Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Dynamic Fair Division with Partial Information

Gerdus Benadè ; , Daniel Halpern ; , Alexandros Psomas

To cite this article:
Gerdus Benadè ; , Daniel Halpern ; , Alexandros Psomas (2025) Dynamic Fair Division with Partial Information. Operations Research

Published online in Articles in Advance 15 Jan 2025

. https://doi.org/10.1287/opre.2023.0608

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Articles in Advance, pp. 1–21
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Dynamic Fair Division with Partial Information

Gerdus Benadè,[a,*] Daniel Halpern,[b] Alexandros Psomas[c]

[a] Boston University, Boston, Massachusetts 02215; [b] Harvard University, Cambridge, Massachusetts 02139; [c] Purdue University, West Lafayette, Indiana 47907

*Corresponding author

Contact: benade@bu.edu, https://orcid.org/0000-0002-0024-0026 (GB); dhalpern@g.harvard.edu, https://orcid.org/0000-0003-3368-7235 (DH); apsomas@cs.purdue.edu, https://orcid.org/0000-0002-7709-5058 (AP)

Received: November 4, 2023
Revised: August 16, 2024; October 25, 2024
Accepted: November 13, 2024
Published Online in Articles in Advance: January 15, 2025

Area of Review: Market Analytics and Revenue Management

https://doi.org/10.1287/opre.2023.0608

Abstract. 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 with a single previous item. For independent but nonidentical agents, we obtain envy-freeness and a constant approximation to Pareto efficiency with high probability. Our results are asymptotically tight. A computational study shows that envy-freeness and efficiency can be achieved on practical time-horizons.

Funding: D. Halpern is supported by the National Science Foundation Graduate Research Fellowship Program [Grant DGE1745303]. A. Psomas is supported in part by an NSF CAREER award (Division of Computing and Communication Foundations) [Grant CCF-2144208], a Google AI for Social Good award, and research awards from Google and Supra.

Supplemental Material: All supplemental materials, including the computer code and data that support the findings of this study, are available at https://doi.org/10.1287/opre.2023.0608.

Keywords: 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, 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 nonidentical agents (where $v_{i,t}$ is drawn from an agent-specific distribution D_i) for the algorithm that maximizes weighted welfare (Bai and Gölz 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, eliciting numerical valuations is often deemed too difficult when low volume, the unpredictability of donations arriving. and the cognitive burden of elicitation may prevent recipients from forming regular habits of reporting valuations, or when it is difficult to compare reports between recipients. Furthermore, interpersonal comparisons of reported utilities are quite controversial (Robbins 1938). Because of this, many real-world settings discussed in the literature involve much simpler forms of eliciting agents' interest than reporting cardinal utilities. 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, before continuing to the next recipient, if necessary; 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, for example, its rank relative to a subset of the past items allocated to this agent, or even just a pairwise comparison with 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 with a subset of prior items allocated to agent i, but otherwise knows nothing about the underlying numerical valuation nor the value distribution. We view this lack of additional knowledge as a key feature of our algorithm, aligning with the Wilson doctrine (Wilson 1985), that mechanisms should not rely on agents' underlying beliefs and value distributions. Developing this algorithm requires balancing exploration and exploitation. We need enough reference items to "estimate" values accurately but not too many to avoid inefficiency. We alternate between these goals with carefully timed phases to achieve the desired properties.

Given this strong positive result, we explore the limits of what we can achieve when further restricting the amount of information available. Indeed, even ranking an item among arbitrary received ones may be too demanding if the reference items were given hundreds of time steps ago. What if each agent can remember only a *single* item previously allocated to them, and the fresh item is compared with just this one item? That is, the algorithm only learns whether the new item is better or worse than the item in memory and may, at that time, choose to replace the item in memory. 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 $1-1/e-\varepsilon$ approximation to the optimal welfare with high probability, for all $\varepsilon > 0$ (Theorem 4). It again requires no extra information about the underlying numerical values or distribution, only making use of the elicited comparisons. In addition, 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 (Theorem 3); therefore, a constant approximation like Algorithm 2 is all we can hope for.

Next, we relax the i.i.d. assumption and show that our algorithms are still effective when agents are correlated or nonidentical. First, we consider agents that agree on a noisy common valuation of each item, so $v_{i,t} = v_t + \epsilon_{i,t}$ for $v_t \sim D^{com}, v_{i,t} \sim D^{noise}$. Algorithm 1 (with some small modifications) 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 $\frac{1+\sqrt{5}}{4} \approx 0.809$ 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 nonidentical agents and Algorithm 2 elicits strictly less information, one might still prefer to use Algorithm 1 because it has significantly shorter exploration phases.

We conclude with a computational study on both synthetic and real-world value distributions. 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 1,000 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 computer science studies what can be achieved by algorithms that only elicit preferences of limited expressiveness. Procaccia and Rosenschein (2006) consider voting rules that receive ordinal information as input but are evaluated on the cardinal utilities underlying the ordinal reports. They define the notion of distortion to measure the worst-case deterioration of an aggregate cardinal objective (e.g., utilitarian social welfare) because of only having access to ordinal information. Recent works prove bounds on the distortion for many problems in social choice, including matching (Filos-Ratsikas et al. 2014, Anshelevich and Sekar 2016, Abramowitz and Anshelevich 2018, Anshelevich and Zhu 2019), voting (Boutilier et al. 2015, Caragiannis et al. 2017, Goel et al. 2017, Anshelevich et al. 2018, Munagala and Wang 2019, Gkatzelis et al. 2020, Kempe 2020, Mandal et al. 2020, Kizilkaya and Kempe 2022, Charikar et al. 2024), 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. (2021). 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, for example, Aziz et al. (2015), Bouveret et al. (2010), Baumeister et al. (2017), and 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 because of 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; Gorokh et al. 2021; Barman et al. 2022; Vardi et al. 2022), and we note that the welfare-maximizing algorithms of Dickerson et al. (2014), Kurokawa et al. (2016), and Bai and Gölz (2022) work in the dynamic setting, even though the 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 a 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, whereas 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 nontrivial 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.

We assume fixed agents and items that arrive over time; however, other models of online allocation have also been studied with the dual objectives of fairness and efficiency. For example, Sinclair et al. (2022) consider a model with a fixed pool of resources where agents arrive over time and a core decision is how much to allocate in this time step versus how much to save for the future.

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 the value from the expert chosen on that day (in the full feedback variant), and the 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 (Dhangwatnotai et al. 2010), where the task is to design mechanisms that perform well in the worst case even compared with the tailor-made mechanism which knows the distributions.

2. Preliminaries

A set of T indivisible items/goods, labeled by $\mathcal{G} = \{1, 2, \ldots, T\}$, needs to be allocated to a set of n agents, labeled by $\mathcal{N} = \{1, \ldots, n\}$. 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 $\boldsymbol{A}$ is a partition of the items into bundles $A_1, \ldots, A_n$, where A_i is the set of items assigned to agent $i \in \mathcal{N}$. Each allocation has an associated utility profile $\boldsymbol{v}(\boldsymbol{A}) = (v_1(A_1), \ldots, v_n(A_n))$.

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 $\mathcal{G}^t = \{1, 2, \ldots, t\}$ 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 cumulative distribution function (CDF) F, that is, $v_{i,t} \sim D$ 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 a 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, that is, $v_{i,t} \sim D_i$ for all $i \in \mathcal{N}$ and $t \in \mathcal{G}$.

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

2.1. 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, that is, $M \subseteq A_i^{t-1}$, $|M| \le 1$. In the other extreme, t is compared with 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 $\sigma_i^{-1}(S, 4) = 2$.

2.2. 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 $D_1, \ldots, D_n$. Let $\mathcal{E}_P(t)$ be the event that some algorithm satisfies property P (e.g., envy-freeness or Pareto optimality (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 $\lim_{t \to \infty} \Pr[\mathcal{E}_P(t)] = 1$.

