent i receives an item is exactly $1/n$ and the expected value conditioned on receiving the item is at least $\mathbb{E}[X_i] \geq E$. Hence, the expected contribution of each item to $v_i(A_i)$ is at least $1/n \cdot E$. A Chernoff bound then tells us \mathcal{E}_1^T holds for agent i with probability at least $1 - \exp(-\frac{ET}{8n})$. A union bound over all agent's tells us this occurs simultaneously for all agents with probability at least $1 - n \exp(-\frac{ET}{8n})$, i.e., with high probability. The claim holds because the intersection of a finite number of high probability events occurs with high probability.

Next, note that for sufficiently large T , since $f(T) \in o(T)$, $f(T) \leq \min_{i \in \mathcal{N}} \frac{(\mathbb{E}[X_i \mid F_i(X_i) \geq 1/2] - \mathbb{E}[X_i])}{8n} \cdot T$ and $f(T) \leq \varepsilon/(4n) \cdot ET$. Fix such a sufficiently large T . We show that conditioned on $\mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T \cap \mathcal{E}_4^T$, both EF and $(1/e - \varepsilon)$ -PO hold. Let $A^{QM} = (A_1^{QM}, \dots, A_n^{QM})$ be the allocation of quantile maximization and $A = (A_1, \dots, A_n)$ be the allocation of Algorithm 2. Beginning with envy-freeness, we have that for all pairs of agents i and j ,

$$\begin{aligned} v_i(A_i) &\geq^{(\mathcal{E}_4^T)} v_i(A_i^{QM}) - f(T) \\ &\geq^{(\mathcal{E}_2^T)} v_i(A_j^{QM}) - f(T) + \min_{i \in \mathcal{N}} \frac{(\mathbb{E}[X_i \mid F_i(X_i) \geq 1/2] - \mathbb{E}[X_i])}{4n} \\ &\geq^{(\mathcal{E}_4^T)} v_i(A_j) - 2f(T) + \min_{i \in \mathcal{N}} \frac{(\mathbb{E}[X_i \mid F_i(X_i) \geq 1/2] - \mathbb{E}[X_i])}{4n} \\ &\geq v_i(A_j), \end{aligned}$$

so the allocation is envy-free. Next, we show that for all agents $\frac{v_i(A_i)}{v_i(A_i^{QM})} \geq 1 - \varepsilon/2$. Since A^{QM} is $(1/e - \varepsilon/2)$ -PO under \mathcal{E}_3^T , this implies that A is a $(1/e - \varepsilon/2)(1 - \varepsilon/2) \geq 1/e - \varepsilon$ approximation to PO as well.

To that end, for each agent i we have

$$\frac{v_i(A_i)}{v_i(A_i^{QM})} = \frac{v_i(A_i^{QM}) - (v_i(A_i^{QM}) - v_i(A_i))}{v_i(A_i^{QM})}$$

$$\begin{aligned}
&= 1 - \frac{(v_i(A_i^{QM}) - v_i(A_i))}{v_i(A_i^{QM})} \\
&\geq^{(\varepsilon_4^T)} 1 - \frac{f(T)}{v_i(A_i^{QM})} \\
&\geq^{(\varepsilon_1^T)} 1 - \frac{f(T)}{ET/(2n)} \\
&\geq 1 - \frac{\varepsilon \cdot ET/(4n)}{ET/(2n)} \\
&= 1 - \varepsilon/2,
\end{aligned}$$

as needed. \square

EC.2. Additional computational results

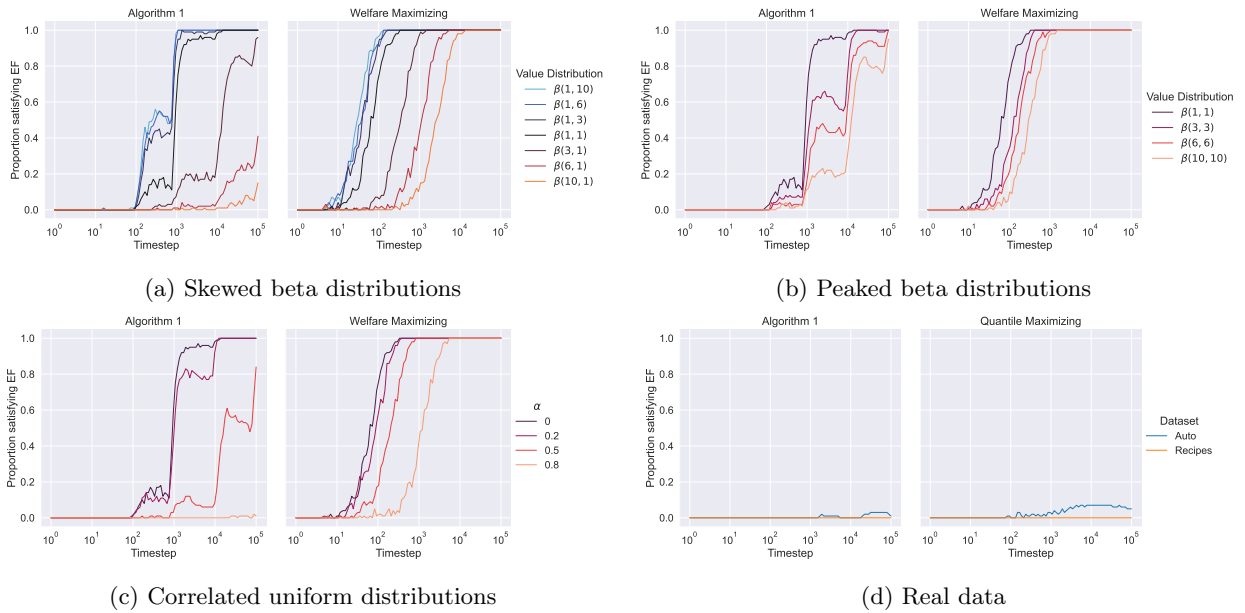


Figure EC.1 Equivalent of Figure 3 with five agents.

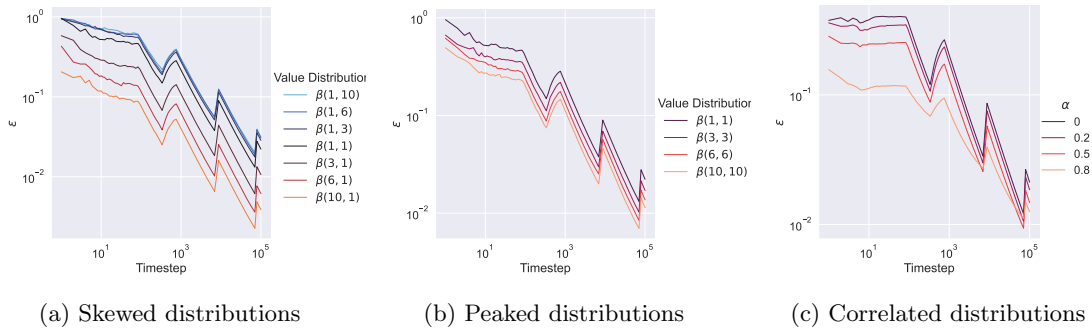




Figure EC.2 Equivalent of Figure 4 with five agents.



(a) Proportion of envy-free runs.

(b) 95th percentile welfare approximations.**Figure EC.3** Equivalent of Figure 5 with five agents.