

Fair Division Under Inaccurate Preferences

Trung Dang* Daniel Halpern† Anuran Makur‡ Alexandros Psomas§
Japneet Singh¶ Paritosh Verma||

Abstract

The fair allocation of scarce resources is a central problem in mathematics, computer science, operations research, and economics. While much of the fair-division literature assumes that individuals have underlying cardinal preferences, eliciting exact numerical values is often cognitively burdensome and prone to inaccuracies. A growing body of work in fair division addresses this challenge by assuming access only to ordinal preferences. However, the restricted expressiveness of ordinal preferences makes it challenging to quantify and optimize cardinal fairness objectives such as envy.

In this paper, we explore the broad landscape of fair division of indivisible items given inaccurate cardinal preferences, with a focus on minimizing envy. We consider various settings based on whether the true preferences of the agents are stochastic or worst-case, and whether the inaccuracies, modeled as additive noise, are stochastic or worst-case. When the true preferences are stochastic, we show that envy-free allocations can be computed with high probability; this is achieved both in the setting with stochastic and worst-case noise. This generalizes a notable result in stochastic fair division, which establishes a similar guarantee, albeit in the absence of any noise. When the true preferences are worst-case, and the noise is bounded, we analyze the maximum envy achieved by the well-known Round-Robin algorithm. This bound is shown to be tight for deterministic algorithms, and applications of this bound involving known results from statistical estimation are provided. Lastly, we consider a setting with worst-case preferences and noise, where the true preferences for each item are revealed upon its allocation. Here, we give an efficient online algorithm that guarantees logarithmic maximum envy with high probability, for bounded noise. This result builds upon and generalizes a known result from algorithmic discrepancy to a setting with noisy input data. The broad range of tools used in our results highlights the theoretical relevance of this natural fair division problem.

1 Introduction

Fair division studies the fundamental problem of fairly allocating resources to agents with heterogeneous preferences over these resources. Mathematical formulations date back to Steinhaus (1948) and Dubins and Spanier (1961), though some of the field’s core questions (and even key algorithms) can be traced back to the Bible and ancient Greece. Most of the existing literature assumes that agents’ preferences are represented by valuation functions, which the agents report to the designer, or the designer can access through exact valuation queries.

While we have vastly deepened our understanding of when fairness is — and when it is not — attainable under these models, when taking this approach to practice, an especially unappealing aspect is that participants must report *exact* numerical values to the algorithm. Motivated by the reality that eliciting cardinal valuations is prone to inaccuracies as well as cognitively burdensome for the participants,

*The University of Texas at Austin. dddtrung@cs.utexas.edu

†Google Research. dhalpern@google.com

‡Purdue University. amakur@purdue.edu

§Purdue University and Google Research. apsomas@purdue.edu

¶Purdue University. japneetsingh776@gmail.com

||Purdue University. paritoshverma97@gmail.com

a growing body of work studies the problem of a designer that only has access to *ordinal information*. This idea originates from the seminal work of Procaccia and Rosenschein (2006) that defines the concept of distortion: the worst-case loss with respect to some aggregate cardinal objective (e.g., utilitarian social welfare) due to limited expressiveness; see Anshelevich et al. (2021) for a survey.

Most works in this area take a worst-case approach; after the algorithm has made its decisions, the underlying cardinal valuations are chosen adversarially (but, of course, consistently with the ordinal information available to the algorithm). This approach gives robust, albeit pessimistic guarantees. While we agree on the assumption that participants being able to submit *precise* numerical values is unrealistic in practice, we also seek a model which takes into account the fact that participants *can* provide useful, but perhaps imperfect, cardinal signals.

Hence, in this paper, we consider a middle ground between the “perfectly accurate cardinal” and “fully ordinal” models. We initiate the study of fair division given access to inaccurate cardinal preferences. Our goal is to design algorithms that minimize the maximum envy between agents (with respect to true preferences), using only inaccurate preferences as input. We investigate various formal realizations of this general fair division problem, such as different ways to model inaccurate cardinal preferences as well as different assumptions about agents’ true preferences.

1.1 Our Contributions

We introduce the problem of fair division with imperfect information. There are n agents with additive preferences, and m indivisible goods/items. Given noisy estimates $\hat{v}_{i,j}$ of the true values $v_{i,j}^*$, for each agent i and item j as input, our goal is to allocate the items in a way that minimizes the envy, with respect to the true values $v_{i,j}^*$. We model inaccuracies as additive noise on the true values. Our results are broadly divided into two categories depending on whether the true values of the agents are stochastic or worst-case.

Stochastic true valuations. These results focus on scenarios in which the true valuations $v_{i,j}^*$ are stochastic. We consider two subsettings, depending on whether the added noise is stochastic or worst-case. In both cases, we establish the existence of envy-free allocations, with high probability.

- First, we consider the case of stochastic true valuations and stochastic noise, i.e., $v_{i,j}^* \sim \mathcal{D}$, $\eta_{i,j} \sim \mathcal{D}'$, where we get to observe inaccurate/noisy values $\hat{v}_{i,j} = v_{i,j}^* + \eta_{i,j}$; here \mathcal{D} , \mathcal{D}' are two unknown, fixed, real-valued distributions. In Theorem 2, we show that the simple procedure of maximizing welfare with respect to the observed values produces envy-free allocations with high probability. This generalizes a result of Dickerson et al. (2014) to the setting with inaccurate preferences.¹ To prove our result, we establish a strict statistical correlation inequality (Theorem 1) that might be of independent interest. We show that for any two real-valued, nondecreasing functions f, g and any real-valued random variable X , we have $\mathbb{E}[f(X) \cdot g(X)] > \mathbb{E}[f(X)] \cdot \mathbb{E}[g(X)]$ if and only if $\text{Var}(f(X)) > 0$ and $\text{Var}(g(X)) > 0$. This is used to establish the inequality $\mathbb{E}[v_{i,j}^* \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}] > \mathbb{E}[v_{i,j}^* \mid i \notin \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$ (Lemma 2), which, in turn, allows us to argue that agents have higher true values ($v_{i,j}^*$) for items they are allocated (compared to the ones they are not allocated), where the allocation decisions are based on the noisy values ($\hat{v}_{i,j}$). Note that the weaker inequality, $\mathbb{E}[v_{i,j}^* \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}] \geq \mathbb{E}[v_{i,j}^* \mid i \notin \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$, follows directly from Chebyshev’s association inequality; unfortunately though, it is insufficient for our result, which necessarily requires a strict inequality.
- For stochastic true valuations and worst-case noise, we consider a general setting where the value vector for each item j is drawn i.i.d. from a bounded non-negative vector-valued distribution \mathcal{D} over $[0, b]^n$; that is, agents’ valuations for a specific item are arbitrarily correlated, but values across items are independent. We prove that, if \mathcal{D} satisfies a mild technical condition (essentially that no two agents are identical up to scaling) then as long as the noise is bounded by a constant, there exists

¹Dickerson et al. (2014) show that envy-free allocations exist with high probability in a setting where $\eta_{i,j} = 0$ for all i, j .

an algorithm that, given noisy estimates, outputs an allocation that is envy-free with high probability (Theorem 3). To prove Theorem 3, we establish the existence of strongly envy-free allocations, with high probability, in the underlying noiseless setting, by adapting an existence result from the literature on cake-cutting Barbanel (2005).

Worst-case true valuations. These results focus on scenarios where the underlying true valuations $v_{i,j}^*$ s are fixed (i.e., worst-case).

- Our first set of results concerns the case of worst-case true valuations and worst-case noise, i.e., we assume that the unknown true values $v_{i,j}^* \in [0, b]$ are fixed, and we have estimates $\hat{v}_{i,j}$ that satisfy $|\hat{v}_{i,j} - v_{i,j}^*| \leq \varepsilon$ for some (possibly unknown) $\varepsilon \geq 0$. Here, we establish a tight bound (Theorems 5 and 6) on the best possible envy guarantee that can be achieved by any deterministic algorithm that only has access to $\hat{v}_{i,j}$. In particular, we show that the standard Round-Robin algorithm² achieves an optimal guarantee.
- We proceed to give two applications of our Round-Robin result.

In the first application (Application 1), we present an end-to-end framework where we use *ordinal information* to estimate values, and then use the estimated values to make allocation decisions. Specifically, we assume access to stochastic ordinal observations; we choose arguably the most popular discrete choice model, the Bradley-Terry-Luce (BTL) model (sometimes referred to as the multinomial logit model). In this model, the probability of an agent ranking one alternative higher than another is given by the logistic function applied to the difference between the valuations of the two options Zermelo (1929); Bradley and Terry (1952); Luce (1959); Plackett (1975); McFadden (1973). Given such ordinal information, we first produce estimates $\hat{v}_{i,j}$ for the true preferences. Then these estimates are plugged into the Round-Robin algorithm to get an allocation with a bounded maximum envy. The fact that Round-Robin outputs allocations that are *balanced*³ in addition to being fair is a key property that enables us to combine BTL and Round-Robin in an end-to-end framework.

The second application (Application 2) deals with stochastic noise with an *unbiased* error. Specifically, we assume that $\hat{v}_{i,j} := v_{i,j}^* + \varepsilon_{i,j}$, where $\varepsilon_{i,j} := \sigma_{i,j} \cdot z_{i,j}$ where $z_{i,j}$ s are drawn i.i.d. from a non-negative (possibly unknown) distribution \mathcal{D} and $\sigma_{i,j}$ s are drawn uniformly at random from the set $\{-1, 1\}$. We focus on small tail distributions \mathcal{D} , and in particular, MHR distributions.⁴ We show that, given $\hat{v}_{i,j}$ s, allocations with at most a constant maximum envy can be computed efficiently. Towards this, we employ known results about the order statistics of MHR distributions and a recent result concerning the concentration of order statistics of MHR distributions. Application 2 is obtained by combining the bound on the largest noise with our aforementioned Round-Robin result.

- Finally, after exploring various aspects of worst-case valuations and bounded noise (the tight Round-Robin guarantee and its applications), we proceed to study a different, more permissive noise model. In particular, we focus on a scenario where, just as before, we have to allocate items to agents based on the noisy estimates $\hat{v}_{i,j}$; however, this time, once any item j is allocated, then the true valuation vector v_j^* corresponding to this item is revealed.⁵ This immediate revelation of the true preference

²The Round-Robin algorithm operates in rounds, where in each round, agents follow an order $1, 2, \dots, n$, and pick, from the set of unallocated/available items, an item having the highest value. Here we assume that these picks are made based on the noisy estimates $\hat{v}_{i,j}$.

³An allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is called balanced if all agents get roughly equal number of items, i.e., $|A_i| - |A_j| \leq 1$ for all $i, j \in [n]$.

⁴A distribution \mathcal{D} has monotone hazard rate (MHR) iff $\frac{1-F(x)}{f(x)}$ is non-increasing, where $F(x)$ and $f(x)$ are its cumulative distribution function and probability density function, respectively.

⁵The revelation of information (preference) that governs the objective (envy), following a decision of the algorithm (allocation of an item) is also a key theme in many well-studied problems like the online convex optimization (OCO), the experts problem, etc.

gives more power to the algorithm designer. In this model, we show that there is a simple, online algorithm that achieves a maximum envy bound of $\mathcal{O}\left(\sqrt{n}\left(\log\frac{mn}{\delta} + m\varepsilon\right)\right)$ with probability at least $1 - \delta$ (Theorem 7), where the noise is bounded by ε , i.e., $|v_{i,j}^* - \hat{v}_{i,j}| \leq \varepsilon$. The algorithm achieving this bound (Algorithm 1) is inspired by (and is a modified version of) an online discrepancy minimization algorithm given by Alweiss et al. (2021).⁶ Notably, if the noise satisfies the upper bound $\varepsilon \leq \frac{1}{m} \log\frac{mn}{\delta}$, then we get an envy as well as a discrepancy bound of $\mathcal{O}\left(\sqrt{n} \cdot \log\frac{mn}{\delta}\right)$ that grows logarithmically in m . This matches the result of Alweiss et al. (2021) in the noiseless setting, and is obtained by a comparatively simpler induction-based analysis. To establish this result, we first bound discrepancy in the online vector balancing problem (Theorem 8), and then bound multicolored discrepancy in an online setting with noisy vectors (Theorem 9). A known reduction from envy minimization to multicolored discrepancy minimization then gives us our final envy minimization result (Theorem 7).

1.2 Additional Related Work

Most literature on fair division assumes that agents’ preferences are captured by valuation functions. An (often implicit) assumption in this literature is that the agents have the ability to report *exact* numerical values to the designer. In practice, eliciting cardinal valuations is error-prone, as well as cognitively burdensome. A growing body of work focuses on settings where access to only *ordinal information* is available. A central notion is distortion Procaccia and Rosenschein (2006): the worst-case loss with respect to some aggregate cardinal objective resulting from the access to limited ordinal information. Bounds on distortion have been studied for problems, including fair division Amanatidis et al. (2016); Halpern and Shah (2021), matching Anshelevich and Sekar (2016); Abramowitz and Anshelevich (2018); Anshelevich and Zhu (2019); Ebadian and Shah (2025), voting Boutilier et al. (2012); Anshelevich et al. (2018) and participatory budgeting Benade et al. (2021) (see Anshelevich et al. (2021) for a survey).

Closer in spirit to our setting are works that elicit ordinal feedback (often via pairwise comparisons) while allowing the underlying cardinal valuations to be chosen adversarially, e.g., Benadè et al. (2025); Amanatidis et al. (2016); Halpern and Shah (2021); Ebadian et al. (2022); Amanatidis et al. (2022). Our results for stochastic true valuations also relate to the well-studied area of stochastic fair division Dickerson et al. (2014); Bai and Gözl (2021); Benadè et al. (2023). Finally, Bai et al. (2022) study a *smoothed* model, where a worst-case utility profile is slightly perturbed at random. In Bai et al. (2022), the perturbed profile constitutes agents’ true valuations; in contrast, we keep true valuations fixed and the perturbed profile is the observed instance.

2 Preliminaries

We consider the problem of dividing a set of m indivisible items among a set of $n \geq 2$ agents. We will represent the set of agents by $[n]$ and the set of items by $[m]$. Each agent $i \in [n]$ has a valuation function $v_i^* : 2^{[m]} \rightarrow \mathbb{R}$, which maps every subset of items $S \subseteq [m]$ to $v_i^*(S)$, the agent’s value for S . In this paper, we focus on the case of *additive* valuation functions. That is, each agent $i \in [n]$ has a value $v_{i,j}^* \in \mathbb{R}$ for each item $j \in [m]$, and the value of agent i for a subset of items $S \subseteq [m]$ is $v_i^*(S) = \sum_{j \in S} v_{i,j}^*$.

Allocations. An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is an n -partition of $[m]$, where A_i denotes the set of items allocated to agent $i \in [n]$. That is, in any allocation \mathcal{A} , all items are allocated, $\bigcup_{i \in [n]} A_i = [m]$, and each item is allocated to only one agent, $A_i \cap A_j = \emptyset$ for all distinct agents $i, j \in [n]$.

Fairness Notions. We are interested in finding allocations that are fair, i.e., each agent is satisfied with the items they receive. Envy-freeness is the central fairness notion in fair division.

⁶Alweiss et al. (2021) study the problem of online discrepancy minimization in the absence of any noise.

Definition 1 (Envy-freeness). *An allocation $\mathcal{A} = (A_1, \dots, A_n)$ is envy-free (EF) if each agent prefers their own bundle over the bundle of any other agent, i.e., for all agents $i, j \in [n]$, we have $v_i^*(A_i) \geq v_i^*(A_j)$.*

Notably, envy-free allocations do not always exist when the items being allocated are indivisible. Hence, our goal in such scenarios is to minimize the maximum envy between any pair of agents. For any two agents $i, j \in [n]$, the envy that agent i has for agent j is defined as $v_i^*(A_j) - v_i^*(A_i)$. Our objective will be to minimize the maximum envy, i.e., $\max_{i \neq j} v_i^*(A_j) - v_i^*(A_i)$. In an envy-free allocation, the maximum envy is non-positive.