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

2.3. Efficiency Notions

An allocation $\boldsymbol{A}$ Pareto dominates an allocation $\boldsymbol{A}'$, denoted $\boldsymbol{A} \succ \boldsymbol{A}'$, when $v_i(A_i) \ge v_i(A_i')$ for all $i \in \mathcal{N}$ and there exists $j \in \mathcal{N}$ with $v_j(A_j) > v_j(A_j')$. An allocation $\boldsymbol{A}$ is *Pareto efficient* or *Pareto optimal* (PO) if there is no feasible (integral) allocation that Pareto dominates it. An allocation $\boldsymbol{A}'$ is in the (one) swap-neighborhood of $\boldsymbol{A}$ when it can be created from $\boldsymbol{A}$ 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 $\boldsymbol{A}$ is *one-swap Pareto optimal* (SPO) if it is undominated by any allocation in its swap-neighborhood. Several notions of approximate Pareto efficiency exist (see, e.g., Ikeda et al. 2001 and Leung et al. 2015); we use the definition by Ruhe and Fruhwirth

(1990) according to which an allocation $\boldsymbol{A}$ is *α-Pareto efficient* when $v(\boldsymbol{A})/\alpha$ is undominated.

The social welfare of an allocation $\boldsymbol{A}$ is $\mathsf{sw}(\boldsymbol{A}) = \sum_{i\in\mathcal{N}} v_i(A_i)$. Let allocation $\boldsymbol{A}^*$ denote a (social) welfare optimal allocation for which $\mathsf{sw}(\boldsymbol{A}^*) \geq \mathsf{sw}(\boldsymbol{A})$ for all feasible allocations $\boldsymbol{A}$. An allocation $\boldsymbol{A}$ provides an α approximation to welfare if $\mathsf{sw}(\boldsymbol{A}) \geq \alpha \cdot \mathsf{sw}(\boldsymbol{A}^*)$. For the notion of approximate efficiency we consider, observe that an α-approximation to welfare implies that the allocation is also α-Pareto efficient.

2.4. Fairness Notions

We focus on a prominent notion of fairness called *envy-freeness*. An allocation $\boldsymbol{A}^T = (A_1^T, \ldots, A_n^T)$ 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 (in fact, several of our lower bounds apply to these stronger algorithms, too). Given quantiles, two 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). Threshold algorithms are natural when the memory length is one, whereas unbounded memory length allows (approximate) quantile maximization.

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, largely unchanged, for the $\frac{n-1}{n}$-threshold algorithm; it can be found in Section EC.1.1 of the Online Appendix.

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 probability $1 - \exp(-\Omega(T))$, where the constant $c = \min_{i\in\mathcal{N}}(\mathbb{E}[V_i | Q_i \geq 1/2] - \mathbb{E}[V_i])/(4n)$.*

Note that c is strictly positive because our distributions are continuous.

Next, we show that in the i.i.d. model, the $\frac{n-1}{n}$-threshold algorithm gives a $1 - \frac{1}{e} - \varepsilon$ approximation to welfare (Lemma 2) with high probability. This approximation is also obtained by a more general result on single threshold algorithms for prophet inequalities of Ehsani et al. (2018), who use the threshold $e^{-1/n}$. Our setting with identical distributions permits a simpler proof, which we provide here for the sake of completeness for threshold $1 - 1/n$, which simplifies some later computations.

Lemma 2. *In the i.i.d. model, the $\frac{n-1}{n}$-threshold algorithm guarantees a $\left(\left(1 - \frac{1}{e}\right) - \varepsilon\right)$ approximation to welfare, with probability $1 - \exp(-\Omega(T))$, for all $\varepsilon > 0$.*

Proof. Let F be the CDF of an arbitrary continuous distribution. Let $\tau = F^{-1}\left(\frac{n-1}{n}\right)$ be the value at the $\frac{n-1}{n}$ threshold. Note that having $Q_i \geq \frac{n-1}{n}$ is equivalent to having $V_i \geq \tau$. We can upper bound the expected maximum value by

$$\begin{aligned}\mathbb{E}[\max_i V_i] &\leq \tau + \mathbb{E}[(\max_i V_i - \tau)_+] \\ &\leq \tau + \sum_i \mathbb{E}[(V_i - \tau)_+] = \tau + n \cdot \mathbb{E}[(V - \tau)_+]\end{aligned}$$

where $(s)_+ := \max(s, 0)$ and V represents a generic draw from $\mathcal{D}$.

The $\frac{n-1}{n}$ threshold algorithm can also be interpreted as follows: pick a random order over the agents and give it to the first one whose value is above τ. We will lower bound the expected welfare generated by each item in this algorithm, ignoring contributions to the welfare when no agent is above the threshold. Fix an arbitrary ordering of the agents. The probability the item is given to the i-th agent is $\Pr[V_i \geq \tau \wedge V_{i'} < \tau \forall i' < i] = \Pr[V_i \geq \tau] \prod_{i' < i} \Pr[V_{i'} < \tau]$ (because values are independent). Conditioned on this event, the value is $\mathbb{E}[V_i | V_i \geq \tau]$. So, the total welfare is

$$\sum_i \mathbb{E}[V_i | V_i \geq \tau] \Pr[V_i \geq \tau] \prod_{i' < i} \Pr[V_{i'} < \tau].$$

Furthermore, $\mathbb{E}[V_i | V_i \geq \tau] = \tau + \mathbb{E}[V_i - \tau | V_i \geq \tau]$. In addition, we can write $\mathbb{E}[V_i - \tau | V_i \geq \tau] \cdot \Pr[V_i \geq \tau] = \mathbb{E}[(V_i - \tau)_+]$. Putting this together, we have that the welfare is

$$\sum_i (\tau \cdot \Pr[V_i \geq \tau] + \mathbb{E}[(V_i - \tau)_+]) \prod_{i' < i} \Pr[V_{i'} < \tau].$$

Now, $\Pr[V \geq \tau] = 1/n$ for $V \sim D$, so we can simplify this to

$$\begin{aligned}&(\tau/n + \mathbb{E}[(V - \tau)_+]) \sum_i (1 - 1/n)^{i-1} \\ &\quad = (\tau/n + \mathbb{E}[(V - \tau)_+]) \cdot \frac{1 - (1 - 1/n)^n}{1 - (1 - 1/n)} \\ &\quad \geq (\tau/n + \mathbb{E}[(V - \tau)_+]) \cdot n \cdot (1 - 1/e) \\ &\quad = (1 - 1/e)(\tau + n \cdot \mathbb{E}[(V - \tau)_+]) \\ &\quad \geq (1 - 1/e)\mathbb{E}[\max_i V_i].\end{aligned}$$

Finally, for any fixed $\varepsilon > 0$, standard Chernoff bounds tell us that with probability $1 - \exp(-\Omega(T))$, the optimal welfare of T items is at most $T \cdot (1 + \varepsilon/2)\mathbb{E}[\max_i V_i]$ whereas the welfare of the threshold algorithm is at least $T \cdot (1-\varepsilon/2)\left(1-\frac{1}{e}\right)\mathbb{E}[\max_i V_i]$. Indeed, the expected optimal welfare is equal to $T \cdot \mathbb{E}[\max_i V_i]$. The standard multiplicative Chernoff bound says that the probability of the sum of i.i.d. variables exceeding $(1+\varepsilon/2)$ times its expectation μ is at most $\exp(-\mu\varepsilon^2/12)$. Plugging in $\mu = T \cdot \mathbb{E}[\max_i V_i]$, we get the desired statement. The statement about the welfare of the threshold algorithm follows similarly. Thus, the algorithm is a

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

approximation to welfare, with probability $1-\exp(-\Omega(T))$. □

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 $F_1, \ldots, F_n$ 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 $A_i^{\text{high}} = \left\{t \in \mathcal{G}^T \mid F_i(v_{i,t}) \geq 1-\frac{1+\varepsilon'}{n}\right\}$ 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 one 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 $\left(F_i(v_{i,t}) \geq 1-\frac{1+\varepsilon'}{n}\right)$ 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

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

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

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

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 $\exp((-\varepsilon'^2(1+\varepsilon')\mathbb{E}_{Q\sim\text{Unif}[1-1/n,1]}[F_i^{-1}(Q)]\cdot T)/3n)$.

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 $\frac{1}{en}\cdot\mathbb{E}_{Q\sim\text{Unif}\left[1-\frac{1}{n},1\right]}[F_i^{-1}(Q)]$. Indeed, consider an item such that the quantile for agent i is $Q_i > 1-1/n$ whereas $Q_j < 1-1/n$ for all agents $j \neq i$. This occurs with probability $\frac{1}{n}\cdot\left(1-\frac{1}{n}\right)^{n-1} \geq 1/en$, and when this occurs, because 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 $\mathbb{E}_{Q\sim\text{Unif}\left[1-\frac{1}{n},1\right]}[F_i^{-1}(Q)]$, because it is independent of the other agent's values. Hence, the expectation is at least $\frac{1}{en}\mathbb{E}_{Q\sim\text{Unif}\left[1-\frac{1}{n},1\right]}[F_i^{-1}(Q)]$. Finally, we again use a multiplicative Chernoff bound to show that

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

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 $A_1, \ldots, A_n$ is $(1/e-\varepsilon)$-PO. Fix an arbitrary allocation $A'_1, \ldots, A'_n$. 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 $|A_i'| \le T/n$. Because A_i' can be at most as valuable as the most valuable T/n items, we have

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

as needed. □

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 $g:[0,1] \to [0,1]$ be the function $g(q) = \mathbb{E}[V \mid V \le F^{-1}(q)]/\mathbb{E}[V]$, 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 $g(\hat{q}) \ge 0.9$. Let $q_2^* = \max(\hat{q}, 0.9)$. Because g is increasing, $g(q_2^*) \ge g(\hat{q}) \ge 0.9$. 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 $V_{1,1} < v_1^*$ and $V_{2,1} > v_2^*$. 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 $V_{1,j} \ge v_1^*$ and $V_{2,j} \le v_2^*$. Notice that under $\mathcal{E}_j$, though agent 1 has higher *expected* quantile than agent 2, agent 2 still has higher actual quantile for the item with constant probability. 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 because, conditioned on $\mathcal{E}^{\text{mistake}}$, $\mathcal{E}_j$ will occur for a large majority of items, leading to a large discrepancy in bundle sizes.

Let $\mathcal{E}^{\text{manyhigh}}(t)$ be the event $\sum_{j=2}^{t} V_{2,j} \cdot \mathbb{I}[\mathcal{E}_j] \ge (t-1) \cdot 0.7 \cdot \mathbb{E}[V]$. In other words, $\mathcal{E}^{\text{manyhigh}}(t)$ occurs when agent 2 has a high value for items j, $2 \le j \le t$, 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 $\sum_{j=2}^{t} V_{2,j} \le (t-1) \cdot 1.1 \cdot \mathbb{E}[V]$. 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, because each $V_{1,j}$ and $V_{2,j}$ are independent, $\Pr[\mathcal{E}_j] \ge 0.9 \cdot 0.9 = 0.81$, and $\mathbb{E}[V_{2,j} \mid \mathcal{E}_j] = \mathbb{E}[V_{2,j} \mid V_{2,j} \le v_2^*]$. Also, from the definition of $g(\hat{q})$ and the choice of q_2^*, $\mathbb{E}[V_{2,j} \mid V_{2,j} \le v_2^*] \ge 0.9 \cdot \mathbb{E}[V]$. It follows that $\mathbb{E}[V_{2,j} \cdot \mathbb{I}[\mathcal{E}_j]] = \mathbb{E}[V_{2,j} \mid \mathcal{E}_j] \cdot Pr[\mathcal{E}_j] \ge 0.729 \cdot \mathbb{E}[V]$. A straightforward Chernoff bound establishes that $\Pr[\mathcal{E}^{\text{manyhigh}}(t)] \ge 3/4$ for t at least $\frac{6}{\mathbb{E}[V]}$.

Let $Y_j = V_{2,j} \cdot \mathbb{I}[\mathcal{E}_j]$ for all j. Then, $\mathbb{E}[Y_j] \ge 0.729 \cdot \mathbb{E}[V]$, and $\mathbb{E}[\sum_{j=2}^{T} Y_j] \ge (t-1) \cdot 0.729 \cdot \mathbb{E}[V]$. We are interested in the probability that $\sum_{j=2}^{t} Y_j$ is at least $(t-1) \cdot 0.7 \cdot \mathbb{E}[V]$, that is, 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 $\delta \in [0,1)$, $\Pr[Y \le (1-\delta)\mathbb{E}[Y]] \le e^{-\mathbb{E}[Y]\delta^2/2}$.

Continuing our derivation:

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

which is at least $3/4$ when $\mathbb{E}[\sum_{j=2}^{t} Y_j](0.89)^2/2$ is at least

$\ln(4)$, or, equivalently, if $t \geq 1 + 2\ln(4)/(0.7 \cdot (0.89)^2 \cdot \mathbb{E}[V])$. Because $2\ln(4)/(0.7 \cdot 0.89^2) < 5$ and $\mathbb{E}[V] < 1$, so $t \geq 6/\mathbb{E}[V]$ 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 $\mathcal{E}^{\text{mistake}}$, because 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 $t \geq 6/\mathbb{E}[V]$.

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, that is, $\Pr[\overline{\mathcal{E}_{\text{SPO}}(t) \cap \mathcal{E}_{\text{EF}}(t)} \,|\, \mathcal{E}^{\text{mistake}} \cap \mathcal{E}^{\text{manyhigh}}(t) \cap \mathcal{E}^{\text{normalval}}(t)] = 1$. To see this, notice that first, because of $\mathcal{E}^{\text{mistake}}$, 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 $0.7 \cdot (t-1) \cdot \mathbb{E}[V] + v_{2,1}$. Third, $\mathcal{E}^{\text{normalval}}(t)$ ensures that agent 2's value for all items is at most $1.1 \cdot (t-1) \cdot \mathbb{E}[V] + v_{2,1}$, 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

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

Therefore, for $t \geq 6/\mathbb{E}[V]$, 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. □

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 (noting that an allocation satisfying a $(1-\varepsilon)$ approximation to welfare is also $(1-\varepsilon)$-PO).

Theorem 2. *In the i.i.d. model, Algorithm 1 achieves envy-freeness and a $(1-\varepsilon)$ approximation to welfare, with probability $1 - \exp(-\Omega(T^{1/10}))$, for all $\varepsilon > 0$.*

Algorithm 1 works in epochs: each epoch k has an exploration/sampling phase, where each agent i receives a predetermined 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* $k = 1 \ldots$ **do**
  Sampling Phase: ($n \cdot k^4$ items)
  Give the j-th item in this phase to agent $j (\text{mod } n)$.
  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, that is, $\hat{F}_i(q)$ 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 $\gamma > 0$, $\Pr[\sup_{q \in [0,1]} |\hat{F}_i(q) - q| > \gamma] \leq 2e^{-2m\gamma^2}$. 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 $\gamma = \sqrt{\ln(2n/\delta)/2m}$.

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 $Q_1, \ldots, Q_n$ 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 $Q_1, \ldots, Q_n$) $\hat{F}_{i^{\max}}(Q_{i^{\max}}) > \hat{F}_j(Q_j)$ for all $j \neq i^{\max}$. This ensures that $i^{\max}$ has the highest empirical rank, and hence receives the item. Let $Q_{(1)}, \ldots, Q_{(n)}$ be the respective order statistics. A key observation is that $Q_{(n)} - Q_{(n-1)} \sim \text{Beta}[1, n]$ (Gentle 2009). The probability density function (PDF) of a Beta$[1, n]$ distribution is $f(x) = nx^{n-1}$ for $x \in [0, 1]$. Because $f(x) \le n$, $\Pr[Q_{(n)} - Q_{(n-1)} < \rho] < n\rho$ for all $\rho > 0$. Plugging in $\rho = 2\gamma$, we have $\Pr[Q_{(n)} - Q_{(n-1)} \le 2\gamma] < 2n\gamma$. 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 $Q_j \le Q_{(n-1)}$ for all $j \ne i^{\max}$, by definition. Using the DKW inequality condition, it follows that $\hat{F}_{i^{\max}}(Q_{i^{\max}}) \ge Q_{i^{\max}} - \gamma > Q_j + \gamma \ge \hat{F}_j(Q_j)$. We conclude that for $\varepsilon > 2n\sqrt{\ln(2n/\delta)/2m}$, ε-accuracy is satisfied with probability at least $1 - \delta$. □

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 probability* $1 - \exp(-\Omega(T^{1/10}))$, *where* $f(T) \in O(T^{15/16})$.