Inaccurate Preference Model. Our paper focuses on the setting where we don't get to observe the true values $v_{i,j}^*$ of the agents. Instead, we assume that for each agent $i \in [n]$ and item $j \in [m]$, we only have access to a noisy estimate $\hat{v}_{i,j}$ of the true values $v_{i,j}^*$. Indeed, the envy between the agents is still based on the true values $v_{i,j}^*$.

3 Stochastic True Valuations

In this section, we focus on the setting where the true valuations of all the agents, $\{v_{i,j}^*\}_{i,j}$, are stochastic. Section 3.1 describes our result for the case when the noise $(\hat{v}_{i,j} - v_{i,j}^*)$ is stochastic and Section 3.2 describes the result for worst-case, bounded noise.

3.1 Stochastic Values and Stochastic Noise

Let \mathcal{D} and \mathcal{D}' be two real-valued distributions. We consider a model where for every agent $i \in [n]$, item $j \in [m]$, we have $v_{i,j}^* \sim \mathcal{D}$, $\eta_{i,j} \sim \mathcal{D}'$, and the algorithm gets to observe values $\hat{v}_{i,j} = v_{i,j}^* + \eta_{i,j}$. We further assume that the distributions \mathcal{D} and \mathcal{D}' are unknown to the algorithm. In this setting, we prove that maximizing welfare with respect to the observed values — a simple and computationally efficient algorithm — outputs an envy-free allocation, with high probability (Theorem 2).

Theorem 2, only requires \mathcal{D} to be a continuous distribution with positive variance; \mathcal{D}' can be an arbitrary (discrete or continuous) distribution. It is easy to see that positive variance for \mathcal{D} is a necessary requirement for envy-freeness. If \mathcal{D} has a zero variance, for e.g., \mathcal{D} takes the value 1 with probability 1 and m is not a multiple of n , then an envy-free allocation does not exist. Continuity of \mathcal{D} is assumed for the ease of exposition; in particular, to avoid tie-breaking in the algorithm.

To prove this result, we first develop a correlation inequality (Theorem 1) which is a refinement of the well-known Chebyshev's association inequality stated below in Lemma 1.

Lemma 1 (Chebyshev's Association Inequality). *Boucheron et al. (2013) Given any two real-valued, nondecreasing functions f, g and a real-valued random variable X , $\mathbb{E}[f(X) \cdot g(X)] \geq \mathbb{E}[f(X)] \cdot \mathbb{E}[g(X)]$.*

The following correlation inequality recovers Lemma 1 and additionally presents the necessary and sufficient condition for a strict inequality to hold. Notably, a strict inequality is crucially needed for the main theorem, Theorem 2. The result holds for arbitrary (discrete, absolutely continuous, singular continuous) random variables; the proof appears in Appendix B.

Theorem 1. *Given any two real-valued, nondecreasing functions f, g and a real-valued random variable X , we have $\mathbb{E}[f(X) \cdot g(X)] > \mathbb{E}[f(X)] \cdot \mathbb{E}[g(X)]$ if and only if $\text{Var}(f(X)) > 0$ and $\text{Var}(g(X)) > 0$.*

Theorem 1 is used to prove Lemma 2, which plays a central role in Theorem 2. At a high-level, Lemma 2 shows that performing welfare maximization on inaccurate values is, in fact, effective in maximizing welfare with respect to the true valuations; this eventually leads to envy-freeness.

Lemma 2. *Let \mathcal{D} and \mathcal{D}' be two probability distributions. Furthermore, suppose $X_1, \dots, X_n, Y_1, \dots, Y_n$ are independent random variables such that each $X_i \sim \mathcal{D}$ and each $Y_i \sim \mathcal{D}'$. Then, we have:*

1. $\mathbb{E}[X_1 \mid X_1 + Y_1 > \max_{i=2}^n X_i + Y_i] \geq \mathbb{E}[X_1 \mid X_1 + Y_1 \leq \max_{i=2}^n X_i + Y_i]$ if \mathcal{D}' has a positive variance.
2. $\mathbb{E}[X_1 \mid X_1 + Y_1 > \max_{i=2}^n X_i + Y_i] > \mathbb{E}[X_1 \mid X_1 + Y_1 \leq \max_{i=2}^n X_i + Y_i]$ if \mathcal{D} has a positive variance.

Proof. Define $\mathcal{E}_>$ to be the event $[X_1 + Y_1 > \max_{i=2}^n X_i + Y_i]$, and \mathcal{E}_\leq to be the event $[X_1 + Y_1 \leq \max_{i=2}^n X_i + Y_i]$.

(1) If \mathcal{D}' has a positive variance, then the random variables $X_i + Y_i$ have a positive variance. Positive variance implies that $\mathbb{P}(\mathcal{E}_>), \mathbb{P}(\mathcal{E}_\leq) > 0$, i.e., the conditional expectations $\mathbb{E}[X_1 \mid \mathcal{E}_>], \mathbb{E}[X_1 \mid \mathcal{E}_\leq]$ are well defined. To prove (1), it suffices to show that

$$\mathbb{E}[X_1 \mid \mathcal{E}_>] \geq \mathbb{E}[X_1] = \mathbb{E}[X_1 \mid \mathcal{E}_>]\mathbb{P}(\mathcal{E}_>) + \mathbb{E}[X_1 \mid \mathcal{E}_\leq]\mathbb{P}(\mathcal{E}_\leq),$$

or equivalently,

$$\mathbb{E}[X_1 \mid \mathcal{E}_>] = \frac{\mathbb{E}[X_1 \mathbb{1}_{\mathcal{E}_>}]}{\mathbb{P}(\mathcal{E}_>)} \geq \mathbb{E}[X_1].$$

This can be written as

$$\mathbb{E}[X_1 \mathbb{1}_{\mathcal{E}_>}] \geq \mathbb{E}[X_1] \mathbb{E}[\mathbb{1}_{\mathcal{E}_>}].$$

Define $h : \mathbb{R} \rightarrow \mathbb{R}$, $h(x) = x$ and $g : \mathbb{R}^2 \rightarrow \{0, 1\}$, $g(x, z) = \mathbb{1}\{x > z\}$. Additionally, let random variable $Z := \max_{i=2}^n (X_i + (Y_i - Y_1))$; note that $\mathbb{1}_{\mathcal{E}_>} = \mathbb{1}[X_1 > Z]$. Then, we want to show that

$$\mathbb{E}[h(X_1)g(X_1, Z)] \geq \mathbb{E}[h(X_1)]\mathbb{E}[g(X_1, Z)].$$

To prove this, we condition on Z . Then, using the association inequality (Lemma 1) we have that

$$\mathbb{E}[h(X_1)g(X_1, Z) \mid Z] \geq \mathbb{E}[h(X_1) \mid Z]\mathbb{E}[g(X_1, Z) \mid Z] = \mathbb{E}[h(X_1)]\mathbb{E}[g(X_1, Z) \mid Z],$$

where we remove the conditioning in $\mathbb{E}[h(X_1) \mid Z]$ due to independence. Then, taking expectations with respect to Z gives us,

$$\mathbb{E}[h(X_1)g(X_1, Z)] \geq \mathbb{E}[h(X_1)]\mathbb{E}[g(X_1, Z)].$$

(2) Based on the previous part, in order to prove (2), we need to establish $\mathbb{E}[h(X_1)g(X_1, Z)] > \mathbb{E}[X_1]\mathbb{E}[g(X_1, Z)]$. Similar to the above, we condition on a fixed value on Z . Theorem 1 implies that conditioned on Z , the strict inequality

$$\mathbb{E}[h(X_1)g(X_1, Z) \mid Z] > \mathbb{E}[h(X_1)]\mathbb{E}[g(X_1, Z) \mid Z], \tag{1}$$

holds if and only if $\text{Var}(h(X_1)) = \text{Var}(X_1) > 0$ and $\text{Var}(\mathbb{1}\{X_1 > Z\}) = \mathbb{P}[X_1 > Z](1 - \mathbb{P}[X_1 > Z]) > 0$; here, the former condition is true as per our assumption on \mathcal{D} , and the latter condition is equivalent to $\mathbb{P}[X_1 > Z] \in (0, 1)$. Hence, for a fixed value of Z , if $\mathbb{P}[X_1 > Z] \in (0, 1)$, then, and only then, we have $\mathbb{E}[h(X_1)g(X_1, Z) \mid Z] > \mathbb{E}[h(X_1)]\mathbb{E}[g(X_1, Z) \mid Z]$. The final step of the proof will be to show that the strict inequality persists after taking expectations with respect to Z . To argue this, we will show that the probability

$$\mathbb{P}_Z[\mathbb{P}_{X_1}[X_1 > Z] \in (0, 1)] > 0,$$

i.e., with positive probability over selecting Z , we get a strict inequality conditioned on Z , thereby ensuring that the final inequality $\mathbb{E}[h(X_1)g(X_1, Z)] > \mathbb{E}[X_1]\mathbb{E}[g(X_1, Z)]$ obtained after taking expectations with respect to Z is strict as well.

Let \mathcal{E} be the event that $Y_1 \geq \max_{i=2}^n Y_i$. Since Y_i s are iid drawn from \mathcal{D}' , we have $\mathbb{P}[\mathcal{E}] = \mathbb{P}[Y_1 \geq \max_{i=2}^n Y_i] \geq \frac{1}{n} > 0$. Additionally, conditioned on \mathcal{E} , we have $Z \leq \max_{i=2}^n X_i$. Hence,

$$\begin{aligned} \mathbb{P}_Z[\mathbb{P}_{X_1}[X_1 > Z] \in (0, 1)] &\geq \mathbb{P}[\mathcal{E}] \cdot \mathbb{P}_Z[\mathbb{P}_{X_1}[X_1 > Z] \in (0, 1) \mid \mathcal{E}] \\ &\geq \mathbb{P}[\mathcal{E}] \cdot \mathbb{P}_Z[\mathbb{P}_{X_1}[X_1 > \max_{i=2}^n X_i] \in (0, 1) \mid \mathcal{E}] \\ &= \mathbb{P}[\mathcal{E}] \cdot \mathbb{P}_{X_2, \dots, X_n}[\mathbb{P}_{X_1}[X_1 > \max_{i=2}^n X_i] \in (0, 1)], \end{aligned}$$

where we drop the conditioning in the second term because X_i s and Y_i s are independent and \mathcal{E} only depends on Y_i s. Finally, note $\mathbb{P}[\mathcal{E}] > 0$, and the second term in the product above is also positive since the variance of \mathcal{D} is positive. This implies that $\mathbb{P}_Z[\mathbb{P}_{X_1}[X_1 > Z] \in (0, 1)]$ is positive and completes the proof. \square

We now state the main result of this section, Theorem 2; its proof is deferred to Appendix B.

Theorem 2. *Let \mathcal{D} and \mathcal{D}' be two fixed distributions such that \mathcal{D} is continuous and has a positive variance (\mathcal{D}' could be arbitrary). Assume that $v_{i,j}^* \sim \mathcal{D}$ for all $i \in [n]$, $j \in [m]$, and we observe the noisy estimates $\hat{v}_{i,j} = v_{i,j}^* + \eta_{i,j}$ where $\eta_{i,j}$ for all $i \in [n]$, $j \in [m]$. Then, running the welfare maximization algorithm on the observed noisy values $\{\hat{v}_{i,j}\}_{i,j}$, computes an envy-free allocation with probability at least $1 - O(1/m)$ if $m/\log m \in \Omega(n)$.*

3.2 Stochastic Values and Worst-case Noise

We now consider a general setting where the value vector for each item j , $v_j^* = (v_{1,j}^*, \dots, v_{n,j}^*)$ is drawn i.i.d. from a bounded vector-valued distribution \mathcal{D} over $[0, b]^n$ for some bound b . This framework allows the agents' valuations for a specific item to be arbitrarily correlated (as the components of a single draw from \mathcal{D}). However, values across items are independent.

We show that if \mathcal{D} satisfies a mild condition — essentially that no two agents are “identical” up to scaling — and if the number of items m is sufficiently large, an efficient algorithm can use noisy estimates to find an envy-free allocation with high probability. The condition is the following: Let $X = (X_1, \dots, X_N) \sim \mathcal{D}$ be a random vector. We require that $\Pr[X_i = c \cdot X_j] = 0$ for all $i \neq j$ and constants $c \in \mathbb{R}$.

Theorem 3. *Let \mathcal{D} be a bounded nonnegative distribution satisfying the non-identical condition. Assume that $v_j^* \sim \mathcal{D}$, independently, for all items j . For each $i \in [n]$, $j \in [m]$, let $\hat{v}_{i,j}$ be a noisy estimate of $v_{i,j}^*$ such that for some $\varepsilon > 0$, we have $\max_{i,j} |v_{i,j}^* - \hat{v}_{i,j}| \leq \varepsilon$. There is a constant $c_{\mathcal{D}}$ depending only on \mathcal{D} and an algorithm such that if $\varepsilon \leq c_{\mathcal{D}}$, given the noisy estimates $\hat{v}_{i,j}$ for each i, j , the algorithm outputs an allocation $\mathcal{A} = (A_1, \dots, A_n)$ such that \mathcal{A} is envy-free with probability at least $1 - \exp(-\Omega_{\mathcal{D}}(m))$.*

The algorithm uses linear programming to find a strongly envy-free fractional allocation with respect to the noisy estimates \hat{v} . It then applies randomized rounding to this fractional allocation to obtain the final integral allocation. The size of the linear program is polynomial in n and m , with coefficients based on \hat{v} . Thus, under the mild assumption that the estimates are polynomial in bit length (which must be the case if they were produced by a polynomial time algorithm), the theorem's algorithm runs in polynomial time.

The proof of Theorem 3 (which appears in Appendix C) relies on establishing the existence of a strongly envy-free allocation in the underlying noiseless setting, which we formalize in the following lemma. This “strong” guarantee provides a buffer that is robust to the two sources of error: the bounded noise ε in the estimates and the error introduced by randomized rounding. This ensures that the final allocation is envy-free with high probability. This existence result may be of independent interest.

Lemma 3. *Let \mathcal{D} be a distribution satisfying the above conditions. Assume that $v_j^* \sim \mathcal{D}$, independently, for all items j . Then, there is a constant $c'_{\mathcal{D}}$ depending only on \mathcal{D} such that with probability at least $1 - \exp(-\Omega_{\mathcal{D}}(m))$, there exists an allocation \mathcal{A}^* such that, for all agents $i, j \in [n]$, $v_i^*(A_i^*) \geq v_i^*(A_j^*) + c'_{\mathcal{D}} m$.*

The proof of Lemma 3 relies on the following theorem about the existence of strongly envy-free allocations in cake cutting. Recall, given a σ -algebra W over a set C , a function $\mu : W \rightarrow \mathbb{R}_{\geq 0}$ is a *non-atomic probability measure* if (a) $\mu_i(A) \geq 0$ for every $A \in W$, (b) $\mu_i(\emptyset) = 0$, (c) given a countable collection of pairwise disjoint elements of W , A_1, A_2, \dots , $\mu_i(\bigcup_j A_j) = \sum_{j=1}^{\infty} \mu_i(A_j)$, (d) if $\mu_i(A) > 0$, then for some $B \subseteq A$ with $B \in W$, $0 < \mu_i(B) < \mu_i(A)$, and (e) $\mu_i(C) = 1$.

Theorem 4 (Theorem 5.5 (b) of Barbanel (2005)). *Fix n non-atomic probability measures μ_1, \dots, μ_n over a set C with σ -algebra W . Suppose that no two measures are identical. Then, there exists a partition (P_1, \dots, P_n) of C such that $\mu_i(P_i) > \mu_i(P_j)$ for all $i \neq j$.*

We are now ready to prove the lemma.

Proof of Lemma 3. Fix such a distribution \mathcal{D} satisfying the conditions. Let $X = (X_1, \dots, X_n) \sim \mathcal{D}$ be a random vector.