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 $k \ge 3n$ 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 $k(t) = \min\{K \in \mathbb{N} \mid \sum_{k=1}^{K} nk^4 + k^8 \ge t\}$ be the function that given an item t returns the epoch item t is in. Notice that $T \ge \sum_{k=1}^{k(T)-1} nk^4 + k^8 \ge (k(T) - 1)^8$, and therefore $k(T) \le 2T^{1/8}$. In any run of the algorithm, we can classify every item $t \le T$ into at least one of the following five 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 one of the first $\lfloor T^{1/10} \rfloor$ epochs, that is, $k(t) \le \lfloor T^{1/10} \rfloor$.
3. Item t was allocated in the sampling phase of epoch $k(t) \ge 3n$.
4. Item t was allocated in the ranking phase of epoch $k(t) \ge 3n$; the epoch was $\varepsilon_{k(t)}$-accurate.
5. Item t was allocated in the ranking phase of epoch $k(t) \ge \lfloor T^{1/10} \rfloor + 1$; 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 nonmaximum quantile for it. We show that the numbers of mistakes in each category are bounded by $3^{10}n^9, T^{1/2}n + T^{9/10}, 2nT^{5/8}, 9nT^{15/16}$, and 0, with probabilities $1, 1, 1, 1 - \exp(-\Omega(T^{7/8}))$, and $1 - \exp(-\Omega(T^{1/10}))$, respectively. This implies that the total number of mistakes is at most the sum of these quantities, which, via a union bound, is $O(T^{15/16})$ with probability $1 - \exp(-\Omega(T^{1/10}))$, via a union bound.

The number of items in the first category is at most

$$\begin{aligned}\sum_{k=1}^{3n-1} k^4 n + k^8 &\le \sum_{k=1}^{3n} (3n)^4 n + (3n)^8 \\ &\le (3n)^5 n + (3n)^9 \le 3^{10} n^9.\end{aligned}$$

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

For the second category, a similar computation gives a bound of

$$\begin{aligned}\sum_{k=1}^{\lfloor T^{1/10} \rfloor} k^4 n + k^8 &\le \lfloor T^{1/10} \rfloor \cdot (\lfloor T^{1/10} \rfloor^4 n + \lfloor T^{1/10} \rfloor^8) \\ &\le T^{1/2} n + T^{9/10}.\end{aligned}$$

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

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

Each item t in the fourth category has probability $\varepsilon_{k(t)}$ of being a mistake. The expected number of mistakes is therefore at most $\sum_{k=3n}^{k(T)} \varepsilon_k k^8 = \sum_{k=3n}^{k(T)} 3nk^{13/2} \le 3nk(T)^{15/2} \le 8nT^{15/16}$. Using Hoeffding's inequality, we get that the number of mistakes is at most $(8n + 1)T^{15/16}$, because a deviation of $T^{15/16}$ occurs with probability at most $\exp(-2T^{15/8}/T) = \exp(-2T^{7/8})$.

For the fifth category, we will union bound over the probability that *any* epoch $k \ge T^{1/10}$ is not ε_k-accurate. This probability is at most

$$\begin{aligned}\sum_{k=\lfloor T^{1/10} \rfloor + 1}^{\infty} \delta_k &= \sum_{k=\lfloor T^{1/10} \rfloor + 1}^{\infty} 2n/e^{2k} \\ &\le 2n \exp(-2T^{1/10}) \cdot \frac{1}{1 - 1/e^2} \\ &\le 3n \exp(-2T^{1/10}).\end{aligned}$$

Hence, with probability at least $1 - 3n\exp(-2T^{1/10})$, there will be zero items in this category. □

Finally, we can prove Theorem 2 as a relatively straightforward consequence of Lemma 5, because 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 V 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 $1/2 \cdot \mathbb{E}[V] \cdot T$, let $\mathcal{E}_2^T$ be the event that quantile maximization is c-strongly EF for $c = \frac{(\mathbb{E}[V \mid F(V) \geq 1/2] - \mathbb{E}[V])}{4n}$, 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 probability $1 - \exp(-\Omega(T^{1/10}))$. Note that Lemmas 1 and 5 tell us $\mathcal{E}_2^T$ and $\mathcal{E}_3^T$ each occur with probability $1 - \exp(-\Omega(T))$ and $1 - \exp(-\Omega(T^{1/10}))$, respectively. For $\mathcal{E}_1^T$, the maximum value for each item is in expectation at least the expected value for a single agent $\mathbb{E}[V]$. Hence, a Chernoff bound tells us $\mathcal{E}_1^T$ occurs with probability at least $1 - \exp(-\mathbb{E}[V]T/8)$. The claim holds via a union bound.

Next, note that for sufficiently large T, because $f(T) \in o(T), f(T) \leq \frac{(\mathbb{E}[V \mid F(V) \geq 1/2] - \mathbb{E}[V])}{8n} \cdot T$ and $f(T) \leq \varepsilon/2 \cdot \mathbb{E}[V] \cdot T$ (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 $\boldsymbol{A}^{QM} = (A_1^{QM}, \ldots, A_n^{QM})$ be the allocation of quantile maximization and $\boldsymbol{A} = (A_1, \ldots, A_n)$ be the allocation of Algorithm 1. Beginning with envy-freeness, we have that for all pairs of agents i and j,

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

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

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

as needed. □

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 0.999-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 $q_1^t, \ldots, q_n^t \in [0,1]$ 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, because 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 one; rescaling (specifically, dividing all values by $2+\varepsilon$) gives a distribution upper bounded by one and does not affect any of our arguments. Consider the value distribution V, with

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

for some small $\varepsilon > 0$ to be fixed later. Intuitively, V is a continuous version of a discrete distribution which takes low value (near zero) with probability $1 - \frac{1}{n}$, medium value (near one) with probability $\frac{2}{3n}$, and high value (near two) with probability $\frac{1}{3n}$. Let F_V 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 $q_1, \ldots, q_n$. First, we partition the agents depending on whether their quantile q_i is above or below $1 - \frac{2n}{3}$. We let $N^{\text{below}} = \left\{ i \in [n] \middle| q_i < 1 - \frac{2n}{3} \right\}$ and

$N^{\text{above}} = \left\{i \in [n] \mid q_i \geq 1 - \frac{2n}{3}\right\}$. Either $|N^{\text{below}}| \geq \lceil n/2 \rceil$ or $|N^{\text{above}}| \geq \lceil n/2 \rceil$; we analyze each case separately. Because $n \geq 3$, 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

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

Because $\lceil n/2 \rceil \cdot (\lceil n/2 \rceil - 1) \geq n^2/8$, we can that 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 two exists. Hence, in this time step, the algorithm has an additive error (compared with 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 $Q_{i^{\max}} \in \left(1 - \frac{1}{n}, 1 - \frac{2}{3n}\right)$ 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 one exists. Hence, in this time step, the algorithm has an additive error (compared with 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 $\frac{1-\varepsilon}{1{,}000} T$ 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.

Finally, we handle the case of two agents. Assume values are drawn from a 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, 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 $T/1{,}000$ because of concentration, so the algorithm cannot guarantee more than a 0.999 approximation, as needed. □

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

Algorithm 2 (Bounded Memory)

for *Epoch* $k = 1 \ldots$ **do**
  Sampling Phase: (k^9 items)
  NotWithinError $\leftarrow \mathcal{N}$
  for *trial* $= 1, \ldots, k^3$ **do**
    for $i \in$ NotWithinError **do**
      Allocate the next item to agent i, and update her memory
    Test $k^6 - |$NotWithinError$|$ number of items (for each agent)
    for $i \in$ NotWithinError **do**
      if *Proportion of test items for agent i is within* $\pm 1/k^2$ *of* $(n-1)/n$ **then**
        NotWithinError $\leftarrow$ 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

Theorem 4. *In the i.i.d. model, given a memory of one item per agent, Algorithm 2 achieves envy-freeness and a* $1 - 1/e - \varepsilon$ *approximation to welfare, with probability* $1 - \exp(-\Omega(T^{1/20}))$, *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 of the Online Appendix.