The key to this lemma is to represent \mathcal{D} as n different valuations over the unit interval $[0, 1]$ and to find a *strongly envy-free allocation* of $[0, 1]$. This allocation will correspond to a strongly envy-free randomized allocation. By taking a sufficiently large number of samples, we can convert this into a strongly envy-free integral allocation using concentration inequalities.

Recall that X_i is bounded in $[0, b]$ for some bound b .

By Lemma 4.22 of Kallenberg (1997), there exists a measurable function $f : [0, 1] \rightarrow [0, b]$ such that $f(U)$ has distribution \mathcal{D} when $U \sim \text{Unif}[0, 1]$. For each i , let f_i denote the i -th coordinate function of f , i.e., $f_i : [0, 1] \rightarrow [0, b]$. Define the following n set functions μ_1, \dots, μ_n on the Borel sigma-algebra $\mathcal{B}([0, 1])$:

$$\mu_i(A) = \int_A f_i d\lambda, \quad \text{for all } A \in \mathcal{B}([0, 1]),$$

where λ is the Lebesgue measure. Since $\mathbb{E}[|X_i|]$ is finite, each f_i is integrable, and μ_i is a measure on $[0, 1]$. It is nonnegative as f_i is nonnegative. Since f_i is the Radon-Nikodym derivative of μ_i with respect to Lebesgue measure, each μ_i is absolutely continuous and hence nonatomic. We have that $\mu_i(\emptyset) = 0$, and it satisfies countable additivity by properties of Lebesgue integration.

For each i , let $\mu'_i(A) = \mu_i(A)/\mu_i([0, 1])$ be its normalized measure. Note that this is well-defined, as if $\mu_i([0, 1]) = 0$, then $\mu_i(A) = 0$ for all Borel A , which implies that $f_i(U) = 0$ with probability 1. This means that $\Pr[X_i = 0 \cdot X_j] = 1$ for $j \neq i$, contradicting the assumption on \mathcal{D} .

The Radon-Nikodym derivative of μ'_i is $f'_i(x) = f_i(x)/\mu_i([0, 1])$. If $\mu'_i = \mu'_j$, their derivatives must be equal almost everywhere (see Theorem 3.8 of Folland (1999)), i.e., $f'_i = f'_j$ almost everywhere. This implies that $f_i(x) = c \cdot f_j(x)$ for $c = \mu_i([0, 1])/\mu_j([0, 1])$ almost everywhere on $[0, 1]$, which again implies that $\Pr[X_i = c \cdot X_j] = 1$, breaking our condition on \mathcal{D} .

Using Theorem 4, there is a partition of $[0, 1]$ among the agents B_1, \dots, B_n such that $\mu'_i(B_i) > \mu'_i(B_j)$ for $i \neq j$. Note that this immediately implies that $\mu_i(B_i) \geq \mu_i(B_j)$ for all $i \neq j$ as well.

Define:

$$\varepsilon = \min_{i \neq j} (\mu_i(B_i) - \mu_i(B_j)),$$

which is positive since it is the minimum over a finite number of strictly positive values.

We now interpret μ_1, \dots, μ_n as valuation functions over $[0, 1]$. We observe:

$$\mathbb{E}[f_i(U) \cdot \mathbb{1}[U \in B_i]] = \int_{B_i} f_i d\lambda = \mu_i(B_i),$$

and thus for any $i \neq j$:

$$\mathbb{E}[f_i(U) \cdot (\mathbb{1}[U \in B_i] - \mathbb{1}[U \in B_j])] = \mu_i(B_i) - \mu_i(B_j) \geq \varepsilon.$$

We now sample m i.i.d. draws $X^1, \dots, X^m \sim \mathcal{D}$ and show that with high probability, we obtain a strongly envy-free allocation. We can couple these draws with uniform draws by sampling $U^1, \dots, U^m \sim \text{Unif}[0, 1]$ i.i.d., and setting

$$X^j = f(U^j).$$

We define an allocation \mathcal{A} by:

$$A_i = \{k \mid U^k \in B_i\}.$$

Set $c'_D = \varepsilon/2$, which so far depends only on the distribution \mathcal{D} and not m . We will show that with probability $1 - \exp(-\Omega_{\mathcal{D}}(m))$, we have:

$$v_i(A_i) - v_i(A_j) > c'_D \cdot m, \quad \forall i \neq j.$$

Fix $i \neq j$. Define:

$$Y_{i,j}^k = X_i^k \cdot (\mathbb{1}[k \in A_i] - \mathbb{1}[k \in A_j]).$$

Observe that $v_i(A_i) - v_i(A_j) = \sum_k Y_{i,j}^k$. We would like to show that for all $i \neq j$, $\sum_k Y_{i,j}^k \geq c'_D \cdot m$. To do so, we bound the complement: $\Pr[\sum_k Y_{i,j}^k < c'_D \cdot m]$. Note that $\mathbb{E}[\sum_k Y_{i,j}^k] \geq m \cdot \varepsilon = 2c'_D \cdot m$. Thus, this is a deviation from the mean of $> c'_D \cdot m$. Furthermore, each $Y_{i,j}^k \in [-b, b]$, $\mathbb{E}[Y_k] \geq \varepsilon$, the Y_k s are independent. We can therefore apply *Hoeffding's inequality* to get:

$$\Pr \left[\sum_k Y_k < c \cdot m \right] \leq \exp \left(-\frac{c^2 m}{2b^2} \right).$$

Applying a union bound over all $n(n-1)$ pairs:

$$\Pr [\exists i \neq j, v_i(A_i) - v_i(A_j) < c \cdot m] \leq n(n-1) \exp \left(-\frac{c^2 m}{2b^2} \right) \in \exp(-\Omega_{\mathcal{D}}(m)),$$

completing the proof. □

The purpose of the non-identical condition is now clearer. If two agents had the same values (up to scaling), then it is impossible that both strongly prefer their own bundles.

4 Worst-Case True Valuations

The focus of this section will be on settings where the true preferences of the agents $v_{i,j}^*$ are fixed/worst-case. In Section 4.1, we detail our analysis of the Round-Robin algorithm. Section 4.2 presents our two applications of the Round-Robin bound. Finally, in Section 4.3, we describe our results on discrepancy and envy minimization.

4.1 Worst-Case Values and Worst-case Noise

We start by showing that, given worst-case true valuations and (bounded) worst-case noise, we can find an allocation with maximum envy at most m/n times the noise (Theorem 5); we prove that this is the best bound possible by giving a matching lower bound (Theorem 6).

The guarantee of Theorem 5 is achieved by the simple Round-Robin process. Here, for each agent $i \in [n]$, we assume the existence of a possibly unknown value $\delta_i \geq 0$, which when added to the estimates $\hat{v}_{i,j}$, results in a value ε close to $v_{i,j}^*$, i.e., $|v_{i,j}^* - (\hat{v}_{i,j} + \delta_i)| \leq \varepsilon$. The theorem holds for any values of $\delta_i \geq 0$, and, in particular, the natural choice of $\delta_i = 0$ for all $i \in [n]$. The reason the envy bound in Theorem 5 is independent of δ_i is a consequence of the fact that round robin produces *balanced* allocations, ensuring that every agent gets roughly the same number of items. While it is natural to assume that $\delta_i = 0$ for all $i \in [n]$, Theorem 5 is stated in terms of the additional parameters δ_i s as the balancedness property will be used in an application we will present in Section 4.2.

Theorem 5 (Upper bound for worst-case noise). *For every $i \in [n]$ and $j \in [m]$, let $\hat{v}_{i,j}$ be the estimate of true value $v_{i,j}^* \in [0, b]$, such that $\max_{i,j} |v_{i,j}^* - \hat{v}_{i,j} - \delta_i| \leq \varepsilon$, for some possibly unknown constants $\varepsilon > 0$, $\delta_i \geq 0$ for each agent $i \in [n]$. Given $\hat{v}_{i,j}$ for each i, j , we can efficiently compute an allocation $\mathcal{A} = (A_1, \dots, A_n)$ such that, for every pair of agents $i, i' \in [n]$, we have $v_i^*(A_{i'}) - v_i^*(A_i) \leq 2\varepsilon \lceil \frac{m}{n} \rceil + b$.*

Proof. Consider the allocation produced by the Round-Robin algorithm, when executed on the estimated values $\hat{v}_{i,j}$. Recall that the Round-Robin algorithm allocates items to agents over a sequence of rounds, according to a fixed, arbitrary order. In each round, the next agent in the order picks an item that maximizes her value ($\hat{v}_{i,j}$, in our case), among the set of remaining items (ties can be broken in an arbitrary manner). The allocation \mathcal{A} produced by Round-Robin satisfies envy-freeness up to one item (EF1) *with respect to the noisy values*. That is, for any pair of agents $i, i' \in [n]$ with $\hat{v}_i(A_{i'}) \neq \emptyset$, we have $\hat{v}_i(A_i) \geq \hat{v}_i(A_{i'} \setminus \{g\})$ for some $g \in A_{i'}$. Furthermore, the allocation \mathcal{A} produced by Round-Robin is *balanced*, i.e., all agents get roughly equal number of items, for each agent $i \in [n]$, $|A_i|$ is either $\lfloor \frac{m}{n} \rfloor$ or $\lceil \frac{m}{n} \rceil$.

Given these facts, we will bound the envy of an agent $i \in [n]$ for agent $i' \in [n]$ with respect to the true (unknown) values as follows. If $A_{i'} = \emptyset$, then there is no envy since $v_i^*(A_{i'}) - v_i^*(A_i) = -v_i^*(A_i) \leq 0$. Otherwise, if $A_{i'} \neq \emptyset$, let $g \in A_{i'}$ be an item such that $\hat{v}_i(A_i) \geq \hat{v}_i(A_{i'} \setminus \{g\})$. Additionally, note that the inequality $|v_{i,j}^* - \hat{v}_{i,j} - \delta_i| \leq \varepsilon$ implies that (i) $v_i^*(A_i) \geq \hat{v}_i(A_i) + \delta_i|A_i| - \varepsilon|A_i|$ and (ii) $v_i^*(A_{i'} \setminus \{g\}) \leq \hat{v}_i(A_{i'} \setminus \{g\}) + \delta_i|A_{i'} \setminus \{g\}| + \varepsilon|A_{i'} \setminus \{g\}|$. Using this we get,

$$\begin{aligned}
v_i^*(A_{i'}) - v_i^*(A_i) & \leq v_i^*(A_{i'} \setminus \{g\}) - v_i^*(A_i) + b && (v_{i,j}^* \leq b \text{ for all } j \in [m]) \\
& \leq \hat{v}_i(A_{i'} \setminus \{g\}) + \delta_i|A_{i'} \setminus \{g\}| + \varepsilon|A_{i'} \setminus \{g\}| - v_i^*(A_i) + b && (\text{Via (i)}) \\
& \leq (\hat{v}_i(A_{i'} \setminus \{g\}) - \hat{v}_i(A_i)) + \delta_i(|A_{i'} \setminus \{g\}| - |A_i|) + \varepsilon(|A_i| + |A_{i'} \setminus \{g\}|) + b && (\text{Via (ii)}) \\
& \leq (\hat{v}_i(A_{i'} \setminus \{g\}) - \hat{v}_i(A_i)) + \varepsilon(|A_i| + |A_{i'} \setminus \{g\}|) + b && (\delta_i \geq 0, |A_i|, |A_{i'}| \in \{\lfloor \frac{m}{n} \rfloor, \lceil \frac{m}{n} \rceil\}) \\
& \leq \hat{v}_i(A_{i'} \setminus \{g\}) - \hat{v}_i(A_i) + 2\varepsilon \lceil \frac{m}{n} \rceil + b && (|A_i|, |A_{i'}| \leq \lceil \frac{m}{n} \rceil) \\
& \leq 2\varepsilon \lceil \frac{m}{n} \rceil + b. && (\hat{v}_i(A_{i'} \setminus \{g\}) \leq \hat{v}_i(A_i))
\end{aligned}$$

This concludes the proof. \square

The next theorem shows that the envy bound achieved by Round-Robin is tight (up to an additive term) for deterministic algorithms.

Theorem 6 (Lower bound for worst-case noise). *For every $i \in [n]$ and $j \in [m]$, let $\hat{v}_{i,j}$ be an estimate of $v_{i,j}^*$, such that $\max_{i,j} |v_{i,j}^* - \hat{v}_{i,j}| \leq \varepsilon$, for some $\varepsilon > 0$. No deterministic algorithm that takes $\hat{v}_{i,j}$ as input, and outputs a complete allocation, can guarantee that the envy between any two agents is strictly less than $2\varepsilon \frac{m}{n}$.*

Proof. Consider any deterministic algorithm, and let $\mathcal{A} = (A_1, \dots, A_n)$ be the allocation returned by this algorithm when $\hat{v}_{i,j} = 1/2$ for all $i \in [n]$ and $j \in [m]$. Let $p \in \arg \max_{i \in [n]} |A_i|$ and $q \in \arg \max_{i \in [n] \setminus \{p\}} |A_i|$ be the agent with the largest and second largest number of items in \mathcal{A} , respectively. We have that $|A_p| + |A_q| \geq \frac{2m}{n}$. We will construct the true values $v_{i,j}^*$ to maximize the envy of agent q for agent p . In particular, we set $v_{q,j}^* = 1/2 + \varepsilon$ for all $j \in A_p$ and $v_{q,j}^* = 1/2 - \varepsilon$ for all $j \in A_q$. This implies that the envy of agent q for agent p , $v_q^*(A_p) - v_q^*(A_q) = |A_p|(1/2 + \varepsilon) - |A_q|(1/2 - \varepsilon) = \varepsilon(|A_p| + |A_q|) + 1/2(|A_p| - |A_q|)$. Using the fact that $|A_p| \geq |A_q|$ and $|A_p| + |A_q| \geq \frac{2m}{n}$, we get $v_q^*(A_p) - v_q^*(A_q) \geq 2\varepsilon \frac{m}{n}$. \square

4.2 Applications of Theorem 5

Estimating Preferences Via a BTL Query Model

We now describe a method for obtaining a reasonably accurate estimate of all agents' preferences using the Bradley-Terry-Luce (BTL) estimation model. In this model, each agent is asked pairwise comparison queries of the form "which item do you like more, k or ℓ ?" and it responds with a binary choice from $\{k, \ell\}$. Following the BTL model, the set of queried item pairs for each agent is determined independently by an Erdős-Rényi random graph with edge probability p , where each unordered pair $\{k, \ell\}$ is included

independently with probability p . For every queried pair, we collect K independent comparison responses (i.e., K repeated queries per observed pair). A detailed formulation of the BTL model is provided in Appendix A. Using the collected responses, we run a standard maximum-likelihood estimator for each agent $i \in [n]$, to recover $\tilde{\theta}_i \in \mathbb{R}^m$, an approximation of θ_i^* , which is a centered version of the true preference vector of agent i . Let $\Theta^* = [\theta_1^*; \dots; \theta_n^*]^T$ and $\hat{\Theta} = [\tilde{\theta}_1; \dots; \tilde{\theta}_n]^T$ denote the concatenated true and estimated preference matrices, then the following lemma provides a bound on the estimation error of each θ_i^* .

Lemma 4 (ℓ^2 and ℓ^∞ -Error Bounds (Chen et al., 2022, Theorem 3.1)). *Assume $p \geq c_0 \frac{\log m}{m}$ for a sufficiently large constant $c_0 > 0$, and that K comparisons are collected per observed pair under an Erdős–Rényi observation graph with parameter p . Let $\theta^* \in \mathbb{R}^m$ satisfies $\|\theta^*\|_\infty \leq b$ and $\theta^{*\top} \mathbf{1}_m = 0$. Then, there exists a polynomial-time estimator that computes $\tilde{\theta} \in \mathbb{R}^m$ by asking $\binom{m}{2} pK$ average number of pairwise comparisons and satisfies the following bounds*

$$\|\tilde{\theta} - \theta^*\|_2 \leq c_2 \cdot \frac{1}{\sqrt{pK}}, \quad \|\tilde{\theta} - \theta^*\|_\infty \leq c_\infty \cdot \sqrt{\frac{\log m}{mpK}},$$

with probability at least $1 - O(1/m^7)$, for some constants $c_2, c_\infty > 0$ only depending on b .