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 analog to Lemma 5: the number of disagreements between Algorithm 2 and the ideal threshold algorithm is sublinear. 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 of the Online Appendix, 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 probability $1-\exp(-\Omega(T^{1/20}))$, where $f(T) \in O(T^{19/20})$.*

6. Agents with Correlated Values

Recall that $v_{i,t} = v_t^{com} + \varepsilon_{i,t}$, 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 because 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 is some intervals $[a^{com}, b^{com}]$ and $[a^{noise}, b^{noise}]$. (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 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 Gölz (2022) and weaker than 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 probability $1-\exp(-\Omega(T^{1/14}))$, 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 zero, that is, are of the form $[0, b^{com}-a^{com}]$ and $[0, b^{noise}-a^{noise}]$. Indeed, translating the values does not change whether 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, because 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 $\gamma > 0$, $\Pr[\sup_v |\hat{F}_i(v) - F^{sum}(v)| > \gamma] \le 2e^{-2m\gamma^2}$, and specifically for $\gamma = \sqrt{\ln(2n/\delta)/2m}$, 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 $V_1^{noise}, \ldots, V_n^{noise}$ 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 $j \ne i^{max}$. 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 $j \ne i^{max}$, 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 $v_2 \ge v_1 \in [0, b^{com} + b^{noise}]$ such that $v_2 - v_1 \le b^{noise}$,*

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

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

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

Because $0 \le x \le b^{com} + b^{noise}$, rearranging shows that $x - b^{noise} \le b^{com}$, and (trivially), $x \ge 0$. 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

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

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

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

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 zero 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 $[c, d] \subseteq [0, b^{com} + b^{noise}]$ of length $\ell := d - c$ is at least as large as integrating $[0, \ell]$, that is, $\int_c^d g(x)\,dx \ge \int_0^\ell g(x)\,dx$. Further, as long as $\ell \le \min(b^{com}, b^{noise})$, 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 $v_2 - v_1 \le \min(b^{com}, b^{noise})$, then

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

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

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

It follows from Inequality (2), $v_2' \le v_2$, and the definition of v_2' that, as required,

$$\begin{aligned} F^{sum}(v_2) - F^{sum}(v_1) &\ge F^{sum}(v_2') - F^{sum}(v_1) \\ &\ge p^2(v_2' - v_1)^2/2 \\ &= \left(\min\left(\frac{b^{com}}{b^{noise}}, 1\right) p(v_2 - v_1)\right)^2 /2. \quad \square \end{aligned}$$

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 $F^{sum}(V_{i^{max}}) - F^{sum}(V_j) \ge 2\gamma$ for all $j \ne i^{max}$. Because $V_{i^{max}}$ and V_j can differ by at most b^{noise}, we immediately get that $F^{sum}(V_{i^{max}}) - F^{sum}(V_j) \ge \left(\min\left(\frac{b^{com}}{b^{noise}}, 1\right) p(V_{i^{max}} - V_j)\right)^2/2$. A sufficient condition for this to be at least 2γ is that

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

and, equivalently,

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

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, because f^{noise} is upper bounded by q,

$$\begin{aligned} F^{noise}(V_{i^{max}}^{noise}) - F^{noise}(V_j^{noise}) &= \int_{V_j^{noise}}^{V_{i^{max}}^{noise}} f^{noise}(x)\,dx \\ &\le q(V_{i^{max}}^{noise} - V_j^{noise}). \end{aligned}$$

Figure 1. Plot of Lower Bound $g(x)$

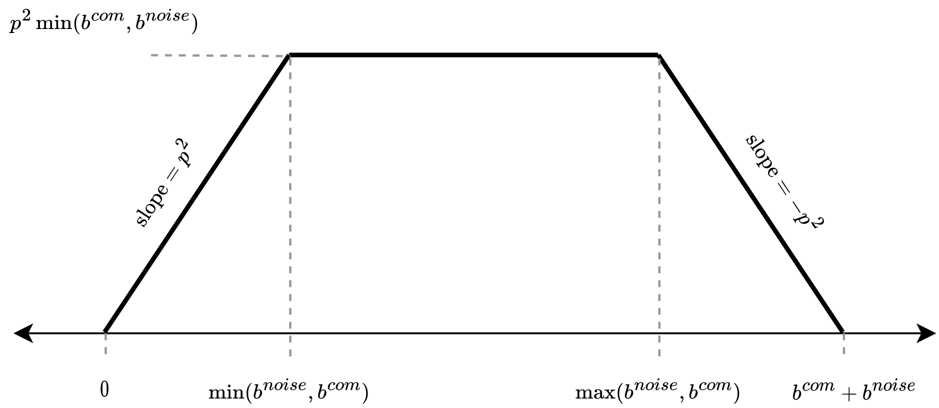

Hence, as long as $F^{noise}(V^{noise}_{i^{max}}) - F^{noise}(V^{noise}_j) \geq \frac{2q\sqrt{\gamma}}{\min\left(\frac{b^{com}}{b^{noise}}, 1\right)p}$, then $V^{noise}_{i^{max}} - V^{noise}_j \geq \frac{2\sqrt{\gamma}}{\min\left(\frac{b^{com}}{b^{noise}}, 1\right)p}$.

$F^{noise}(V^{noise}_j)$ is distributed uniformly on $[0,1]$ for each i, and these random variables are independent across agents. $F^{noise}(V^{noise}_{i^{max}})$ 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 $\Pr[\forall j \neq i^{max}, F^{noise}(V^{noise}_{i^{max}}) - F^{noise}(V^{noise}_j) > \rho] \geq 1 - n\rho$. Plugging in $\rho = \frac{2q\sqrt{\gamma}}{\min\left(\frac{b^{com}}{b^{noise}}, 1\right)p}$, we get that this holds with probability at least $1 - \frac{2nq\sqrt{\gamma}}{\min\left(\frac{b^{com}}{b^{noise}}, 1\right)p}$. Therefore, as long as $\varepsilon > \frac{2nq}{\min\left(\frac{b^{com}}{b^{noise}}, 1\right)p} \cdot \left(\frac{\ln(2n/\delta)}{2m}\right)^{1/4}$, ε-accuracy holds. □

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 probability $1 - \exp(-\Omega(T^{1/14}))$. 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 $C = 3nq/\left(\min\left(\frac{b^{com} - a^{com}}{b^{noise} - a^{noise}}, 1\right)p\right)$. For each epoch k, choose $\delta_k = 2n/e^{2k}$ and set $\varepsilon_k = C \cdot k^{-5/4}$. Now

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

The constant C is unimportant; as it does not depend on T or k, it will disappear as the time step grows large. The primary difference from Theorem 2 is that, because of 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 $k^{-3/2}$), and use $T^{1/14}$ instead of $T^{1/10}$.

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 time step T. Because of the different sampling lengths, it now holds that $k(T) \leq 2T^{1/12}$.

We now count the number in the first $\lfloor T^{1/14} \rfloor$ instead of $\lfloor T^{1/10} \rfloor$ and get

$$\sum_{k=1}^{k(T)} nk^6 + k^{12} \leq nT^{1/2} + T^{13/14}.$$

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

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

Because of Hoeffding's inequality, a deviation of more than $T^{47/48}$ does not occur with probability $1 - \exp(-\Omega(T^{47/48}))$, so the number of mistakes here is $O(T^{47/48})$ with the corresponding probability.

Union bounding over the number of items in non-ε_k-accurate exploit phases is again at most

$$\sum_{k=\lfloor T^{1/14} \rfloor + 1}^{\infty} \delta_k \leq 2n \sum_{k=\lfloor T^{1/14} \rfloor + 1}^{\infty} \frac{2n}{e^{2k}} = \exp(-\Omega(T^{1/14})).$$

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

6.1. Correlated Values with Bounded Memory

One may wonder if these positive results for Algorithm 1 also carry over using bounded memory; that is, could some modification of Algorithm 2 perform well in the common-noise noise model? The answer comes down to what guarantees we would like.

Achieving envy-freeness is relatively straightforward. Indeed, any q-threshold algorithm achieves envy-freeness as long as, with positive probability, some agents are above the threshold and some below: by symmetry, all agents have an equal chance of receiving each item, and their value conditioned on receiving an item is strictly higher than their value conditioned on not receiving it. All thresholds satisfy this property assuming interval support on the distributions, and minimal modification to the proof of Theorem 4 is required to show Algorithm 2 is envy-free.

Getting a welfare guarantee appears to be more challenging. There exist distributions where the $\frac{n-1}{n}$-threshold algorithm does not achieve a constant approximation. For example, consider a common distribution that takes value n with probability $2/n$ and zero otherwise, and a noise distribution that takes value $n-1$ with probability $1/n$ and zero otherwise. The $\frac{n-1}{n}$ quantile is at least n. However, this implies that whether a value is above the threshold is completely determined by the common draw. Therefore, either all agents are above the threshold or all are below, so items will be given to random agents. Random agents have an expected value of at most three. On the other hand, the expected maximum of n draws from the noise distribution is $\Theta(n)$. This example can be extended to have continuous distributions (by slightly spreading out the mass around the points) and to interval support (draw uniformly from $[0,n]$ with some tiny probability). Given the connection between threshold algorithms and prophet inequalities, we could hope to import results from threshold algorithms for correlated distributions (for example, from Immorlica et al. 2020). Unfortunately, all existing results require knowing the distribution (e.g., set a threshold of $\mathbb{E}[\max_i V_i]/2$). This conflicts with a fundamental feature of our algorithms: they do not need to know anything about exact values or underlying distributions, just ordinal relationships. Of course, our setup is not as general as theirs, and hence, with new techniques, positive results might be possible. We leave this as an interesting direction for future work.

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 two nonidentical agents, there is no algorithm that is EF and c-PO with probability p, for* $c > \frac{1+\sqrt{5}}{4} \approx 0.809$ *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 (V_1, V_2) there exists a $T^* = T^*(V_1, V_2)$ 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 $\mathcal{E}_j^{\mathcal{A},t}$ 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, $\Pr[\mathcal{E}_j^{\mathcal{A},t}] \geq 2/3 + \varepsilon$ 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 one with probability z^2 and zero 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, $\Pr[\overline{\mathcal{E}}_1^T] = \Pr[|\frac{1}{T}\sum_{t=1}^T Z_1^t \cdot Z_2^t - z^2| \geq \delta] \leq 2\exp(-2T\delta^2)$. It follows that for $T \geq \log(2/\varepsilon)/(2\delta^2)$, $\Pr[\overline{\mathcal{E}}_1^T] \leq \varepsilon$. Similarly, for $T \geq \log(2/\varepsilon)/(2\delta^2)$, it holds that $\Pr[\overline{\mathcal{E}}_2^T] \leq \varepsilon$, and $\Pr[\overline{\mathcal{E}}_3^T] \leq \varepsilon$. Consider an arbitrary $T > T_{\max} = \max\{T_0, T_1, T_2, \log(2/\varepsilon)/(2\delta^2)\}$. Applying a union bound,

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

It follows that $\Pr[\mathcal{E}_0^{\mathcal{A},T} \cap \mathcal{E}_1^{\mathcal{A},T} \cap \mathcal{E}_2^{\mathcal{A},T} \cap \mathcal{E}_1^T \cap \mathcal{E}_2^T \cap \mathcal{E}_3^T] > 0$. 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, because $\mathcal{A}$ 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 $H_B = \{t \in [T] : q_1(t) \geq 1 - z$ 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 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, that is, $|A_2^T \cap H_B| \geq |H_B|/2$. 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 $\hat{A}_2 = \overline{H}_1$.

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

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

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

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

By $\mathcal{E}_0^{\mathcal{A}}$, A^T is envy-free on I_0. It follows that $|A_1^T| \geq (1-w)|A_2^T|$. Because $|A_1^T| + |A_2^T| = T$, 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

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

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

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

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. □

Algorithms 1 and 2 are envy-free with high probability, even in the non-i.i.d. model, because envy-freeness is not an "interagent" property. Our last result shows that they also give a constant approximation to Pareto efficiency, by combining Lemma 3 with Lemmas 5 and 7. The proof of Theorem 7 can be found in Section EC.1.5 of the Online Appendix.

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 probability $1 - \exp(-\Omega(T^{1/10}))$ and $1 - \exp(-\Omega(T^{1/10}))$, respectively, for all $\varepsilon > 0$.*

Although the formal guarantees in Theorem 7 are similar for the two algorithms despite Algorithm 2 using a memory size of one, Algorithm 1 has the benefit of much shorter epoch lengths (and better guarantees in the i.i.d. case).

8. Computational Study

The purpose of this section is twofold. First, although 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 rates. Second, although Algorithm 1 was designed to be amenable for theoretical analysis, we compare it with 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 sample with 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 distribution skews 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 are 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 agent-specific values $\varepsilon_{it} \sim U(0,1)$ for each agent i and item t. We then set 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 values bid by real food banks on actual donations using artificial currency over the course of a year, similar to part of what is analyzed in Prendergast (2017, 2022) and Altmann (2023). We interpret these bids as a proxy for correlated values for each donation. To generate an instance, we first restrict to sets of five food banks that bid together on at least 20 distinct donations. We sample such a set of five, treat their bids on a common donation as a correlated value distribution, and draw all item values from this distribution. A small number of bids (under 4%) are negative, which, in the original context, meant that the organization expected to receive artificial currency to accept the donation; in our context we interpret these as zero-valued.

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 Online 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. The performance on the real data is very much in line with what is observed for the simpler (small α) correlated distributions, where envy-freeness takes roughly 1,000 time steps to establish. In general, Algorithm 1 kept pace reasonably well alongside its "ideal" counterpart, 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 distributions had worse approximations, and increasing correlation led 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

Figure 2. (Color online) Density Functions of the Beta Distributions on Which We Test

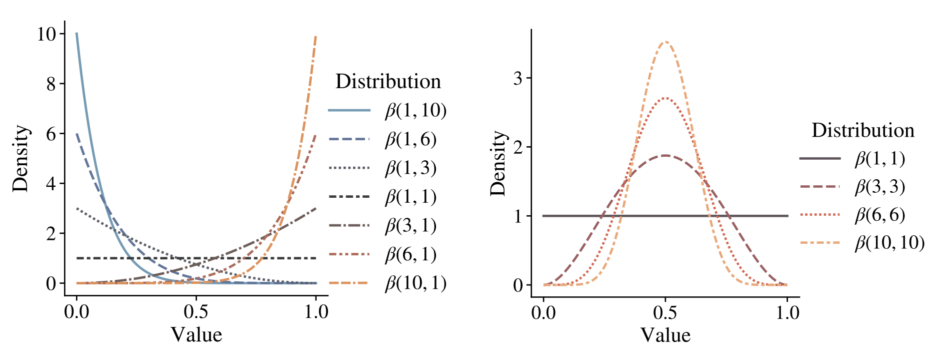

Note. The skewed distributions are on the left, peaked distributions on the right.

Figure 3. (Color online) Proportion of Envy-Free Runs

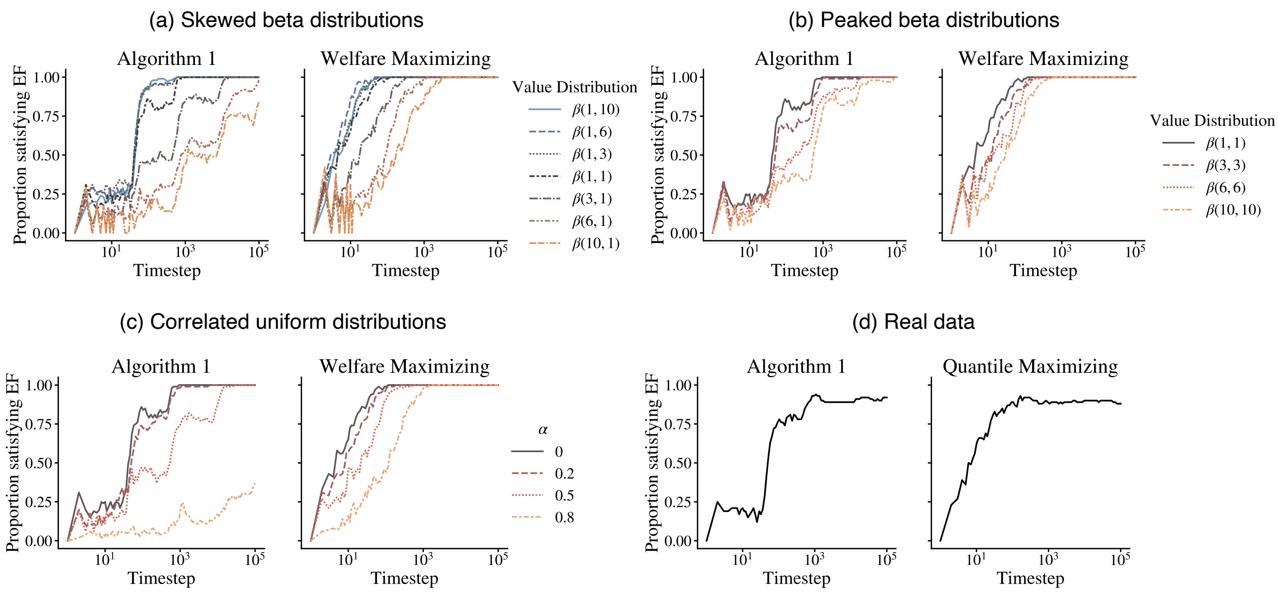

Note. For each data set, the left graph shows the results for Algorithm 1, whereas the right graph shows the performance of the "ideal" welfare-maximizing algorithm (or quantile-maximizing in the case of real data where the underlying value distributions are heterogenous).