The estimated preferences obtained from Lemma 4 are given to the Round-Robin algorithm to compute an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ with a bounded envy. The guarantees of the end-to-end application combining the BTL model with Round-Robin are stated in Application 1.

Before stating Application 1, we note that the combination of BTL estimation with Round-Robin is made possible only because of the balancedness property that Round-Robin exhibits. By design, the BTL model is invariant to linear shifts.⁷ Consequently, we cannot recover the average value of all the items, for any agent using BTL. On the other hand, maximum envy between agents does depend on this average value, and cannot be bounded without it. Using Round-Robin addresses this problem by making envy between any pair of agents independent of the average value of all the items — this is a direct consequence of the balancedness property, which ensures that agents get a roughly equal number of items.⁸

Application 1 (Inference and Fair Division). *Assume there are m indivisible items and n additive agents with $n < m$, where $v_{i,j}^* \in [0, 1]$ is the value of agent i for item j . Let $V^* \in [0, 1]^{n \times m}$ be the valuation matrix where $V_{i,j}^* = v_{i,j}^*$. If the users respond to BTL-based pairwise queries, then there is an algorithm that asks $\binom{m}{2} pK$ average number of pairwise-comparison queries to each agent $i \in [n]$, and with probability at least $1 - O(1/m^6)$, produces an allocation $\mathcal{A} = (A_1, \dots, A_n)$ in polynomial time that has bounded maximum envy. Formally, the obtained allocation \mathcal{A} satisfies*

$$\max_{i,j} v_i(A_i) - v_i(A_j) = O\left(\sqrt{\frac{m \log m}{n^2 p K}}\right),$$

where K is the number of queries per-pair and p is the observation graph probability satisfying $p = \Omega\left(\frac{\log m}{m}\right)$.

Proof of Application 1. We begin by transforming the true valuation matrix $V^* \in [0, 1]^{n \times m}$ into a preference score matrix Θ^* via a centering operation. Specifically, for each agent $i \in [n]$ and item $j \in [m]$, we define

$$\theta_{i,j}^* = v_{i,j}^* - \frac{1}{m} \sum_{l=1}^m v_{i,l}^* \quad (2)$$

⁷Formally, the pairwise comparison probability in BTL, stated in Equation (9), depends only on the difference in preference values and not the mean.

⁸This independence is evident from the fact that the envy bound in Theorem 5 does not depend on the constants δ_i .

Centering ensures that $\Theta^* \mathbf{1}_m = \mathbf{0}_n$. Furthermore, as $v_{i,j}^*$ s lie in $[0, 1]$, we have $\theta_{i,j}^* \in (-1, 1)$.

Next, for every row of Θ^* , that is each vector $\theta_i^{*\top}$, we invoke Lemma 4 on θ_i^* with $b = 1$. Under the conditions specified in that theorem, each invocation returns an estimate $\tilde{\theta}_i$ satisfying

$$\|\tilde{\theta}_i - \theta_i^*\|_\infty = O\left(\sqrt{\frac{\log m}{mpK}}\right),$$

with probability at least $1 - O(1/m^7)$. Therefore, by taking $\hat{\Theta} = [\tilde{\theta}_1; \dots; \tilde{\theta}_n]^\top$, we have

$$\max_{i,j} |\hat{\Theta}_{i,j} - \Theta_{i,j}^*| = O\left(\sqrt{\frac{\log m}{mpK}}\right) \quad (3)$$

Through a union bound, we get that Equation (3) holds with probability at least $1 - O(n/m^7)$, which is at least $1 - O(1/m^6)$ since $n < m$. Finally, we define the estimated valuation matrix as $\hat{V} := \hat{\Theta}$, and for each agent i , we define $\delta_i := \frac{1}{m} \sum_{j=1}^m v_{i,j}^* \in [0, 1]$, which represents the mean valuation of agent i . Using this along with Equation (2) and (3), for any $i \in [n]$ and $j \in [m]$, we get that

$$\left| \hat{v}_{i,j} - \left(v_{i,j}^* - \frac{1}{m} \sum_{l=1}^m v_{i,l}^* \right) \right| = |\hat{v}_{i,j} - \Theta_{i,j}^*| = |\hat{\Theta}_{i,j} - \Theta_{i,j}^*| = O\left(\sqrt{\frac{\log m}{mpK}}\right).$$

This implies that, for each $i \in [n]$ and $j \in [m]$, an estimate $\hat{v}_{i,j}$ that satisfies $|v_{i,j}^* - \hat{v}_{i,j} - \delta_i| = |\hat{v}_{i,j} - v_{i,j}^* + \delta_i| = O\left(\sqrt{\frac{\log m}{mpK}}\right)$ for the value of $\delta_i = \frac{1}{m} \sum_{l=1}^m v_{i,l}^* \in [0, 1]$. Note that while $\hat{v}_{i,j}$ is known, δ_i is unknown. Regardless, we can apply Theorem 5 to compute an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ in polynomial time that has bounded maximum envy with respect to the true values, i.e., $\max_{i,j} v_i^*(A_j) - v_i^*(A_i) = O\left(\varepsilon \cdot \frac{m}{n}\right) = O\left(\sqrt{\frac{m \log m}{n^2 p K}}\right)$. This concludes the proof. \square

Worst-Case Values and Stochastic Noise

We consider the following simple model for stochastic noise: $\hat{v}_{i,j} = v_{i,j}^* + \varepsilon_{i,j}$, for $\varepsilon_{i,j} = \sigma_{i,j} \cdot z_{i,j}$. We assume that $z_{i,j}$ s are drawn i.i.d. from a non-negative distribution \mathcal{D} , and $\sigma_{i,j}$ s are drawn uniformly at random from the set $\{-1, 1\}$. That is, $z_{i,j}$ is the magnitude of the noise, and $\sigma_{i,j}$ denotes whether the noise underestimates or overestimates the value. The next theorem shows that, when \mathcal{D} is not heavy tailed, then we can efficiently compute an allocation with small envy (with respect to the true valuations). Formally, our condition on \mathcal{D} is that \mathcal{D} has monotone hazard rate (MHR): $\frac{1-F(x)}{f(x)}$ is a non-increasing function, where $F(x)$ and $f(x)$ are its cumulative distribution function and probability density function, respectively.

Application 2. For every $i \in [n]$ and $j \in [m]$, let $z_{i,j}$ s be i.i.d. draws from a non-negative MHR distribution \mathcal{D} such that $\mathbb{E}[\mathcal{D}] \leq \frac{n}{m \log(nm)}$, and $\sigma_{i,j}$ s be uniformly random draws from the set $\{-1, 1\}$. Let $\hat{v}_{i,j} = v_{i,j}^* + \varepsilon_{i,j}$, for $\varepsilon_{i,j} = \sigma_{i,j} \cdot z_{i,j}$. Given $\hat{v}_{i,j}$ for each i, j , we can efficiently compute an allocation $\mathcal{A} = (A_1, \dots, A_n)$ such that, for every pair of agents $i, i' \in [n]$, we have $v_i^*(A_{i'}) - v_i^*(A_i) \leq 10$, with probability at least $1 - \frac{1}{(nm)^{3/5}}$.

Proof of Application 2. We use the following two lemmas about MHR distributions, where $\mathcal{D}_{k:n}$ denotes the distribution of the k -th lowest of n i.i.d. samples from \mathcal{D} .

Lemma 5 (Babaioff et al. (2017); Lemma 5.3). For any non-negative MHR distribution \mathcal{D} and $m \leq n$, it holds that $\frac{\mathbb{E}[\mathcal{D}_{m:m}]}{\mathbb{E}[\mathcal{D}_{n:n}]} \geq \frac{H_m}{H_n}$, where H_k is the k -th harmonic number.

Lemma 6 (Azizzadenesheli et al. (2026); Lemma 3). For any non-negative MHR distribution \mathcal{D} , and $n \geq 4$, $\Pr[\mathcal{D}_{n:n} < 2\mathbb{E}[\mathcal{D}_{n:n}]] \geq 1 - \frac{1}{n^{3/5}}$.

Let $\varepsilon_{\max} = \max_{i,j} |\varepsilon_{i,j}|$ be the magnitude of the largest noise in an instance. ε_{\max} follows the same distribution as $\mathcal{D}_{nm:nm}$. Therefore, letting $\mu_{\max} = \mathbb{E}[\mathcal{D}_{nm:nm}]$, by Lemma 6, $\Pr[\varepsilon_{\max} < 2\mu_{\max}] \geq 1 - \frac{1}{(nm)^{3/5}}$. However, by Lemma 5, $\mu_{\max} = \mathbb{E}[\mathcal{D}_{nm:nm}] \leq H_{nm}\mathbb{E}[\mathcal{D}] \leq \log(nm)\mathbb{E}[\mathcal{D}]$. Therefore,

$$\Pr[\varepsilon_{\max} < 2 \log(nm) \mathbb{E}[\mathcal{D}]] \geq 1 - \frac{1}{(nm)^{3/5}}.$$

Let $\beta = 2 \log(nm) \mathbb{E}[\mathcal{D}] \leq \frac{2n}{m}$, and condition on the event that $\varepsilon_{\max} < \beta$. Then, applying Theorem 5 for $\delta_i = 0$ for all agents i , we can find an allocation $\mathcal{A} = (A_1, \dots, A_n)$ such that $\max_{i,i'} v^*(A_{i'}) - v^*(A_i) \leq 2\beta \left\lceil \frac{m}{n} \right\rceil + 1 \leq 10$. \square

4.3 Discrepancy Minimization with Inaccurate Inputs

This section focuses on a variant of the multicolored discrepancy minimization problem. Upper bounds on multicolored discrepancy are derived, as it is known that these bounds directly translate to the problem of envy minimization Alweiss et al. (2021); Halpern et al. (2025). We begin with a brief refresher on discrepancy minimization.

In the vector balancing problem, given a set of vectors u_1, u_2, \dots, u_m satisfying $\|u_i\|_{\infty} \leq 1$, the goal is to compute signs $s_1, \dots, s_m \in \{-1, 1\}$ to minimize the *discrepancy*, defined as $\|\sum_{i=1}^m s_i u_i\|_{\infty}$. The multicolored discrepancy minimization problem is a generalization of the vector balancing problem in which we have to assign one color, among a set of $[k] = \{1, 2, \dots, k\}$ colors, to each vector, to minimize the *multicolored discrepancy*, i.e., $\max_{i,j \in [k]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_{\infty}$, where for each $\ell \in [k]$, S_{ℓ} is the set of vectors that are assigned the color ℓ .

To see how envy minimization reduces to multicolored discrepancy minimization, consider the following. If we let each agent $i \in [n]$ correspond to a color (i.e., we set $k = n$) and each item $j \in [m]$ correspond to a vector $v_j^* = (v_{1,j}^*, v_{2,j}^*, \dots, v_{n,j}^*)/\sqrt{n}$ (ensuring $\|v_j^*\|_{\infty} \leq 1$), then we get an instance of the multicolored discrepancy minimization problem. Moreover, for any allocation $\mathcal{A} = (A_1, \dots, A_n)$, the maximum envy $\max_{i,j \in [n]} v_i(A_j) - v_i(A_i)$ is always at most the multicolored discrepancy $\max_{i,j \in [n]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_{\infty}$ of the corresponding problem, where the set S_i contains the vectors that correspond to the bundle of items A_i assigned to agent i . Hence, $\max_{i,j \in [n]} v_i(A_j) - v_i(A_i) \leq \max_{i,j \in [n]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_{\infty}$ and any upper bound on multicolored discrepancy translates to an upper bound on maximum envy.

We study envy minimization for a model in which we get to observe the inaccurate preferences $(v_j^* + \eta_j) \in \mathbb{R}^n$ for each item $j \in [m]$. We then have to allocate this item to an agent, and after the allocation, the true preference vector v_j^* is revealed to us. We will assume that $v_{i,j}^* \in [-1, 1]$ and $\eta_{i,j} \in [-\varepsilon, \varepsilon]$ are both worst-case.⁹ In this model, we give an efficient *online* algorithm that computes an allocation whose maximum envy is upper bounded by $\mathcal{O}(\sqrt{n}(\log \frac{mn}{\delta} + m\varepsilon))$ with probability at least $1 - \delta$.

Theorem 7. *Let $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m \in \mathbb{R}^n$ be a sequence of items (valuation vectors) that arrive one at a time, where $v_{i,j}^* \in [-1, 1], \eta_{i,j} \in [-\varepsilon, \varepsilon]$ for all $i \in [m], j \in [n]$. When each item $v_i^* + \eta_i$ arrives, we need to allocate it to one of the n agents, after which the true valuation vector v_i^* is revealed to us. In this model, there is an efficient online algorithm that computes an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ that guarantees the following maximum envy bound $\max_{i,j \in [n]} v_i^*(A_j) - v_i^*(A_i) \leq \frac{27\sqrt{n}}{C} \log \frac{5nm}{\delta} + 6m\sqrt{n}\varepsilon$, with probability at least $1 - \delta$, for $C > 0$ being a universal constant following Definition 2.*

To prove Theorem 7, we first develop an online algorithm for vector balancing problem with noisy input vectors (Section 4.3.1). Using this result, we obtain a bound for the corresponding multicolored discrepancy problem (Section 4.3.2). Finally, via the reduction from envy minimization to multicolored discrepancy minimization explained earlier, we directly obtain Theorem 7.

⁹This is without loss of generality, as the bounds we obtain can be scaled by a factor of b if $v_{i,j}^* \in [-b, b]$ and $\eta_{i,j} \in [-\varepsilon b, \varepsilon b]$.

4.3.1 Vector balancing with noisy input

We consider the following weighted vector-balancing problem with inaccurate vectors as input. Formally, given a sequence of inaccurate vectors $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m$ as input, we want to find signs $s_1, s_2, \dots, s_m \in \{p, p-1\}$ where $p \in [1/2, 2/3]$ is a fixed parameter, to minimize $\|\sum_{i=1}^m s_i v_i^*\|_\infty$. We show that $\|\sum_{i=1}^m s_i v_i^*\|_\infty \leq \mathcal{O}(\sqrt{n}(\log \frac{mn}{\delta} + m\varepsilon))$ with probability at least $1 - \delta$ is achieved by a modified version of the online algorithm given by Alweiss et al. (2021), which achieves a similar logarithmic bound on discrepancy, in the absence of noise. Our algorithm (Balance Under Noise) is stated as Algorithm 1.

ALGORITHM 1: BALANCE UNDER NOISE

Input : Parameter $p \in [1/2, 2/3]$ and a sequence of vector $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m$ where $\|v_t\|_\infty \leq 1, \|\eta_t\|_\infty \leq \varepsilon$ for all $t \in [m]$.

Output: Signs $s_1, s_2, \dots, s_m \in \{p, p-1\}$.

- 1 Initialize $w_1 = \mathbf{0} \in \mathbb{R}^n$ and $s_1, s_2, \dots, s_m = 0$.
 - 2 Define function $r(x) := \min\{1, \max\{0, x\}\}$ for all $x \in \mathbb{R}$.
 - 3 **for** $t \leftarrow 1, 2, \dots, m$ **do**
 - 4 Set $c_t \leftarrow \frac{9n}{4C} \log \frac{5nm}{\delta} + \varepsilon nt$
 - 5 Randomly sample $s_t \sim \begin{cases} p, & \text{with probability } r\left((1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right), \\ p-1, & \text{with probability } 1 - r\left((1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right). \end{cases}$
 - 6 Set $w_{t+1} \leftarrow w_t + s_t v_t^*$
 - 7 **end**
 - 8 **return** signs s_1, s_2, \dots, s_m .
-

Algorithm 1 maintains a running discrepancy vector, $w_t = \sum_{j=1}^{t-1} s_j v_j^*$. When a new inaccurate vector $v_t^* + \eta_t$ arrives, it simply samples the sign s_t from the following distribution: p with probability $r\left((1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right)$ and $p-1$ otherwise (Line 5). The function $r(\cdot)$ defined in Line 2, ensures that the probability value is valid, i.e., lies between zero and one. Note that the probability is computed based on the inaccurate vector $v_t^* + \eta_t$. Once the sign s_t has been sampled, the true vector v_t^* is revealed, and w_{t+1} is updated (Line 6).