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 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 leads to a marginal boost in performance. Finally, we try running the algorithm *without* resetting the sampled items at the end of every epoch. Namely, at first agents were given 20 items each, and then all future items were 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 reasonably 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

Figure 4. (Color online) Ninety-Fifth Percentile Welfare Approximations

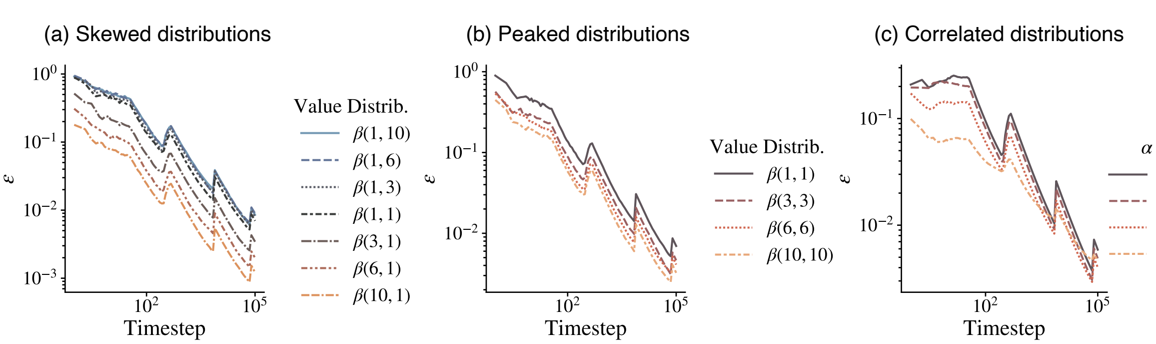

Note. If at time step 1,000 the algorithm has $\varepsilon = 0.02$, then on 95% of runs, at the 1,000th time step, the algorithm achieved at least 98% of the optimal welfare.

Figure 5. (Color online) Comparison of Alternative Algorithms in Various Settings

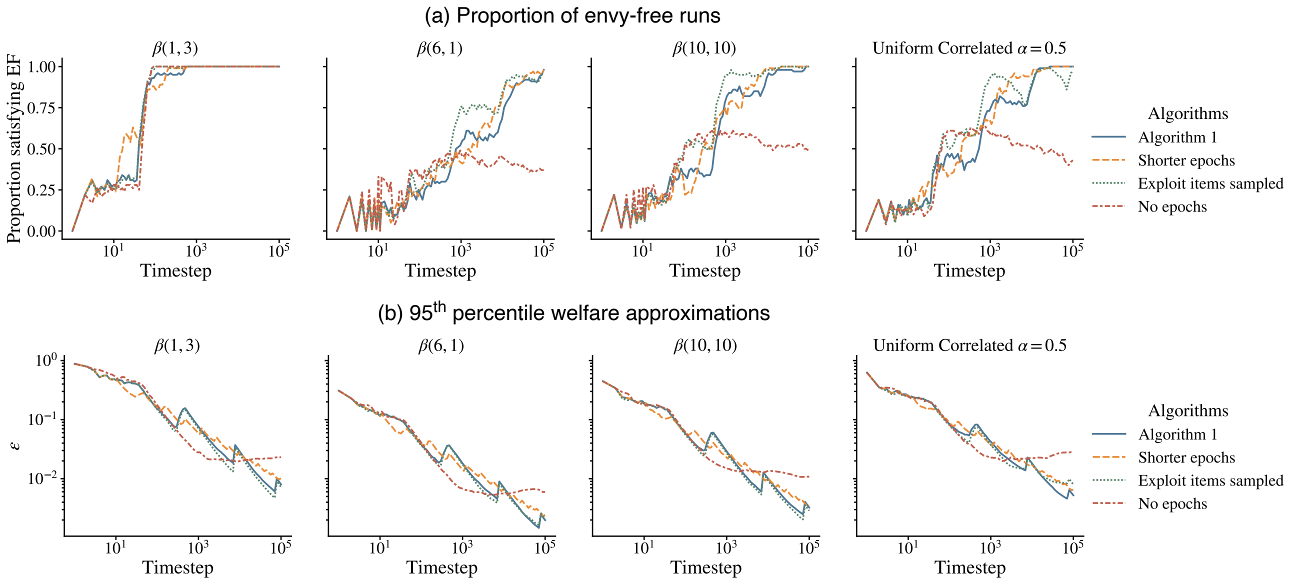

to EF allocations and welfare appears to cap out well below 99% of optimal without hope of further improvements from additional items.

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. Finally, we assume that agents provide ordinal information while having underlying cardinal utilities. One could, instead, explore fairness notions like stochastic dominance (SD) envy-freeness which are defined directly on ordinal preferences. Positive results in this setting may be more challenging and, for example, require that the number of items is divisible by the number of agents.

Acknowledgments

The authors are grateful to Canice Prendergast for generously sharing data.

References

Abramowitz B, Anshelevich E (2018) Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. *Proc. 32nd AAAI Conf. Artificial Intelligence (AAAI-18)*, 894–901.

Aleksandrov MD, Aziz H, Gaspers S, Walsh T (2015) Online fair division: Analysing a food bank problem. Qiang Yang Q, Wooldridge MJ, eds. *Proc. 24th Internat. Joint Conf. Artificial Intelligence (IJCAI 2015)* (AAAI Press, Palo Alto, CA), 2540–2546.

Altmann SM (2023) Choice, welfare, and market design: An empirical investigation of Feeding America's choice system. Working paper, Wadham College, University of Oxford, Oxford, UK.

Amanatidis G, Birmpas G, Markakis E (2016) On truthful mechanisms for maximin share allocations. Brewka G, ed. *Proc. 25th Internat. Joint Conf. Artificial Intelligence (IJCAI '16)* (AAAI Press, Washington, DC), 31–37.

Anshelevich E, Sekar S (2016) Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. *Proc. 30th AAAI Conf. Artificial Intelligence (AAAI)*, 390–396.

Anshelevich E, Zhu W (2019) Tradeoffs between information and ordinal approximation for bipartite matching. *Theory Comput. Systems* 63(7):1499–1530.

Anshelevich E, Filos-Ratsikas A, Shah N, Voudouris AA (2021) Distortion in social choice problems: The first 15 years and beyond. Zhou ZH, ed. *Proc. 30th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org, San Mateo County, CA), 4294–4301.

Anshelevich E, Bhardwaj O, Elkind E, Postl J, Skowron P (2018) Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264:27–51.

Aziz H, Gaspers S, Mackenzie S, Walsh T (2015) Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.

Bai Y, Gölz P (2022) Envy-free and pareto-optimal allocations for agents with asymmetric random valuations. De Raedt L, ed. *Proc. 31st Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org, San Mateo County, CA), 53–59.

Barman S, Khan A, Maiti A (2022) Universal and tight online algorithms for generalized-mean welfare. Honavar V, Spaan MTJ, eds. *Proc. 36th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 4793–4800.

Baumeister D, Bouveret S, Lang J, Nguyen NT, Nguyen TT, Rothe J, Saffidine A (2017) Positional scoring-based allocation of indivisible goods. *Autonomous Agents Multi-agent Systems* 31(3):628–655.

Benade G, Kazachkov AM, Procaccia AD, Psomas CA (2018) How to make envy vanish over time. Tardos É, Elkind E, Vohra R, eds. *Proc. 10th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 593–610.

Benade G, Nath S, Procaccia AD, Shah N (2021) Preference elicitation for participatory budgeting. *Management Sci.* 67(5):2813–2827.

Bogomolnaia A, Moulin H, Sandomirskiy F (2022) On the fair division of a random object. *Management Sci.* 68(2):1174–1194.

Boutilier C, Caragiannis I, Haber S, Lu T, Procaccia AD, Sheffet O (2015) Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.