To establish the discrepancy bound, we will show that the high-dimensional random vector w_t behaves like a subgaussian.¹⁰ We begin by introducing relevant definitions and a useful lemma.

Definition 2 (Vershynin (2018)). *Every σ -subgaussian random vector $X \in \mathbb{R}^n$ satisfies the following bounds*

1. $\mathbb{E}[\exp(\langle X, u \rangle^2 / \sigma^2)] \leq 2$ for all $u \in \mathcal{S}^{n-1}$ and¹¹
2. If $\mathbb{E}[X] = 0$, then $\mathbb{E}[\exp(\langle X, \theta \rangle)] \leq \exp(C\sigma^2 \|\theta\|_2^2)$ for all $\theta \in \mathbb{R}^n$, where $C > 0$ is a universal constant, independent of the random vector.

Lemma 7 (Hoeffding's Lemma). *Let X be any random variable that lies in the range $[a, b]$. Then, for any constant $\lambda \in \mathbb{R}$, we have $\mathbb{E}[\exp(\lambda X)] \leq \exp(\lambda \mathbb{E}[X] + \frac{\lambda^2 (b-a)^2}{8})$.*

We will now define an event that will play a crucial role in the analysis of Algorithm 1. Let Fail_t be the event that during the start of iteration t , w_t satisfies the following condition,

$$\text{Fail}_t := \mathbb{1}[\langle w_t, v_t^* + \eta_t \rangle > 4(1-p)c_t] \vee \mathbb{1}[\langle w_t, v_t^* + \eta_t \rangle < -4pc_t] \vee \mathbb{1}[\|w_t\|_\infty > c_t/\sqrt{n}] \quad (4)$$

¹⁰In Lemma 8, we will show that w_t satisfies the definition of subgaussianity, albeit with an additional multiplicative factor.

¹¹The set $\mathcal{S}^{n-1} := \{x \in \mathbb{R}^n \mid \|x\|_2 = 1\}$ denotes the unit sphere in n dimension.

That is, Fail_t is true if either $\|w_t\|_\infty$ is large or $(1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t} \notin [0, 1]$. Note that if the latter condition doesn't hold, i.e., $(1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t} \in [0, 1]$, then we have $r\left((1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right) = (1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}$. We additionally define Fail_t^c be the complement event of Fail_t . The following lemma proves a concentration result for the random vector w_t , showing that it satisfies the definition of subgaussianity (Definition 2), with an extra multiplicative factor.

Lemma 8. *For all $t \in [m]$, the random vector $w_t \in \mathbb{R}^n$ satisfies the following properties:*

1. for all $\theta \in \mathbb{R}^n$, we have $\mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \theta \rangle) \mid \cap_{\ell=1}^{t-1} \text{Fail}_\ell^c] \leq (1 - \frac{\delta}{m})^{1-t} \exp(\sigma_t^2 \|\theta\|_2^2)$ where $\sigma_t^2 = c_t = \frac{9n}{4C} \log \frac{5nm}{\delta} + \varepsilon nt$.
2. $\mathbb{P}[\text{Fail}_t^c \mid \cap_{\ell=1}^{t-1} \text{Fail}_\ell^c] \geq 1 - \frac{\delta}{m}$.

Proof. We will prove (1) and (2) property by induction on t . Since $w_1 = \mathbf{0}$, both (1) and (2) are trivially satisfied for $t = 1$. Now assuming that (1) and (2) are true for w_t , we now prove that both of these hold for w_{t+1} as well; this will complete the inductive argument.

Proving property (1). For any $\theta \in \mathbb{R}^n$ consider,

$$\begin{aligned} & \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\ &= \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_t + s_t v_t^*, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \quad (w_{t+1} = w_t + s_t v_t^*) \\ &= \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \theta \rangle) \cdot \mathbb{E}_{s_t} [\exp(s_t \langle v_t^*, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \end{aligned} \quad (5)$$

To proceed further, we will now invoke Hoeffding's lemma (Lemma 7) by setting $X = s_t$, $\lambda = \langle v_t^*, \theta \rangle$, $b = p$, and $a = p - 1$. This gives us

$$\begin{aligned} \mathbb{E}_{s_t} [\exp(s_t \langle v_t^*, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] &\leq \exp\left(\langle v_t^*, \theta \rangle \mathbb{E}[s_t] + \frac{\langle v_t^*, \theta \rangle^2 (p - (p-1))^2}{8}\right) \\ &= \exp\left(\frac{-\langle v_t^*, \theta \rangle \langle w_t, v_t^* + \eta_t \rangle}{4c_t} + \frac{\langle v_t^*, \theta \rangle^2}{8}\right), \end{aligned} \quad (6)$$

where the final equality uses the fact that given the event Fail_t^c , we have $r\left((1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right) = (1-p) - \frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}$, and therefore, $\mathbb{E}[s_t] = -\frac{\langle w_t, v_t^* + \eta_t \rangle}{4c_t}$ (Line 5). Substituting Equation (6) in Equation (5) gives us,

$$\begin{aligned} & \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\ &\leq \mathbb{E}_{s_1, \dots, s_{t-1}} \left[\exp(\langle w_t, \theta \rangle) \cdot \exp\left(\frac{-\langle v_t^*, \theta \rangle \langle w_t, v_t^* + \eta_t \rangle}{4c_t} + \frac{\langle v_t^*, \theta \rangle^2}{8}\right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \\ &= \mathbb{E}_{s_1, \dots, s_{t-1}} \left[\exp\left(\langle w_t, \theta \rangle - \frac{\langle v_t^*, \theta \rangle \langle w_t, v_t^* + \eta_t \rangle}{4c_t}\right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \exp\left(\frac{\langle v_t^*, \theta \rangle^2}{8}\right) \\ &= \mathbb{E}_{s_1, \dots, s_{t-1}} \left[\exp\left(\left\langle w_t, \theta - \frac{\langle v_t^*, \theta \rangle (v_t^* + \eta_t)}{4c_t} \right\rangle\right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \exp\left(\frac{\langle v_t^*, \theta \rangle^2}{8}\right) \end{aligned} \quad (7)$$

Note that our inductive hypothesis, in particular property (1) for w_t does not apply to Equation (7) since the event being conditioned on is $\cap_{\ell=1}^t \text{Fail}_\ell^c$ and not $\cap_{\ell=1}^{t-1} \text{Fail}_\ell^c$. Nevertheless, using the inductive hypothesis, we can prove the following useful claim whose proof appears in Appendix D.

Claim 1. For all $\beta \in \mathbb{R}^n$, we have $\mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \leq (1 - \frac{\delta}{m})^{-t} \exp(\sigma_t^2 \|\beta\|_2^2)$.

By setting $\beta = \theta - \frac{\langle v_t^*, \theta \rangle (v_t^* + \eta_t)}{4c_t}$ in Claim 1 and substituting in Equation (7) we get

$$\begin{aligned} & \left(1 - \frac{\delta}{m}\right)^t \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\ & \leq \exp\left(\sigma_t^2 \cdot \left\| \theta - \frac{\langle v_t^*, \theta \rangle (v_t^* + \eta_t)}{4c_t} \right\|_2^2\right) \exp\left(\frac{\langle v_t^*, \theta \rangle^2}{8}\right) \\ & = \exp\left(\sigma_t^2 \cdot \left(\|\theta\|_2^2 + \frac{\langle v_t^*, \theta \rangle^2}{16c_t^2} \|v_t^* + \eta_t\|_2^2 - \frac{\langle v_t^*, \theta \rangle (\langle v_t^*, \theta \rangle + \langle \eta_t, \theta \rangle)}{2c_t}\right)\right) \exp\left(\frac{\langle v_t^*, \theta \rangle^2}{8}\right) \\ & = \exp(\sigma_t^2 \|\theta\|_2^2) \exp\left(\langle v_t^*, \theta \rangle^2 \left(\frac{\sigma_t^2}{16c_t^2} (\|v_t^*\|_2^2 + \|\eta_t\|_2^2 + 2\langle v_t^*, \eta_t \rangle) - \frac{\sigma_t^2}{2c_t} + \frac{1}{8}\right)\right) \exp\left(\frac{-\sigma_t^2 \langle v_t^*, \theta \rangle \langle \eta_t, \theta \rangle}{2c_t}\right). \end{aligned}$$

To simplify the above expression, we will use the fact that $\langle a, b \rangle \leq \|a\|_2 \|b\|_2$, $\|v_t^*\|_2 \leq \sqrt{n}$, $\|\eta_t\|_2 \leq \varepsilon \sqrt{n} \leq \sqrt{n}$, along with $\sigma_t^2 = c_t$. Substituting these gives us,

$$\begin{aligned} & \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\ & \leq \left(1 - \frac{\delta}{m}\right)^{-t} \exp(\sigma_t^2 \|\theta\|_2^2) \cdot \exp\left(\langle v_t^*, \theta \rangle^2 \left(\frac{4n}{16c_t} - \frac{1}{2} + \frac{1}{8}\right)\right) \cdot \exp\left(\frac{-\langle v_t^*, \theta \rangle \langle \eta_t, \theta \rangle}{2}\right). \end{aligned}$$

Note that for a small value of δ , the term $\frac{4n}{16c_t} = 1/(\frac{9}{C} \log \frac{5nm}{\delta} + 4\varepsilon t)$ is negligible in comparison to the other terms. Hence, $\frac{4n}{16c_t} - \frac{1}{2} + \frac{1}{8} < 0$ and we get

$$\begin{aligned} \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] & \leq \left(1 - \frac{\delta}{m}\right)^{-t} \exp(\sigma_t^2 \|\theta\|_2^2) \exp\left(\frac{-\langle v_t^*, \theta \rangle \langle \eta_t, \theta \rangle}{2}\right) \\ & \leq \left(1 - \frac{\delta}{m}\right)^{-t} \exp(\sigma_t^2 \|\theta\|_2^2) \exp\left(\frac{\varepsilon n \|\theta\|_2^2}{2}\right) \\ & \leq \left(1 - \frac{\delta}{m}\right)^{-t} \exp((\sigma_t^2 + \varepsilon n) \|\theta\|_2^2) \\ & = \left(1 - \frac{\delta}{m}\right)^{-t} \exp(\sigma_{t+1}^2 \|\theta\|_2^2), \end{aligned}$$

where the last equality follows from the fact that $\sigma_{t+1}^2 = \frac{9n}{4C} \log \frac{5nm}{\delta} + \varepsilon n(t+1)$. This establishes property (1) for w_{t+1} .

Under the assumption that $\delta \leq 1/2$, an immediate consequence of property (1), which we will use in proving property (2) is that,

$$\begin{aligned} \mathbb{E}_{s_1, \dots, s_t} [\exp(\langle w_{t+1}, \theta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] & \leq \left(1 - \frac{\delta}{m}\right)^{-t} \exp(\sigma_{t+1}^2 \|\theta\|_2^2) \\ & \leq \left(1 - \frac{\delta t}{m}\right)^{-1} \exp(\sigma_{t+1}^2 \|\theta\|_2^2) \\ & \leq (1 - \delta)^{-1} \exp(\sigma_{t+1}^2 \|\theta\|_2^2) \quad (t \leq m) \\ & \leq 2 \exp(\sigma_{t+1}^2 \|\theta\|_2^2) \quad (8) \end{aligned}$$

where the second inequality follows from the fact that $(1+x)^a \geq 1+xa$ if $x \geq -1, a \geq 1$, and final inequality follows from $\delta \leq 1/2$.

Proving property (2). We will show that $\Pr[\text{Fail}_{t+1} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \leq \delta/m$, which will directly prove property (2) for iteration $t+1$.

Note that Fail_{t+1} happens only if in iteration $t+1$, $\langle w_{t+1}, v_{t+1}^* + \eta_{t+1} \rangle > 4(1-p)c_{t+1}$ or $\langle w_{t+1}, v_{t+1}^* + \eta_{t+1} \rangle < -4pc_{t+1}$ or $\|w_{t+1}\|_\infty > c_{t+1}/\sqrt{n}$ (Equation (4)). Since $p \in [1/2, 2/3]$, we have $4p > 4(1-p)$. We will utilize this fact to simplify as follows,

$$\begin{aligned}
& \Pr[\text{Fail}_{t+1} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\
&= \Pr[\langle w_{t+1}, v_{t+1}^* + \eta_{t+1} \rangle > 4(1-p)c_{t+1} \text{ or } \langle w_{t+1}, v_{t+1}^* + \eta_{t+1} \rangle < -4pc_{t+1} \text{ or } \|w_{t+1}\|_\infty > c_{t+1}/\sqrt{n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\
&\leq \Pr[|\langle w_{t+1}, v_{t+1}^* + \eta_{t+1} \rangle| > 4(1-p)c_{t+1} \text{ or } \|w_{t+1}\|_\infty > c_{t+1}/\sqrt{n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \quad (4p > 4(1-p)) \\
&= \Pr[|\langle w_{t+1}, (v_{t+1}^* + \eta_{t+1})/(4(1-p)\sqrt{n}) \rangle| > c_{t+1}/\sqrt{n} \text{ or } \|w_{t+1}\|_\infty > c_{t+1}/\sqrt{n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\
&\leq \sum_{u \in S} \Pr[|\langle w_{t+1}, u \rangle| > c_{t+1}/\sqrt{n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c],
\end{aligned}$$

where set $S := \{(v_{t+1}^* + \eta_{t+1})/(4(1-p)\sqrt{n}), e_1, e_2, \dots, e_n\}$ and e_1, \dots, e_n are the basis unit vectors along the axes. Note that $\forall u \in S$, we have $\|u\|_2 \leq 3/2$: for all $i \in [n]$, $\|e_i\|_2 = 1$ and $\|(v_{t+1}^* + \eta_{t+1})/(4(1-p)\sqrt{n})\|_2 \leq 2\sqrt{n}/(4(1-p)\sqrt{n}) \leq 2/(4(1-2/3)) = 3/2$, where the last inequality uses the fact that $p \leq 2/3$. We proceed to further evaluate the above probability expression,

$$\begin{aligned}
& \Pr[\text{Fail}_{t+1} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\
&\leq \sum_{u \in S} \Pr[|\langle w_{t+1}, u \rangle| > c_{t+1}/\sqrt{n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c] \\
&\leq \sum_{u \in S} \Pr \left[\langle w_{t+1}, u/\|u\|_2 \rangle^2 > \frac{c_{t+1}^2}{(3/2)^2 n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \quad (\|u\|_2 \leq 3/2, \forall u \in S) \\
&= \sum_{u \in S} \Pr \left[\frac{\langle w_{t+1}, u/\|u\|_2 \rangle^2}{\sigma_{t+1}^2/C} > \frac{c_{t+1}C}{(3/2)^2 n} \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \quad (\sigma_{t+1}^2 = c_{t+1}, \text{ for } C < 1/8) \\
&= \sum_{u \in S} \Pr \left[\exp \left(\frac{\langle w_{t+1}, u/\|u\|_2 \rangle^2}{\sigma_{t+1}^2/C} \right) > \exp \left(\frac{c_{t+1}C}{(3/2)^2 n} \right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \\
&\leq \sum_{u \in S} \mathbb{E} \left[\exp \left(\frac{\langle w_{t+1}, u/\|u\|_2 \rangle^2}{\sigma_{t+1}^2/C} \right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] / \exp \left(\frac{c_{t+1}C}{(3/2)^2 n} \right) \quad (\text{Markov's inequality})
\end{aligned}$$

Now, we will make use of the following lemma, whose proof has been deferred to Appendix D.¹²

Lemma 9. *For a random vector $w \in \mathbb{R}^n$, if there exists a $\lambda > 0$ such that $\forall \theta \in \mathbb{R}^n$, $\mathbb{E}[\exp(\langle w, \theta \rangle)] \leq 2 \exp(C\lambda^2 \|\theta\|_2^2)$, where $C < 1/8$ is a fixed constant, then $\forall u \in \mathcal{S}^{n-1}$, we have $\mathbb{E}[\exp(\langle w, u \rangle^2/\lambda^2)] \leq 3$.*

From Equation (8), we know that the random vector w_{t+1} satisfies the condition in Lemma 9 with $\lambda^2 = \sigma_{t+1}^2/C$. Hence, we get,

$$\mathbb{E} \left[\exp \left(\frac{\langle w_{t+1}, u/\|u\|_2 \rangle^2}{\sigma_{t+1}^2/C} \right) \mid \cap_{\ell=1}^t \text{Fail}_\ell^c \right] \leq 3.$$

¹²The proof of this lemma closely follows Vershynin (2018) Proposition 2.6.1, differing only in multiplicative constants. We include the proof for completeness.

Using the above bound, and plugging in $c_{t+1} = \frac{9n}{4C} \log \frac{5nm}{\delta} + \varepsilon n(t+1)$ in the running chain of inequalities,

$$\begin{aligned} \Pr[\text{Fail}_{t+1} \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] &\leq \sum_{u \in S} 3 / \exp \left(\log \frac{5nm}{\delta} + \frac{4C\varepsilon(t+1)}{9} \right) \\ &= (n+1) \cdot \frac{3\delta}{5nm \exp(4C(t+1)\varepsilon/9)} \quad (|S| = n+1) \\ &\leq \frac{\delta}{m} \quad (4C(t+1)\varepsilon/9 > 0, n \geq 2) \end{aligned}$$

This completes the proof. \square

Lemma 8 directly leads to our main theorem about vector balancing.

Theorem 8. *During the execution of Algorithm 1, the random vector $w_t = \sum_{i=1}^{t-1} s_i v_i^*$ satisfies $\|w_t\|_{\infty} \leq \frac{9\sqrt{n}}{4C} \log \frac{5nm}{\delta} + \varepsilon t\sqrt{n}$, for all $t \in [m]$ with probability at least $1 - \delta$.*

Proof. We will prove that with probability at least $1 - \delta$, the event $\cup_{\ell=1}^m \text{Fail}_{\ell}$ doesn't happen during the execution of Algorithm 1, i.e., Equation (4) is never satisfied. The theorem then follows, since Equation (4) not being satisfied means that the inequality $\|w_t\|_{\infty} \leq c_t/\sqrt{n} = \frac{9\sqrt{n}}{4C} \log \frac{5nm}{\delta} + \varepsilon t\sqrt{n}$ holds for all t .

Using chain rule of probability we have $\Pr[\cap_{t=1}^m \text{Fail}_{t}^c] = \prod_{t=1}^m \Pr[\text{Fail}_{t}^c \mid \cap_{\ell < t} \text{Fail}_{\ell}^c]$. From Lemma 8, we know that $\Pr[\text{Fail}_{t}^c \mid \cap_{\ell < t} \text{Fail}_{\ell}^c] \geq 1 - \delta/m$. Hence, we get $\Pr[\cap_{t=1}^m \text{Fail}_{t}^c] = \prod_{t=1}^m \Pr[\text{Fail}_{t}^c \mid \cap_{\ell < t} \text{Fail}_{\ell}^c] \geq (1 - \delta/m)^m \geq 1 - \delta$. This completes the proof. \square

4.3.2 Multicolored discrepancy and Envy minimization.

We now generalize the result of Theorem 8 to the multicolored discrepancy problem (Theorem 9). The idea is based on a known reduction that has been utilized in prior work Alweiss et al. (2021); Halpern et al. (2025).

We begin by first describing the algorithm that achieves the multicolored discrepancy bound of Theorem 9. Consider a complete and full binary tree \mathcal{T} having exactly k leaves and root node r .¹³ For each internal node u of \mathcal{T} , we use n_{ℓ}^u and n_r^u to denote the number of leaves in its left and right subtree, respectively. We run a copy of Algorithm 1 at each internal node u of \mathcal{T} by setting the parameter p to $p_u = n_{\ell}^u / (n_{\ell}^u + n_r^u)$.

When each vector $v_j^* + \eta_j$ arrives, we pass it down along the tree from the root to one of the leaves; $v_j^* + \eta_j$ is then assigned a color equal to the index of the leaf it eventually arrives at. In particular, first $v_j^* + \eta_j$ is passed through the copy of Algorithm 1 running at the root r of \mathcal{T} . If $v_j^* + \eta_j$ is assigned a sign of p_r by Algorithm 1 running at r , then it is passed down to the next node in the right subtree, if $v_j^* + \eta_j$ is assigned a sign of $p_r - 1$, then it is passed down to the next node in the left subtree. Supposed the next node is u , then we again check whether $v_j^* + \eta_j$ is assigned a sign of p_u , $p_u - 1$ at u , and pass it down to the right, left subtree respectively. This continues till the vector $v_j^* + \eta_j$ reaches a leaf node, i.e, it gets assigned a color equal to the index of the leaf node. When this happens, the true vector v_j^* is revealed. Using v_j^* , the value of the vector w_j^u (superscript u identifies the node) used in Algorithm 1 is updated at each internal node through which $v_j^* + \eta_j$ passes; this update happens in Line 6. This process of assigning vectors to leaves is repeated for all the vectors. This completes the description of the multicolored discrepancy algorithm.

The guarantee of this algorithm is stated in Theorem 9, its proof appears in Appendix D. Notably, while we run Algorithm 1 in every internal node of \mathcal{T} having a depth of tree $O(\log k)$, the final multicolored discrepancy bound we obtain is only a constant factor larger than Theorem 8.

¹³A binary tree is *full* if every internal node has either 0 or 2 children. A binary tree is *complete* if all levels except possibly the last one are completely filled.

Theorem 9. Let $k \geq 2$ be an integer and $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m \in \mathbb{R}^n$ be a sequence of vectors that arrive one at a time, where $\|v_i^*\|_\infty \leq 1, \|\eta_i\|_\infty \leq \varepsilon \leq 1$ for all $i \in [m]$. When each vector $v_i^* + \eta_i$ arrives, we need to assign it a color from $\{1, 2, \dots, k\}$ after which the vector v_i^* is revealed to us. In this model, there is an efficient algorithm that achieves a multicolored discrepancy bound of $\max_{i,j \in [k]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_\infty \leq \frac{27\sqrt{n}}{2C} \log \frac{5nmk}{\delta} + 6\epsilon t\sqrt{n}$, with probability at least $1 - \delta$, where S_a is the set of vectors that are assigned the color $a \in [k]$ after all the vectors have arrived.

Finally, we obtain Theorem 7 by setting $k = n$ in Theorem 9 and recalling that any upper bound on multicolored discrepancy is also an upper bound for the corresponding envy-minimization problem, as highlighted at the start of Section 4.3.

Theorem 7. Let $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m \in \mathbb{R}^n$ be a sequence of items (valuation vectors) that arrive one at a time, where $v_{i,j}^* \in [-1, 1], \eta_{i,j} \in [-\varepsilon, \varepsilon]$ for all $i \in [m], j \in [n]$. When each item $v_i^* + \eta_i$ arrives, we need to allocate it to one of the n agents, after which the true valuation vector v_i^* is revealed to us. In this model, there is an efficient online algorithm that computes an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ that guarantees the following maximum envy bound $\max_{i,j \in [n]} v_i^*(A_j) - v_i^*(A_i) \leq \frac{27\sqrt{n}}{C} \log \frac{5nm}{\delta} + 6m\sqrt{n}\varepsilon$, with probability at least $1 - \delta$, for $C > 0$ being a universal constant following Definition 2.

Acknowledgments

Anuran Makur was supported in part by the NSF CAREER Award under Grant CCF-2337808. Alexandros Psomas and Paritosh Verma were supported in part by an NSF CAREER Award under Grant CCF-2144208, and a research award from the Herbert Simon Family Foundation. Paritosh Verma also wants to acknowledge the support of the Bilsland Dissertation Fellowship. Trung Dang was funded in part by NSF award CCF-2217069.

References

- Ben Abramowitz and Elliot Anshelevich. 2018. Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- Ryan Alweiss, Yang P Liu, and Mehtaab Sawhney. 2021. Discrepancy minimization via a self-balancing walk. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 14–20.
- Georgios Amanatidis, Georgios Birmipas, Aris Filos-Ratsikas, and Alexandros A Voudouris. 2022. A few queries go a long way: Information-distortion tradeoffs in matching. *Journal of Artificial Intelligence Research* 74 (2022), 227–261.
- Georgios Amanatidis, Georgios Birmipas, and Evangelos Markakis. 2016. On truthful mechanisms for maximin share allocations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 31–37.
- Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264 (2018), 27–51.
- Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A. Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *International Joint Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:232075911>
- Elliot Anshelevich and Shreyas Sekar. 2016. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.

- Elliot Anshelevich and Wennan Zhu. 2019. Tradeoffs between information and ordinal approximation for bipartite matching. *Theory of Computing Systems* 63, 7 (2019), 1499–1530.
- Kamyar Azizzadenesheli, Trung Dang, Aranyak Mehta, Alexandros Psomas, and Qian Zhang. 2026. Reward Selection with Noisy Observations. In *International Conference on Algorithmic Learning Theory*. PMLR.
- Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. 2017. Posting prices with unknown distributions. *ACM Transactions on Economics and Computation (TEAC)* 5, 2 (2017), 1–20.
- Yushi Bai, Uriel Feige, Paul Gözl, and Ariel D Procaccia. 2022. Fair allocations for smoothed utilities. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 436–465.
- Yushi Bai and Paul Gözl. 2021. Envy-free and Pareto-optimal allocations for agents with asymmetric random valuations. *arXiv preprint arXiv:2109.08971* (2021).
- Julius B Barbanel. 2005. *The geometry of efficient fair division*. Cambridge University Press.
- Gerdus Benadè, Daniel Halpern, and Alexandros Psomas. 2025. Dynamic fair division with partial information. *Operations Research* (2025).
- Gerdus Benadè, Daniel Halpern, Alexandros Psomas, and Paritosh Verma. 2023. On the existence of envy-free allocations beyond additive valuations. *arXiv preprint arXiv:2307.09648* (2023).
- Gerdus Benade, Swaprava Nath, Ariel D Procaccia, and Nisarg Shah. 2021. Preference elicitation for participatory budgeting. *Management Science* 67, 5 (2021), 2813–2827.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2013. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199535255.001.0001>
- Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. 2012. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 197–214.
- Ralph Allan Bradley and Milton E. Terry. 1952. Rank Analysis of Incomplete Block Designs. I. The Method of Paired Comparisons. *Biometrika* 39, 3/4 (December 1952), 324–345.
- Pinhan Chen, Chao Gao, and Anderson Y Zhang. 2022. Partial recovery for top-k ranking: optimality of MLE and suboptimality of the spectral method. *The Annals of Statistics* 50, 3 (2022), 1618–1652.
- John Dickerson, Jonathan Goldman, Jeremy Karp, Ariel Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 28.
- Lester E Dubins and Edwin H Spanier. 1961. How to cut a cake fairly. *The American Mathematical Monthly* 68, 1P1 (1961), 1–17.
- Soroush Ebadian, Rupert Freeman, and Nisarg Shah. 2022. Efficient Resource Allocation with Secretive Agents.. In *IJCAI*. 272–278.
- Soroush Ebadian and Nisarg Shah. 2025. Every Bit Helps: Achieving the Optimal Distortion with a Few Queries. (2025).
- Gerald B Folland. 1999. *Real analysis: modern techniques and their applications*. Vol. 40. John Wiley & Sons.

- Daniel Halpern, Alexandros Psomas, Paritosh Verma, and Daniel Xie. 2025. Online Envy Minimization and Multicolor Discrepancy: Equivalences and Separations. In *Proceedings of the 26th ACM Conference on Economics and Computation*. 188–188.
- Daniel Halpern and Nisarg Shah. 2021. Fair and Efficient Resource Allocation with Partial Information. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 224–230. <https://doi.org/10.24963/ijcai.2021/32> Main Track.
- Olav Kallenberg. 1997. *Foundations of modern probability*. Vol. 2. Springer.
- Robert Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical Analysis*. John Wiley & Sons Inc., New York, NY, USA.
- Daniel McFadden. 1973. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics (Economic Theory and Mathematical Economics)*, Paul Zarembka (Ed.). Academic Press, New York, NY, USA, 105–142.
- Robin L. Plackett. 1975. The Analysis of Permutations. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* 24, 2 (1975), 193–202.
- Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*. Springer, 317–331.
- Hugo Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- Roman Vershynin. 2018. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press.
- Ernst Zermelo. 1929. Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift* 29, 1 (December 1929), 436–460.

A Preliminaries for the BTL model

This appendix provides the formal setup for agents’ preference elicitation via the BTL model. The material here expounds on the estimation procedure described in Section 4.2.

We consider a natural extension of the well-studied BTL model Bradley and Terry (1952); Plackett (1975); Luce (1959), where each agent has a separate set of BTL preferences. Specifically, for n agents and m items, the model assumes that each agent $i \in [n]$ has an associated (unknown) preference score $\theta_{i,j}^* \in [-b, b]$ for every item $j \in [m]$. Under the BTL model, the probability of agent i preferring item j over item k depends only on the difference in their latent scores and is given by:

$$\mathbb{P}(\text{agent } i \text{ prefers item } j \text{ over } k) = \frac{\exp(\theta_{i,j}^*)}{\exp(\theta_{i,j}^*) + \exp(\theta_{i,k}^*)} = \psi(\theta_{i,j}^* - \theta_{i,k}^*), \quad (9)$$

where $\psi(t) = 1/(1 + \exp(-t))$ denotes the sigmoid function.

We collect the preference scores as a matrix $\Theta^* \in \mathbb{R}^{n \times m}$ such that $\Theta_{i,j}^* = \theta_{i,j}^*$. Our goal is to estimate the unknown preference matrix Θ^* from noisy pairwise comparison queries collected from the agents. Since the preference probabilities in Eq. (9) are invariant to translation, we impose the following identifiability constraint

$$\Theta^* \mathbf{1}_m = \mathbf{0}_n. \quad (10)$$

For each agent, comparisons are observed only on a subset of item pairs. This comparison structure is modeled by an agent-specific undirected graph \mathcal{E}_i . Specifically, for each agent $i \in [n]$, an edge $(j, k) \in \mathcal{E}_i$ if and only if agent i provides a comparison between items j and k . We assume that the observation graph of each agent is drawn independently from an Erdős-Rényi random graph model with parameter p .

We assume that for every agent $i \in [n]$ and for each pair $(j, k) \in \mathcal{E}_i$ with $k > j$, we observe K independent comparisons. Let $y_{j,k}^{(i,l)} \in \{0, 1\}$ denote the outcome of the l -th comparison between items j and k by agent i , where $y_{j,k}^{(i,l)} = 1$ indicates a preference for item j over item k . Under the BTL model, these outcomes are generated according to

$$y_{j,k}^{(i,l)} = \begin{cases} 1, & \text{with probability } \psi(\theta_{i,j}^* - \theta_{i,k}^*), \\ 0, & \text{otherwise} \end{cases}, \quad (11)$$

For convenience, we set $y_{k,j}^{(i,l)} = 1 - y_{j,k}^{(i,l)}$ for all $(k > j)$ with $(j, k) \in \mathcal{E}_i$. We define the empirical preference probabilities $y_{j,k}^{(i)} = \frac{1}{K} \sum_{l=1}^K y_{j,k}^{(i,l)}$. Given the observed empirical preference probabilities $\{y_{j,k}^{(i)}\}_{i,j,k}$, our goal is to estimate the latent preference matrix Θ^* .