Bouveret S, Endriss U, Lang J (2010) Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. Coelho H, Studer R, Wooldridge MJ, eds. *Proc. 19th Eur. Conf. Artificial Intelligence (ECAI)* (IOS Press, Amsterdam), 387–392.

Caragiannis I, Nath S, Procaccia AD, Shah N (2017) Subset selection via implicit utilitarian voting. *J. Artificial Intelligence Res.* 58:123–152.

Cesa-Bianchi N, Lugosi G (2006) *Prediction, Learning, and Games* (Cambridge University Press, Cambridge, UK).

Charikar M, Ramakrishnan P, Wang K, Wu H (2024) Breaking the metric voting distortion barrier. *Proc. 35th Annual ACM-SIAM Sympos. Discrete Algorithms (SODA)* (SIAM, Philadelphia), 1621–1640.

Dhangwatnotai P, Roughgarden T, Yan Q (2010) Revenue maximization with a single sample. *Proc. 11th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 129–138.

Dickerson JP, Goldman J, Karp J, Procaccia AD, Sandholm T (2014) The computational rise and fall of fairness. Brodley CE, Stone P, eds. *Proc. 28th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 1405–1411.

Dvoretzky A, Kiefer J, Wolfowitz J (1956) Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Ann. Math. Statist.* 27(3):642–669.

Ehsani S, Hajiaghayi M, Kesselheim T, Singla S (2018) Prophet secretary for combinatorial auctions and matroids. *Proc. 29th Annual ACM-SIAM Sympos. Discrete Algorithms (SODA)* (SIAM, Philadelphia), 700–714.

Filos-Ratsikas A, Frederiksen SKS, Zhang J (2014) Social welfare in one-sided matchings: Random priority and beyond. Lavi R, ed. *Proc. 7th Internat. Sympos. Algorithmic Game Theory (SAGT)* (Springer, Berlin, Heidelberg), 1–12.

Friedman E, Psomas CA, Vardi S (2015) Dynamic fair division with minimal disruptions. *Proc. 16th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 697–713.

Friedman E, Psomas CA, Vardi S (2017) Controlled dynamic fair division. *Proc. 18th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 461–478.

Gentle JE (2009) *Computational Statistics* (Springer, New York).

Gkatzelis V, Halpern D, Shah N (2020) Resolving the optimal metric distortion conjecture. Irani S, ed. *Proc. 61st Sympos. Foundations Comput. Sci. (FOCS)* (IEEE, Piscataway, NJ), 1427–1438.

Gkatzelis V, Psomas A, Tan X (2021) Fair and efficient online allocations with normalized valuations. Leyton-Brown K, Mausam, eds. *Proc. 35th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 5440–5447.

Goel A, Krishnaswamy AK, Munagala K (2017) Metric distortion of social choice rules: Lower bounds and fairness properties. Daskalakis C, Babaioff M, Moulin H, eds. *Proc. 18th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 287–304.

Gorokh A, Banerjee S, Iyer K (2021) From monetary to nonmonetary mechanism design via artificial currencies. *Math. Oper. Res.* 46(3):835–855.

Halpern D, Shah N (2021) Fair and efficient resource allocation with partial information. Zhou Z, ed. *Proc. 30th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org, San Mateo County, CA), 5456–5463.

He J, Procaccia AD, Psomas A, Zeng D (2019) Achieving a fairer future by changing the past. Kraus S, ed. *Proc. 28th Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org, San Mateo County, CA), 343–349.

Hoi SC, Sahoo D, Lu J, Zhao P (2021) Online learning: A comprehensive survey. *Neurocomputing* 459:249–289.

Ikeda K, Kita H, Kobayashi S (2001) Failure of Pareto-based MOEAs: Does non-dominated really mean near to optimal? *Proc. 2001 Congress Evolutionary Comput.*, vol. 2 (IEEE, New York), 957–962.

Immorlica N, Singla S, Waggoner B (2020) Prophet inequalities with linear correlations and augmentations. *Proc. 21st ACM Conf. Econom. Comput. (EC)* (ACM, New York), 159–185.

Kash I, Procaccia AD, Shah N (2014) No agent left behind: Dynamic fair division of multiple resources. *J. Artificial Intelligence Res.* 51:579–603.

Kempe D (2020) Communication, distortion, and randomness in metric voting. Conitzer V, Sha F, eds. *Proc. 34th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 2087–2094.

Kizilkaya FE, Kempe D (2022) Plurality veto: A simple voting rule achieving optimal metric distortion. *Proc. 31st Internat. Joint Conf. Artificial Intelligence (IJCAI)* (ijcai.org, San Mateo County, CA), 349–355.

Kroer C, Peysakhovich A, Sodomka E, Stier-Moses NE (2021) Computing large market equilibria using abstractions. *Oper. Res.* 70(1):329–351.

Kurokawa D, Procaccia AD, Wang J (2016) When can the maximin share guarantee be guaranteed? Schuurmans D, Wellman MP, eds. *Proc. 30th AAAI Conf. Artificial Intelligence (AAAI)* (AAAI Press, Palo Alto, CA), 523–529.

Leung S, Lui E, Pass R (2015) Voting with coarse beliefs. *Proc. 6th Innovations Theoret. Comput. Sci. Conf. (ITCS)* (ACM, New York), 61.

Mandal D, Shah N, Woodruff DP (2020) Optimal communication-distortion tradeoff in voting. Biró P, Hartline JD, Ostrovsky M, Procaccia AD, eds. *Proc. 21st ACM Conf. Econom. Comput. (EC)* (ACM, New York), 795–813.

Manurangsi P, Suksompong W (2021) Closing gaps in asymptotic fair division. *SIAM J. Discrete Math.* 35(2):668–706.

Massart P (1990) The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *Ann. Probab.* 18(3):1269–1283.

MEANS database (2023) Organization website and FAQs. https://meansdatabase.org/.

Munagala K, Wang K (2019) Improved metric distortion for deterministic social choice rules. Karlin AR, Immorlica N, Johari R, eds. *Proc. 20th ACM Conf. Econom. Comput. (EC)* (ACM, New York), 245–262.

Nguyen NT, Nguyen TT, Rothe J (2017) Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. Larson K, Winikoff M, Das S, Durfee EH, eds. *Proc. 16th Internat. Conf. Autonomous Agents Multi-agent Systems (AAMAS)* (ACM, New York), 262–271.

Prendergast C (2017) How food banks use markets to feed the poor. *J. Econom. Perspect.* 31(4):145–162.

Prendergast C (2022) The allocation of food to food banks. *J. Political Econom.* 130(8):1993–2017.

Procaccia AD, Rosenschein JS (2006) The distortion of cardinal preferences in voting. Klusch M, Rovatsos M, Payne TR, eds. *Proc. 10th Internat. Workshop Cooperative Inform. Agents (CIA)* (Springer, Berlin, Heidelberg), 317–331.

Robbins L (1938) Interpersonal comparisons of utility: A comment. *Econom. J. (London)* 48(192):635–641.

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. Leskovec J, Grobelnik M, Najork M, Tang J, Zia L, eds. *Proc. Web Conf. 2021* (ACM, New York), 857–865.

Sinclair SR, Banerjee S, Yu CL (2022) Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. *SIGMETRICS Performance Evaluation Rev.* 50(1):95–96.

Vardi S, Psomas A, Friedman E (2022) Dynamic fair resource division. *Math. Oper. Res.* 47(2):945–968.

Wilson R (1985) *Game-Theoretic Analyses of Trading Processes* (Defense Technical Information Center, Fort Belvoir, VA).

Zeng D, Psomas A (2020) Fairness-efficiency tradeoffs in dynamic fair division. *Proc. 21st ACM Conf. Econom. Comput. (EC)* (ACM, New York), 911–912.

Gerdus Benadè is an assistant professor of information systems at Boston University. He is interested in computational social choice, fair division, and discrete optimization, especially real-world applications like participatory budgeting, gerrymandering, and food rescue.

Daniel Halpern is a PhD candidate in the EconCS group at Harvard University. He works on problems at the intersection of algorithms, economics, and artificial intelligence, including large-scale opinion aggregation and AI alignment.

Alexandros Psomas is an assistant professor in the Department of Computer Science at Purdue University. He is broadly interested in algorithmic economics, artificial intelligence, computational social choice and mechanism design, machine learning, and algorithmic game theory.