For each agent $i \in [n]$, we estimate the agent's parameters $\tilde{\theta}^{(i)}$ by solving the following marginal log-likelihood problem using only agent i 's preference data (cf. Chen et al. (2022)). This problem takes the form

$$\tilde{\theta}^{(i)} = \arg \min_{\theta: \theta^T \mathbf{1}_m = 0} \left\{ -\frac{1}{mp} \sum_{(j,k) \in \mathcal{E}_i} y_{j,k}^{(i)} \log \psi(\theta_j - \theta_k) + (1 - y_{j,k}^{(i)}) \log(1 - \psi(\theta_j - \theta_k)) \right\}. \quad (12)$$

This convex optimization yields n vectors $\tilde{\theta}^{(i)} \in \mathbb{R}^m$, one for each agent. Stacking these vectors as rows produces the initial matrix estimate $\hat{\Theta} = [\tilde{\theta}^{(1)T}; \tilde{\theta}^{(2)T}; \dots; \tilde{\theta}^{(n)T}] \in \mathbb{R}^{n \times m}$. The estimation error bounds for this estimator are quantified in Lemma 4.

B Missing proofs from Section 3.1

Theorem 1. *Given any two real-valued, nondecreasing functions f, g and a real-valued random variable X , we have $\mathbb{E}[f(X) \cdot g(X)] > \mathbb{E}[f(X)] \cdot \mathbb{E}[g(X)]$ if and only if $\text{Var}(f(X)) > 0$ and $\text{Var}(g(X)) > 0$.*

Proof. Note that if $\text{Var}(g(X)) = 0$ or $\text{Var}(f(X)) = 0$, then, almost surely, $g(X) = \mathbb{E}[g(X)]$ or $f(X) = \mathbb{E}[f(X)]$ respectively. Both of these, in turn, imply the equality $E[f(X) \cdot g(X)] = E[f(X)] \cdot E[g(X)]$. Hence, the condition $\text{Var}(f(X)) > 0$ and $\text{Var}(g(X)) > 0$ is necessary.

To prove its sufficiency, we consider the difference $\mathbb{E}[f(X) \cdot g(X)] - E[f(X)] \cdot E[g(X)]$, which, can be equivalently stated as $E[f(X) \cdot g(X) - f(X) \cdot \mathbb{E}[g(X)]] = E[f(X)(g(X) - \mathbb{E}[g(X)])]$. In particular, we will show that if $\text{Var}(f(X)) > 0$ and $\text{Var}(g(X)) > 0$, then $E[f(X)(g(X) - \mathbb{E}[g(X)])] > 0$ as well. This will complete the proof.

Define $\mathcal{E}_>$ and $\mathcal{E}_<$ to respectively denote the events $g(X) > \mathbb{E}[g(X)]$ and $g(X) < \mathbb{E}[g(X)]$. Since $\text{Var}(g(X)) > 0$, both these events can happen, i.e., have a positive probability. To proceed, we will use the observation that the random variable $g(X) - \mathbb{E}[g(X)]$ can be written as $\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)]) - \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))$. Hence, we want to show that the difference stated below is positive,

$$E[f(X)(g(X) - \mathbb{E}[g(X)])] = \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])] - \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))].$$

Note that $0 = \mathbb{E}[g(X) - \mathbb{E}[g(X)]] = \mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])] - \mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]$. Therefore, we have $\mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])] = \mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]$. The fact that $\text{Var}(g(X)) > 0$, further implies that,

$$C := \mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])] = \mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))] > 0.$$

To prove that $\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])] - \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))] > 0$, we will instead prove that $\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]/C - \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]/C > 0$. Indeed, the latter implies the former. Now, using the definition of C , we can write,

$$\begin{aligned} & \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]/C - \mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]/C \\ &= \frac{\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]} - \frac{\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]} \end{aligned}$$

In the above expression, define $u := \frac{\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]}$ and $\ell := \frac{\mathbb{E}[f(X) \cdot \mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]}$ to be the left and right terms, respectively. Additionally, let $x_u := \inf\{x : g(X) > \mathbb{E}[g(X)]\}$ and $x_\ell := \sup\{x : g(X) < \mathbb{E}[g(X)]\}$. The following two observations will be crucial for the subsequent proof, where we will prove $u > \ell$.

- (i) $f(x_u) \geq f(x_\ell)$. Since g is nondecreasing, we have $x_\ell \leq x_u$, which, along with the fact that f is nondecreasing implies that $f(x_u) \geq f(x_\ell)$.
- (ii) $u \geq f(x_u)$ and $f(x_\ell) \geq \ell$. Note that u , by definition, is a weighted mean (or a convex combination) of $f(X)$, for all values of X that satisfy $g(X) > \mathbb{E}[g(X)]$, i.e., all values of X that are at least x_u . Hence, we have $u \geq f(x_u)$. Similarly, ℓ is a weighted mean of $f(X)$ for values of X that satisfy $E[g(X)] > g(X)$, i.e., values that are at most x_ℓ . This gives us the inequality $f(x_\ell) \geq \ell$.

Note that together (i) and (ii) imply,

$$u \geq f(x_u) \geq f(x_\ell) \geq \ell, \tag{13}$$

i.e., we get the weak inequality $u \geq \ell$. However, we want to show that $u - \ell > 0$. To establish the strict inequality, we will show that at least one of the following three strict inequalities must be true: $f(x_u) > f(x_\ell)$, $u > f(x_u)$, or $f(x_\ell) > \ell$. Either of these three strict inequalities along with Equation (13) will imply that $u - \ell > 0$ and will complete the proof.

Indeed, if we have $f(x_u) > f(x_\ell)$, then we're done. Henceforth, we focus on the case when $f(x_u) = f(x_\ell)$. Given that $f(x_u) = f(x_\ell)$ and $\text{Var}(f(X)) > 0$, it must be the case that $\text{Var}(f(X) \mid \mathcal{E}_>) > 0$ or $\text{Var}(f(X) \mid \mathcal{E}_<) > 0$ — this follows from the law of total variance.

We will show that this implies that either $u > f(x_u)$ or $f(x_\ell) > \ell$. This will rely on the following crucial property. The weighted mean u , of $f(X)$, for values of X satisfying $g(X) > \mathbb{E}[g(X)]$, has the property that each value of $X = \hat{x}$ (where $g(\hat{x}) > \mathbb{E}[g(X)]$) has a strictly positive weight of $\frac{g(\hat{x}) - \mathbb{E}[g(X)]}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_>}(g(X) - \mathbb{E}[g(X)])]}$ in the weighted mean. Hence, if we have $\text{Var}(f(X) \mid \mathcal{E}_>) > 0$, then we must have $u > f(x_u)$. Similarly, the weighted mean ℓ , of $f(X)$, for values of X satisfying $g(X) < \mathbb{E}[g(X)]$, is such that each $X = \tilde{x}$ with $\mathbb{E}[g(\tilde{x})] > g(\tilde{x})$ has a strictly positive weight of $\frac{\mathbb{E}[g(X)] - g(\tilde{x})}{\mathbb{E}[\mathbb{1}_{\mathcal{E}_<}(\mathbb{E}[g(X)] - g(X))]}$ in the weighted mean ℓ . Therefore $u > \ell$: either $f(x_u) > f(x_\ell)$ is true, else $\text{Var}(f(X) \mid \mathcal{E}_<) > 0$ or $\text{Var}(f(X) \mid \mathcal{E}_>) > 0$, which implies that either $u > f(x_u)$ or $f(x_\ell) > \ell$. The equivalence is proved. \square

Theorem 2. *Let \mathcal{D} and \mathcal{D}' be two fixed distributions such that \mathcal{D} is continuous and has a positive variance (\mathcal{D}' could be arbitrary). Assume that $v_{i,j}^* \sim \mathcal{D}$ for all $i \in [n]$, $j \in [m]$, and we observe the noisy estimates $\hat{v}_{i,j} = v_{i,j}^* + \eta_{i,j}$ where $\eta_{i,j}$ for all $i \in [n]$, $j \in [m]$. Then, running the welfare maximization algorithm on the observed noisy values $\{\hat{v}_{i,j}\}_{i,j}$, computes an envy-free allocation with probability at least $1 - O(1/m)$ if $m/\log m \in \Omega(n)$.*

Proof. To allocate the items to the agents, we run the social welfare maximization algorithm on the noisy estimates, i.e., we allocate each item $j \in [m]$ to the agent $i \in \arg \max_{i \in [n]} \{\hat{v}_{i,j}\}$ (tie-breaking randomly). We observe that this algorithm satisfies the following two properties: first, for each item $j \in [m]$, we have $\mathbb{P}(i \in \arg \max_{i \in [n]} \{\hat{v}_{i,j}\}) = 1/n$, i.e., every agent has an equal, $1/n$, chance of getting each item —

this follows from the fact that $\hat{v}_{i,j}$ are continuous random variables, and are independent and identically distributed.

Second, we have that $\mathbb{E}[v_{i,j}^* \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}] > \mathbb{E}[v_{i,j}^* \mid i \notin \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$, i.e., agents have a higher *true* value for an item they receive based on the *observed* noisy values. This is the most crucial property of our analysis, and it follows from Lemma 2 since \mathcal{D} is continuous and has a positive variance. With these two properties, we will show that this algorithm results in envy-free allocations with high probability if $m/\log m = \Omega(n)$.¹⁴

Define indicator random variables $X_{i,j} := \mathbb{1}[i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$, which are true if and only if agent i received item j . Define $\mathcal{B}_{i,\ell} = \mathbb{1}[\sum_{j=1}^m v_{i,j} X_{i,j} < \sum_{j=1}^m v_{i,j} X_{\ell,j}]$ to be the event that agent i envies agent ℓ after the execution of the algorithm. We want to argue that the probability $\mathbb{P}(\cup_{i \neq \ell} \mathcal{B}_{i,\ell})$ is tiny. Towards this, we begin by analyzing the $\mathbb{P}[\mathcal{B}_{i,\ell}]$.

For all $j \in [m]$, define $\mu_j := \mathbb{E}[v_{i,j} \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$ and $\mu'_j := \mathbb{E}[v_{i,j} \mid i \notin \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}]$. Note that as items behave identically, we have $\mu_1 = \mu_2 = \dots = \mu_m$ and $\mu'_1 = \mu'_2 = \dots = \mu'_m$, so we define $\mu = \mu_1$ and $\mu' = \mu'_1$. Furthermore, for fixed \mathcal{D} and \mathcal{D}' , we know that $\mu > \mu'$ from part (2) of Lemma 2, since \mathcal{D} and \mathcal{D}' satisfy the required conditions.

Note that for the event $\mathcal{B}_{i,\ell}$ to happen, either the event $\mathcal{B}_{i,\ell}^1 := \mathbb{1}[\sum_{j=1}^m v_{i,j} X_{i,j} < \frac{m}{n} \left(\frac{\mu + \mu'}{2}\right)]$ or $\mathcal{B}_{i,\ell}^2 := \mathbb{1}[\sum_{j=1}^m v_{i,j} X_{\ell,j} > \frac{m}{n} \left(\frac{\mu + \mu'}{2}\right)]$ must happen; indeed, if $\mathcal{B}_{i,\ell}^1$ and $\mathcal{B}_{i,\ell}^2$ do not happen, then $\mathcal{B}_{i,\ell}$ cannot happen. Hence, to show that $\mathcal{B}_{i,\ell}$ happens with small probability, we will show that $\mathcal{B}_{i,\ell}^1$ and $\mathcal{B}_{i,\ell}^2$ happen with small probability.

Note that $\mathbb{E}[\sum_{j=1}^m v_{i,j} X_{i,j}] = m \cdot \mathbb{P}(X_{i,j} = 1) \cdot \mathbb{E}[v_{i,j} \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}] = \frac{m\mu}{n}$ and $\mathbb{E}[\sum_{j=1}^m v_{i,j} X_{\ell,j}] = m \cdot \mathbb{P}(X_{\ell,j} = 1) \cdot \mathbb{E}[v_{i,j} \mid i \in \arg \max_{k \in [n]} \{\hat{v}_{k,j}\}] = \frac{m\mu'}{n}$. Using this, we can rewrite

$$\mathbb{P}(\mathcal{B}_{i,\ell}^1) = \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{i,j} < \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{i,j}\right] \left(1 - \frac{\mu - \mu'}{2\mu}\right)\right)$$

and

$$\mathbb{P}(\mathcal{B}_{i,\ell}^2) = \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{\ell,j} > \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{\ell,j}\right] \left(1 + \frac{\mu - \mu'}{2\mu'}\right)\right).$$

Define $\delta = \min\{\frac{\mu - \mu'}{2\mu}, 1\}$, and note this is a fixed quantity, as μ and μ' are fixed. Since $\sum_{j=1}^m v_{i,j} X_{i,j}$ is a nonnegative random variable, we get that

$$\mathbb{P}(\mathcal{B}_{i,\ell}^1) = \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{i,j} < \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{i,j}\right] \left(1 - \frac{\mu - \mu'}{2\mu}\right)\right) = \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{i,j} < \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{i,j}\right] (1 - \delta)\right).$$

Additionally, the fact that $\mu > \mu'$ implies that

$$\mathbb{P}(\mathcal{B}_{i,\ell}^2) = \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{\ell,j} > \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{\ell,j}\right] \left(1 + \frac{\mu - \mu'}{2\mu'}\right)\right) \leq \mathbb{P}\left(\sum_{j=1}^m v_{i,j} X_{\ell,j} > \mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{\ell,j}\right] (1 + \delta)\right).$$

Now, we can use the Chernoff bounds on the sum of independent-and-[0, 1]-bounded variables $v_{i,j} X_{i,j}$ and $v_{i,j} X_{\ell,j}$ to obtain the following upper bound on $\mathbb{P}(\mathcal{B}_{i,\ell})$,

¹⁴The subsequent argument will follow a similar analysis presented in Dickerson et al. (2014).

$$\begin{aligned}
\mathbb{P}(\mathcal{B}_{i,\ell}) &\leq \mathbb{P}(\mathcal{B}_{i,\ell}^1 \cup \mathcal{B}_{i,\ell}^2) \\
&\leq \mathbb{P}(\mathcal{B}_{i,\ell}^1) + \mathbb{P}(\mathcal{B}_{i,\ell}^2) && \text{(union bound)} \\
&\leq \exp\left(-\mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{i,j}\right] \frac{\delta^2}{2}\right) + \exp\left(-\mathbb{E}\left[\sum_{j=1}^m v_{i,j} X_{\ell,j}\right] \frac{\delta^2}{3}\right) && \text{(Chernoff bound)} \\
&= \exp\left(-\frac{m\mu}{n} \cdot \frac{\delta^2}{2}\right) + \exp\left(-\frac{m\mu'}{n} \cdot \frac{\delta^2}{3}\right) < 2 \exp\left(-\frac{m\mu'}{n} \cdot \frac{\delta^2}{3}\right).
\end{aligned}$$

For m where $m \geq \frac{9n \log m}{\mu' \delta^2}$, we get $\mathbb{P}(\mathcal{B}_{i,\ell}) \leq 2 \exp(-3 \log m) = 2/m^3$. Hence, via a union bound, we get that the probability of the resulting allocation not being envy-free is $\mathbb{P}(\cup_{i \neq \ell} \mathcal{B}_{i,\ell}) \leq \frac{2n^2}{m^3} \in O(1/m)$. This completes the proof. \square

C Missing proofs from Section 3.2

We now prove the full theorem, again restated here for convenience.

Theorem 3. *Let \mathcal{D} be a bounded nonnegative distribution satisfying the non-identical condition. Assume that $v_j^* \sim \mathcal{D}$, independently, for all items j . For each $i \in [n]$, $j \in [m]$, let $\hat{v}_{i,j}$ be a noisy estimate of $v_{i,j}^*$ such that for some $\varepsilon > 0$, we have $\max_{i,j} |v_{i,j}^* - \hat{v}_{i,j}| \leq \varepsilon$. There is a constant $c_{\mathcal{D}}$ depending only on \mathcal{D} and an algorithm such that if $\varepsilon \leq c_{\mathcal{D}}$, given the noisy estimates $\hat{v}_{i,j}$ for each i, j , the algorithm outputs an allocation $\mathcal{A} = (A_1, \dots, A_n)$ such that \mathcal{A} is envy-free with probability at least $1 - \exp(-\Omega_{\mathcal{D}}(m))$.*

Proof of Theorem 3. Let $\mathcal{A}^* = (\mathcal{A}_1^*, \dots, \mathcal{A}_n^*)$ be the allocation that minimizes the maximum envy with respect to the true values, $\max_{i,j \in [n]} v_i^*(\mathcal{A}_j^*) - v_i^*(\mathcal{A}_i^*)$. By Lemma 3, this is at most $-c'_{\mathcal{D}} \cdot m$ with probability $1 - \exp(-\Omega_{\mathcal{D}}(m))$, where $c'_{\mathcal{D}}$ is defined in the lemma. Condition on this event.

Let $\hat{\mathcal{A}} = (\hat{A}_1, \dots, \hat{A}_n)$ be the allocation that minimizes the maximum *observed* envy, $\max_{i,j \in [n]} \hat{v}_i(\hat{A}_j) - \hat{v}_i(\hat{A}_i)$. Since \mathcal{A}^* is a feasible allocation, β can be no larger than the observed envy gap of \mathcal{A}^* , which can be at most $\varepsilon \cdot m$ larger than its true gap. Thus, $\beta \leq -c'_{\mathcal{D}} \cdot m + \varepsilon \cdot m$.

We will choose $c_{\mathcal{D}} = c'_{\mathcal{D}}/4$, so $\varepsilon \leq c'_{\mathcal{D}}/4$. Plugging this in, we get that $\beta \leq -(3/4) \cdot c'_{\mathcal{D}} m$.

Now consider the following linear program that minimizes the maximum envy (with respect to observed values) over all fractional allocations.

$$\begin{aligned}
&\min \alpha \\
\alpha &\geq \sum_{j=1}^m \hat{v}_{i,j} x_{i',j} - \hat{v}_{i,j} x_{i,j} && \forall i \in [n], i' \in [n], i' \neq i \\
&\sum_{i=1}^n x_{i,j} = 1 && \forall j \in [n] \\
&x_{i,j} \geq 0 && \forall i \in [n], j \in [m]
\end{aligned}$$

The integral allocation $\hat{\mathcal{A}}$ is a feasible solution for this LP, therefore, the value of optimal solution of this LP, x^* , has maximum envy at most β . Let $\mu_{i,i'} = \sum_{j=1}^m x_{i',j}^* \hat{v}_{i,j}$. Let $\tilde{\mathcal{A}} = (\tilde{A}_1, \dots, \tilde{A}_n)$ be the integral allocation we get by allocating item j to agent i with probability $x_{i,j}^*$, and let $\mu_{i,i'} = \mathbb{E}[\hat{v}_i(\tilde{A}_{i'})]$. Note that $\mu_{i,i'} - \mu_{i,i} \leq \beta \leq -(3/4)c'_{\mathcal{D}} m$.

The expected observed value from agent i for bundle $\tilde{A}_{i'}$ is $\mathbb{E}[\hat{v}_i(\tilde{A}_{i'})] = \mu_{i,i'}$. The value $\hat{v}_i(\tilde{A}_{i'})$ is the sum of m independent random variables (one for each item), each bounded in the range $[0, \hat{v}_{i,j}] \subseteq [0, b + \varepsilon]$.

Since $\varepsilon \leq c_{\mathcal{D}}$, the range is $[0, b + c_{\mathcal{D}}]$. By Hoeffding's inequality, for any $\delta > 0$:

$$\Pr[|\hat{v}_i(\tilde{A}_{i'}) - \mu_{i,i'}| \geq \delta] \leq 2 \exp\left(\frac{-2\delta^2}{m(b + c_{\mathcal{D}})^2}\right).$$

Applying a union bound over all $\leq n^2$ pairs, the probability that any observed envy deviates from its mean by more than δ is

$$\Pr[\forall i, i' : |\hat{v}_i(\tilde{A}_{i'}) - \mu_{i,i'}| \leq \delta] \geq 1 - 2 \exp\left(\frac{-4\delta^2 \ln(n)}{m}\right).$$

We will choose $\delta = m \cdot c'_{\mathcal{D}}/4$. This ensures that except with probability $\exp(-\Omega(m))$, $\hat{v}_i(\tilde{A}_{i'}) - \hat{v}_i(\tilde{A}_i) \leq \beta + 2\delta \leq -m \cdot c'_{\mathcal{D}}/4$.

Finally, note that $v_i^*(\tilde{A}_{i'}) - v_i^*(\tilde{A}_i)$ can deviate from the observed gap by at most $m \cdot \varepsilon \leq m \cdot c'_{\mathcal{D}}/4$. This ensures that \tilde{A} is envy-free. A union bound over the two events yields the theorem. \square

D Missing proofs from Section 4.3

Claim 1. For all $\beta \in \mathbb{R}^n$, we have $\mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] \leq (1 - \frac{\delta}{m})^{-t} \exp(\sigma_t^2 \|\beta\|_2^2)$.

Proof. Using the law of total probability we get,

$$\begin{aligned} \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] &= \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] \cdot \mathbb{P}[\text{Fail}_t^c \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] + \\ &\quad \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \text{Fail}_t, \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] \cdot \mathbb{P}[\text{Fail}_t \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] \\ &\geq \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] \cdot \mathbb{P}[\text{Fail}_t^c \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] \\ &\geq \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] \left(1 - \frac{\delta}{m}\right), \end{aligned}$$

where the final inequality follows from our induction hypothesis, i.e., $\mathbb{P}[\text{Fail}_t^c \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] \geq 1 - \frac{\delta}{m}$. Hence, we get $\mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^t \text{Fail}_{\ell}^c] \leq (1 - \frac{\delta}{m})^{-1} \mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c]$. Substituting the bound $\mathbb{E}_{s_1, \dots, s_{t-1}} [\exp(\langle w_t, \beta \rangle) \mid \cap_{\ell=1}^{t-1} \text{Fail}_{\ell}^c] \leq (1 - \frac{\delta}{m})^{1-t} \exp(\sigma_t^2 \|\beta\|_2^2)$, from the induction hypothesis, completes the proof. \square

Lemma 9. For a random vector $w \in \mathbb{R}^n$, if there exists a $\lambda > 0$ such that $\forall \theta \in \mathbb{R}^n$, $\mathbb{E}[\exp(\langle w, \theta \rangle)] \leq 2 \exp(C\lambda^2 \|\theta\|_2^2)$, where $C < 1/8$ is a fixed constant, then $\forall u \in \mathcal{S}^{n-1}$, we have $\mathbb{E}[\exp(\langle w, u \rangle^2 / \lambda^2)] \leq 3$.

Proof. Using the Taylor series expansion we have,

$$\mathbb{E}[\exp(\langle w, u \rangle^2 / \lambda^2)] = 1 + \sum_{i=1}^{\infty} \frac{\mathbb{E}[\langle w, u \rangle^{2i}]}{i! \cdot \lambda^{2i}} \quad (14)$$

To bound the series above, we will bound the moments $\mathbb{E}[\langle w, u \rangle^{2i}]$ for all even $i \geq 1$. Towards this note that,

$$\begin{aligned}
\mathbb{E}[\langle w, u \rangle^{2i}] &= \int_0^\infty \mathbb{P}[\langle w, u \rangle^{2i} > x] dx \\
&= \int_0^\infty \mathbb{P}[\exp(\langle w, \alpha u \rangle) > \exp(\alpha x^{1/2i})] dx && \text{(for any } \alpha > 0) \\
&\leq \int_0^\infty \frac{\mathbb{E}[\exp(\langle w, \alpha u \rangle)]}{\exp(\alpha x^{1/2i})} dx && \text{(Markov's inequality)} \\
&\leq 2 \int_0^\infty \frac{\exp(C\lambda^2\alpha^2)}{\exp(\alpha x^{1/2i})} dx && (\mathbb{E}[\exp(\langle w, \alpha u \rangle)] \leq 2 \exp(C\lambda^2\alpha^2)) \\
&= 2 \int_0^\infty \exp\left(-\frac{x^{1/i}}{4C\lambda^2}\right) dx && \text{(setting } \alpha = x^{1/2i}/(2C\lambda^2)) \\
&= 2 \int_0^\infty (4C\lambda^2)^i \cdot i \cdot y^{i-1} \exp(-y) dy && \text{(substituting } y = \frac{x^{1/i}}{4C\lambda^2}) \\
&= 2i(4C\lambda^2)^i \int_0^\infty y^{i-1} \exp(-y) dy \\
&= 2i(4C\lambda^2)^i \cdot \Gamma(i) && (15)
\end{aligned}$$

where the final equality follows from the definition of the well-known gamma function $\Gamma(z) := \int_0^\infty x^{z-1} e^{-x} dx$. Substituting the bound in Equation (15) into Equation (14), we get

$$\begin{aligned}
\mathbb{E}[\exp(\langle w, u \rangle^2/\lambda^2)] &= 1 + \sum_{i=1}^\infty \frac{\mathbb{E}[\langle w, u \rangle^{2i}]}{i! \cdot \lambda^{2i}} \\
&\leq 1 + \sum_{i=1}^\infty \frac{2i(4C\lambda^2)^i \cdot \Gamma(i)}{i! \cdot \lambda^{2i}} \\
&\leq 1 + 2 \sum_{i=1}^\infty (4C)^i && (\Gamma(i) = (i-1)!) \\
&\leq 1 + 2 \sum_{i=1}^\infty 2^{-i} = 3 && (C < 1/8)
\end{aligned}$$

This completes the proof. \square

Theorem 9. *Let $k \geq 2$ be an integer and $v_1^* + \eta_1, v_2^* + \eta_2, \dots, v_m^* + \eta_m \in \mathbb{R}^n$ be a sequence of vectors that arrive one at a time, where $\|v_i^*\|_\infty \leq 1, \|\eta_i\|_\infty \leq \varepsilon \leq 1$ for all $i \in [m]$. When each vector $v_i^* + \eta_i$ arrives, we need to assign it a color from $\{1, 2, \dots, k\}$ after which the vector v_i^* is revealed to us. In this model, there is an efficient algorithm that achieves a multicolored discrepancy bound of $\max_{i,j \in [k]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_\infty \leq \frac{27\sqrt{n}}{2C} \log \frac{5nmk}{\delta} + 6\varepsilon t\sqrt{n}$, with probability at least $1 - \delta$, where S_a is the set of vectors that are assigned the color $a \in [k]$ after all the vectors have arrived.*

Proof. Now we will show that this algorithm described in Section 4.3.2, achieves a multicolored discrepancy bound of $\max_{i,j \in [k]} \|\sum_{v \in S_i} v - \sum_{v \in S_j} v\|_\infty \leq \frac{27\sqrt{n}}{2C} \log \frac{4nmk}{\delta} + 6\varepsilon t\sqrt{n}$ with probability at least $1 - \delta$. Here, for any node u , S_u is the set of vectors that pass through node u during the execution of the multicolored discrepancy algorithm. To assist with the analysis, for each node u , we define π_u to be the fraction of total leaves that are in the subtree rooted at u , i.e., the number of leaves in the subtree of u divided by k . We now state a key claim.

Claim 2. *With probability at least $1 - \delta$, for all nodes u of \mathcal{T} , we have $\|\sum_{v \in S_u} v - \pi_u \sum_{v \in S_r} v\|_\infty \leq 3\left(\frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m\sqrt{n}\right)$*

Before proving this claim, we will show that it suffices to establish the desired multicolored discrepancy bound. Given Claim 2, for all leaves $i, j \in [k]$, we will have

$$\begin{aligned}
\left\| \sum_{v \in S_i} v - \sum_{v \in S_j} v \right\|_\infty &= \left\| \left(\sum_{v \in S_i} v - \frac{1}{k} \sum_{v \in S_r} v \right) - \left(\sum_{v \in S_j} v - \frac{1}{k} \sum_{v \in S_r} v \right) \right\|_\infty \\
&\leq \left\| \sum_{v \in S_i} v - \pi_i \sum_{v \in S_r} v \right\|_\infty + \left\| \sum_{v \in S_j} v - \pi_j \sum_{v \in S_r} v \right\|_\infty \quad (\pi_i = \pi_j = 1/k) \\
&\stackrel{\text{(Claim 2)}}{\leq} 6 \cdot \left(\frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n} \right) = \frac{27\sqrt{n}}{2C} \log \frac{4nmk}{\delta} + 6\varepsilon m \sqrt{n}.
\end{aligned}$$

The above bound holds for all $i, j \in [k]$ with probability at least $1 - \delta$, giving us the desired multicolored discrepancy bound. As the final step, we prove Claim 2.

Proof of Claim 2. From Theorem 8, we know that for each internal node u in \mathcal{T} , $\|w_m^u\|_\infty \leq \frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon t \sqrt{n}$ with probability at least $1 - \delta/k$. Since \mathcal{T} has exactly k leaves, it has $k - 1$ internal nodes. Therefore, through a union bound, we get that with probability at least $1 - \delta$, all internal nodes u will have $\|w_m^u\|_\infty \leq \frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n}$. Additionally, note that if u_ℓ and u_r are respectively the left and right children of u , then we have $\|w_m^u\|_\infty = \|p_u \sum_{v \in S_{u_r}} v - (1 - p_u) \sum_{v \in S_{u_\ell}} v\|_\infty \leq \frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n}$.

We will prove the claim by induction, starting from the root node and going down the tree \mathcal{T} . For the base case note that for the root node r , we have $\|\sum_{v \in S_r} v - \pi_u \sum_{v \in S_r} v\|_\infty = \|\sum_{v \in S_r} v - 1 \cdot \sum_{v \in S_r} v\|_\infty = 0$, i.e., the bound is trivially satisfied. Now, assuming that this bound is satisfied for an internal node u , we will show that it will be satisfied for both its left child u_ℓ and right child u_r . We will prove the argument only for the left child u_ℓ ; the argument for the right child u_r will follow similarly.

As per our induction hypothesis, we have $\|\sum_{v \in S_u} v - \pi_u \sum_{v \in S_r} v\|_\infty \leq 3(\frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n})$. Additionally, the copy of Algorithm 1 running at a , guarantees us that $\|w_m^u\|_\infty = \|(1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \sum_{v \in S_{u_r}} v\|_\infty \leq \frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n}$; this follows from Theorem 8. Using this we can bound,

$$\begin{aligned}
&\left\| \sum_{v \in S_{u_\ell}} v - \pi_{u_\ell} \sum_{v \in S_{u_r}} v \right\|_\infty \\
&= \|(p_u + 1 - p_u) \sum_{v \in S_{u_\ell}} v - \pi_{u_\ell} \sum_{v \in S_{u_r}} v\|_\infty \\
&= \|(p_u + 1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \pi_u \sum_{v \in S_r} v\|_\infty \quad (\pi_{u_\ell} = p_u \pi_u) \\
&= \|p_u (\sum_{v \in S_u} v - \sum_{v \in S_{u_r}} v) + (1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \pi_u \sum_{v \in S_r} v\|_\infty \quad (\sum_{v \in S_\ell} v = \sum_{v \in S_a} v - \sum_{v \in S_r} v) \\
&= \|p_u (\sum_{v \in S_u} v - \pi_u \sum_{v \in S_r} v) + (1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \sum_{v \in S_{u_r}} v\|_\infty \\
&\leq p_u \left\| \sum_{v \in S_u} v - \pi_u \sum_{v \in S_r} v \right\|_\infty + \|(1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \sum_{v \in S_{u_r}} v\|_\infty \\
&\leq \frac{2}{3} \left\| \sum_{v \in S_u} v - \pi_u \sum_{v \in S_r} v \right\|_\infty + \|(1 - p_u) \sum_{v \in S_{u_\ell}} v - p_u \sum_{v \in S_{u_r}} v\|_\infty \quad (p_u \in [\frac{1}{2}, \frac{2}{3}]) \\
&\leq \frac{2}{3} \cdot 3 \left(\frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n} \right) + \frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} + \varepsilon m \sqrt{n} \\
&= 3 \left(\frac{9\sqrt{n}}{4C} \log \frac{4nmk}{\delta} \right).
\end{aligned}$$

This completes the induction step and therefore the proof.

□

□